queue for the traversal not only depends on the cumulated costs to reach a vertex in the graph but also on the expected costs to reach the goal from this vertex. Bound by Minkowski's inequality, whereas $||x + y||_p \leq ||x||_p + ||y||_p$ (known as triangle inequality for $p = 2$), A* prunes the search space in comparison with Dijkstra's Algorithm. A sound heuristic for the remaining cost estimation is the geographical distance that is always lower than the road-based distance.

In the case of static cost functions, contraction hierarchies [223] are a data structure that speeds up the A* algorithm and enables trip calculation in large traffic networks. Instead of searching the shortest path directly within the traffic network, contraction hierarchies reduce the search space to the most important connections. In a preprocessing step, these important segments are identified (based on the topology), and the network is extended by edges between these important links.

In contrast to regular road networks, public transportation data enhances a spatial graph with temporal data by adding timetable information. A trip T serves a sequence of stops $stops(T) = (s_1, \ldots, s_n)$, $s_i \in S$. T connects two stops s_a and s_b if and only if $stop(T, s_a) < stop(T, s_b)$. If multiple trips contain the exact same sequence of stops, they form a line [47].

A common approach is to model the dynamic into the graph G and then to apply Dijkstra's algorithm. This results in a time-extended and time-dependent model. In the time-extended model, every transit node is split into multiple vertices for each event (arrival, transit, and departure). The time-dependent model assigns every transit node one vertex, and arcs encode temporal constraints.

A data structure and algorithm, `Transfer Patterns`, introduced by Hannah [48] is considered state of the art in public transport routing. Based on the assumption that during a day, there are only a few optimal routes from stop s to stop t that differ only in the time they take place. In a preprocessing phase, optimal routes are computed as a sequence of transfer stations, neglecting the time component and information about intermediate stations. For each origin and target destination a directed acyclic graph is saved, containing all routes starting with the destination and containing all intermediate stations until the origin is reached.

In a realistic route planning scenario, various delays occur amongst the public transport vehicles. In contrast to vehicular traffic, trams and trains cannot overtake each other, and vehicles in transit networks wait for connections (e.g., connecting trains). This causes delays to propagate differently than vehicular traffic jams. In addition, two modes of transportation may share the same physical resource (e.g., buses or trams riding on a vehicular street). Thus, two forms of delays in transit networks are distinguished in literature: 1) a vehicle is late due to its own reasons, and 2) other vehicles are late caused by the former [462].

Several models for transit delays are reported in the literature. The work in [175] assumes independence. By contrast, [232] allows delays to cumulate. Sophisticated models incorporate dependencies among the vehicles into the delay [276]. In [439], the delays are analyzed visually.

In a trip planning application, real-time predictions of delay are the main benefit, as future delays may influence the route choice. Thus, we highlight two recent works on delay prediction and delay recognition: [217] applies queueing theory and assumes delays to aggregate, and [774] detects delays and unexpected vehicle movement in real time from the GPS traces.

In this work, we do not focus on the prediction but assume that we have information on delays of vehicles (in the commonly used GTFS real-time data format) either from vehicle observations or predictions.

With such dynamics, the trip computation becomes more difficult. Though a previous publication [47] states that transfer patterns are delay robust, this only holds as long as no new transfers are enabled by the delay. In the likely case that novel transfers are enabled, the existing transfer patterns do not represent this information and cannot result in the optimal transit route.

Transfer Patterns were introduced in [48]. The method comprises a data structure and an algorithm for fast transit route computation. In a preprocessing step, all possible connections are pre-computed and stored in a compressed format. For each public transport line, a table is stored, denoting in the columns the stops along the line. In this way, it holds the maximal possible route without changes.

Our approach to the dynamics of the transit information is to incorporate potential delay information already in the pre-computation phase, and add additional transfer possibilities to the DAGs created during transfer pattern creation.

As we aim to apply the transit route computations in an industrial context, we extend the capabilities of the existing open-source platform OpenTripPlanner (OTP). Our dynamic transfer patterns outperform the algorithms previously available in OTP A* [258] and RAPTOR [165] by an order of magnitude [411].

4.1.4 Control and Planning of Individual Actions

Urban areas are increasingly subject to congestions. Most navigation systems and algorithms that avoid these congestions consider drivers independently and can, thus, cause novel congestions at unexpected places. The precomputation of optimal trips (Nash equilibrium) could be a solution to the problem but due to its static nature is of no practical relevance. By contrast, we describe an approach to avoid traffic jams with dynamic self-organizing trip planning.

In [412], we apply reinforcement learning to learn dynamic weights for routing from the decisions and feedback logs of the vehicles. In order to compare the routing regime against others, the validation uses an open simulation environment (LuST) that allows the reproduction of the traffic in Luxembourg for with varying penetration rates. All of these experiments reveal that the performance of the traffic network is increased, and the occurrence of traffic jams is reduced by applying our routing regime.

Traffic closely resembles a bandit feedback learning environment (see [33] for an introduction to bandit learning). Bandit learning is a reinforcement learning task where the behavior of some blackbox (e.g., a bandit) should be learned only by the feedback we observe. Several actions can be taken (in the bandit problem, this equals drawing an arm). However, only the result of the actions can be observed, and it is unknown what would have happened otherwise. Vehicles serve as agents that move in a road network. The actions are represented by the roads a vehicle can choose at an intersection. Once a road is chosen, a reward will be assigned for that particular road depending on its actual state. The reward for all other roads that could have been chosen remains unknown. This lack of fully labeled data makes a supervised learning approach particularly complex.

The *Policy Optimizer for Exponential Models* algorithm (POEM) [666] is able to learn solely based on the reward values provided by the environment. Additionally, POEM does not perform online learning but rather uses logged data. This abstraction is known from bandit problems, which seek to optimize a reward from the sole information gained after turning the arm of the bandit. This presents a more robust approach, since a learned model can be thoroughly tested before deployment. The system will also not evolve over time, which could lead to unpredictable behavior. This is particularly undesirable in the context of vehicle routing.

In [666], POEM assigns a structured output to an arbitrary input based on its probability of being correct. Therefore, before applying POEM to congestion avoidance, a suitable mapping of the routing problem to a policy h_0, along with an input space X and output space Y, must be modeled. Additionally, a cardinal loss feedback mapping δ is required, which serves as the reward function for all selected input/output combinations.

The input space X was chosen as $X := [0, 1]^m$. Here, each $\vec{x} = (x_1, \ldots, x_m)^T \in X$ represents a feature vector of (normalized) sensor measurements for a road segment. For instance, a road's density, occupancy, mean speed, vehicle count, or waiting time can be used. Any value not in $[0, 1]$ was scaled using min-max scaling.

The output space must be a set of suitable, structured outputs. As POEM should be applied to the problem of congestion control, a single label indicating whether a road is congested or not already provides adequate results. Thus, let $Y := \{(0), (1)\}$, where (0) indicates a road is not congested and (1) corresponds to congestion.

The policy $h_0(Y \mid \vec{x})$ is a probability distribution over the output space. In other words, it assigns a probability to each output $\vec{y}$ given any input $\vec{x}$ based on how likely $\vec{y}$ is to be correct under conditions $\vec{x}$. Hence, predictions are made by sampling $\vec{y} \sim h_0(Y \mid \vec{x})$. The goal of POEM is then to improve this policy. Initially, no such policy exists for the constructed input and output spaces. This is a common problem when applying POEM. Therefore, a default policy is used (compare [666]). Let $h_0(\vec{y} \mid \vec{x}) := 0.5$, meaning both labels are assigned a probability of 0.5 for all $\vec{x}$.

Lastly, in order to improve an existing policy, POEM requires a cardinal loss feedback mapping $\delta : X \times Y \to \mathbb{R}$. This was achieved by applying one of the following

two primitive congestion detection methods to the sensor readings: the primitive density congestion metric, $\delta_{density}$, assumed a road to be congested when its density was greater than one-seventh of its jam density [124]. The primitive mean speed congestion metric δ_{speed} would assume a road as congested when its mean speed was less than ten kilometers per hour of its allowed maximum speed.

In order not only congestion but also reduce it, vehicles must receive frequent information updates about the current state of the road network. Then, POEM is used to predict the next state of the road network. This information will be used by vehicles to bypass roads which are deemed congested. Thus, those results must also be applied in a routing algorithm, such as Dijkstra or A*.

Let $G = (V, E, c, q)$ be a graph representing a road network. Here, c and q are the default cost and heuristic functions. Additionally, assume all vehicles have knowledge about a congestion labeling policy $h \in \mathcal{H}_{lin} \cup \{h_0\}$ [666]. When using dynamic routing, vehicles will receive updates about roads at regular intervals $T \in \mathbb{N}$. The update can then be written as $u_T : E \nrightarrow X$.[3] Then, when a vehicle receives update u_T, it is able to predict how likely a road is to be congested during interval $T + 1$ using h.

The described model receives sensor information only about whole road segments, rather than individual lanes, which might be problematic, as congestion does not always arise on every lane equally. That challenging situation is most likely to occur at junctions where each lane will allow a vehicle to go in a different direction. We address this problem by aggregating sensor data for each connected edge pair (for the use of a line graph of G, see [257]). Additionally, the resulting data allows more precise congestion detection as individual turning lanes are separated in the model.

In order to bypass arising congestion, a vehicle must recalculate its route with respect to the newly received update u_T. This is achieved by increasing the weight of an edge that is likely congested:

$$p^0_{(e_1,e_2)} := h((0) \mid 0.5u_T(e_1) + 0.5u_T(e_2)) \tag{4.4}$$

$$c' : E^2 \to \mathbb{R}_+,\ (e_1, e_2) \nrightarrow \frac{c(e_2)}{p^0_{(e_1,e_2)}} \tag{4.5}$$

The denominator shows the previously mentioned aggregation of sensor data. For notational simplicity, c' is defined for all elements of E^2. However, in practice only a subset of E^2 is used where e_1 is incident or equal to e_2.

The function c' calculates the new weight of an edge e_2 depending on the preceding edge that was reached. For instance, a vehicle on edge $e_1 = (u, v)$ would calculate the weight for edge $e_2 = (v, w)$ using $c'(e_1, e_2)$. A vehicle that starts its route on edge e_2 would use $c'(e_2, e_2)$.

Essentially, c' divides the default weight of an edge by its probability of not being congested in interval $T + 1$. This means the weight of an edge will remain almost

3 Here, it is assumed that updates are received equally for all edges.

unchanged when no congestion is expected. The increase will conversely depend on how likely congestion is to arise.

Finally, it was assumed that sensor data updates are available for every road. In real-world road networks, permanently installed sensors are much more scarcely distributed throughout the network. This problem can be partly alleviated by directly implementing sensors in the vehicle (e.g., using navigation applications provided by smartphones, or self-driving cars). However, some roads will still remain uncovered. Here, u_T can map to $\{0\}^m$. For the previously defined features in X (a road's density, occupancy, mean speed, vehicle count, and waiting time), its dimension m would be equal to 5. This will cause h to assign a probability of 0.5 to both labels (as defined by $\mathcal{H}_{lin}$ in [666]). Another solution might be to map u_T to the average of all sensor readings in an interval. Thus, uncovered roads would reflect the average state of a road network.

Logging

For POEM, no interactive control over actions is required, as it was specifically designed to learn using logged data. Hence, with respect to the previously defined setting, POEM requires a dataset:

$$D := \{(\vec{x}_i, \vec{y}_i, \delta_i, p_i) \mid i \in \mathbb{N}_{\leqslant n}\},\ p_i = h(\vec{y}_i \mid \vec{x}_i)\ . \tag{4.6}$$

This dataset will be created during the logging phase. All edges are assigned weights using c', and routes are calculated using an implementation of A*, which produces the shortest routes for any admissible heuristic. Additionally, POEM is initially applied using the default policy h_0, which will scale all weights equally by a factor of two. The scaling will not affect A*, meaning no route changes will occur, which in turn simplifies learning on previously collected data.

The data itself can either be collected by each vehicle or by a centralized authority monitoring each vehicle. For both approaches, a data entry cannot be created before any feedback is available. Thus, intermediate results must be cached.

First, the aggregated feature vector $\vec{x}_i$ is logged. The respective label $\vec{y}_i$ with its corresponding probability p_i are then determined using:

$$\vec{y}_i = \begin{cases} (0),\ h((0) \mid \vec{x}_i) > 0.5 \\ (1),\ h((1) \mid \vec{x}_i) > 0.5 \\ random((0),(1)),\ otherwise \end{cases} \tag{4.7}$$

Here, $random((0),(1))$ means a label is chosen randomly and uniformly distributed. Lastly, the feedback is logged using either $\delta_{density}$ or δ_{speed}. The respective results will inherently depend on the previously chosen label.

The deployment of our self-organizing routing algorithm in an urban area could be done in two ways. One option is to use the data of an existing stationary traffic information system (e.g., a SCATS [324] system) and feed it into a navigation platform

that can be used by the citizens. The other option is to turn vehicles directly into sensors and retrieve segment-wise statistics on travel time, density, and traffic flow directly from the navigation app. In the latter case, one might be worried about individual privacy because mobility statistics are recorded centrally. However recent work [410] provides an approach to protect individual privacy, known as homeomorphic encryption. This approach encrypts the data such that it still allows for the analysis of the crypto-text, but just the result can be decrypted. In the following, we will test these two deployment settings using stationary and moving sensors and compare them to Nash equilibrium and uninformed routing.

For the comparability of experiments with different routing algorithms, it is essential to guarantee the same traffic demand (i.e., origin/destination pairs) over time. For repeatability of the same origin/destination setting, we perform analysis with a microscopic traffic simulator called SUMO [349]. The simulator models individual vehicles on a microscopic level, so it controls also acceleration and deceleration, and is largely applied in traffic simulation and applications. It allows us to control traffic demand and provides us with complete knowledge of the performance of the street network and the routing performance. In contrast to arbitrary toy experiments, we aim at modeling sound traffic scenarios. We use an open simulation scenario in the city of Luxembourg [143], which enables the reproduction of 24 hours in the city's mobility.

The common procedure of SUMO is to generate the route of each vehicle before the simulation starts, which is why its live routing capabilities are rather limited. However, SUMO provides the *Traffic Control Interface* (TraCI), a network interface that allows full control over the current simulation. We used this to implement a Java application (SUMO-CA) that simulates a central authority. In order to calculate vehicle routes, SUMO-CA loads a road network and converts it to a directed, weighted multi-graph. When running a simulation, SUMO-CA will receive and parse sensor measurements in regular intervals. This information is utilized to predict the next state of the road network using POEM. Finally, those results are used to update vehicle routes.

Unless stated otherwise, each experiment will start at 7:45 (simulation time) and runs over a period of roughly 35 minutes, or exactly 2048 seconds. The reason why this particular window was chosen is that roads generally are more susceptible to congestion during rush hour. Additionally, a size of 2048 seconds allows rerouting intervals to be easily scaled using a factor of two. Finally, in order to create more realistic jams on arterial roads, SUMO was set to scale the original demand by a factor of 1.3.

Evaluating vehicle detours is problematic. Neither absolute nor relative differences will adequately represent measured detours. The reasoning behind this is that long routes will allow longer, absolute detours, whereas, short routes will allow longer, relative detours. Hence, a different metric is required. We propose the usage of the weighted relative detour as follows.

Let $y_A, y_B \in \mathbb{R}_+^\star$ be arbitrary measurements of one vehicle when algorithms A and B are applied, respectively. Then weighted relative detour $diff_{rw}$ will then calculate the

relative difference, while weighting it using the absolute difference.

$$diff_{rw}(y_A, y_B) := |y_A - y_B| \star \frac{y_A - y_B}{y_A + y_B} \tag{4.8}$$

Various charts in [412] present the evaluation results. The baseline is an uninformed Uniform Cost Search (UCS), where each road was assigned its static default weight, and every vehicle chooses its path individually by $A^{\star}$. In this case, congestions are likely to appear. In addition, a Nash equilibrium (NASH) is shown as a baseline.

4.1.5 Discussion

Throughout this work, we highlighted various models to estimate traffic predictions under different model assumptions and properties. The models represent different aspects of traffic at various granularities.

As an example, modeling car-to-car or vehicle-to-infrastructure interactions in inner cities requires different spatio-temporal granularity and thus different model assumptions from a macroscopic daily average traffic flow prediction.

In general, traffic is a chaotic system and the commonly applied Markov assumption is often violated in practice. Future traffic does not only depend on a fixed number of previous observations. Consider, for example, a semaphore in traffic system (a traffic light, a barrier, or a large public parking). In these situations, it is easy to see that, though following certain patterns, traffic is chaotic.

In Google Maps, the inherent structure of traffic data is currently modeled by Graph Neuronal Networks [172, 546]. However p-adic models are also a promising technique at fine granularities to represent the chaotic behavior. Since it is important for production-ready systems that dynamic predictions are tractable, condiditonal sum-product networks [597] are also an interesting future research direction.

We also applied the algorithm used for self-organizing control of navigation plans to control the charging prizes of electric mobility [543]. In this application, we observed that the explored states of the system might be bad for the system provider. As an example, in our experiments, reduced and even negative energy prizes could provide a useful incentive to prevent grid burnouts. However, the total financial risk needs to be bounded. Such constraints could be incorporated into reinforcement learning using stabilities.

4.2 Privacy-Preserving Detection of Persons and Classification of Vehicle Flows

Marcus Haferkamp
Benjamin Sliwa
Christian Wietfeld

Abstract: In some places, the continuously increasing road traffic will soon exhaust the capacity of existing traffic infrastructure unless appropriate measures are taken. Especially in urban environments with a high density of residential and commercial properties, the infrastructure is highly utilized or overloaded during peak hours. Since structural measures are often not possible or only at great expense, a practical solution to counter this issue is to optimize the infrastructure utilization and the control of traffic flows. For this purpose, the widely installed Internet of Things (IoT)-powered Intelligent Traffic Systems (ITS) can be used, which enable automated detection and high-precision classification of different road users and thus transform the infrastructure into a data-driven Cyber-Physical System (CPS).

Although various sensor systems have been proposed, they fulfill only subsets of the requirements, including accuracy, cost-efficiency, privacy preservation, and robustness. One approach that meets those requirements is a novel radio-based sensor system, of which we present two variants in this contribution. The system's fundamental idea is to exploit radio-based fingerprints of road users—multi-dimensional and characteristic attenuation patterns of several radio links—for detection and classification. One of the presented system variants additionally evaluates high-precision channel information extracted from Wireless LAN (WLAN) Channel State Information (CSI) or Ultra-Wideband (UWB) Channel Impulse Response (CIR) data. The proposed solution benefits from increased robustness against a wide range of interferences, e. g., poor visibility due to bad weather conditions. Moreover, the system exclusively uses embedded microcontroller units (MCUs) and radio technologies, allowing compact and cost-efficient installations in rural and dense downtown areas.

We have performed comprehensive field measurement campaigns and machine learning-enabled analyses that confirm the presented approach's high suitability for different requirements and application scenarios. In this regard, we have evaluated multiple applications, including the comparatively simple detection of road users and the fine-grained classifications of several vehicle classes. For instance, the proposed systems achieve more than 99 % for binary classification and 93.83 % for differentiating seven vehicle types.

4.2.1 Introduction

Following the current trend, it is expected that a large part of the existing transport systems will reach their capacity limits in the (near) future. Possible reasons include the approval of new types of personal transport (e. g., e-scooters) and the shift to alternative means of transportation. There are essentially two options to counter this problem without stricter regulatory measures for road users: structural measures to expand current capacities and the efficient utilization of existing infrastructures through optimized traffic flow control. However, the former step is often not an option due to high financial costs and additional long-term restrictions caused by the construction measures. Instead, more efficient traffic flow control is a realistic undertaking thanks to sensor information provided by the vehicles themselves and/or to low-cost IoT components, especially in smart cities. Such systems also collect high-precision and vehicle type-specific information, paving the way for novel and more advanced optimization methods (e. g., type-specific lane assignment or routing and smart parking). For this purpose, the systems must always provide up-to-date and area-wide precise traffic information, which is collected, among other ways, by a sensor network installed over a large area. Next to high accuracy, these systems must also meet other requirements. They should function reliably in challenging weather and traffic conditions while protecting road users' privacy and be energy- and cost-efficient to operate. In some countries, compliance with all these points is a prerequisite for being approved for large-scale installations in road traffic. For instance, some sensor solutions are unsuitable for this use case because of their characteristic weaknesses—e. g., privacy concerns when using camera-based sensors. An increasing number of vehicles is also equipped with GNSS (Global Navigation Satellite Systems) and mobile network connectivity, which provide detailed information about the current position of a vehicle in real time. In the work of Niehöfer et. al. (for example [476, 477]) it has been shown within the CRC that the accuracy of the vehicle position can be enhanced through in-depth system simulation to provide lane-specific positioning information of vehicles. Yet, any system that collects such location information about individual vehicle tracks raises privacy concerns.

Therefore, this contribution presents a novel Wireless Sensor Network (WSN) for detecting and classifying different types of road users, which identifies those based on characteristic inferences of the signal strength of a radio signal (fingerprint). Initially, the wrong-way driver warning system [250] has leveraged the idea of inferring the travel direction of passing vehicles based on the time sequence of the radio links' attenuation. We have enhanced this approach to determine vehicles of certain types utilizing class-specific fingerprints induced by their shapes and materials. We use supervised learning techniques to extract such class-specific similarities from the channel information for the evaluation. Specifically, in this contribution, we present two generations of the novel detection and classification system. The focus here is on the first generation, which correlates the Received Signal Strength Indicator (RSSI) of

many diagonal and transverse radio links. It can reliably infer different vehicle types based on this information [252, 621].

Furthermore, we present a modular and more compact system design that adds further high-precision channel information using WLAN CSI and UWB CIR radio technologies in addition to coarser information [251]. Since our research's focus has primarily been in the context of the initial system design, this section is devoted to the original system. Also, it provides a brief outlook of the successor system.

Figure 4.1 shows the presented IoT-powered sensor system's intended information flow and its use in a smart city context. Here, the communication modules acting as sensors record fingerprints of passing vehicles and preprocess this raw data and the classification task. One could use such exact traffic information in two different application scenarios. In *on-site* applications (e. g., parking-lot balancing, wrong-way driver detection), the acquired data is evaluated immediately on-site and serves as a trigger for further actions (e. g., warnings of wrong-way drivers). By contrast, *global* applications aggregate this locally relevant information to enable predictions and recommendations for larger areas. Finally, the widely deployed sensor systems can dynamically adjust their predictions by periodically verifying the prevailing traffic situation.

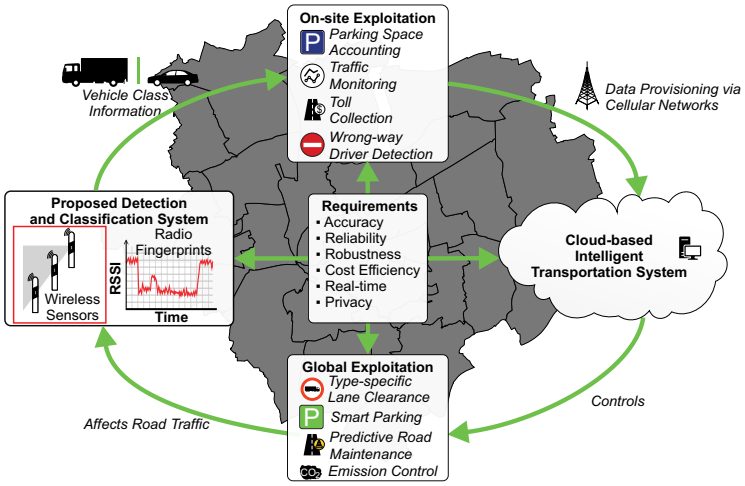

Fig. 4.1: Overall system vision: Embedding of the proposed IoT-based sensor system in a smart city environment. All sensor deployments are *locally* exploited for on-site applications and contribute their data to the *global* data-driven ITS applications ©[2020] IEEE. Reprinted, with permission, from [634].

4.2.2 Related Work

In this section, we provide an overview of existing systems and technologies for vehicle detection and classification. Figure 4.2 shows the abstract process flow and the main logical components of such systems, starting with the gathering of vehicle traces up to the final classification task. The sensor technology generates accurate traces as continuous and high-rate data streams, of which only a part is relevant for the subsequent process steps. To determine suitable sequences within these traces, a detection stage is typically connected afterward, reducing the overall workload. Based on these selected sequences, relevant (statistical) features are then extracted, which serve as input for supervised classification procedures using well-defined taxonomies.

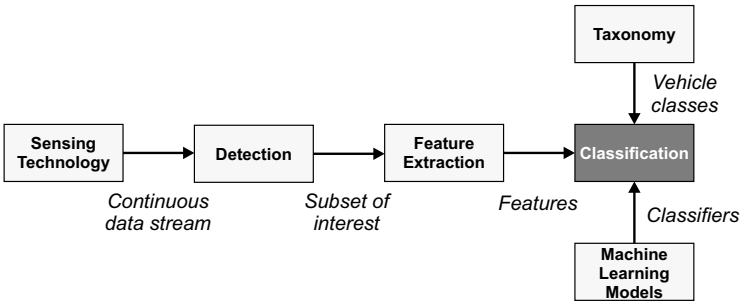

Fig. 4.2: Abstract system model of vehicle classification systems, ranging from data acquisition using different sensor technologies to the final classification tasks ©[2020] IEEE. Reprinted, with permission, from [634].

4.2.2.1 Taxonomies for Classification of Vehicles

The Federal Highway Administration (FHWA) proposes a 13-class scheme for classifying vehicles mainly based on the number of axles [327]. However, this approach's disadvantage is that the number of axles does not indicate the vehicles' exact dimensions. For example, accurate vehicle length information is essential for providing reliable parking space balancing or parking guidance systems.

An alternative taxonomy is the Nordic System for Intelligent Classification of vehicles (NorSIKT) [698], which is used in Scandinavian countries and, with its hierarchical approach, provides different gradations.

The *ISO 3833-1977* standard, the *2007/46/EC* Directive of the European Parliament, and the European New Car Assessment Programme (*Euro NCAP*) provide different schemes for classifying vehicles.

Nonetheless, many academic approaches use individual classification schemes to evaluate the performance of the proposed systems. Following this example, we have developed an adapted scheme with different refinement degrees in this work (cf. Section 4.2.4.1) and applied it to the final performance evaluation.

4.2.2.2 Sensor Technologies Used in Vehicle Detection and Classification Systems

This section provides an overview of established sensor technology used in vehicle classification detection and classification systems. Although much of the referenced work provide performance evaluation in terms of typical metrics—mainly classification accuracy—we want to note that comparing different solutions is difficult. Reasons for this include diverging taxonomies, different approaches for data preparation, and a variety of ML methods—e. g., Support Vector Machine (SVM) [147], Random Forest (RF) [102], k-Nearest-Neighbor (kNN), or Artificial Neural Network (ANN) [236]—used for analysis in the respective works. The sensor systems used can be broadly classified as *intrusive* or *non-intrusive*, respectively.

Intrusive systems represent the *classic* solution approach and are typically embedded in the road surface gathering technology-specific measured parameters. The used sensor technology directly affects the type and scope of measures required for installation or maintenance. While a minimally invasive cutting of the road surface is sufficient for some systems, more extensive and costly measures are necessary for other approaches. Representatives of this system category include Weigh in Motion (WIM) systems, inductive loop detector (ILD)—using one [145] or more inductive loops [365]—fiber Bragg grating sensors [674], vibration sensors [747], and piezoelectric sensors [519].

Non-intrusive systems include acoustic sensor systems, inertial sensors, camera-based approaches, and radio-based solutions. *Acoustic sensor systems* identify road users based on the emitted sounds. The fundamental challenge for these systems is the extraction of the relevant signal component from the dominant noise caused by the traffic noise. However, studies of acoustic sensor systems have shown that their use is of limited value due to comparatively low classification accuracies [225]. By adding other sensor technologies, the precision of these systems is significantly increased [157].

Different types of *inertial sensors*, such as accelerometers, gyroscopes, or magnetometers, are often combined on an inertial measurement unit (IMU). For vehicle detection and classification, the IMUs are either installed directly on the road's surface or at its side. One approach is detecting the number of axles of a passing vehicle, from which the vehicle class is deduced. For example, such systems achieve accuracies of 98.98 % for detection and 97 % for length-based classification [40].

Camera-based systems use pattern recognition and image processing techniques and are widely used due to their high precision. Apart from the detection of road users and the classification of vehicle types, the available high resolutions also allow a reliable recognition of vehicle makes [612], which can be problematic due to regulatory requirements to protect the privacy of road users. Most of these photosensitive systems use ambient light, so these systems' performance varies significantly with the day or visibility conditions. Using a Convolutional Neural Network (CNN), the approach presented in [178] achieves an accuracy of 95.7 % in daylight and 88.8 % in darkness for the classification of six types of vehicles.

In contrast to RSSI, the evaluation of WLAN CSI provides frequency-specific channel information. Orthogonal Frequency-Division Multiplexing (OFDM)-based radio technologies such as IEEE 802.11 use this information to estimate a channel's interference—e. g., multipath propagation—and reconstruct the original symbols. Depending on the number of transmitting and receiving antennas and the channel bandwidth, between 64 and 512 subcarriers are sent in a data packet's training fields. The receiving unit can infer the radio channel's interference by comparing amplitude and phase information of the expected and the received subcarrier sequence. Apart from reconstructing the original symbols, a variety of applications can exploit such detailed information. In addition to vehicle classification [732], localization and tracking of people behind walls and doors [9], as well as privacy-preserving monitoring by healthcare applications [316] are possible. Another technology, UWB, is predestined for the precise measurement of a radio channel because of its high robustness against interference due to its support for large channel bandwidths and its ability to determine accurate channel impulse responses. Although the primary use of UWB is in the area of localization—and recently also as an additional security measure for radio keys—it can also be used for activity detection [599] and vehicle detection classification [251]. *Radio-based approaches* assume that different vehicle types, due to specific shapes and installed materials, characteristically attenuate a radio signal. These attenuation patterns—symbolically referred to as fingerprints—can distinguish between different vehicle classes. Several radio technologies such as Bluetooth [58] or IEEE 802.15.4-based variants [250, 252] are suitable for radio-based methods, provided that the transceiver modules allow access to specific indicators of signal strength. A common approach is to use the RSSI, which is a coarse measure for assessing the received signal strength and depends heavily on the Signal-to-Noise Ratio (SNR) of the radio signal. Since these systems operate in the 2.4 GHz radio range, they exhibit high robustness to poor weather conditions due to rain and snowfall [150, 522].

4.2.3 Radio Fingerprinting-Based Vehicle Detection and Classification

This section describes the two variants of the proposed radio-based systems for vehicle detection and classification, including all essential components. Although both systems follow similar approaches with the evaluation of radio fingerprints , there are differences concerning the hardware components and the data processing, which we discuss in separate sections. First, the original system, which evaluates the signal strength information (RSSI) of multiple transverse and diagonal radio links, is discussed in detail (cf. Section 4.2.3.1). Subsequently, we highlight the significant differences and innovations of the current system approach that leverages high-precision WLAN CSI and UWB CIR channel information in Section 4.2.3.2.

4.2.3.1 RSSI-Based Vehicle Detection and Classification

The system setup initially used for the detection of wrong-way drivers [250], consisting of a total of six radio nodes integrated into delineators—three transmitter and three receiver units each—is shown in Figure 4.3. The system setup uses a constant longitudinal spacing of $\Delta_{lon} = 5m$.

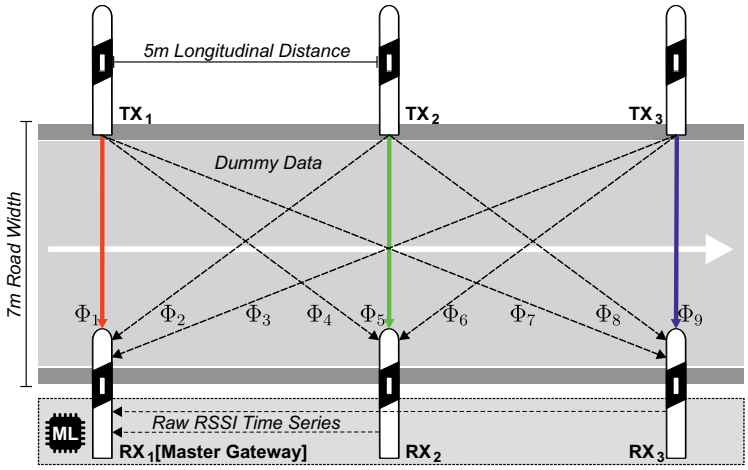

Fig. 4.3: Schematic system overview. Each delineator post contains an RF transceiver module. In total, the system uses nine different radio links ©[2020] IEEE. Reprinted, with permission, from [634].

All nodes use low-cost, off-the-shelf MCU with IEEE 802.15.4 radio modules equipped with omnidirectional antennas and operate with a transmit power of 2.5 dBm in the 2.4 GHz frequency band.

For continuously measuring the RSSI of all radio links, the corresponding transmitter modules periodically transmit pseudo data every 8ms. The system uses a coordinated channel access scheme utilizing tokens to avoid interference between the radio links. Then the receiving nodes send the signal strength information they measure to the master gateway, which aggregates the raw data and synchronizes it for further processing. Figure 4.4 illustrates an example of time-varying radio fingerprints gathered for all radio links for a passing car (top) and a truck (bottom), respectively.

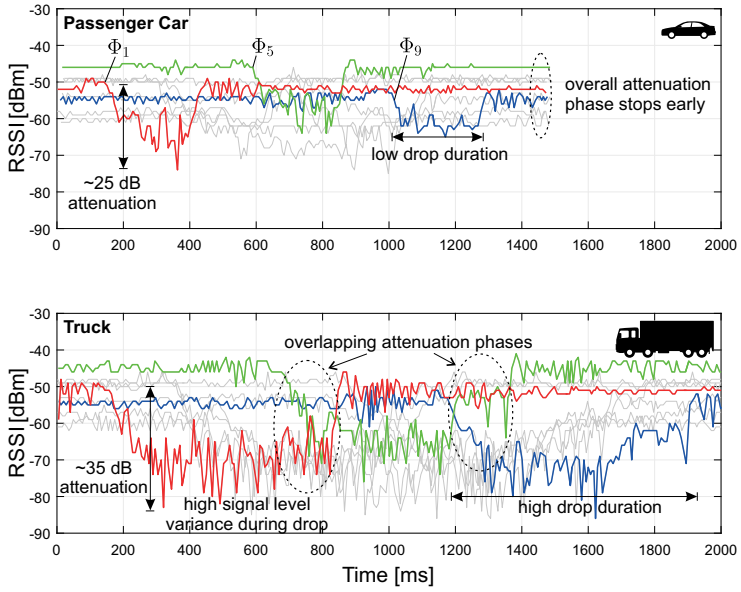

Fig. 4.4: Example multi-dimensional radio fingerprints for a passenger car and a truck. The colorized signals refer to the transverse radio links; the gray signals correspond to diagonal ones ©[2020] IEEE. Reprinted, with permission, from [634].

Figure 4.5 illustrates the entire data processing process for the RSSI-based classification system. First, the RSSI time series of all nine radio links Φ_i are recorded as vehicles pass through. Our approach then forwards the time signals to the data preprocessing block consisting of filtering using a moving average filter and subsequent normalization. These steps are relevant for minimizing the influence of scattered outliers—e. g., multi-path effects—and enabling high compatibility with various machine learning methods (*feature scaling*). Another process block realizes the detection of relevant subsets from the preprocessed time series. The system uses an automated thresholding approach to determine the individual start point t_{start} and endpoint t_{end} for each time series.

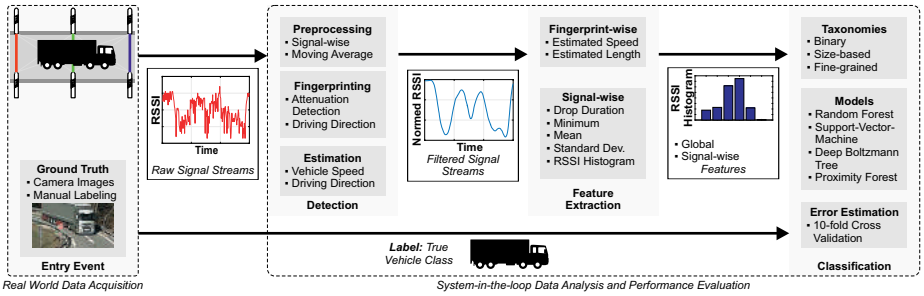

Fig. 4.5: System architecture model and data preprocessing pipeline used for ML-based vehicle classification using radio fingerprints ©[2020] IEEE. Reprinted, with permission, from [634].

The sequences tailored in this way serve as input for the subsequent process steps. The driving speed estimation serves as an additional feature for the classification process. With the help of the known longitudinal distance Δ_{lon} between the individual delineators, the system can estimate the average speed $\tilde{v}$ of the passing vehicles utilizing the temporal difference of the attenuation of the transverse links Φ_1, Φ_5, and Φ_9 using the following equation:

$$\tilde{v} = \frac{1}{3}\left(\frac{d(1,5)}{\Delta t(1,5)} + \frac{d(1,9)}{\Delta t(1,9)} + \frac{d(5,9)}{\Delta t(5,9)}\right) \tag{4.9}$$

where $\Delta t(i,j) = t_{start}(j) - t_{start}(i)$ and $d(i,j)$ is the longitudinal distance between the traversal radio links $\Phi(i)$ and $\Phi(j)$. Negative velocities $\tilde{v} < 0$ refer to an opposite direction, which indicates a wrong-way driver in the case of one-way streets. Similarly, we use Equation 4.10 to determine an approximation of the vehicle length:

$$\tilde{l} = \frac{|\tilde{v}|}{3}(\tau(1) + \tau(5) + \tau(9)) \tag{4.10}$$

where $\tau(i) = t_{end}(i) - t_{start}(i)$ denotes the duration of the attenuation of radio link $\Phi(i)$. The system also considers 90 different indicators—ten features for each of the nine radio links. These represent statistical variables such as mean value, standard deviation, minimum or maximum. In this way, a dimensional reduction is performed, since instead of several hundred signal strength values, the system only needs to process ten features per radio link.

4.2.3.2 Using CSI and CIR Data for Vehicle Detection and Classification

Like the previously presented system, the current modular system proposal also relies on the assumption that it is possible to reliably distinguish different road users by analyzing the characteristic interference they induce to a radio channel. Figure 4.6 illustrates the novel system approach's structure, which uses the radio technologies *WiFi 4* (IEEE 802.11n) and *Ultra-Wideband* (IEEE 802.15.4a). In addition to comparatively

coarse signal strength information, these technologies also measure a wealth of exact channel parameters.

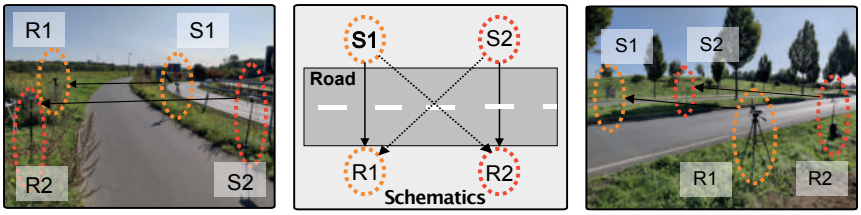

Fig. 4.6: Schematic system overview of the novel system approach leveraging WLAN CSI and UWB CIR data for bicycle detection (left) and motorized vehicle classification (right).

For measuring WLAN CSI, the system uses MCUs based on *Espressif ESP32* with WLAN transceivers connected to directional antennas operating with a transmit power of 20 dBm in the 2.4 GHz frequency band. For a continuous sampling of the radio channel, high-rate dummy packets are exchanged between the respective transmitting and receiving nodes. Each received packet contains CSI information for channel estimation. To reduce protocol overhead and thus increase overall system performance, the system uses unidirectional User Datagram Protocol (UDP) data transmissions. Thanks to an Application Programming Interface (API), the MCUs allow the accessing of CSI and thus amplitude and phase information from various subcarriers. In general, the CSI can contain other fields than Legacy Long Training Field (LLTF) such as High Troughput Long Training Field (HT-LTF) or Space-Time Block Code High Throughput Long Training Field (STBC-HT-LTF), which depends on the supported transmission modes of all WLAN modules involved as well as the channel characteristics. The system currently uses WLAN nodes only for high-rate measurement; the data preparation and ML steps have so far only been performed on more powerful computers.

UWB can accurately determine a radio channel's channel impulse responses thanks to very short signal pulses, allowing further insights regarding a radio channel, e. g., whether a line-of-sight (LOS) path is available or how many significant signal paths exist. The presented system setup uses a custom-made Printed Circuit Board (PCB), combining a *Decawave DWM1000* UWB transceiver module and an *ARM Cortex M3* MCU [687]. Like the WLAN nodes, the systems currently uses the UWB nodes to measure channel impulse responses. This high-resolution channel data is continuously transferred to computers for further processing via USB. Figure 4.7 demonstrates example WLAN CSI and UWB CIR traces.

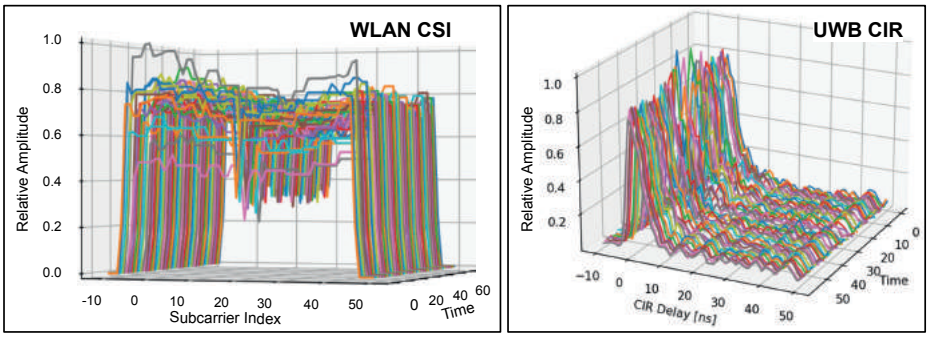

Fig. 4.7: Example WLAN CSI and UWB CIR traces. Each colorized line indicates a complete measurement sample including either multiple subcarriers (WLAN CSI) or CIR buffer sample data (UWB CIR).

4.2.4 Evaluation Methodology

This section presents the methodology used to evaluate both system variants. In this respect, we explain the system setups used for the field measurements, including essential parameters, the taxonomies adopted for the classification, and the ML models for performance evaluation.

4.2.4.1 Field Measurements

For data acquisition, we installed live systems in different environments for both system variants. The original RSSI-based classification system was installed and tested at a rest area on the A9 Autobahn as part of an official test site of the German Federal Ministry of Transport and Digital Infrastructure (shown in Figure 4.8, right). In total, the traces of 2605 vehicles were recorded and then manually labeled using camera images. The main parameters of the RSSI-based system can be found in Table 4.1.

The novel system proposal, which also uses high-resolution WLAN CSI and UWB CIR channel data, was tested at two locations. Traces of cyclists were recorded at a cycle path (Figure 4.8, left), while those of motorized vehicles, especially those similar to passenger cars, were collected at a busy single-lane road (Figure 4.8, right). Thus, the latter setting is similar to that used for the evaluation of the RSSI-based predecessor system. Table 4.1 lists the essential parameters for the novel system proposal.

Tab. 4.1: System parameters of the original RSSI-based system approach and the system evolution using WLAN CSI and UWB CIR.

Parameter	Radio Technology		
	WLAN CSI	UWB	RSSI
Transmission power	20 dBm	-14.31 dBm	2.5 dBm
Operating frequency	2.4 GHz	6.5 GHz	2.4 GHz
Sampling frequency	80 Hz	40 Hz	125 Hz
Antenna type	directional	omnidirectional	omnidirectional
Antenna gain	5-7 dBi	—	—
Number of radio links	1	1	9

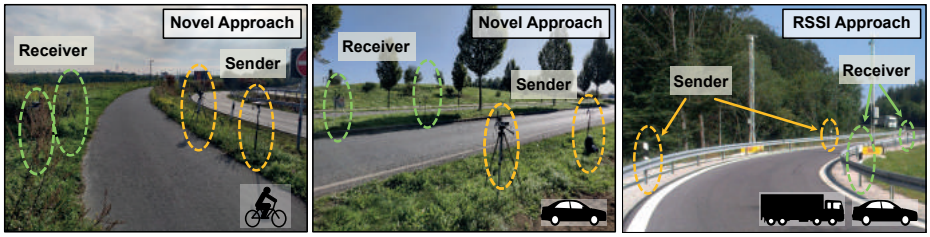

Fig. 4.8: Experimental live deployments of the original RSSI-based system approach (right) and the novel CSI- and CIR-based system evolution (left, middle) on three different settings for gathering real-world vehicle traces.

We used multiple taxonomies for the ML-based performance evaluation of the presented system variants, illustrated by Figure 4.9 for both the original RSSI-based (left) and the novel classification system (right). Defining different taxonomies was necessary because we tested both systems at diverse locations characterized by divergent traffic flows. Specifically, we evaluated the performance of the systems using taxonomies of varying complexity, which we briefly explain in the following:

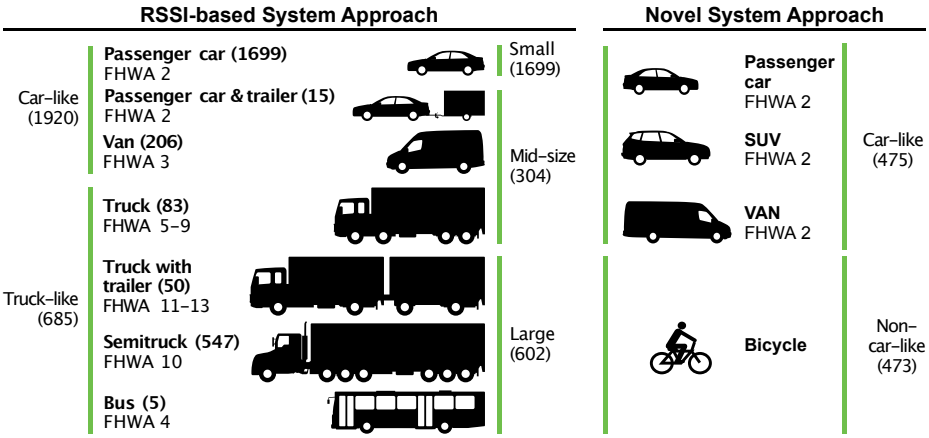

Fig. 4.9: Overview of the vehicle classes and sample counts for the different taxonomies used for evaluating the RSSI-based system approach (left) and the novel system design (right) ©[2020] IEEE. Reprinted, with permission, from [634].

Binary This category distinguishes between *car-like* and *truck-like* or *non-car-like* subclasses. While we classified car-like and truck-like vehicles for the RSSI-based system, we investigated the detection accuracy of the novel system with regard to cyclists using a binary classification with traces of cyclists as well as LOS (*idle*). No object was in the system during the LOS measurements, so fingerprints of the LOS radio channel were measured.

Cyclist vs. Motorized Vehicles Because the dataset of traces for different road users was not large enough, we performed the detection and classification of three classes: *car-like*, bicycle (*non-car-like*), and *idle*.

Size-based This was a 3-type classification of vehicles by vehicle length (only for the RSSI-based system).

Body style-based Here we use a fine-grained classification of seven vehicle types (only for the RSSI-based system).

In the *body style-based taxonomy*, the fine-grained classification task results in an increased overlap of vehicle classes with similar shapes (e. g., bus and semi-truck), increasing the overall classification inaccuracy. Nevertheless, we considered this complex taxonomy for the performance evaluation to illustrate the RSSI-based system approach's strengths and limitations.

4.2.4.2 Machine Learning-Based Classification

We have used several established and state-of-the-art machine learning models to detect and classify vehicles, which we compare and explain below. The following models were used to evaluate the performance of both system variants:

Artificial Neural Networks (ANNs) Inspired by the human nervous system, ANNs have received keen attention for different scientific applications in the context of deep learning. From an implementation perspective, these models are realized by multiple matrix multiplications directly affecting its resulting memory footprint.

Random Forests (RFs) Typical representatives of ensemble learning methods are RFs, which leverage the fact that most instances are assumed to be correct (*wisdom of the crowd*). Random subsets of features and training data are used for training each tree incorporated in an RF. Thanks to their binary decision-making, RFs allow for a resource-efficient implementation using simple if/else statements. By adjusting parameters such as limiting the number of allowed trees or the maximum depth for all trees, both processor and memory utilization can be controlled conveniently.

Support Vector Machines (SVMs) SVMs aim to separate data points in a multi-dimensional space through a hyperplane such that for each feature, the members of each class are separated as precisely as possible, which is achieved by minimizing a specific objective function.

In addition, we have used the following ML models for evaluating the performance of the original RSSI-based system approach:

Deep Boltzmann Trees (DBTs) Belonging to deep learning models, DBTs benefit from the fact that users have neither to extract features nor define transformation functions because they automatically derive differentiable functions from the given dataset. As an inherent downside, DBTs also require the user to select proper hyperparameters and a sufficient amount of training data due to the mass of trainable weights.

Proximity Forests (PFs) Like RFs, PFs belong to ensemble learning models, but instead of CART trees, they utilize proximity trees. Proximity trees use associated data points from the training set and implement–as its name suggests–a proximity-based approach where an object follows the branch with the highest similarity regarding a parametrized similarity measure.

4.2.5 Real-World Validation

This section presents and discusses the results for both proposed vehicle classification system approaches. Because we have developed and tested both systems independently—i. e., in different locations with divergent road users—we cover the results in separate sections, starting with the original RSSI-based vehicle detection and classification system.

4.2.5.1 Radio-Based Detection and Classification System

This subsection covers the results gained for the original RSSI-based system approach for vehicle detection and classification. At first, we describe how we have evaluated its detection performance, i.e., how accurately the system can determine a passing road user. To this end, we fed the raw traces of 2605 vehicles into a system-in-the-loop evaluation setup, allowing for flexible parameter tuning of the detection algorithm. Due to its relatively high system complexity in multiple diagonal and cross-radio links, the system also facilitates speed estimation and wrong-way driver detection (see Figure 4.10). We simulated the latter task by virtually inverting the order of the radio links spanned between the different nodes. Accordingly, an estimated negative speed indicates a wrong-way driver. The histogram shows a noticeable distribution for the dataset, implying a sound detection of the driving direction for passing vehicles. Since the number of detected vehicles matches captured vehicle traces, detection accuracy is 100 %. Nonetheless, we want to note that further real-world measurements are necessary to confirm the results of our virtual detection of wrong-way drivers.

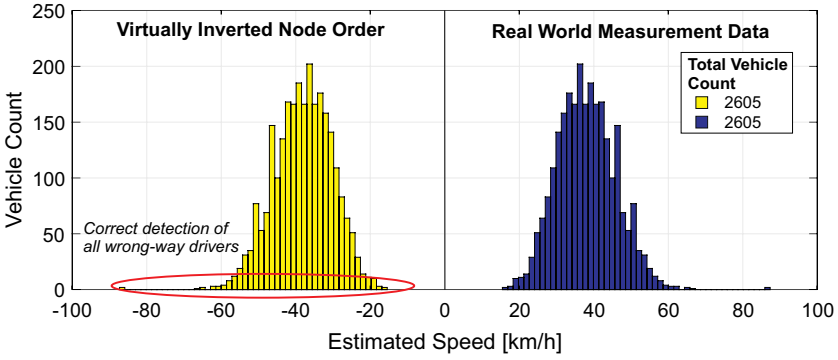

Fig. 4.10: Histograms of the speed estimations for the real world data and the virtually inverted node sequence. The vehicle count matches the total of captured vehicles traces and all wrong-way drivers are detected ©[2020] IEEE. Reprinted, with permission, from [634].

Next, we want to provide and discuss the results of the vehicle classification. We have utilized 10-fold cross-validation with 1/9 data split in each fold, i. e., 90 % of the data is used for training and the remaining 10 % for testing. After ten iterations, the statistical deviations of those folds are derived and used for performance evaluation. Figure 4.11 illustrates the classification accuracies for different machine learning models and the considered vehicle taxonomies (cf. Figure4.9). The 99 % classification accuracy, a typical minimum requirement for some applications, is illustrated as a horizontal dashed line.

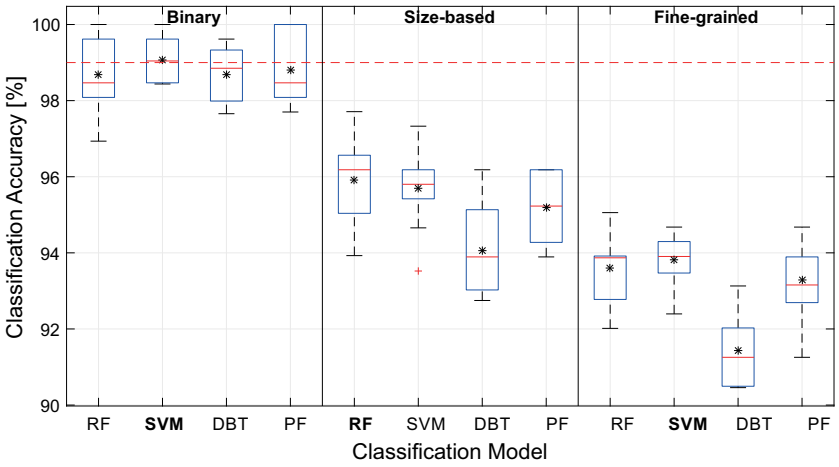

Fig. 4.11: Comparison of the overall classification accuracies for the considered machine learning models and considered vehicle taxonomies ©[2020] IEEE. Reprinted, with permission, from [634].

All evaluated models can exceed the given 99 % threshold for some runs but otherwise fall below it for the binary taxonomy. Only the SVM achieves a mean accuracy that matches this threshold. For the more complex vehicle taxonomies, the overall accuracies of all models decrease significantly: 93 % to 98 % for the size-based taxonomy and 90 % to 95 % concerning the fine-grained task. The apparent deviations between the models' performances result from their different learning strategies. While the stochastic nature of RF induces more significant standard errors in cross-validation, the DBT obtains lower performance levels than the remaining models because it calculates a probability measure for the given data.

Finally, the class-specific classification accuracy for the three considered vehicle taxonomies, i. e., binary, size-based, and fine-grained, is depicted in Figures 4.12, 4.13, 4.14. Starting with the binary taxonomy, which is the most simple classification task differentiating car-like and truck-like vehicles, the main challenge for all models is to classify truck-like cars correctly. We can validate this assumption by interpreting the classification results for mid-sized vehicles, as shown in Figure 4.13: all models have similar standard error values, whereas they perform notably better for small- and large-sized cars. For the fine-grained taxonomy, the multitude of similarly shaped vehicles and the underrepresentation of traces for specific vehicle types (e. g., bus) leads to lower classification accuracies due to larger standard deviations.

4.2.5.2 Vehicle Classification Using WLAN CSI and UWB CIR

For the new modular classification system, we present and discuss the classification results in this subsection. As previously mentioned, we have conducted multiple measurement campaigns for gathering traces of both cyclists and motorized vehicles. Be-

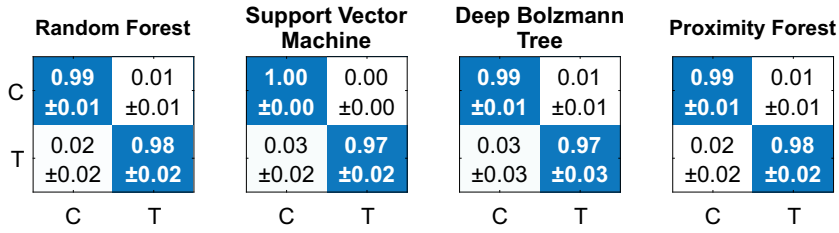

Fig. 4.12: Normalized confusion matrices for binary vehicle classification. *C*: car-like, *T*: truck-like ©[2020] IEEE. Reprinted, with permission, from [634].

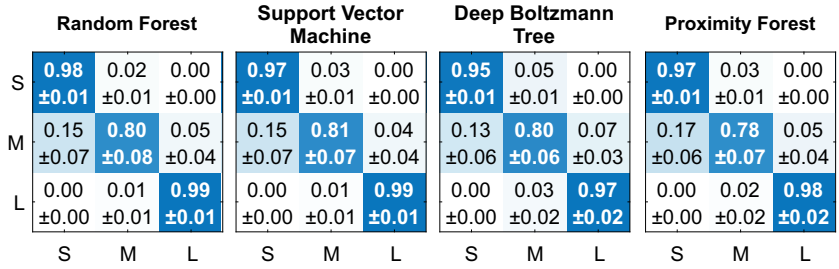

Fig. 4.13: Normalized confusion matrices for size-based vehicle classification. *S*: small, *M*: medium, *L*: large ©[2020] IEEE. Reprinted, with permission, from [634].

cause our measurements have focussed on capturing traces induced by cyclists, we have performed most of the analysis on detecting these. We start discussing the results for cyclist detection. Then we provide the performance results for a multi-type classification task with regard to cyclists and different motorized vehicle types.

For evaluating the bicycle detection performance, we have considered a binary classification task with the classes bicycle and non-bicycle (idle). Table 4.2 lists the maximum classification results achieved for separately analyzing different channel parameters gathered from WLAN CSI and UWB CIR data using ANN, RF, and SVM. Regarding WLAN CSI, the RSSI is the dominant channel indicator leading to the highest classification accuracy for all models. A possible explanation is that the WLAN transceiver modules evaluate multiple channel parameters to extract a significant measure for the link quality. Similarly, there is also a single channel parameter for UWB—the quotient of the estimated first path signal power and the channel impulse response power—leading to the highest classification accuracies. In particular, using this quotient and ANN, we could reach 100 % accuracy for detecting cyclists.

Next, we present the results for the multi-type vehicle classification applied for cyclists and different motorized vehicles. Specifically, we have conducted this task for a total of three classes, i.e., idle, bicycle, and car-like vehicles. Table 4.3 shows the classification results for WLAN CSI and UWB CIR data using ANN, RF, and SVM. Contrary

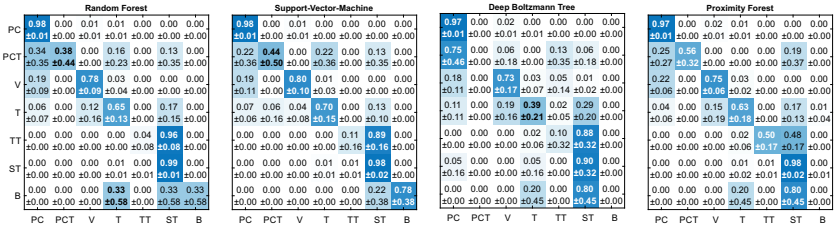

Fig. 4.14: Normalized confusion matrices for body style vehicle classification. *PC*: passenger car, *PCT*: passenger car with trailer, *V*: van, *T*: truck, *TT*: truck with trailer, *ST*: semitruck, *B*: bus ©[2020] IEEE. Reprinted, with permission, from [634].

Tab. 4.2: Bicycle Detection: Results for WLAN CSI and UWB using ANN, RF, and SVM with a 10-fold CV.

Model	Score	WLAN CSI		UWB	
		Value [%]	Param.	Value [%]	Param.
ANN	Accuracy	99.27±0.57	R (f2)	100±0	FC (f0)
	Precision	99.35±0.52	R (f2)	100±0	FC (f0)
	Recall	99.24±0.61	R (f2)	100±0	FC (f0)
	F-Score	99.30±0.56	R (f2)	100±0	FC (f0)
RF	Accuracy	99.45±0.54	R (f0)	99.83±0.26	FC (f1)
	Precision	99.48±0.52	R (f0)	99.84±0.25	FC (f1)
	Recall	99.45±0.51	R (f0)	99.8±0.26	FC (f1)
	F-Score	99.46±0.51	R (f0)	99.83±0.26	FC (f1)
SVM	Accuracy	99.32±0.51	R (f2)	99.83±0.26	FC (f0)
	Precision	99.38±0.47	R (f2)	99.84±0.24	FC (f0)
	Recall	99.30±0.53	R (f2)	99.82±0.27	FC (f0)
	F-Score	99.34±0.50	R (f2)	99.83±0.26	FC (f0)

f: Filter size, *FC*: Ratio of first path signal power and CIR power, *R*: RSSI

to the cyclist detection task, there are multiple predominant channel indicators for each system.

For WLAN CSI, the RSSI seems to be inadequate for achieving the highest accuracy. Instead, the subcarrier's amplitude values of different training fields are more relevant for this task: LLTF when using ANN, STBC-HT-LTF for RF, and several subcarriers in the case of SVM. By contrast, there are two crucial parameters when using UWB: the amplitudes of all raw CIR accumulator data (A) and the amplitudes of accumulator sample 15 (A15). By comparing the resulting classification accuracies for WLAN CSI and UWB CIR data, we can identify a considerable performance gap of about 5 % for the benefit of the former radio technology. However, we note that we have gathered the traces for cyclists and motorized vehicles in different environments with diverging system dimensions, impacting the transmissions between the UWB transceiver modules equipped with internal PCB antennas.

Tab. 4.3: Multi-type vehicle classification: Results for WLAN CSI and UWB using ANN, RF, and SVM with a 10-fold cross-validation.

Model	Score	WLAN CSI		UWB	
		Value [%]	Param.	Value [%]	Param.
ANN	Accuracy	98.23±0.67	L (f4)	92.38±1.30	A (f2)
	Precision	98.52±0.49	L (f5)	93.53±1.46	A (f2)
	Recall	98.31±0.63	L (f4)	93.30±1.34	A (f2)
	F-Score	98.39±0.71	L (f3)	93.41±1.38	A (f2)
RF	Accuracy	98.67±0.62	S (f0)	92.96±1.67	A (f0)
	Precision	98.83±0.59	S (f0)	93.74±1.74	A (f2)
	Recall	98.84±0.60	S (f1)	93.28±1.79	A (f2)
	F-Score	98.8±0.61	S (f0)	93.51±1.75	A (f2)
SVM	Accuracy	96.95±1.66	H_{SC52} (f0)	91.17±2.03	A_{15} (f0)
	Precision	97.86±1.24	H_{SC52} (f0)	92.13±1.85	A_{15} (f0)
	Recall	97.46±0.43	L (f4)	90.48±2.74	A_{15} (f0)
	F-Score	97.39±1.44	H_{SC52} (f0)	91.29±2.25	A_{15} (f0)

A: Amplitudes of all CIR accumulator samples, A_{15}: Amplitudes of CIR accumulator sample 15, f: Filter size, H_{SC52}: HT-LTF sub-carrier 52 amplitudes, L: LLTF sub-carrier amplitudes, S: STBC-HT-LTF sub-carrier amplitudes

4.2.6 Conclusion

This section presented two variants of novel radio-based systems that exploit different indicators of radio channels for accurate vehicle detection and classification. While the original system approach leverages relatively rough attenuation patterns of wireless signals (*RSSI fingerprints*), the evolved modular system approach uses exact channel parameters provided by the radio technologies WLAN CSI and UWB. Compared with existing detection and classification solutions, the proposed system variants are privacy-preserving, robust against challenging weather conditions, accurate, and cost-efficient. We have analyzed the suitability of both systems in comprehensive measurement campaigns in different environments: on a rest area on a highway, a busy one-lane road in a rural setting, and on a cycle path. The presented results approve the high performance of those system approaches for a set of differently challenging applications ranging from simple detection tasks of road users to a fine-grained classification of multiple vehicle types.

In future work, we want to investigate the applicability of different radio technologies (e. g., mmWave) within our detection and classification system. Moreover, we will obtain additional vehicle traces in challenging urban environments (e. g., in a downtown setting) to evaluate and strengthen system performance.

4.2.7 Acknowledgments

In addition to the CRC 876, part of this work has been supported by the German Federal Ministry of Transport and Digital Infrastructure (BMVI) in the course the PuLS project (03EMF0203B).

4.3 Green Networking and Resource Constrained Clients for Smart Cities

Pascal Jörke, Christian Wietfeld

Abstract: The Internet of Things (IoT) will enable a variety of new use cases by linking billions of IoT devices. Introducing new use cases each day, IoT devices will be found everywhere in the future. With a new generation of resource-constrained clients, communication networks have to face new challenges such as high communication ranges, small data transmission efficiency, and large scalability. With the Narrowband Internet of Things (NB-IoT) and enhanced Machine Type Communication (eMTC), cellular communication solutions have been adapted to these new challenges. Including mechanisms for larger communication ranges as well as lower power consumption, NB-IoT and eMTC aim to fulfil the requirements defined by new massive Machine Type Communication (mMTC) use cases. While performance is often only optimized on the lower layers, upper layers including transmission and application protocols need to be addressed by reducing overhead and enabling efficient small data transmissions in order to deliver good performance for resource-constrained clients.

This section describes the achievements in evaluating the performance of Low-Power Wide-Area Network (LPWAN) solutions for resource-constrained clients in terms of energy efficiency, spectral efficiency, and latency. Therefore, new cellular IoT features for power saving and coverage extension are explained in detail, while taking the costs for the scalability of the networks into account. With this knowledge, a performance analysis of resource-constrained LPWAN clients with different coverage conditions is provided.

Although both NB-IoT and eMTC use the same power-saving techniques as well as repetitions to extend the communication range, the analysis reveals a different performance in the context of data size, rate, and coupling loss. While eMTC comes with a 4 % better battery lifetime than NB-IoT when considering 144 dB coupling loss, the NB-IoT battery lifetime has 18 % better performance in 164 dB coupling loss scenarios. The overall analysis shows that in coverage areas with a coupling loss of 155 dB or less, eMTC performs better, but requires much more bandwidth. Taking the spectral efficiency into account, NB-IoT is in all evaluated scenarios the better choice and more suitable for future networks where the number of devices connected is expected to be close to or go beyond the network capacity.

While communication is possible with coupling losses up to 164 dB, the results show that the overall performance is very limited with decreasing signal quality. Although being designed for extended coverage, the mobile network operators should continuously improve the signal quality for both uplink and downlink directions. When increasing

the number of base stations is not feasible, alternative signal quality improvement solutions should be addressed. In this context, the coverage and link quality improvement of cellular IoT networks with multi-operator and multi-link strategies was evaluated as a case study, using the smart city Dortmund, Germany. The results show that the link quality can be improved by up to 13.6 dB, which enables shorter time-on-air for resource-constrained devices and thus drastically improves the energy and spectral efficiency.

4.3.1 Introduction

Waste bins with fill-level sensors, distributed environmental sensors monitoring the overall air quality in large cities, and beehive sensors regulating the hive temperature and transmitting the hive weight are just some use cases that integrate small sensor devices. The IoT enables countless new use cases. While some sensors have fixed power sources, others need to be independent of fixed power sources (e.g. smart waste bins) and therefore must rely on batteries, or, for an even better battery life, energy harvesting [262]. In large-scale scenarios, such as smart waste management, the operational costs need to be as low as possible and therefore the clients have to rely on a single battery for years, but still provide large communication ranges. In the past few years, many new communication solutions have addressed the requirements of low power consumption and wide area communication and are therefore called Low-Power Wide-Area Networks (LPWANs). A promising solutions is Long-Range Wide-Area Network (LoRaWAN), which is used by many public utilities, because it broadcasts in the license-free spectrum and is easy to set up. An alternative in licensed frequency bands is NB-IoT and eMTC, which were derived from the LTE communication technology. Both NB-IoT and eMTC can be deployed in existing LTE networks and therefore provide a fast and easy rollout in many countries.

The next section will give a short overview of the relevant characteristics when considering solutions in the license-free and licensed spectrums.

Clients in the License-Free Spectrum LoRaWAN is an easy-to-use communication solution for IoT. Designed for small data transmissions, LoRaWAN uses a proprietary communication protocol with small overhead and short time-on-air. To further reduce overhead and power consumption, the channel access is based on unslotted Aloha, directly transmitting data when available. With no channel access mechanism and deployment in the license-free spectrum, collisions are inevitable [93], making LoRaWAN unreliable in large scaled networks. Therefore, LoRaWAN is well suited for use cases with minimum Quality of Service (QoS) requirements, where the loss of packets is acceptable. Additionally, duty-cycle limitations need to be taken into account. Duty-cycle limitations are used in license-free spectrum to restrict the maximum transmit time of a device, e.g. 1 %, which means that a device can only transmit data for 36 s each

hour, affecting the maximum transmit interval. More details on clients in license-free spectrum can be found in Section 5.1.

Clients the in Licensed Spectrum When communication needs to be more reliable, communication solutions in licensed spectrum are the better choice. Derived from LTE, NB-IoT and eMTC use central scheduling for available frequency and time resources. Therefore, collisions on the air interface are prevented and the scalability of the network itself is mostly limited by the given frequency bandwidth. The price to pay for a scheduled transmission is the increased overhead for synchronization and control traffic, which affects the spectral and energy efficiency of resource-constrained clients.

4.3.2 Design Objectives of Resource-Constrained IoT Clients

With an exclusive spectrum available, NB-IoT and eMTC (often summarized as cellular IoT) are not limited by duty cycles and can be used with various application protocols and in many use cases. While eMTC relies on an existing LTE network and reuses LTE synchronization, NB-IoT can also be deployed as a stand-alone network [398]. Since NB-IoT uses only 180 kHz bandwidth (compared with 1.08 MHz in eMTC), it can also be deployed in guard bands–the unused bandwidth between two LTE networks–and usually used avoid interference.

Since clients may be distributed over a large area or even in challenging communication environments such as basements, cellular IoT solutions have to provide an extended network coverage, while still enabling low power consumption. The following section will give a short overview on power-saving and range-extending mechanisms that are introduced by cellular IoT.

4.3.2.1 Low Power Consumption

With current cellular communication solutions, the battery lifetime is often limited to a maximum of several weeks. By contrast, cellular IoT has the design objective of 10 years on a single 5 Wh battery [5]. Most IoT devices are designed to transmit small amounts of data on a hourly, daily, or weekly basis, which means the device is mostly in an idle state. Therefore, the new mechanisms for low power consumption focus on energy efficiency in idle mode.

Figure 4.15 depicts a typical NB-IoT transmission cycle. In the connected state, the device transmits its data and waits for a response. When no more data is transmitted or received, the device enters Extended Discontinuous Reception (eDRX). The eDRX mode extends the DRX (Discontinuous Reception) cycle to allow a device to remain longer in a power-saving state between paging occasions [398] and thus to further reduce the power consumption. The device remains synchronized and periodically available for

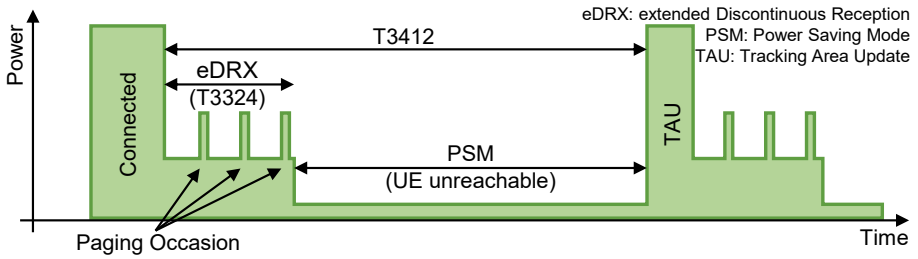

Fig. 4.15: NB-IoT transmission cycle with eDRX and PSM. ©[2018] IEEE. Reprinted, with permission, from [305].

mobile-terminated services. When the eDRX timer T3324 expires, the device will switch into the Power Saving Mode (PSM).

When using PSM, the device enters a power-saving state in which it reduces its power consumption to a bare minimum [398]. In PSM, the device remains registered to the network and maintains its connection configurations. As soon as the device leaves PSM, it does not need to attach to the network; rather, it reestablishes the previous connection, which leads to a reduced signaling overhead and optimized device power consumption. However, the device is unreachable for the network as long as it remains in PSM because it does not listen to the paging time windows. Mobile terminated services have to be suspended until the device reconnects to the network for mobile originated events. Tracking Area Updates (TAU) also trigger the device to end PSM and reestablish the connection to the network. While performing a TAU, the device listens to paging time windows and queued downlink transmissions.

4.3.2.2 Extended Coverage

For a comparism of signal ranges, the Maximum Coupling Loss (MCL) is often used, since it defines the maximum signal attenuation at which the receiver is still able to decode the signal. While the eMTC design objective defines an MCL of 155.7 dB [4], NB-IoT aims to extend the MCL to 164 dB. Figure 4.16 provides corresponding basement penetration ranges for the different MCL objectives.

Besides small bandwidths, cellular IoT solutions use repetitions for an increased energy per bit and therefore improved signal decoding. Therefore, eMTC introduces Coverage Enhancement (CE) Modes A and B. CE Mode A is mandatory and supports up to 32 repetitions while CE Mode B is optional and defines up to 2048 repetitions. NB-IoT also supports up to 2048 repetitions, though it does not divide the number of repetitions in different CE Modes, making all repetition options mandatory to all devices. Table 4.4 gives a detailed overview of the maximum number of repetitions for each NB-IoT and eMTC signal [398].

As similar data is transmitted, the application data rate decreases drastically with each repetition and devices consume more power compared with a transmission with-

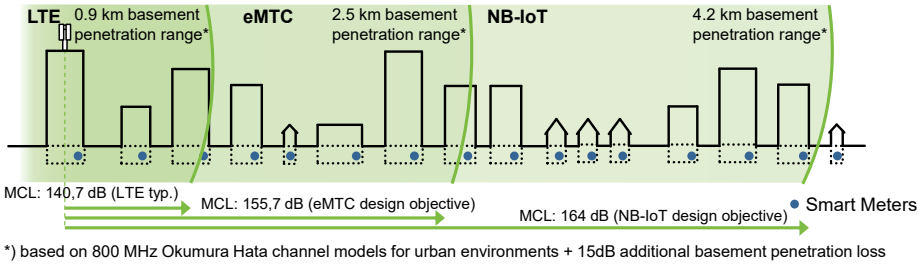

Fig. 4.16: Coverage Enhancement in NB-IoT and eMTC

Tab. 4.4: Maximum number of repetitions in eMTC and NB-IoT.

eMTC	Max. repetitions		NB-IoT	Max. repetitions
	CE Mode A	CE Mode B		
PDSCH	32	2048	NPDSCH	2048
MPDCCH	16	256	NPDCCH	2048
PRACH	32	128	NPRACH	128
PUSCH	32	2048	NPUSCH	128
PUCCH	8	32		

out repetitions. While new mechanisms like eDRX and PSM aim to increase the energy efficiency and battery lifetimes of cellular IoT devices, the extended coverage will lead to a significant reduction of energy efficiency. To ensure that the cellular IoT design objective is still achievable, both NB-IoT and eMTC must be subjected to a performance analysis.

4.3.2.3 Application Protocols for IoT

While LPWAN solutions in a license-free spectrum such as LoRaWAN often lack end-to-end Internet Protocol (IP) support due to the large protocol overhead, both NB-IoT and eMTC are able to transmit IP traffic such Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) messages. Due to the reduced transmission capacity, the packet size and number of message sequences should be as low as possible even with the coverage extension mechanism. For a decent system performance, the choice of a suitable application protocol is essential.

Message Queuing Telemetry Transport (MQTT) is a TCP-based IoT communications protocol, designed for Machine to Machine (M2M) data transmissions in low bandwidth environments [41]. It uses a centralized broker to which clients can publish data, while other clients can subscribe to data updates. In addition to transmission protection through TCP, MQTT introduces three Quality of Service (QoS) levels. QoS level 0 is a simple, low-overhead method of sending a message. The client simply connects to

the broker and publishes the message without MQTT acknowledgement by the broker. The data is still acknowledged on lower layers by TCP. QoS level 1 guarantees that the message is transmitted successfully to the broker. The broker sends an acknowledgement back to the sender. In case of a transmission loss of either the data or the acknowledgement, the data is retransmitted by the sender. QoS level 2 is the highest level of service. It comprises a sequence of four messages between the sender and the broker, and a handshake to confirm that the main message has been sent and that the acknowledgment has been received. When the handshake is completed, both sender and receiver are sure that the message was received exactly once, though this approach creates the highest message overhead.

MQTT for Sensor Networks (MQTT-SN) is an UDP-based, optimized version of the IoT communications protocol MQTT, designed specifically for efficient operation in large, low-power IoT sensor networks [645]. Like MQTT, it uses a centralized broker. Besides QoS levels 0 to 2, a new QoS level is introduced, which further reduces the message overhead. Publishing messages with a QoS level of -1 doesn't require an initial connection setup and broker registration; rather, it only transmits a single publish message with all required data. While QoS -1 comes with the lowest overhead, it does not provide acknowledgements and other responses from the broker.

Constrained Application Protocol (CoAP) is a third important IoT application protocol. CoAP is an UDP-based, specialized web transfer protocol for use with resource-constrained nodes and constrained networks in the IoT [600]. The protocol is designed for M2M applications. Unlike MQTT and MQTT-SN, it transmits data directly between a client and a server. Since no connection needs to be established at first, CoAP comes with a low message overhead. When using confirmed transmissions, all data that is transmitted between server and client is confirmed by application acknowledgements. The data is re-sent until it is acknowledged or the maximum number of retries is reached.

When high QoS is required, MQTT or MQTT-SN should be used for data transmission, since both protocols provide multiple layers of a protected transmission. When energy efficiency is essential and data are transmitted only from point to point, CoAP is the better choice. Figure 4.17 gives an overview of message overheads for these IoT application protocols.

For the highest energy efficiency, the number of transmitted messages should be as low as possible. When no application acknowledgement is required, both CoAP Non-Confirmable and MQTT-SN QoS -1 are applicable. In case of acknowledgement is needed on the application layer, CoAP is the most energy efficient choice since both MQTT and MQTT-SN require a previous connection setup, before transmitting and acknowledging user data.

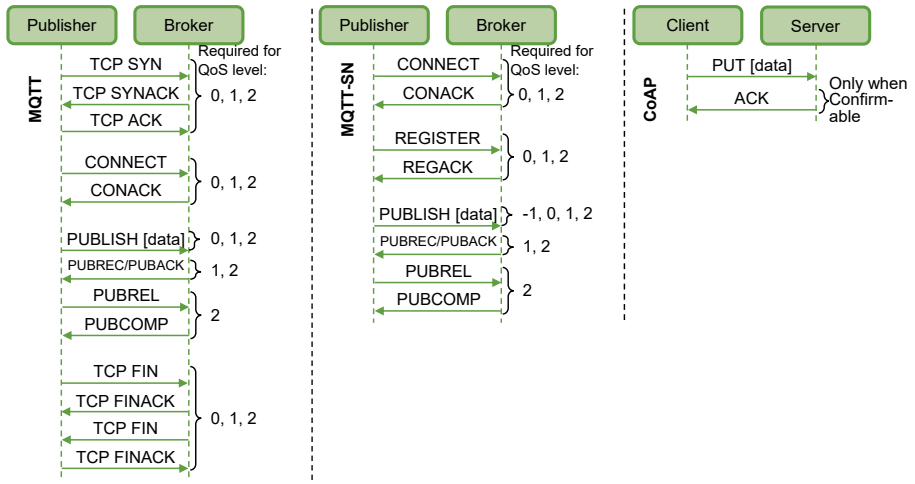

Fig. 4.17: Comparison of message sequences for different IoT application protocols.

4.3.3 Performance Analysis of Resource-Constrained LPWAN Clients

While features like eDRX and PSM aim for lower power consumption and thus longer battery lifetime, coverage enhancement provides a much higher power consumption for UEs by introducing longer transmission and reception intervals. Both features are required to fulfill the challenges of 10-years battery lifetimes, 164 dB maximum coupling loss, and a maximum latency for a single data transmission of 10 seconds. In this section, performance studies of cellular IoT solutions will be analyzed.

4.3.3.1 Power Consumption Analysis of NB-IoT and eMTC in Challenging Smart-City Environments

In Section 4.3.2.1, two new power-saving states for NB-IoT and eMTC devices are introduced. Both states provide a reduced power consumption compared with current GSM or LTE devices. Besides PSM and eDRX, devices can enter three additional power states: Connected, Tail, and TAU. In the Connected power state, random access, data transmission, and reception are performed. After data transmission and reception, the device remains for a predefined time in a Tail state, also called a data inactivity timer, where it remains connected on the RRC communcation layer for additional data exchange. Then, it switches to the eDRX power state, where the device wakes up only for paging occasions. Finally, the device reduces its power consumption to a bare minimum in the PSM state. It periodically wakes up for TAU and checks if downlink transmissions are queued, since these messages can't be received while in PSM. Figure 4.18 gives an overview of the state machine that is used to determine the energy consumption of UEs.

To compare different cellular IoT solutions and assess if all IoT requirements can be fulfilled by NB-IoT and eMTC, the authors in [305] provide a performance analysis of

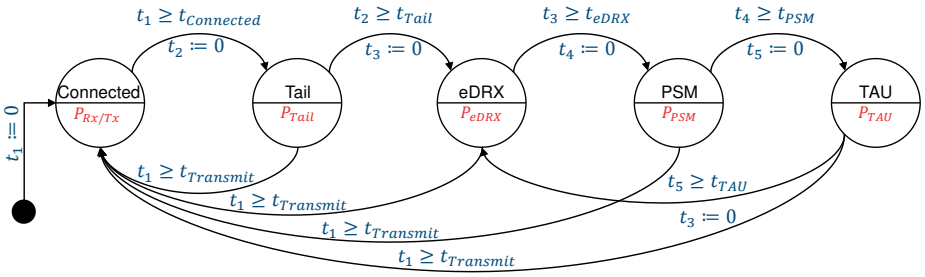

Connected:	Data transmission or reception	PSM:	Power Saving Mode
Tail:	Idle state before entering eDRX	TAU:	Tracking Area Update
eDRX:	Extended Discontinuous Reception Window	P:	Electrical Power

Fig. 4.18: State machine for power consumption evaluation. ©[2018] IEEE. Reprinted, with permission, from [305].

NB-IoT and eMTC latency, data rate, battery lifetime, and spectral efficiency for three different coverage classes. The results are shown in Figure 4.19.

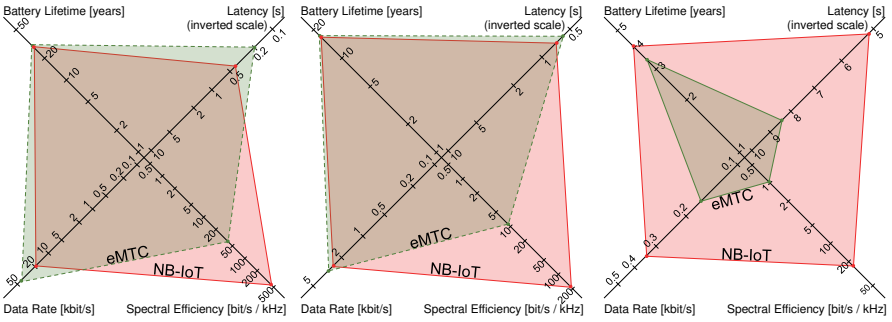

(a) 144 dB coupling loss **(b)** 154 dB coupling loss **(c)** 164 dB coupling loss

Fig. 4.19: Comparison of NB-IoT and eMTC devices for 84 bytes of acknowledged uplink data every 24 hours in different coverage conditions. Note that the axis scales vary between the three figures. ©[2018] IEEE. Reprinted, with permission, from [305].

The results show that eMTC performs slightly better than NB-IoT at coupling losses of 144 dB and 154 dB. While the eMTC performance gain is rather small, it uses 6 times more bandwidth than NB-IoT, making transmissions less spectral-efficient. When it comes to cell edges such as basements, where the coupling loss can increase to up to 164 dB, NB-IoT clearly outperforms eMTC. While both cellular technologies use the same power saving mechanisms, NB-IoT needs fewer repetitions to transmit data, which reduces the time on air and therefore extends the time in PSM between transmissions.

Although eMTC performs better than NB-IoT in good coverage conditions, the difference in data rate, latency, and battery lifetime performance is rather small. When it comes to poor coverage conditions as well as spectral efficiency, NB-IoT is recommended over eMTC.

4.3.3.2 Coverage and Link Quality Improvement of Cellular IoT Networks with Multi-Operator and Multi-Link Strategies

Section 4.3.3.1 has given an overview of the performance of cellular IoT technologies in different coverage scenarios. With increasing signal attenuation the overall performance decreases drastically. Therefore, the signal quality should always be as good as possible. Instead of expanding the networks by installing new base stations, multi-operator strategies (such as National Roaming) can provide better coverage and link quality for LTE and cellular IoT technologies, by allowing cellular devices to use networks from different Mobile Network Operators (MNOs) as well.

The authors in [306] evaluated the potential of coverage and link quality improvement in terms of multi-operator strategies in the Smart City Dortmund as a case study. By extracting the number and locations of all cellular base stations in Dortmund and applying empirical path loss models for urban environments, we performed a comprehensive coverage analysis. The results are given in Figure 4.20.

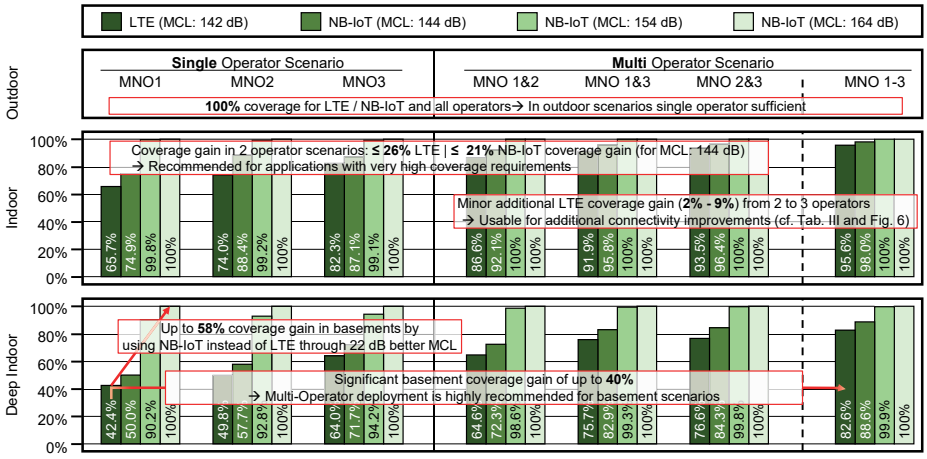

Fig. 4.20: Results of the coverage analysis for outdoor, indoor and deep indoor scenarios and different cellular communication technologies in an urban environment. ©[2019] IEEE. Reprinted, with permission, from [306].

While all MNOs can provide full LTE and NB-IoT coverage outdoors, indoor, and basement coverage, deep indoor coverage from LTE decreases to 66 % and 42 %, respectively.

In multi-operator scenarios, the deep indoor coverage increases by up to 40 %, which makes multi-operator deployment highly recommended in LTE scenarios.

Due to its extended coverage, NB-IoT can provide full coverage in all scenarios when using a maximum coupling loss of 164 dB. If the maximum coupling loss is limited to 154 dB for a better performance (see Section 4.3.3.1), NB-IoT can still provide full coverage when multi-operator strategies are used.

Even in scenarios with full coverage, multi-operator strategies can be reasonable by improving the average link quality by up to 13.6 dB (Figure 4.21). In Figure 4.19, a 10 dB improvement of link quality can already increase the battery lifetime of an NB-IoT device from 4 to 18 years and decrease the latency from 5 s to 0.8 s. Therefore, multi-operator strategies are highly recommended for link-quality improvement.

Average Signal Power Gain for LTE and NB-IoT	MNO 1&2		MNO 1&3		MNO 2&3		MNO 1-3 (National Roaming)		
	MNO 1	MNO 2	MNO 1	MNO 3	MNO 2	MNO 3	MNO 1	MNO 2	MNO 3
Outdoor	9.5 dB	8.0 dB	10.2 dB	8.3 dB	9.0 dB	7.8 dB	11.4 dB	10.0 dB	9.0 dB
Indoor	11.2 dB	9.1 dB	12.3 dB	10.1 dB	10.2 dB	9.0 dB	13.6 dB	11.0 dB	10.6 dB
Deep Indoor	10.4 dB	8.9 dB	11.0 dB	9.1 dB	9.6 dB	8.4 dB	12.6 dB	10.9 dB	10.0 dB

Fig. 4.21: Results of the coupling loss reduction potential for different coupling loss scenarios and cellular communication technologies. ©[2019] IEEE. Reprinted, with permission, from [306].

4.3.4 Conclusion

Energy efficiency is an important factor in the IoT. Many use cases rely on sensors that can last at least 10 years on a single battery. With new communication technologies such as NB-IoT and eMTC, cellular solutions respond to the new challenges that are introduced by the IoT. But optimized communication technologies alone are not sufficient. Energy efficiency needs to be addressed on all layers, from the choice of an appropriate application protocol that produces as low an overhead as possible to link-quality improvement strategies that obviate a high number of repetitions on the air interface. In view of the results of the performance and coverage analysis, a good device placement is of great importance for resource-constrained devices. In the future, extensive measurements of latency and energy consumption can be used to derive an ML-based predictive model for latency and energy performance by using passive measurement parameters such as signal strength and signal quality. Additionally, the number and size of transmitted messages can be reduced by ML-based model-predictive communication, as introduced in [30]. If green networking and resource constraints are taken into account from the very beginning of an application's design, the performance of the system from a user and network perspective can be significantly increased.

4.3.5 Acknowledgments

In addition to the CRC 876, part of this work has been supported by the Federal Ministry for Economic Affairs and Climate Action (BMWK, former Federal Ministry for Economic Affairs and Energy (BMWi)) in the course of the ”Ideal” project, funded under grant agreement no. 03ET7557A.

4.4 Vehicle to Vehicle Communications: Machine Learning-Enabled Predictive Routing

Cedrik Krieger
Benjamin Sliwa
Christian Wietfeld

Abstract: Vehicular communication is used to exchange safety-related status information, enable efficiency-oriented mobility planning, and share other user data, such as video streams, among a locally restricted area and the vehicles and nearby infrastructure within it. Without the need for cellular coverage, the particular devices, or agents, organize themselves in a distributed fashion without a central coordination unit. This ability not only allows the realization of Intelligent Transportation Systems (ITSs) that will have a major impact on the cities of the future, but it also enables spontaneously deployed networks, that cover the task of on-demand network provisioning for events. A well-known example is the support of rescue units that can utilize Unmanned Aerial Vehicles (UAVs) for remote sensing and delegate exploration tasks and reduce the risk of endangering human personnel. These applications all have high requirements and need a robust and reliable communication behavior. As Mobile Ad-hoc NETworks (MANETs) are not managed centrally, data needs to be routed efficiently from the sender to the receiver, whereas link losses and unnecessary hops need to be avoided. Established protocols rely on simple distance measurements and try to minimize the sender-to-receiver distance. In challenging networks, these simple approaches can not cope with the complexity of the task. Therefore, more advanced techniques integrate more information and provide a higher grade of generalization. Comprehensive simulations have shown that the utilization of cross-layer knowledge and the prediction of future network states enable reliable and robust reinforcement learning-based routing algorithms, that achieve high performances under different conditions. Moreover, this technology outperforms established routing protocols by up to 51 % in all considered studies.

4.4.1 Introduction: Direct Agent Communication in Ad-Hoc Networks

Self-organizing networks, where the nodes communicate directly, are described as ad-hoc networks. Routes are built not only to directly reachable neighbors, but also to more distinct nodes with which communication is only possible by invoking (multiple) intermediate hops. This way, all agents create a mesh. Here, the sub-class of Mobile Ad-hoc NETworks (MANETs) is of particular interest, as they explicitly specify the possible movement of nodes, which leads to frequent changes in the network topology and

thus requires suitable solutions for providing a robust and reliable communication. As MANETs are infrastructure-less, efficient routing protocols are required to cover a trade-off between overhead and the provision of valid routes. Whereas too much traffic used for coordination would reduce the overall throughput, outdated information leads to the loss of packets. The nodes' mobility is a primary factor of impact on the network topology. MANETs can have a hybrid consistency of mobile nodes like those of UAVs and cars, but it also has nodes with very low or stationary mobility like those of pedestrians and Road-Side Units (RSUs). Besides mobility, varying channel conditions due to multi-path propagation, especially in urban environments, signal attenuation, and shadowing, have a significant impact on the node's reachability and can harm the end-to-end routing performance. As these varieties of influxes offer challenging conditions for MANETs, established routing protocols that rely on considerably simple metrics, such as distance vectors represented by hop counts, are not able to fulfill the requirement for a robust and reliable communication. Therefore, the integration of further information and the enhancement of routing metrics are motivated in order to assess occurring network situations adequately and increase the overall performance.

4.4.2 Related Work: Evolution of Mobility-Predictive Ad-Hoc Routing Protocols

A classification of the developments of ad-hoc routing protocols is given in Figure 4.22. The proposed protocols originate from the Better Approach To Mobile Ad-hoc Networking (B.A.T.M.A.N.) [472] project located in the Freifunk community. B.A.T.M.A.N. III was originally developed to tackle scalability problems in the established routing protocol Optimized Link State Routing (OLSR) by distributing the topology knowledge among local entities and thus obviate the need to calculate the full network graph on every node, which is an expensive task and especially unsuitable for resource-constrained systems. Subsequent B.A.T.M.A.N. versions relocate their point of operation from Internet Protocol (IP)-based routing in layer 3 to layer 2 in order to provide an network protocol-independent routing approach and have a more direct impact on the packets. As B.A.T.M.A.N. is intended for real-use cases, kernel implementations are available, but, simulation models for scientific research are often omitted. Therefore, a simulation model of the B.A.T.M.A.N V protocol version for the well-known discrete event simulator Objective Modular NETwork testbed in C++ (OMNET++) [701] has been developed and validated by field experiments in [627]. However, as the overall goal of the Freifunk community is to provide mesh-based Internet access within cities, B.A.T.M.A.N. implementations contain overhead to fulfill this task, such as Host Network Announcements (HNAs), is separate from the actual routing process. While the first extension B.A.T.Mobile [623] forks from the main branch and hauls those measures, the consecutive protocol Predictive Ad-hoc Routing fueled by Reinforcement learning and Trajectory knowledge (PARRoT) [635] omits additional overhead for network provisioning. Thus, PARRoT, which gained additional influences of reinforcement learning-based routing,

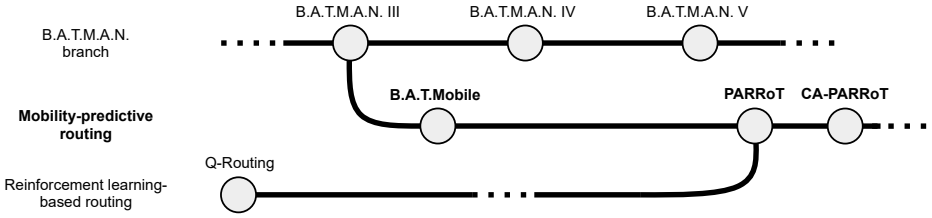

Fig. 4.22: Evolution graph of routing protocols.

concentrates on IP-based routing for an assessment of the concepts. The de-capsulated development of novel approaches intends to merge the newer routing mechanism back into the latest B.A.T.M.A.N. version. The latter protocols, B.A.T.Mobile and PARRoT, follow a mobility-predictive routing approach and are explained in more detail in the following. The most recent development is Context-Adaptive PARRoT (CA-PARRoT) [586], which can be regarded as an extension to PARRoT and follows a hybrid machine learning approach.

Routing protocols can be classified into *reactive* and *proactive* protocols. The first initiate a route-building process on-demand. Well-known examples are Ad-hoc On-demand Distance Vector (AODV) and DYnamic MANET On-demand routing protocol (DYMO). The latter maintain routing tables that are used for lookups when necessary and are updated periodically. Destination-Sequenced Distance Vector (DSDV) and OLSR are widely known proactive protocols. Greedy Perimeter Stateless Routing in wireless networks (GPSR) is a geo-based routing approach that considers mobility and communication as a dependent task. The route building is done by minimizing the geo-distance between sender and destination node with each hop. Extensive summaries about existing protocols are found in [486] and [470]. An empirical analysis of used protocols in vehicular networks is provided in [119].

Recent developments in the machine learning field have also had an impact on routing algorithms. The authors of [671] use a centralized Artificial Neural Network (ANN) to enable a Software Defined Network (SDN) approach for latency minimization in vehicular networks. However, reinforcement learning-based approaches allow routing entities to make autonomous decisions in a decentralized manner. In [99] Q-routing is proposed as an integration of autonomous routing decisions based on learned latency. The authors of [481] extend this with mobility-based metrics and take into account the swarm coherence of agents. A summary of channel and propagation models is given in [704]. The authors of [60] present an approach of learning from stochastic channel parameters. Intelligent routing algorithms and developments for ad-hoc networks might also have an impact on future network generations [17].

4.4.3 Approaches: Enhance Routing by Prediction and Machine Learning

In this subsection, the routing approaches of B.A.T.Mobile, PARRoT, and CA-PARRoT are presented. B.A.T.Mobile leverages cross-layer knowledge from the mobility domain to predict the agents' future trajectory and integrate this information in the routing decision. Following the anticipatory mobile networking paradigm [107], an overall increase in terms of robustness and reliability was achieved. PARRoT, as a follow-up protocol, takes the mobility prediction approach and integrates it into a reinforcement learning process that utilizes abstract metrics [481] and represents the routing process [99]. CA-PARRoT extends this with a mechanism to compensate short-term influences and introduces a hybrid machine learning approach where the routing decision still relies on reinforcement learning, but a machine learning component is used to classify a Radio Environment Prototype (REP) and select an appropriate parametrization to achieve the best possible end-to-end performance.

4.4.3.1 B.A.T.Mobile: Leveraging Cross-Layer Knowledge

B.A.T.Mobile introduces a multi-factorial mobility prediction that classifies available information from the mobility control into three levels of assumed accuracy. An iterative prediction process is performed that always chooses the most accurate prediction method for the current step. A prediction width τ is divided into N iteration steps to predict the future position $\tilde{\mathbf{p}}(t + \tau)$. The considered types of information, named in descending order of their accuracy, are:

- The steering vector σ_i describes the current heading to the position of the next iteration step. This information is only available in the very first iteration.
- The currently targeted waypoint $\mathbf{w}(t)$. If the agent's position is in a specific range of this waypoint, it is considered reached, and the next target is used for the remaining prediction process.
- The history of N_e previous positions. Every iteration step appends a new estimate to this list. It is used as a fallback solution to enable a mobility prediction even if no advanced information is available.

The prediction result $\tilde{\mathbf{p}}(t+\tau)$ is then integrated into the periodically broadcasted routing messages of the underlying B.A.T.M.A.N. routing protocol and replaces the former Transmission Quality (TQ) metric for routing decisions.

4.4.3.2 PARRoT: Transition to an Autonomous Routing Process

PARRoT inherits the mobility prediction of B.A.T.Mobile. The routing metric is not solely built on relative mobility; rather, it is gathered by a Q-learning process. Agents share their current and predicted positions, which are then used by the receiving agent to reconstruct their neighbor's trajectory. Further, this is set into a relationship with the agent's own trajectory, and the future availability of a link between these two agents is

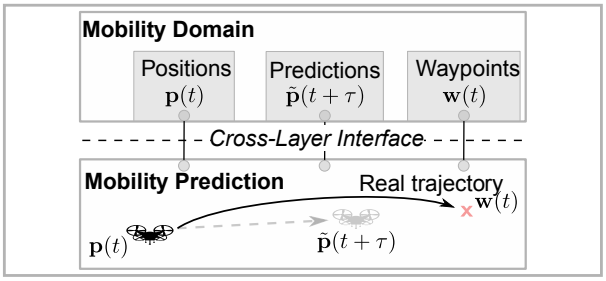

Fig. 4.23: Interconnection of the routing protocol and the mobility domain for multi-factorial mobility predictions. Reprinted from [585].

expressed by the metric $\Phi_{\text{LET}}(d, j)$, representing the Link Expiry Time (LET). Every agent evaluates its own cohesion by comparing its current set of neighbors with a previous set. This produces the metric $\Phi_{\text{Coh}}(j)$, which is shared through the routing messages. The reinforcement learning process is performed by reverse route building. An originator creates a routing message, referred to as chirps in PARRoT, sets the reward value V to 1.0, and broadcasts it. Recipient agents carry out the learning process, considering the originator, which will be the destination d in reverse route building, and the adjacent agent j, from whom the message was received, and process it to

$$Q(d, j) = Q(d, j) + \alpha \left[\gamma_0 \cdot \Phi_{\text{LET}}(d, j) \cdot \Phi_{\text{Coh}}(j) \cdot V_j - Q(d, j)\right] . \tag{4.11}$$

Here, the learning rate α controls the impact of how new routing messages that deliver the reward V_j affect the learned knowledge. The basic discount factor γ_0 guarantees a path degradation. This is of particular interest when all other metrics become 1.0, e.g. in a static network. This measure then forces the shared reward to be decreased and prevents routing loops. The agents maintain a Q-table, where a quality indicator $Q(d, j)$ exists for every destination/gateway pair over which, a chirp has been received. This table is then utilized to feed the known routing table, which enables the packet forwarding process.

Figure 4.24 shows the datagram of a PARRoT routing message that contains identification data, age information, and position information.

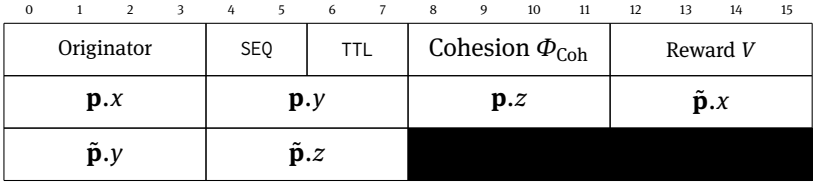

Fig. 4.24: Byte datagram of a PARRoT routing message with a total length of 40 bytes.

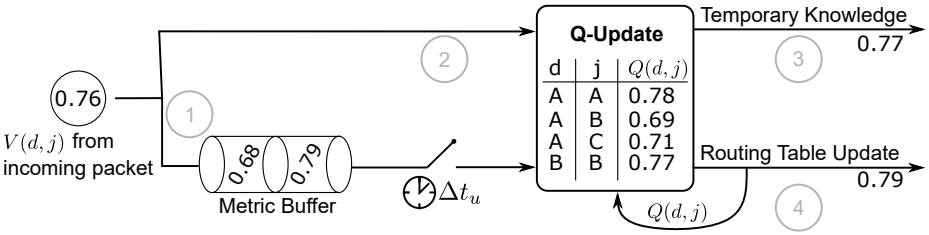

Fig. 4.25: Timer-based update procedure to avoid the short-term effects of having too much impact on learned knowledge. ©[2021] IEEE. Reprinted, with permission, from [586].

4.4.3.3 CA-PARRoT: Introducing a Machine Learning-Enabled Context Adaption

CA-PARRoT proposes further enhancements to PARRoT in order to increase robustness under distinct conditions and reduce the pre-configuration effort. As known from comprehensive simulations, PARRoT has shown a negative impact for time-variant environments as short-term effects harm the accuracy of learned knowledge. B.A.T.Mobile, in turn, proves to be more robust because it buffers incoming information to smoothen the update process. Thus, a timer-based update procedure is also introduced to CA-PARRoT, as shown in Figure 4.25. It is divided into four phases, where an incoming value is pushed to a metric buffer in *phase 1*. The same value leads to an immediate Q-update in *phase 2*, yielding a temporary knowledge in *phase 3* that is used to decide the forwarding of the current routing message. After an elapsed time Δt_u, the best candidate is read from the metric buffer and triggers a Q-update whose result is persisted in the stored knowledge and is also used to update the routing tables in which the packet forwarding is managed (*phase 4*).

Besides the refactored update process, a machine learning component is introduced to obtain a parametrization after the classification of the current environment. For this purpose, different radio environmental prototypes are considered, for which a parameter optimization is carried out in advance, and the best parametrization is provided to the routing protocol. Figure 4.26 shows the adaption approach, which starts with an initial monitoring of the signal strengths of incoming routing messages within an evaluation window. Statistical features are extracted and used to classify a prototype and select its parameters such as the learning rate α, the basic discount factor γ_0, and newly introduced λ and ω, which are used to exponentially weight partial Q-learning metrics.

Random Forests (RFs), Support Vector Machines (SVMs), and Artificial Neural Networks (ANNs) are provided as classification methods. Figure 4.27 shows the classification accuracy that is obtained through a 10-fold cross-validation with the Lightweight Machine Learning for IoT Systems with Resource Limitations (LIMITS) [633] framework, which also enables a model export.

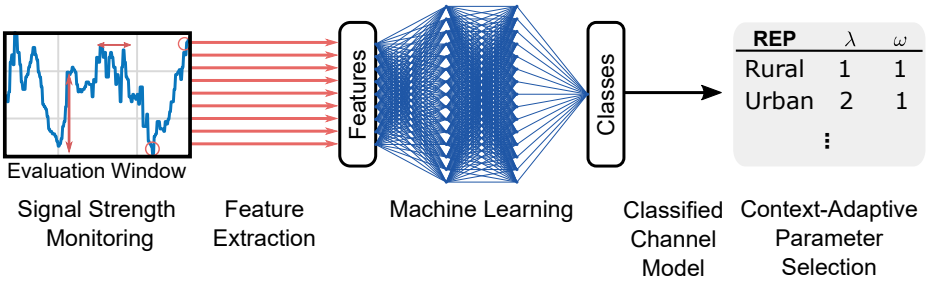

Fig. 4.26: Context-adaption approach for classifiying *radio environment prototypes* (REPs) and selecting parameters accordingly. ©[2021] IEEE. Reprinted, with permission, from [586].

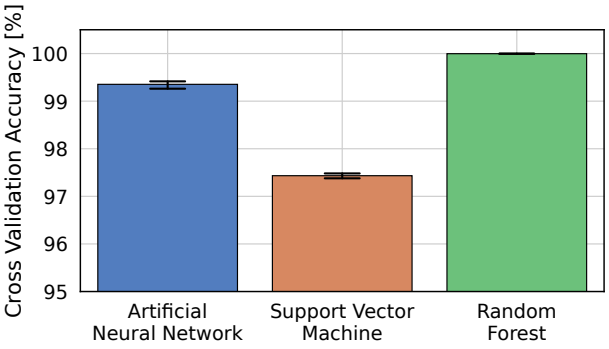

Fig. 4.27: Cross-validation accuracy for the classification of radio environment prototypes. ©[2021] IEEE. Reprinted, with permission, from [586].

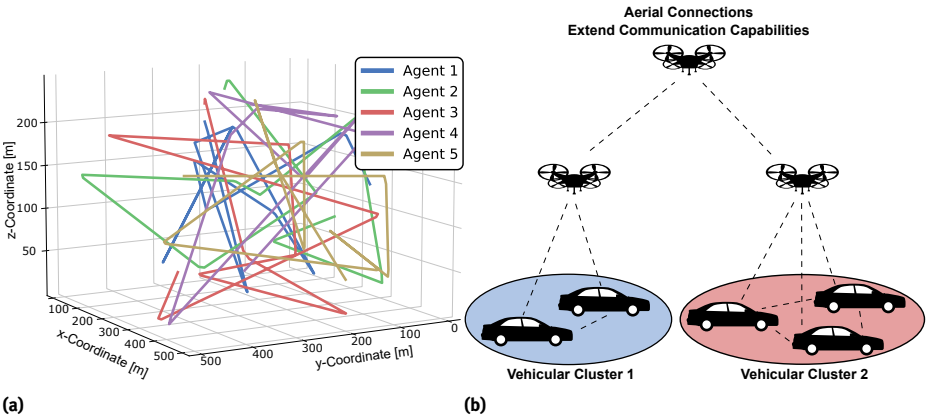

Fig. 4.28: (a) shows the trajectories of 5 agents, following a random waypoint mobility pattern in the three-dimensional reference playground. (b) shows an application-driven mobility, where UAVs fly over clusters of ground vehicles. ©[2021] IEEE. Reprinted, with permission, from [635].

4.4.4 Performance Evaluation of MANET Routing Protocols

In order to carry out a performance evaluation among different routing protocols, a reproducible scenario setup is required that provides a comparable frame for all candidates and makes the overall performance dependent on the evaluated protocol. As Key Performance Indicators (KPIs), the *Packet Delivery Ratio* (PDR) and the mean latency are considered as end-to-end metrics.

A reference scenario is constructed, where ten agents establish a video stream between two agents with a Constant Bit Rate (CBR) of 2 Mbit/s. Besides the communication aspect, the agent mobility has a major impact on the evolution and characteristics of the network topology. Figure 4.28 (a) shows a generic Random WayPoint (RWP) mobility in the three-dimensional playground of the reference scenario. The agents move with a constant speed of 50 km/h and immediately choose their next target when they reach their current one. RWP mobility is considered as a benchmark mobility. Figure 4.28 (b) shows an example of an application-driven mobility, where UAVs hover clusters of ground vehicles to extend their communication capabilities [624]. This mobility pattern involves hybrid types of vehicles that possess inherently different characteristics—a particular challenge for efficient routing protocols.

Another crucial aspect is the choice of a channel model. In rural environments, a Line-Of-Sight (LOS) connection between two agents is usually given. With a lack of objects in the playground, multipath propagation can be neglected, and the free-space path loss L is proportional to the exponentially weighted distance d with $L \propto d^{\eta}$, where η is the attenuation coefficient. This simple but generic model is utilized in the reference scenario for performance evaluation. For other environments, more complex path loss models need to be considered. In urban areas, the impact of objects leads to a higher

importance of multipath propagation, as the LOS path is superposed by reflected signal paths. Therefore, the Nakagami model presents a stochastic impact on the Received Signal Strength (RSS) in addition to the free-space path loss. The empirical analysis of [119] points out commonly used channel models and well-established reference protocols in vehicular research with which the novel protocols are compared.

4.4.5 Results

In this subsection, a comprehensive simulative performance analysis is presented. At first, a scalability analysis is carried out to access the behavior of routing protocols for different types of networks. Afterward, the end-to-end performances are evaluated in scenario studies, that present potential fields of application. To assess an upper bound, an optimal PDR for free-space conditions is introduced, which represents a post-processed analysis of the agents' positions and the theoretical availability of routes. Thus, it is considered a mobility-constrained upper bound.

4.4.5.1 Scalability Analysis

Number of Agents in the Network As seen in Figure 4.29 (a), the number of agents has a significant impact on the PDR of routing protocols. First, for a low number of agents and a low density in the playground, an unneglectable PDR limitation due to mobility can be observed. B.A.T.Mobile, PARRoT, and CA-PARRoT outperform all established protocols for higher density networks. (CA-)PARRoT's course of the PDR is close to the mobility constraint and only shows an impact of routing overhead, which increases for larger scales. Thus, the results show a good scalability, proving (CA-)PARRoT to be suitable for high-density networks.

Impact of Speed on the End-to-End Performance MANETs are characterized by their mobility, which can mean high agent speeds in many cases. Figure 4.29 (b) shows the performance for a range of slow-moving agents up to highly mobile scenarios. High speeds require routing protocols to adapt to the network topology very quickly. For speeds over 100 km/h, most routing protocols fail to provide reliable routes, which causes the PDR to drop and appear undisclosed. Only (CA-)PARRoT, which is affected by the higher requirements of increasing speeds, is capable of providing high PDRs by anticipating the network topology and compensating link losses.

4.4.5.2 Scenario Studies

UAV Communication in a Rural 3D-Playground As seen in Figure 4.30, the novel routing protocols, B.A.T.Mobile, and (CA-)PARRoT, outperform the established reference protocols. The study considers a rural three-dimensional playground with an air-to-air communication between two UAVs. As the communication range does not cover the

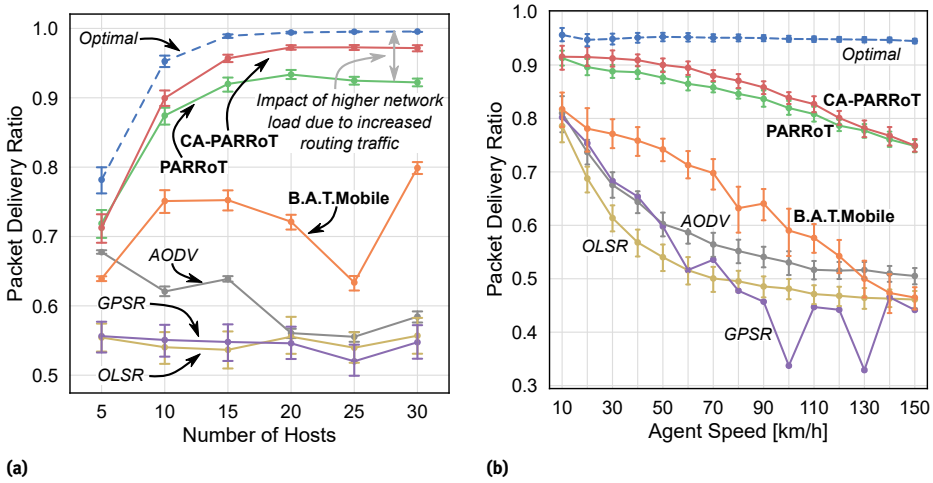

Fig. 4.29: Scalability analysis of routing protocols. PARRoT and CA-PARRoT outperform the considered reference protocols significantly and, thus, provide a robust communication behavior even for challenging conditions.

whole playground, a mobility constrained optimum is calculated, as a routing path is not necessarily available for every point of time.

All established protocols fail to provide a reliable PDR above the 70 % mark, which allows CA-PARRoT to outperform them by 48 % in means. CA-PARRoT achieves a 3 % and 19 % higher PDR than PARRoT and B.A.T.Mobile, respectively, and is thus the best-performing mobility-predictive protocol, with only a 5 % gap to the theoretical upper bound. The proposed protocols, therefore, show gradual improvements with every development stage. Also, considering the latency, CA-PARRoT performs best and shows a 21 % reduced latency compared with OLSR, which is the lowest latency established protocol in this analysis.

Challenging Conditions in Urban Areas Figure 4.31 shows the end-to-end performance for urban radio conditions. B.A.T.Mobile and PARRoT outperform the reference protocols for both KPIs under consideration. PARRoT's performance is a bit weaker than that of B.A.T.Mobile due to the immediate impact of incoming routing packets, where each packet is used for updates, and short-term effects compromise the learning accuracy. The proposed CA-PARRoT is able to avoid this behavior and achieves the highest reliability of all considered routing protocols. Also, in terms of latency, CA-PARRoT shows up to 51 %, lower values compared with established protocols and a 9 %, lower latency compared with that of the B.A.T.Mobile. In general, higher latencies can be observed. These are caused by spontaneous link losses that enforce more buffering in the MAC layer.

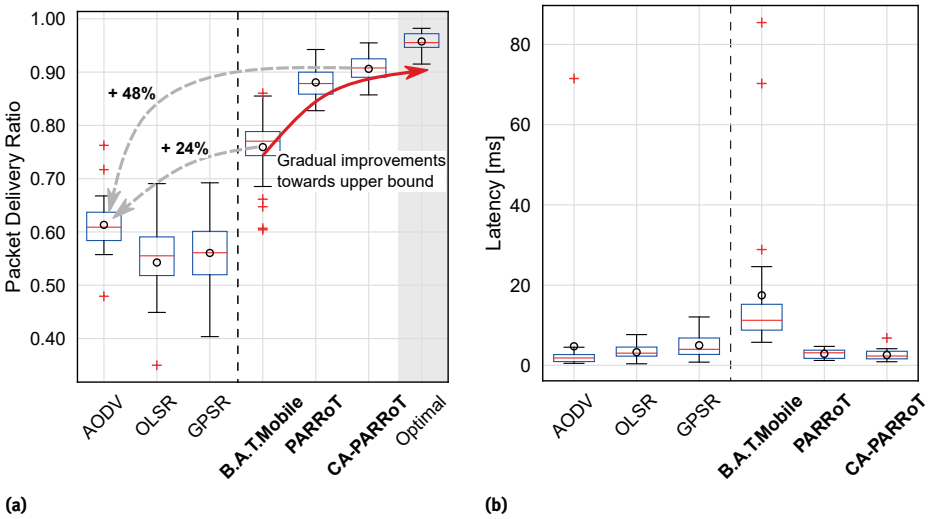

Fig. 4.30: The mobility-predictive routing protocols outperform the established references in terms of reliability, expressed by the PDR. The proactive integration of trajectory knowledge establishes a robust route finding and advances the performance towards the mobility-constrained optimum. ©[2021] IEEE. Reprinted, with permission, from [586].

Performance of MANET Routing in Application-Driven Scenarios The previously presented analysis respects different scalability and radio propagation influences but uses a generic random waypoint mobility model. As real-world applications of MANETs may have mobility in coincidence with the corresponding task, two examples are studied in the following.

Aerial Cluster Hovering

UAVs are used to hover over clusters of cars. In this analysis, ten UAVs are deployed to cover a total of 50 cars, of which ten are equipped with communication interfaces. The remaining 40, therefore, impact only the cluster selection and mobility of other cars, but not the communication system. Figure 4.32 (a) shows the PDR of this scenario. The incremental position updates reduce the accuracy of the mobility prediction and cause B.A.T.Mobile to have high performance losses. However, PARRoT also uses mobility prediction, but is still able outperform all other protocols due to a considered cohesion-aware metric in the learning algorithm.

Distributed Dispersion Detection

To explore plumes, random mobilities are not effective. The distributed dispersion detection (DDD) is a mobility model that is aware of maintaining the cohesion of the UAV swarm during exploration and, therefore, provides a high amount of available

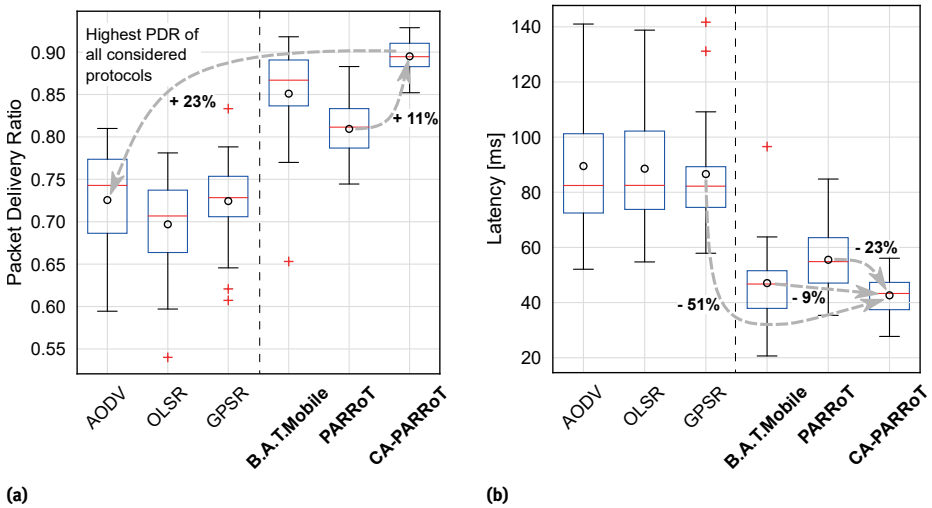

Fig. 4.31: End-to-End performances in a three dimensional playground with urban radio conditions. ©[2021] IEEE. Reprinted, with permission, from [586].

routes. This is also reflected in Figure 4.32 (b), as all routing protocols show a similar, high performance. Nevertheless, B.A.T.Mobile and PARRoT are able to reduce negative outliers and provide a more reliable communication.

4.4.6 Conclusion

The results show that the proactive integration of mobility-domain knowledge enables a significantly more robust behavior, which can outperform established routing approaches in a vast variation of challenging conditions. The utilization of machine learning and reinforcement learning adds an additional gain and robustness, as the comparison between B.A.T.Mobile and (CA-)PARRoT has shown. The more the routing protocols are aware of their environment, e.g., its mobility and radio conditions, the higher the achievable robustness becomes. Therefore, high KPIs could be observed, even in high-scale scenarios. Intelligent routing algorithms are a key component for the realization of efficient infrastructure-less device-to-device communication, not only for WiFi-based networks but also for other technologies such as cellular approaches. The lack of a centralized unit promises lower latency and enables the network participants to learn and adapt to their situation based on local observations, reducing the need for communication expensive status updates.

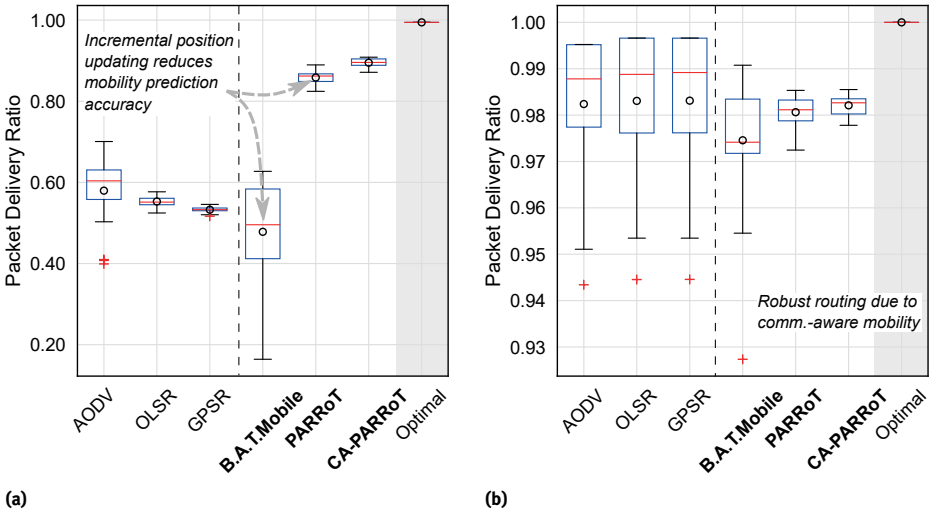

Fig. 4.32: Achieved packet delivery ratios for application-driven scenarios.

4.4.7 Acknowledgments

In addition to the CRC 876, part of this work has been supported by the Ministry of Economic Affairs, Innovation, Digitalization and Energy of the state of North Rhine–Westphalia in the course of the Competence Center 5G.NRW under grant number 005–01903–0047, and in the course of the project Plan Play under grant number 005-2008-0047.

4.5 Modelling of Hybrid Vehicular Traffic with Extended Cellular Automata

Michael Schreckenberg
Tim Vranken

Abstract: Vehicular traffic is a complex system with multiple challenges. For example, highways and urban traffic networks, different vehicle types with varying maximum velocities (cars, trucks, public transportation, etc.), and varying driving behaviors each impacts traffic flow uniquely and strongly. This means that in order to minimize the number of congestions and the average travel times, it is necessary to analyze, model, and simulate traffic in multiple different scenarios.

In the following, we will introduce cellular automaton models for different scenarios. These cellular automaton models aim to reproduce macroscopic traffic phenomena through microscopic simulations. With the help of these simulations, we are able to analyze, understand, and predict traffic in the given scenarios. Lastly, based on the predictions, we can attempt to simulate the same scenarios with small adjustments in order to maximize traffic flow and minimize travel time.

To this end, we will start with introducing and analyzing highway traffic. Here we will focus on applying real-life weather data and how it impacts traffic flow. Next, we will investigate, where the limited space and regular interruption of urban traffic flow by traffic lights and intersections result in new and additional constraints. Lastly, communicating and automated vehicles will be introduced into the simulations. The different reaction times, behaviors, and the human-robot interaction are expected to result in new challenges that have to be investigated and predicted.

4.5.1 Introduction

The topic of vehicular traffic is gaining more and more attention amid rapidly growing numbers of vehicles on the road and increasing amounts of traffic congestion, leading to longer average travel times and more fuel consumption. The road capacity cannot be increased indefinitely through the addition of new lanes, which means that other methods to increase road capacity or use it more efficiently need studying.

In order to use the road more efficiently, traffic and congestion have to be understood better. To this end, we will first analyze and model traffic behavior on highways in Section 4.5.2 with two main goals. The first is to understand the creation of jams and their influence on the traffic flow more deeply. For this, we will use the model by Nagel and Schreckenberg [467]. The second is to model the asymmetric lane changing

required in multiple countries around the world. For this, we will use the one-lane model by Lee et al. [379]. Asymmetric lane changing means that agents are supposed to drive on the most outer (right) lane unless they overtake a slower vehicle.

After analyzing some general traffic flow behavior, we take a look at two more specific problems, the influence of weather on highway capacity, and the missing data of traffic detectors for real-time applications.

While these analyses allow better traffic flow predictions on highways, the situation in urban and inner-city networks is more complicated due to a higher degree of interactions of crossing flows and a regular interruption of the traffic caused by traffic lights. Therefore, in Section 4.5.3 we adapt the model by Lee et al. to include driving behavior before traffic lights and in intersections. Furthermore, traffic flow, based on empirical data from Düsseldorf's inner city, is analyzed. There, different analytical real-time routing methods are applied to minimize traffic jams and the average travel time of the agents. These analyses help identify traffic bottlenecks that have a high impact on the creation of traffic jams and the increase of the average travel time. However, space in urban and inner-city networks is often very restricted, so it is not possible to increase the capacity of these bottlenecks by building more lanes. Due to this, in a follow-up simulation, we changed one of these bottlenecks in a way that two lanes (one leading in and the other leading out of the city) at the chosen intersection were dynamically changed to either lead into or out of the city. This way, the road capacity could be dynamically adjusted to fit the changing demand of commuters.

Lastly, in Section 4.5.4, we simulate heterogeneous traffic with automated and human-driven vehicles. It is expected that automated vehicles will reduce or even eliminate traffic jams at 100 % penetration. However, it will take multiple decades until 100 % is achieved, and the impact of automated vehicles in heterogeneous traffic is unclear due to the different behaviors of automated and human-driven vehicles. Due to this, we adapt our model to simulate the different behaviors of automated, communicating automated, and human-driven vehicles. The goal is to simulate heterogeneous traffic where the three different vehicle types mix and then predict how this will impact traffic flow and road capacities.

4.5.2 Highway Traffic Data Aggregation

The analysis of highway traffic flow can be divided into two topics. General traffic behavior and the real-time traffic situation will be analyzed in this section. For that, we will first model and simulate how traffic jams are created to understand better how and when free-flowing traffic transitions out of free flow due to traffic jams. After that, empirical data is used to create realistic lane-changing behavior while considering asymmetric lane changing rules, as they are applied in countries like Germany or France.

After that, the impact of real-time weather data on jam creation and travel velocities will be analyzed, and the results will be used to increase traffic-flow predictions. The prediction of real-time traffic requires real-time data, which sometimes can be missing due to failing detectors or communication channels. Therefore, the last part of this section will deal with the problem of how to replace missing data accurately.

4.5.2.1 Traffic Jam Analyses

The goal of traffic research is often to prevent traffic jams, or reduce their lifetimes. To this end, Bette et al. [59] analyzed traffic jams using the Nagel-Schreckenberg model [467]. For that, the traffic density was determined at which free-flowing traffic transitions into jammed traffic based on a stability criterion. Afterward, the ratio of jammed cars was separated into different mechanisms, the jamming rate, jam lifetime, and jam size. It was shown that small jams already occur at very low densities and that the increasing life-time of these jams at higher densities is what leads to the transition of the traffic flow from free-flowing to jammed traffic. Furthermore, exponents that control the scaling of all three jam mechanisms close to the critical density have been derived from random walk arguments.

4.5.2.2 Asymmetric Lane-Changing Rules

Lane changing in many countries is asymmetric because drivers are required by law to drive on the most outer (right) lane as long as they do not overtake a slower vehicle. This asymmetric driving behavior creates multiple differences compared with highway traffic without overtaking restrictions. For that reason, empirical data from two countries with such asymmetric rules (Germany and France) have been considered in [248]. A multi-lane cellular automaton model with asymmetric lane changing rules has been created and calibrated based on this empirical data.

This model is based on the one-lane model by Lee et al. [379], where agents have different driving behaviors and a maximal deceleration capability. These two points together allow the model to reproduce accidents due to miss behavior. Because of this, the lane-changing rules have to fulfill three functionalities. Firstly, they have to be safe, which means that the distance to both the preceding and to the following agents has to be large enough, depending on the current velocity. Secondly, the agents should change to an outer lane (one on their right) as soon as the lane change is safe and they do not have to decelerate, while they change to an inner lane (one on their left) only if it is safe and they can accelerate or prevent a forced deceleration. Lastly, the agents have to be prevented from overtaking a slower driving vehicle on a more outer lane.

After adding rules to ensure all three points, the model is able to reproduce empirical lane usage for two- and three-lane highways. Furthermore, a higher number of lanes can be simulated easily after one parameter of the model is re-calibrated accordingly.

Tab. 4.5: Maximum vehicular traffic flows obtained from the whole empirical dataset. Here, w represents the mean of all available water film sensors at a certain point of time.

Surface	Range of water film thickness w	Maximum flow
Dry	$w \leq 0.15$ mm	1960 vehs/(h, lane)
Damp	$0.15 \text{ mm} < w < 0.9$ mm	1720 vehs/(h, lane)
Wet	$w \geq 0.9$ mm	1320 vehs/(h, lane)

4.5.2.3 Influence of Weather Data on Traffic Predictions

Vehicle-2-X communication has grown rapidly within the past two decades, which has increased the availability of extended Floating Car Data (xFCD) that can be applied in the field of traffic information and improvement. One possible integration of this additional real-time data is the inclusion of weather data in traffic-flow predictions [246]. In order to identify the current weather on the road, the water film thickness is taken by local weather stations. Vehicles equipped with xFCD are able to gather and communicate this data through the use of rain-sensing windscreen wipers, which react to water spray. While this vastly increases the available data, floating car data can be more unreliably due to the limited radio spectrum available to transmit this data. An efficient way of communication is developed in Section 5.2 to transmit data reliably and efficiently. An analyses of the correlation between the average velocities of passenger cars and water film thickness on the road showed a strong negative correlation of up to -0.4 for rush-hour traffic. This means that the incorporation of weather data into traffic information systems is expected to be exceptionally beneficial for commuters. Furthermore, Table 4.5 shows that an increased water film thickness also decreases the minimal road capacity above which the traffic flow becomes unstable and can transition away from free-flow, which indicates that the inclusion of weather data improves traffic-jam predictions.

After the impact of water film thickness on the traffic flow was analyzed, the new insights were added to the previously discussed asymmetric multi-lane version of the Lee et al. model. For that, an additional dallying parameter $p(w)$, which depends on the water film thickness on the road w, was introduced. A higher value of this parameter increases the probability of an agent decelerating even if the leading vehicle is far enough away for its velocity to be safe. The results in Figure 4.33 show a good agreement with the empirical data for the roughly 6 km-long Autobahn section chosen for the study. This also shows that an accurate traffic prediction needs reliable real-time data in order to work. However, empirical data is often not reliable enough to ensure constant real-time updates, and data can be missing. This missing data then has to be replaced by approximations in real time.

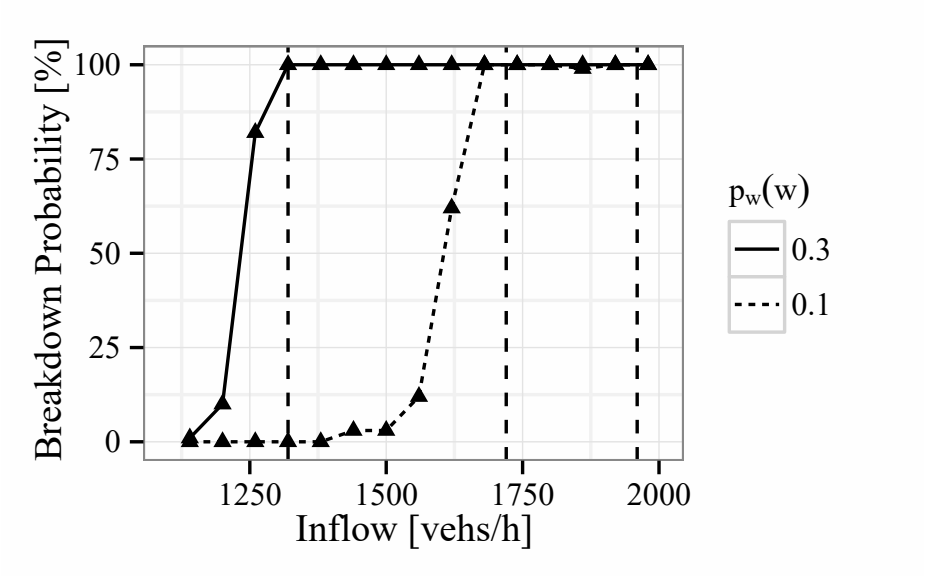

Fig. 4.33: Simulated breakdown probability for two different degrees of surface wetness.

4.5.2.4 Replacing Missing Data

Previously, in the exponential smoothing prediction method [136], missing data of a traffic detector was replaced analyzing the historical data of this detector. This has two major downsides. Firstly, the historical data ideally has to be from the same day and time in the past 30 weeks of all D detectors. If, for example, a detector stops communicating its findings on a Monday at 12 o'clock, then the data of 12 o'clock of the past 30 Mondays has to be considered. This means that all data points of all D detectors have to be saved for at least 30 weeks in order to make accurate predictions on missing data points. The second problem is that the traffic situation ten weeks before, for example, does not have to be the same as today. Accidents or road works, for example, could shift the traffic flow from one street to another, which would strongly increase the error of such predictions. Because of this, a new method to replace missing data has been introduced in [247].

In this new method, 60 minutes of historical data of the surrounding N detectors from the preceding week is taken. This data is then used to train a Poisson Dependency Network (PDN) [249] (which is a form of a Poisson model explained in Section 4.1.2.6). This PDN shows how strongly the traffic-flow data at one detector point correlates with the data of the detector that has missing data. Then, in a final step, the real-time data of the other d detectors is inserted into the PDN to fill the missing data point.

In order to test the method, empirical data was taken, and a prediction of it was made as if it were missing. Then the empirical data was compared with its prediction in Figure 4.34. One can see that the PDN is closer to the actual data than the exponential

smoothing prediction that just used the past N weeks of traffic data of the "missing" detector. A general test at different times of the day and different days of the week showed that the PDN not only uses less historical data but also predicts missing traffic data more accurately. A more in-depth analysis of the problem of traffic flow prediction is given in Section 4.1.

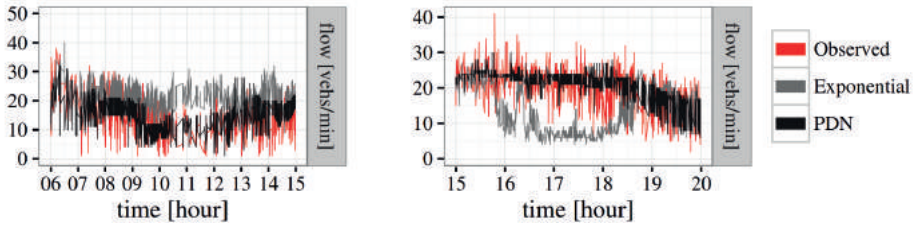

Fig. 4.34: Example empirical data together with exponential smoothing and PDN predictions based on the data.

4.5.3 Urban Traffic Simulations

Inner-city traffic is more complicated than highway traffic due to more interactions of vehicles with different travel directions and intersections with or without traffic lights. In order to analyze and simulate inner-city traffic, the Lee et al. model, which was used as a basis for highway traffic simulations in the previous section, was modified with additional rules in [710] to reproduce empirical intersection traffic data. In another work, different methods to dynamically optimize inner-city traffic through different routing methods and a dynamical application of lanes were analyzed [709]. There, we were able to show that while traffic flow is often above the network capacity, one can decrease the number of traffic jams and the average travel time through more dynamic routing methods and the use of the infrastructure.

4.5.3.1 Cellular Automaton Model

The simulation of urban traffic with a model based on that by Lee et al. [379] has to include a couple of complex situations. Different intersections, for example, can be very different in their structure and serve different purposes. Some intersections can have lanes on which one is only allowed to turn left, while another intersection with the same number of lanes and roads allows turning left or keeping going straight from the most inner (left) lane. Furthermore, agents have to ensure that they arrive on the right lane, depending on their route before they arrive at the traffic lights. Finally, vehicles that turn left within an intersection sometimes have to take into account that the crossing

traffic has a green light at the same time, and they are only allowed to turn if they do not disturb this ongoing traffic flow.

After these points were included within a cellular automaton model in [710], the model was calibrated so that the time requirement $t_B \approx 1.6 \pm 0.6$ s is in accordance with empirical data. To calculate the time requirement t_B, the number of cars that arrive at a red traffic light in front of an intersection within the simulation is set higher than the number of vehicles that can pass this traffic light once it turns green. When the traffic light turns green, the time each agent requires to move into the intersection after the preceding agent moved into it is taken. This time is the time requirement t_B and, after 3 – 4 s of a green traffic light, it averages out to around $t_B = 2.23 \pm 0.04$ s. This is in accordance with the empirical data.

Afterward, an intersection was modeled in that the traffic lights that control left-turning traffic have green at the same time as the crossing traffic flow. Because turning agents decelerate before they turn, the time requirement is already higher than for traffic that does not turn, which means that the lane capacity is lower. Furthermore, this time requirement increases strongly in correlation with the crossing traffic flow. This means that in order to minimize traffic jams and travel times in urban traffic, one not only has to consider the traffic flow from one point to another but also all other traffic flows within the network.

4.5.3.2 Inner-City Traffic Optimization

In [709], inner-city traffic was analyzed and simulated with the help of empirical data from Düsseldorf's inner-city. Based on the analyses of the traffic flow, different routing methods were used to guide the traffic through the application of real-time data. The analyses showed that the traffic capacity of roads inside the city is not enough to cover demand at all times. However, a routing method that aims to make maximal use of the road capacity rather than route vehicles depending on their travel time would improve the traffic flow significantly and could reduce the average travel time by up to 23 %. The downside of optimizing the traffic depending on the network capacity rather than the travel times is that if the traffic flow is below the road capacity, vehicles will take roads with longer travel times than necessary, which increases fuel consumption and travel times unnecessarily. Due to these findings, a new routing method was developed that considers both the road capacity and the travel time of each agent individually. As one can see in Figure 4.35, the new routing method (green) recreates shorter travel times than the network optimization method (red) at low traffic volumes while also reducing the travel times over routing methods where each agent uses the route with the shortest travel time (blue and purple) at high travel volumes.

After it was shown that a more efficient routing method would decrease the average travel time and make more use of the given road capacity, we tested how the road capacity itself could be dynamically optimized. To this end, one of the busiest intersections leading into and out of the inner city was identified. This intersection

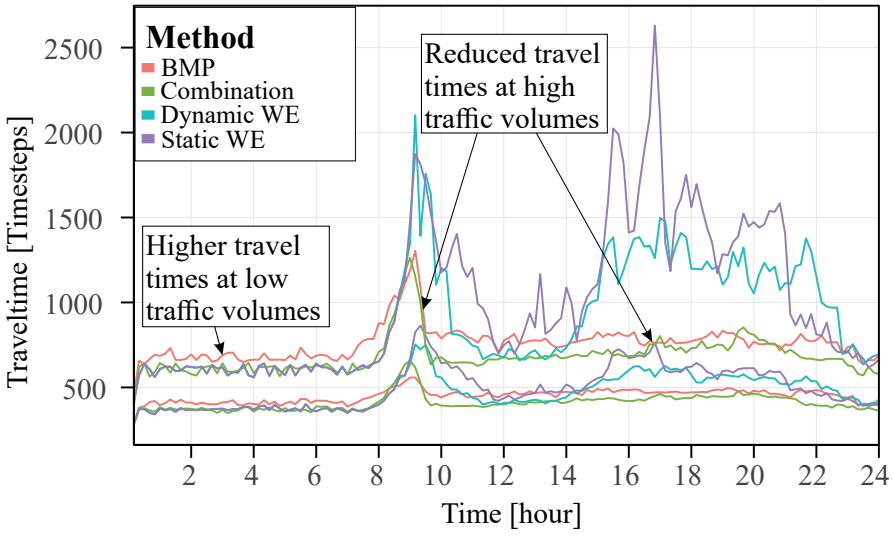

Fig. 4.35: Travel times for different traffic assignment methods. The lower lines represent the average travel time taken over 10 minutes while the upper line represents the 95 % confident interval upper bound.

connects the city ring with one of the main roads leading through the city. One of the lanes on the city ring, as well as one of these on the road leading through the city, were changed to dynamic lanes. Traffic into and out of Düsseldorf is very unbalanced at different times of the day due to the high number of commuters. The dynamic lanes were able to reduce the commute time into the city in the morning while also decreasing the travel time needed to leave the city in the afternoon by increasing the road capacity where it was needed and decreasing it where it was not needed. Through this dynamic change done to a single intersection, the average travel time could be reduced by over 10 % without changing the way vehicles currently choose their routes. Note that these travel-time reductions aren't necessarily the optimal reductions. The goal is rather to understand the network and find its bottlenecks. How to find the optimal routes for each vehicle to reduce the global average travel time is analyzed in Section 4.1.4. There, a new method that applies a reinforcement learning algorithm is simulated, and the results are compared against others.

4.5.4 Automated Vehicular Traffic Flow

Even though dynamic routing and dynamic shifting of the road capacity can reduce the average travel time, this is not a permanent fix for the increasing number of traffic jams

since the number of vehicles is expected to continue increasing in the near future. By contrast, automated vehicles in an 100 % automated traffic are expected to increase road capacity, reduce travel times, fuel consumption, and accidents significantly. However, it is not clear yet how automated vehicles will be introduced into the traffic. Their impact on the traffic dynamics in heterogeneous traffic, when they mix with human-driven vehicles, has been widely discussed. Traffic will be heterogeneous for multiple decades due to an average lifespan of a vehicle of around ten years [42].

In order to predict heterogeneous traffic behavior, a new cellular automaton model was introduced in [708]. One of the big challenges of cellular automaton automated vehicle traffic is that Automated Vehicles (AVs) and Communicating Automated Vehicles (CAVs) have a reduced reaction time compared with human-driven vehicles (HVs). For CAVs, this reaction time could go as low as the time it takes to communicate, which is currently around 0.1 s but could go even lower with 5G, which is expected to become an important method for future connected and automated vehicles. Currently, communication does not always take a fixed time length but instead varies depending on the limited available radio spectrum. If the full spectrum is used, communication can take a lot longer than 0.1 s, depending on the means of communication. An in-depth analysis of the problems of communication and how it is realized is given in Section 5.2. For the remainder of this section, we will assume a stable communication with a 0.1 s communication time. Therefore, the first step toward simulating AVs or CAVs with cellular automaton models was to reduce the time-step length of the model to 0.1 s per time step.

4.5.4.1 Reduced Time-Step Length

The cellular automaton model introduced in [708] is based upon the Pottmeier et al. [508] accident-free version of the Lee model [379]. Human-driven vehicle agents in this model can judge their situation optimistically or pessimistically. If they judge it optimistically, they do not expect their leading vehicle to decelerate strongly (only dawdle). In this case, they can follow the leading vehicle with less than the minimum safety distance, something that is often found in empirical data. If they judge their situation pessimistically, they expect the leading vehicle to decelerate at any moment. They do not follow with the minimum safety distance but even apply an additional safety distance, depending on their velocity.

The newly introduced model reproduces this behavior while also having a 0.1 s long time step and keeping an average reaction time of 1 s for the HVs. For that, multiple changes were made to the calculation of the safe velocity, dawdling, and the judgment of the situation. However, Figure 4.36 shows that the resulting 0.1 s long time-step model presented in [708] is still able to reproduce realistic human-driven vehicle traffic comparable with the model presented by Lee et al. [379] and the modified version of that by Pottmeier et al., which are both known to reproduce empirical traffic well [508]. As one can see, the main difference between the three models is at a density of $20 - 45\ \frac{\text{Veh}}{\text{km}}$

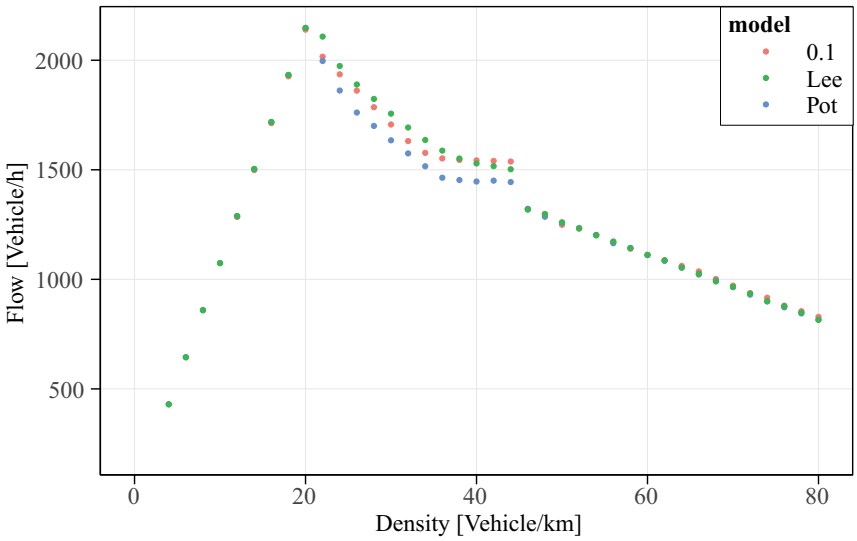

Fig. 4.36: Fundamental diagram for the three different human-driven vehicle cellular automaton model.

in the so-called synchronized traffic phase [317]. Pottmeier et al.'s modification of the Lee model only differs from the Lee model in the calculation of the judgment of the situation. Agents in this model are less likely to be optimistic, which prevents accidents but also increases the average vehicle following time and so reduces the traffic flow. The reduced time-step length also uses this curbed optimism due to which the traffic flow is initially below that of the Lee model. However, due to changes to the dawdling, the velocity distribution is more uniform in the 0.1 s time-step length model, which strengthens the synchronization and increases the traffic flow towards the end of this traffic phase.

Overall the differences between the three models are smaller than the fluctuations observed in empirical traffic, and they all reproduce realistic traffic flow.

4.5.4.2 Heterogeneous Automated Vehicle Traffic

AV and CAV agents have multiple differences compared with HV agents. Their three most important differences are reaction time, dawdle, and behavior calculation. An AV has a reaction time of 0.5 s and a (CAV) one of at least 0.1 s compared with the 1 s of HVs. Neither CAVs nor AVs dawdle at all, while HVs dawdle with a probability of up to 37 % [379]. Lastly, CAVs and AVs do not judge their situation optimistically or pessimistically, but instead, they always follow the leading vehicle with at least the minimum safety distance. If the leading vehicle is human-driven, then they follow with more than the

minimum safety distance to be able to react to unpredictable human behavior without creating strong deceleration waves.

These differences between AVs, CAVs, and HVs were defined and included in the model in [708] before heterogeneous traffic flow was simulated. The simulation results are shown in Figure 4.37, together with the theoretically predicated capacity increase [215] for heterogeneous traffic. One can see that automated vehicular traffic is expected to increase the traffic capacity compared with homogeneous human-driven. The effect is even stronger for communicating automated vehicles. A reduced reaction time means that the minimum safety distance (and so the average vehicle following time) is lower, which allows higher traffic flow at similar densities. Furthermore, the reduced reaction time also means that these automated agents overreact less to human dawdling, which reduces deceleration waves. This effect is strengthened because they do not apply optimistic or pessimistic behavior states. If an HV agent changes its behavior from optimistic to pessimistic, it needs a higher distance to its leading vehicle even if the velocities would not change. This means that the agent has to decelerate more than AVs or CAVs that only have to decelerate as much as the leading vehicle to keep up the same distance.

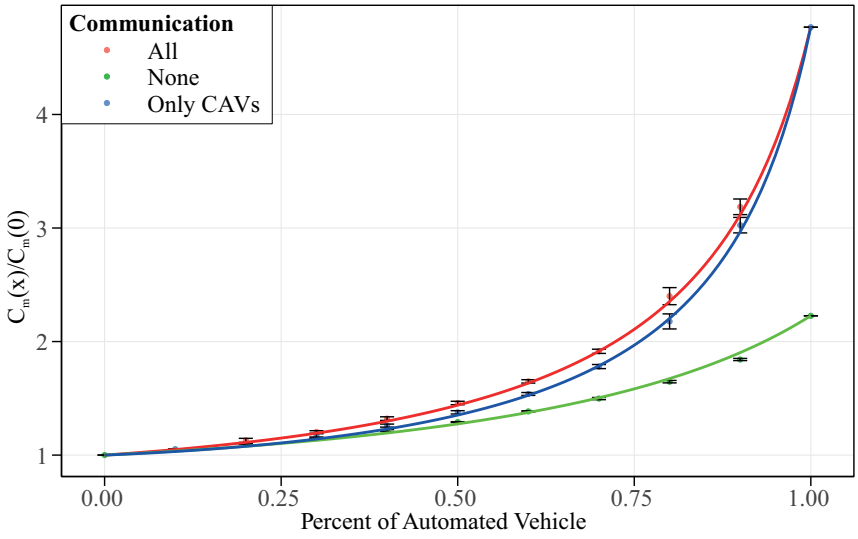

Fig. 4.37: Road capacity increase over homogeneous human-driven vehicular traffic depends on the percentage of automated or communicating automated vehicles.

However, while these results show an improved road capacity in every heterogeneous traffic situation compared with purely human-driven traffic, the model was also able

to reproduce an increased rear-end-collision risk due to human-robot interactions. Human drivers tend to drive with less than the minimum safety distance if they judge the traffic situation optimistically and do not expect to have to decelerate within the next couple of seconds. Before, when they followed another HV agent or even an AV agent, the remaining safety distance the agent used was just enough to prevent accidents if the traffic situation suddenly changed. Now, however, because the CAVs are able to react with only one-tenth of the time the average human needs, the velocity difference between a decelerating CAV and the following HV is so much larger than before that the following HV is not able to prevent an accident once the agent has reacted.

This shows that the different behaviors of HVs and CAVs have the potential to increase the rear accident risk. However, the results shown in Figure 4.37 are simulated after this accident risk was prevented through a change in the behavior of CAVs. If an HV agent n follows a CAV agent $n + 1$ with less than the minimum safety distance, then the CAV agent increases its distance to its leading agent $n + 2$. This way, if agent $n + 2$ decelerates, then agent $n + 1$ can decelerate after driving this additional distance, which gives agent n the time to react to the brake lights of agent $n + 2$, thus preventing accidents.

4.5.4.3 Conclusion

In this section, we highlighted different analyses of vehicle traffic scenarios on highways, in urban areas, and among heterogeneous traffic where automated and conventional vehicles are mixed. The main goal of those analyses is to understand and predict traffic better. Furthermore, the already existing cellular automaton vehicle model introduced by Lee et al. [379] was modified for multiple occasions.

We were able to show that traffic jams already form at very low densities and that the increase in their duration leads to a transition of free-flowing to jammed traffic. While a critical density from that on free flow could transition to jammed traffic flow, it was also shown in other works that this critical density is not fixed and can vary. On highways, the local weather in the form of rain (measured by water film thickness) was identified as an important influence on traffic flow, while urban traffic is mostly dominated by traffic lights. Both problems were analyzed through the use of real-life data as well as simulations of a modified Lee model.

Through the results of those analyses, the accuracy of traffic forecasting could be improved as long as enough real-time data was given. Unfortunately, the reliability of empirical detectors is often not fully guaranteed, and detectors can malfunction or not communicate data. For such cases, we have introduced a new method to fill missing data gaps that uses less historical data and is more accurate. Lastly, we modified the Lee model to predict the impact automated and communicating automated vehicles will have in future traffic when mixing with conventional vehicles.

4.6 Embedded Crowdsensing for Pavement Monitoring and its Incentive Mechanisms

Maximillian Machado
Ran Ran
Liang Cheng

Abstract: Due to its pervasiveness and convenience, crowdsensing is regarded as an effective method to collect specific data. This section surveys projects that take advantage of embedded crowdsensing to collect pavement condition data and describe how crowdsensing platforms conduct road damage detection using deep neural networks with images captured with smartphones. Before such discussion, we explore how to motivate users to participate in low platform-cost crowdsensing tasks. Our research models the pavement crowdsensing problem and designs new incentive mechanisms based on a platform-driven greedy algorithm. Through extensive simulations, the performance of the incentive mechanisms is evaluated and compared in different scenarios in terms of the platform cost and the overall task completion time. The best of them can reduce the total completion time by half compared with the reverse auction incentive mechanism. We conclude this contribution with future work discussions.

4.6.1 Introduction

As road infrastructure increases in size and complexity, innovative solutions must be developed to cope with road degradation. With the current methods used to collect road condition information, covering 4.18 million miles of road in the U.S [555] is a difficult undertaking. One viable solution is to create a crowdsensing platform where smartphone users collect road pavement data. In this case, users could automatically detect and report poor road conditions using embedded cameras, accelerometers, and 4G/5G networks. However, such a network would require an active user base and a means of maintaining crowdsensing participation. Thus, crowdsensing schemes must design an incentive structure for continuous user activity. A crowdsensing platform utilizes an incentive mechanism to motivate user participation, produce diverse data pools, and generate quality information.

Later in this section, we will formally express the incentive mechanisms as one of the key components of crowdsensing platforms. The mechanisms designed in this section are evaluated through experimental procedure, and vary based on the reward distribution assigned to sensing tasks. Favorable mechanisms are those that exhibit both low *platform cost* and *total operation time* when completing sensing tasks. Note, these criteria are not the only metrics for measuring incentive mechanism efficacy.

However, *platform cost* and *total operation time* metrics are sufficient for providing a solution framework to the purposed crowdsensing problem.

Our crowdsensing platform is parameterized with nine unique incentive mechanisms, and aims to collect pavement condition data for varying percentages of a fixed area with economic feasibility in mind. All mechanisms are based on a platform-driven greedy algorithm that motivates users to select sensing tasks that can provide the highest net profit margin for the participant. Eight of the nine incentive mechanisms are uniquely defined in this section while an additional incentive mechanism is motivated from recent related work in crowdsensing literature. Our results provide guidance in selecting the best incentive mechanism in different settings of pavement crowdsensing.

Here are the key contributions of the section:

- The incentive mechanisms we design can effectively avoid the cost explosion problem as users choose their sensing tasks before starting to work on them. Thus, sensing tasks can only be committed to by one participant at a time. Cases where a sensing task should be reexamined–possibility due to poor readings–can be addressed by modeling repeated sensing tasks.
- Our mechanisms enables users to select sensing tasks that offer the highest net profit margin based on a greedy algorithm.
- The total operation time of our approach is reduced compared with that of the task-reverse-auction incentive mechanism. The results highlight this claim and provide solutions for crowdsensing given a target area within a limited budget.

The rest of this discussion is organized as follows: survey the related work, introduce the research problem and its model, present our incentive mechanism solutions, construct simulations for evaluating the incentive mechanisms, discuss the evaluation results, consider augmented machine learning techniques, and conclude with final statements.

4.6.2 Incentive Mechanisms

4.6.2.1 Existing Monetary Incentive Mechanisms

Zhang [763] and Jaimes [299] both assort incentive mechanisms by the types of incentives. In Jaimes [299], monetary and non-monetary incentives are compared. Non-monetary mechanisms [158, 166] rely on the continued participation of users due to intrinsic motivations. Monetary mechanisms [15, 347, 356, 380, 381, 735, 743, 762, 766] rely on the direct backing of fiat money or indirect backing of fiat money through alternative currencies. According to a survey paper [763], monetary incentives will be more likely to motivate users to complete the sensing tasks than non-monetary incentives. Therefore, a monetary mechanism is more fitting for crowdsensing and will be considered in our discussion. However, picking the correct monetary incentive mechanism poses additional challenges.

When deciding on which monetary mechanism to use, one must define the optimization criteria that best suits the crowdsensing scheme. Examples of such criteria include economic feasibility, area coverage, data quality, fairness, and time duration. Producing a platform that optimizes one or more of these categories is non-trivial because the typical problem framework of a crowdsensing scheme comes down to exponential time complexity problems or typical game theoretical model challenges. In the case of economic feasibility, a platform must be designed carefully in order not to allow users too much control over the price of their service. This is known as the cost explosion problem, and is one of many challenges that must be addressed when considering appropriate monetary mechanisms. The following three monetary incentive mechanisms are well studied solutions to various crowdsensing schemes.

- The task-reverse-auction incentive mechanisms [738, 762, 766] allows for a set of users to bid on the set of tasks posted by the platform. Each bid represents a promise to finish a task provided that the platform will pay the user the bid value. Naturally, the user who bids the lowest price wins and gets the opportunity to perform the sensing task.
- In the case of the data-reverse-auction incentive mechanisms [356, 381], a set of users auction their sensing data for the posted set of tasks and their prices per already finished tasks. Then, the platform selects the data that satisfies its criteria and pays the users their bid price.
- The platform-centric model [743] treats the crowdsensing problem as a Stackelberg game. The reward of the task is changed until the platform and users reach a Nash equilibrium.

4.6.2.2 Examination of Three Typical Incentive Mechanisms

Most of the existing incentive mechanisms fall under game theoretical models. In this subsection, we examine the three aforementioned mechanisms.

The main problem of the task-reverse-auction approach [738, 762, 766] is that, because of untruthful bids, the auction style does not always select the nearest user to complete the sensing tasks [766]. In this situation, the user who is far away from a sensing task can win the auction. Further distances result in a longer travel time for users. Thus, the task-reverse-auction incentive mechanisms need more time to complete all the sensing tasks than our incentive mechanisms.

For data-reverse-auction incentive mechanisms [356, 381], while multiple users collect the data for one sensing task, only one user's data can be accepted by the platform. In other words, other users' data is wasted. As a result, this type of incentive mechanisms increases costs for car fuel, personal free time, etc. For our incentive mechanisms, users can select the sensing task before they go to collect the data. Thus, the cost explosion problem can be avoided.

The platform-centric model [743] assumes that the platform has no upper bound on budget budget. Therefore, it can find an optimal solution to giving the platform the

highest-quality data available. In practice, the platform usually has a limited budget and may not be able to obtain the game theory equilibrium reward. Our incentive mechanisms provide heuristics to obtain data while working under tighter budget constraints, and can be considered more consistent with an online crowdsensing scheme.

4.6.3 Incentive Mechanism Research Problem and Its Model

For our research problem, the platform needs to motivate the users of the platform to collect the road pavement data constraint to a budget and target area. In this case, our research objective is to design an appropriate incentive mechanism to help the platform achieve an area coverage target with a low cost and total operation time. Based on the comparison results of incentive mechanisms, the platform can choose the best incentive mechanism with the lowest budget for different area coverage targets.

Our model of the research problem contains three entities: the environment, the sensing task, and the user. Each entity can be described by its behavior and/or its relationship with other entities:

- The environment entity based on the Manhattan model; it is a grid of cells without loss of generality for incentive mechanism studies. The grid has a uniform cost distribution for traveling across adjacent cells, and no missing cells within. The environment represents the types of roads that users may encounter and the varying costs of traveling with different pavement conditions. Lastly, user entities can transfer their position only to one orthogonal cell per unit of time; users cannot move diagonally.
- The sensing task entity contains information on the location of interest and the monetary incentive associated with user participation. The sensing tasks specify roads where pavement sensing is needed.
- The user entity represents users participating in the crowdsensing scheme. As users continue to collect and report data for rewards, they accumulate monetary rewards and endure operation costs.

4.6.4 Incentive Mechanism Solutions

Modularity and scalability are critical features needed in designing a crowdsensing framework for deploying and testing incentive mechanisms. It would be difficult to swap incentive mechanisms and evaluate them without these features. Our crowdsensing platform and incentive mechanism designs are guided by the evaluation metrics described in this section.

4.6.4.1 Notations

The symbols we use in this report are shown in Table 4.6. Two important variables in our model are s_j and u_i. They represent identification numbers of the sensing tasks and users. The tasks s_j and the users u_i have attributes $< u_i, R_{ij}, x_j, y_j >$ and $< s_j, a_i, C_{ij}, x_i, y_i >$, respectively. For users, if s_j is 0 or -1, then the user is currently not participating because the user has not selected a sensing task or has dropped out. For sensing tasks, if u_i is 0 then the sensing task has not been assigned to a user. In addition, if a sensing task has a reward equal to 0, then its reward has been claimed.

Tab. 4.6: Common symbols.

Symbols	Meanings
a_i	Accumulated reward of user u_i
Avg_j	Average distance from task s_j to all users
B	Budget for the platform
BR	Base reward
b	The side length of the grid
C_{ij}	The travel cost for u_i to complete s_j
CR	The reward of the task that offers MP
$d_{j,uc}$	Distance from s_j to uc
$d_{j,tc}$	Distance from s_j to tc
IM	Incentive mechanism
k_i	The ranking number for u_i
MP	Maximum profit for user u_i
NPM	Net profit margin
P	Area coverage percentage
P_{ij}	Profit for u_i of sensing task s_j
PC	The platform cost
R_{ij}	Reward of the sensing task s_j for user u_i
(S) s_j	(Set of) Sensing task/ID
S_a	The set of available tasks
SID	The index of task selected by user u_i
s_r	The percentage of trials succeed
T	Threshold for net profit margin
tc	The center of locations of sensing tasks
t_f	Total operation time
(U) u_i	(Set of) User/ID
uc	The center of locations of users
x_i	x-coordinate
y_i	y-coordinate

4.6.4.2 Evaluation Metrics

The purpose of the evaluation metrics is to differentiate the incentive mechanisms and to guide the design of the crowdsensing solutions. The simulations for incentive mechanism evaluations consist of an extensive number of trials. In each trial, we initialize the tasks and users at the beginning and the simulation runs until all tasks are completed or all users drop out. The details of the evaluation metrics are described as follows:

- The total operation time t_f represents the duration of a trial. In one trial, a timer starts from time 0 and ends at the time t_f when all sensing tasks are completed or all users drop out. While two incentive mechanisms may have an equal success rate s_r, one incentive mechanism might have less total operation time t_f. This implies that users have been incentivized to select and perform tasks efficiently.
- The platform cost in Equation 4.12 is the amount of money that the platform pays the users through sensing task rewards. The *surplus* is the portion of the budget that is not used by the end of a trial. A lower platform cost reflects the ability of incentive mechanisms to reduce the cost of sensing task rewards.

$$PC = B - surplus. \tag{4.12}$$

4.6.4.3 Platform-Driven Greedy Algorithm

The platform-driven greedy algorithm that we use to design our incentive mechanisms is shown in Algorithm 4. The idea of this algorithm is to select an available task that gives the maximum profit to the user. Thus, this platform-driven greedy algorithm computes the gain of task s_j to user u_i by Equation 4.13.

$$P_{ij} = R_{ij} - C_{ij} \tag{4.13}$$

in which R_{ij} is determined by the incentive mechanisms. We will describe more details of R_{ij} in the following subsection. After this algorithm finds out the task s_i which can provide the maximum profit for user u_i, the user u_i needs to check if the net profit margin of the task s_i is greater than the threshold T. If positive, the user u_i selects the task; otherwise, the user u_i drops out.

4.6.5 Incentive Mechanisms

We will cover nine unique incentive mechanisms, each with unique characteristics. The task-reverse-auction (*TRA*) incentive mechanism has been discussed in the literature [738, 762, 766]. It is known that the task-reverse-auction incentive mechanism cannot guarantee that all tasks are completed within a short total operation time in untruthful bid scenarios [766]. Our incentive mechanism design has a goal of reducing the total operation time. Thus, we will compare their total operation times in Section VI. The other eight incentive mechanisms are described as follows.

Algorithm 4: Platform-driven greedy algorithm

Input: u_i, S_a, T where $u_i =< s_j, a_i, C_{ij}, x_i, y_i >$
Output: Updated $u_i.s_j$

if $S_a == \emptyset$ **then**
  $u_i.s_j = -1$; // user u_i drops out as no task is available
  return;
end
$MP = -\infty$, $CR = -\infty$;
for s_j *in* S_a **do**
  $P_{ij} = R_{ij} - C_{ij}$;
  if $P_{ij} \geq MP$ **then**
    $MP = P_{ij}$;
    $CR = R_{ij}$;
    $s = s_j$;
  end
end
if $u_i.a_i == 0$ **then**
  $NPM = 100 \times \frac{MP+CR}{CR}$;
else
  $NPM = 100 \times \frac{MP+u_i.a_i}{u_i.a_i}$;
end
if $NPM < T$ **then**
  $u_i.s_j = -1$ // u_i drops out as no task gives ample profit;
  return;
end
$u_i.a_i = u_i.a_i + CR$;
$u_i.s_j = s$;
return;

Static Uniform (SU) Incentive Mechanism In the static uniform incentive mechanism [524], the incentives of sensing tasks are fixed values that are uniformly distributed and have the value R_{ij} calculated by Equation 4.14. In this case, R_{ij} is set to the base reward BR.

$$R_{ij} = \frac{B}{|S|} = BR \tag{4.14}$$

Dynamic Relative (DR) Incentive Mechanism The incentives change their values R_{ij} based on the distance from currently unavailable users and the user u_i to the sensing task s_j. This incentive mechanism ranks the currently unavailable users and user u_i by their distance to the sensing task s_j in an increasing order. Then, the value of incentive for the sensing task s_j can be calculated by Equation 4.15.

$$R_{ij} = \begin{cases} BR & k_i = 1 \\ BR(1 - \frac{1}{2}\frac{k_i}{|U|}) & k_i \geq 2 \end{cases} \tag{4.15}$$

Dynamic/Static User-Centric (DUC/SUC) Incentive Mechanisms First, the center of user locations is calculated by Equation 4.16. Then we compute the distance $d_{s,uc}$ from the task s to the user center using Equation 4.17. The value R_{ij} is inversely proportional to the distance as shown in Equation 4.18.

- Static case: rewards of sensing tasks are computed only once at the beginning of each trial.
- Dynamic case: like the static case, but the calculation repeats whenever a user is about to select a sensing task.

$$(x_{uc}, y_{uc}) = (\frac{\sum_{i \in U} x_i}{|U|}, \frac{\sum_{i \in U} y_i}{|U|}) \tag{4.16}$$

$$d_{s,uc} = |x_s - x_{uc}| + |y_s - y_{uc}| \tag{4.17}$$

$$R_{ij} = BR(1 - \frac{1}{2}\frac{d_{s,uc}}{b * 2}) \tag{4.18}$$

Dynamic/Static Task-Centric (DTC/STC) Incentive Mechanisms This mechanism first computes the center of the locations of sensing tasks, i.e. tc, by Equation 4.19. It then calculates the distance $d_{s,tc}$ from the sensing task s to the sensing task center by Equation 4.20. Finally it derives the value R_{ij} by Equation 4.21, which is inversely proportional to the distance.

- Static case: rewards of sensing tasks are computed only once at the beginning of each trial.
- Dynamic case: like the static case, but the calculation repeats whenever a user is about to select a sensing task.

$$(x_{tc}, y_{tc}) = (\frac{\sum_{s \in S} x_s}{|S|}, \frac{\sum_{s \in S} y_s}{|S|}) \tag{4.19}$$

$$d_{s,tc} = |x_s - x_{tc}| + |y_s - y_{tc}| \tag{4.20}$$

$$R_{ij} = BR(1 - \frac{1}{2}\frac{d_{s,tc}}{b * 2}) \tag{4.21}$$

Dynamic/Static Pit (DPIT/SPIT) Incentive Mechanisms In this pit-based incentive mechanism, we use all the users' coordinates to calculate an average distance to the sensing task s by Equation 4.22. Then, we compute the incentive R_{ij} of the sensing task s by Equation 4.23.

- Static case: rewards of sensing tasks are computed only once at the beginning of each trial.
- Dynamic case: we need to recalculate the incentives when a user is about to select a sensing task.

$$avg_s = \frac{\sum_i (|x_s - x_i| + |y_s - y_i|)}{|U|} \tag{4.22}$$

$$R_{ij} = \frac{BR}{2}(1 + \frac{avg_s}{b * 2}) \tag{4.23}$$

4.6.6 Incentive Mechanism Simulation

This section describes the parameters and the processes that have been used in the simulations for the performance study of incentive mechanisms.

4.6.6.1 Parameters

The parameter tuple for each trial is $< B, P, IM >$. After simulations, the evaluation metric tuple $< t_f, PC >$ will be averaged across the total number of simulation trials. In our simulation, the unit of time and money are time unit and fiat unit. Here is the description of the parameters of the experiments:

- Budget B represents the quantity of money that allows the platform to use in a trial. For this experiment, 100 data points were collected in the interval $B \in [100.00, 1090.00]$ with 10.00 spacing between each data point.
- Area coverage percentage P represents the percentage of the area that requires sensing data. As with the budget, 100 trials were conducted such that $P \in [20.0\,\%, 79.4\,\%]$ with 0.6 % spacing between each data point. This interval represents a wide range of possible target percentages for pavement crowdsensing. Note that we round down the area coverage percentage when calculating the number of tasks.
- The final parameter is the incentive mechanism IM used in the trial. The different IM calculate rewards of tasks differently.

4.6.6.2 Simulation Execution

Given < B, P, IM >, the construction phase initializes the numbers of cells, users, and sensing tasks in the following order:

- For each cell, any references to users or sensing tasks are cleared.
- For all users, s_j, a_i, C_{ij}, x_i, and y_i are initialized. s_j is set to 0. Each user would be placed in a cell randomly without overlapping.
- For all sensing tasks, u_i, R_{ij}, x_j, and y_j are initialized. u_i is set to 0. Each sensing task will be randomly placed in a cell with no overlap between other sensing tasks.

In the execution phase, available users start their turns by selecting and committing to a sensing task based on Algorithm 4. Then, the user will update its s_j. In turn, the user information associated with the sensing task s_j will be updated to reflect that the user u_i now performs task s_j. If no suitable sensing task is found, then the user drops out of the trial for all future turns. Unavailable users are the ones who have not dropped out and commit their turns by moving towards their sensing tasks. If the user lands on the sensing task, then a_i increases by R_{ij}. If the user is not on the sensing task, then the user must wait another turn to move closer. In both cases, C_{ij}, x_i, and y_i are updated to reflect the current user location.

4.6.7 Incentive Mechanism Evaluation Results

Incentive mechanisms are evaluated and compared in three scenarios corresponding to low, medium, and high area coverage percentages for pavement crowdsensing with different numbers of users. The platform cost is used to order the incentive mechanisms based on their performance data, as shown in the following figures. The minimal budgets shown in the figures are the lowest budgets that can realize a 100 % success rate for the targeted area coverage percentage. It means that any budgets higher than this value allow the platform to achieve a 100 % success rate for the targeted area coverage.

4.6.7.1 Platform Cost Comparison

In this section, we discuss the comparison of incentive mechanisms in terms of the platform cost.

- Given 25 % area coverage, Figure 4.38 shows that the SU and DR incentive mechanisms have the lowest platform costs when the platform has 3 users and 15 users, respectively. Apart from this, the static and dynamic pit incentive mechanisms consistently rank among the top three incentive mechanisms in all scenarios.
- Given 50 % area coverage, Figure 4.39 shows that static and dynamic pit incentive mechanisms still have the best performances of the platform cost in all scenarios. Even though the DTC incentive mechanism achieves the lowest platform cost

in *50 % area coverage with 45 users*, this observation does not conflict with the previous statement.

– Given 75 % area coverage, Figure 4.40 shows that SPIT and DPIT always have the lowest platform cost regardless of how many users the platform has.

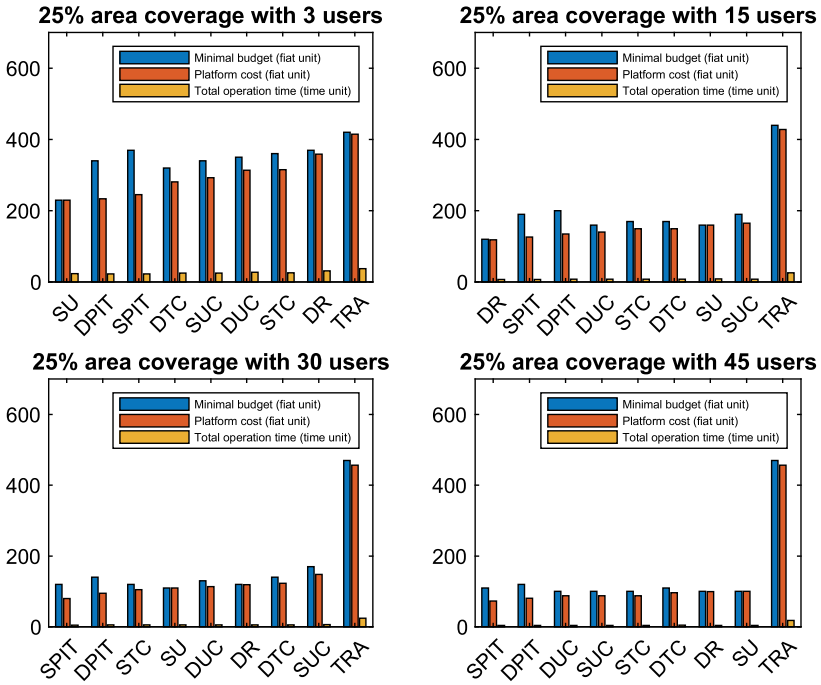

Fig. 4.38: Incentive mechanism comparison: 25 % area coverage.

Based on the observations described above, we can conclude that SPIT and DPIT are the two incentive mechanisms with the lowest platform cost.

4.6.7.2 Total Operation Time Comparison

In this subsection, we discuss the comparison of incentive mechanisms in terms of the total operation time. From Figs. 4.38, 4.39, and 4.40, the total operation time of the task-reverse-auction (TRA) incentive mechanism is nearly twice the total operation times of ours. Additionally, the total operation time of the TRA incentive mechanism becomes longer as the number of participatory users increases while the total operation times of our incentive mechanisms would decrease in the same situation. This result proves that our incentive mechanisms have much less total operation time than the Task-Reverse-Auction (TRA) incentive mechanism.

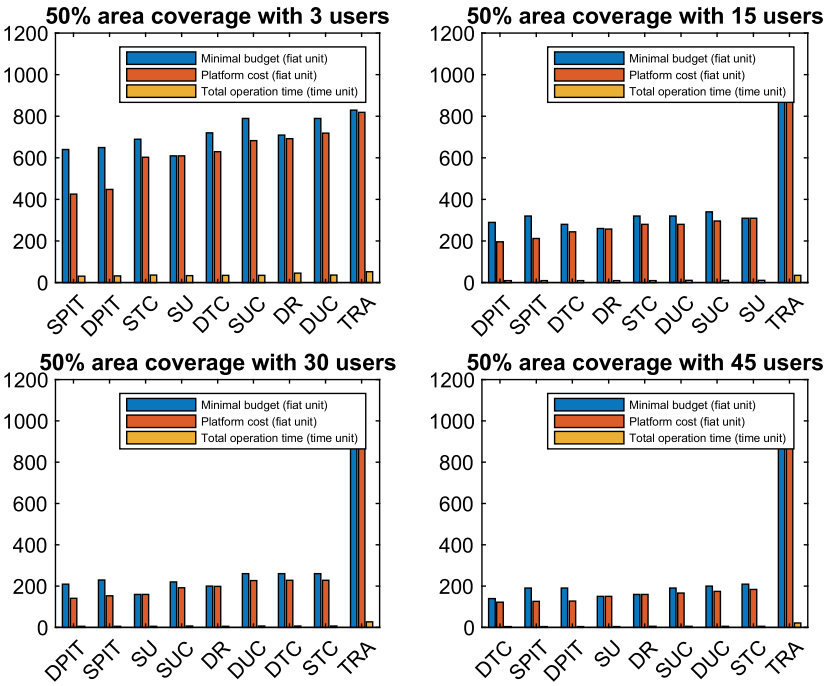

Fig. 4.39: Incentive mechanism comparison: 50 % area coverage.

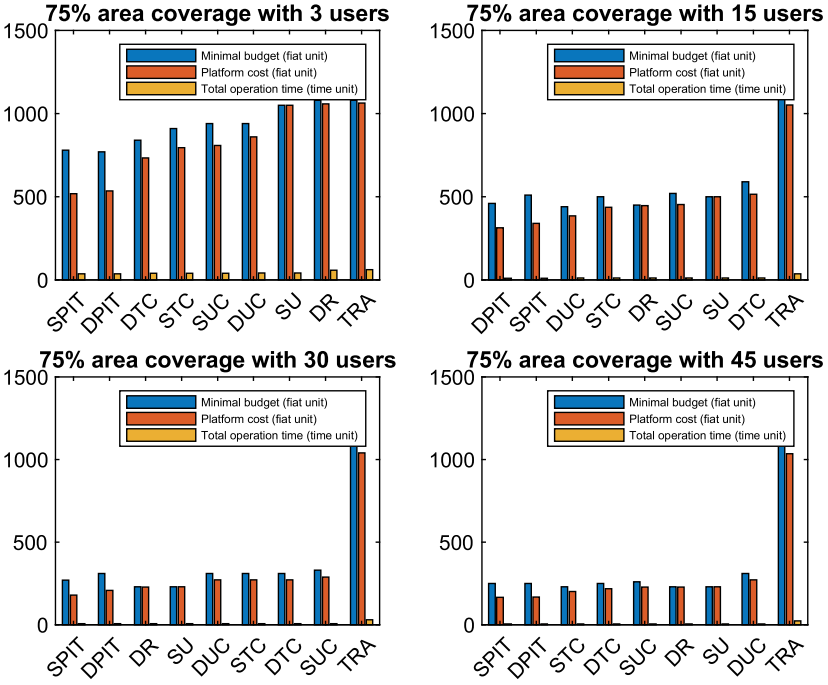

Fig. 4.40: Incentive mechanism comparison: 75 % area coverage.

4.6.8 Machine Learning Augmentation

Based on computer vision, crowdsensing and machine learning can be applied to pavement distress monitoring. For example, we have used a state-of-the-art machine learning model for detecting pavement damages based on images captured by the Android-phone camera and classifying them into eight types with corresponding confidence [738]. The types include (i) Liner crack, longitudinal, wheel mark part, (ii) Liner crack, longitudinal, construction joint part, (iii) Liner crack, lateral, equal interval, (iv) Liner crack, lateral, construction joint part, (v) Alligator crack, (vi) Rutting, bump, pothole, separation, (vii) White line blur, and (viii) Crosswalk blur. We have chosen this machine learning model because it "achieved recalls and precisions greater than 75 % with an inference time of 1.5s on a smartphone." [738]

Moreover, machine learning has been used for incentive mechanisms in embedded crowdsensing applications. For example, neural network and clustering algorithm have been applied for user grouping in [427] and the resulting incentive mechanisms can reduce the social cost, overpayment ratio, and grouping time. In the following sections, we discuss general considerations of machine learning augmentation, supervised learning, and unsupervised learning for incentive mechanisms.

4.6.8.1 General Considerations

We have witnessed great strides in the development of machine learning and its applications in recent years. Work in the spaces of image classification, text generation, language translation, and generative adversarial networks have produced results that could only be described as magic to the untrained eyes. Furthermore, being able to harness the full potential of these techniques, including subsequent derivatives, will be the aim of research for the foreseeable future. Such tantalizing thoughts act as motivation to incorporate different machine learning algorithms within the framework of crowdsensing, and our discussion is no exception. The crowdsensing scheme can be augmented using both supervised and unsupervised learning.

- Supervised learning focuses on mapping input data to target classes. Typically the input data will be sampled from a dataset where particular data points must be categorized. One common supervised learning task is image recognition. For example, one may use a convolutional neural network that receives input data in the form of road pictures and produces output data in a string distinguishing the road condition [738].
- Unsupervised learning focuses on clustering datasets such that embedded classes may be revealed. These class embeddings may reveal subsets of data and help highlight underlying relationships introspectively. One common unsupervised learning task is dimensionality reduction. In this case, we may consider a dataset with copious amounts of features. Using a clustering algorithm, such as spectral clustering, we may be able to reduce the principal features required to character-

ize data points uniquely. In other words, we may only need a proper subset of a dataset's features to represent the grouping of a data point. These transformations may lead to lossless decomposition of our data. Thus, we can reduce memory stress for supervised learning tasks.

4.6.8.2 Supervised Learning for Predicted Budget

One apparent shortcoming between all covered works in this contribution is the predicted budget. Having incentive mechanisms work under a limited budget is only reasonable if the budget has been methodically selected. In this case, a supervised learning problem is clearly established. Either a classification of the simulation parameters or a forecast of user costs can be used to determine a predicted budget. In the first case, the input data would include basic simulation parameters, such as the number of users and percentage area coverage, and the output would be the predicted budget. In the second case, the simulation would provide a seed state, including coordinates of users and sensing tasks, as input and calculate the final board statistics in terms of overall operation costs. The overall operation costs would be correlated with the predicted budget. Previous simulations would act as the data needed to construct these models in either case.

4.6.8.3 Unsupervised Learning for Incentive Mechanisms

Although the different incentive mechanisms studied in this contribution showed various levels of effectiveness in finishing the sensing tasks, a question of interest in crowdsensing applications is how their performances might differ if the incentives follow a non-uniform or random distribution. One may model the environment of the crowdsensing scheme as scattered normal distributions where sensing tasks may cluster in different neighborhoods. This scenario is realistic in rural living where communities may be sparse apart but dense around some centroid. The covered incentive mechanisms require a new component to scale incentives across clustered communities effectively. A clustering algorithm could locate the centroid of sensing task clusters and calculate rewards relative to these neighborhoods for the unsupervised learning task. In this case, the incentive mechanisms could scale to any size of environment given sufficient resources.

4.6.9 Conclusion and Future Work

In this contribution, we proposed eight incentive mechanisms based on a platform-driven greedy algorithm to help the crowdsensing platform motivate users to collect pavement condition data. Since our incentive mechanisms allow users to select the sensing tasks based on a platform-driven greedy algorithm before they start to collect the data, they can avoid the cost explosion problem observed in the data-reverse-

auction incentive mechanisms. From the simulation results, we find that SPIT and DPIT are the incentive mechanisms that have the lowest platform cost. Compared with the task-reverse-auction incentive mechanism, our incentive mechanisms reduce the total operation time by half. Our future research includes large-scale simulations and real-life experiments by extending our prototype pavement crowdsensing system. Lastly, we discussed machine learning augmentations for embedded crowdsensing applications and different incentive mechanisms.

4.6.10 Acknowledgment

The contents of this section reflect the views of the authors, who are responsible for the facts and the accuracy of the information presented herein. This document is disseminated in the interest of information exchange. The project is funded, partially or entirely, by a grant from the U.S. Department of Transportation's University Transportation Centers Program. However, the U.S. government assumes no liability for the contents or use thereof. This work was sponsored by a grant from the Center for Integrated Asset Management for Multimodal Transportation Infrastructure Systems (CIAMTIS), a U.S. Department of Transportation University Transportation Center, under federal grant number 69A3551847103, and the Commonwealth of Pennsylvania's PITA program. The authors are grateful for the support.

5 Communication Networks

5.1 Capacity Analysis of IoT Networks in the Unlicensed Spectrum

Stefan Böcker
Christian Arendt
Christian Wietfeld

Abstract: The ongoing digitalization and the steadily increasing number of distributed sensor devices and Internet-of-Things (IoT) systems implies a massive increase of subscribers. At the same time, the amount of available frequency spectrum resources remains static. In this respect, current 5G networks are already aiming for large-scale connectivity with an ambitious node density of 1 000 000 devices per square kilometer in the area of *massive Machine Type Communication* (mMTC). A huge number of potential technology solutions are available, but a comprehensive networking solution based on one technology seems unlikely. Among typical cellular IoT technologies, these challenging 5G mMTC requirements are also addressed by a growing number of unlicensed technologies enabling a simple, cost-effective network operation independent of licensed operators.

In this context, the potentials of Low-Power Wide Area Networks (LPWAN) technologies, as an additional technology option in unlicensed frequency bands are analyzed. Specifically, this work aims to analyze the suitability of LoRaWAN to contribute to given 5G requirements for specific mMTC applications in large-scale deployments. The performance evaluation illustrates that LoRaWAN is attractive due to high communication ranges up to multiple kilometers, enabling a high coverage even with a small number of cells. The evaluation also finds that the technology has a high potential to contribute to 5G mMTC application areas, especially for non-time-critical sensor applications.

To further increase the reliability of LPWAN systems, especially for critical services, different approaches to increase spectral efficiency are discussed. In addition to purely scheduling-based approaches, a data-driven analysis of the spectral power density to predict and avoid technology-independent interferences is presented. This is used to increase the robustness of LPWAN systems by centrally deriving communication profiles that address and bypass the predicted interference characteristics.

Apart from intelligently scheduling data transmission, another way for increasing efficiency is to reduce the amount of data that has to be transferred in the first place. Nowadays, initial generations of connected IoT devices and applications enabled by Cellular-IoT (CIoT) and LPWAN technologies are deliberately kept simple and based on equidistant, regular communication intervals. By contrast, we illustrate an Artificial

https://doi.org/10.1515/9783110785982-005

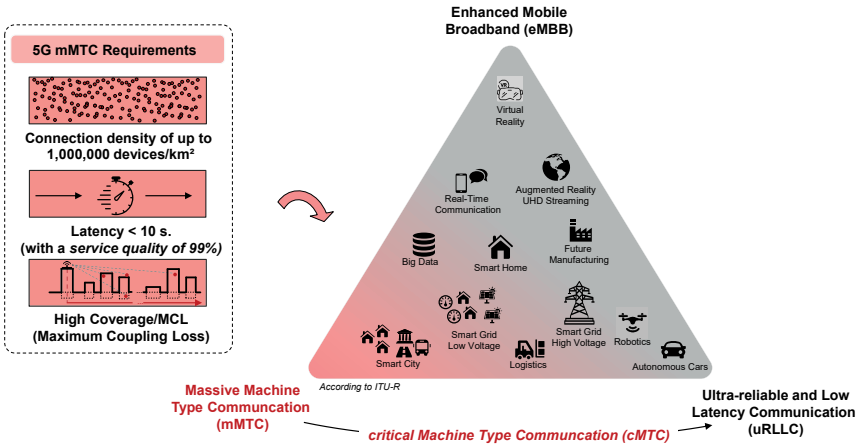

Fig. 5.1: Achieving the 5G scalability targets for IoT environments

Intelligence (AI)-based model-predictive communication approach taking advantage of knowledge about the underlying data of a sensor system. An AutoRegressive Integrated Moving Average (ARIMA) model is used in order to depict the behavior of an applications sensor data, leaving only values deviating from the model to be transmitted.

5.1.1 Introduction

Wireless connectivity has become a ubiquitous part of daily life. While wireless networks were originally developed to connect people, cellular networks have evolved to enable Machine-to-Machine (M2M) communication. In a wide variety of application areas, such as smart cities, energy systems, or production and logistics, devices are linked to each other to enable fully autonomous operation without human intervention. In this context, the 5G specification defines an mMTC requirement profile that aims for an ambitious scalability target of 1 million subscribers per square kilometer, while maintaining a maximum latency of 10 s (see Figure 5.1). At the same time this scalability target is linked to a boundary condition, that considers a pre-defined Poisson arrival process traffic pattern for non-full buffer systems with a payload of 32 B.

While Section 4.3 presents the performance evaluation of the NB-IoT technology as a current 3GPP solution to address 5G mMTC requirements, this contribution covers the research challenge to identify complementary technologies operated in unlicensed frequency bands to contribute to tight 5G mMTC requirement profiles. To this end, this work first discusses potentials and limitations of the Long-Range Wide-Area Network (LoRaWAN) technology, as a representative of LPWAN solutions, in order to subse-

quently introduce optimizations for further performance enhancements of unlicensed technology solutions with specific respect to mMTC applications.

5.1.2 Opportunities and Challenges of the Licensed and the Unlicensed Spectrum for IoT Environments

First, this section introduces the necessary LoRaWAN fundamentals from which the opportunities and challenges are derived. In addition, the impact of regulatory policies on the performance of the LoRaWAN technology is discussed.

LoRaWAN is an LPWAN specification for wireless battery-powered systems in a regional, national, or global range. It is based on the LoRa Modulation technique and mainly operated in the Short-Range Device (SRD) band at around 868 MHz in Europe and 915 MHz in the US. LoRaWAN enables a wide-range communication even in urban scenarios and provides a very good deep indoor penetration [203],[304]. The definition of the spreading factor (SF) permits the trade-off between efficient and very robust communication, whereby the data rates vary from 0.25 kbit/s (SF=12) to 5.5 kbit/s (SF=7).

Because LoRaWAN is operated in unlicensed frequency bands, the channel access must comply with regulatory frequency band requirements that ensure that all participants have equal access to frequency resources. For the underlying SRD band, the European Commission in cooperation with ETSI allows mitigation techniques such as Listen Before Talk (LBT), detect and avoid (DAA), and duty-cycle limitations [192], whereby LoRaWAN relies on a fairly simple pure ALOHA channel access and implements the duty cycle limitations to meet regulatory ETSI requirements. Thus, peak physical data rates are further reduced by a factor of more than 99 % due to the regulatory impact of a given duty cycle of 1 % and a MAC overhead with minor impact. The resulting average throughput ranges from 1.5 to 48 bit/s (as shown in Figure 5.2).

Consequently, throughput limitations are mostly driven by the idle time (time off) following the transmission time per packet (time-on-air), which is required to meet duty cycle limits [446]. The LoRaWAN specification defines three mandatory channels: 868.1 MHz, 868.3 MHz, and 868.5 MHz, additional resources are optional (see Figure 5.3). To reduce interference, channels are cycled in a pseudo-random approach.

As shown in Figure 5.3, the duty-cycle constraints apply to each SRD sub-band and may vary between different sub-bands, i.e. a dedicated LoRaWAN downlink communication is deployed at 869.525 MHz with a duty cycle of 10 % and additional limiting to data-rate class 0 (DR0).

5.1.3 Capacity Limits of LPWAN Technologies in Unlicensed Band Operation

The determination of capacity limits is based on a performance evaluation derived from an analytical model [446] that has been enhanced fundamentally for the underlying

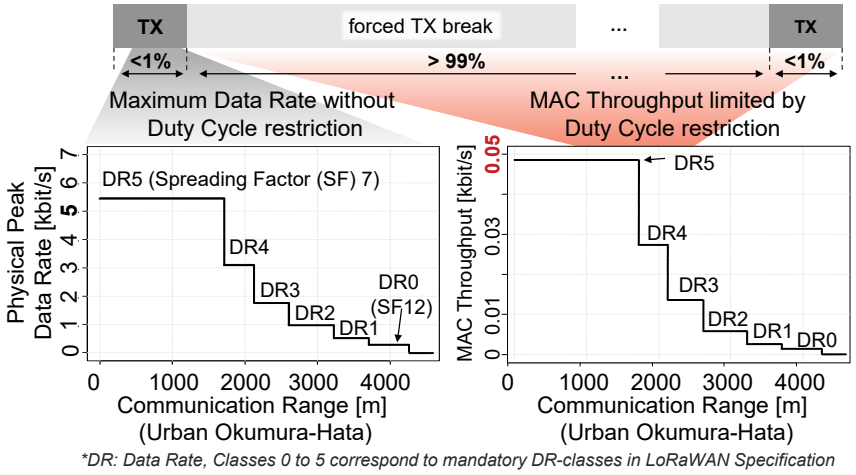

Fig. 5.2: Impact of regulatory limitations in unlicensed band deployments ©[2019] IEEE. Reprinted, with permission, from [93].

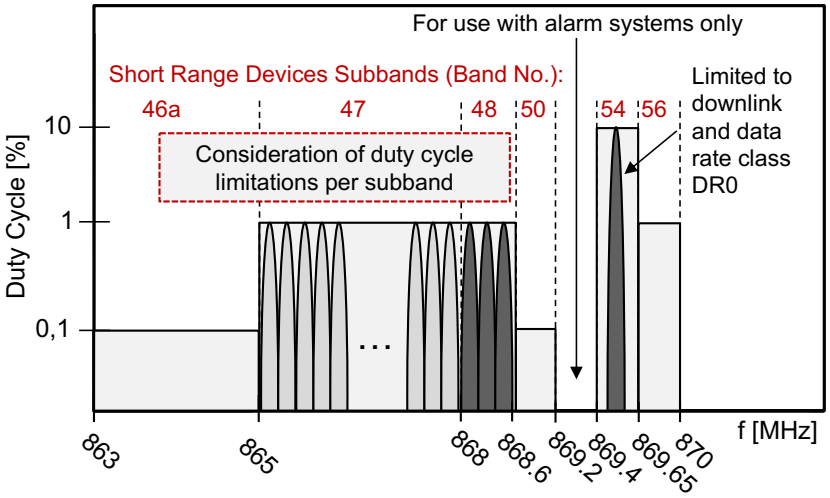

Fig. 5.3: European LoRaWAN channel frequencies in the ISM Band (EU863-870) ©[2019] IEEE. Reprinted, with permission, from [93].

work. The resulting analytical model enables demand-based derivation of key performance parameters, such as data rate and coverage area. The main extension aims at deriving latency bounds for different scalability scenarios and large-scale deployments, enabling the simultaneous determination of service guarantees.

Because downlink communication between a LoRaWAN gateway and distributed LoRaWAN nodes is not interfered with by uplink communication, performance limits

can be easily determined by evaluating a maximum capacity of a single LoRaWAN link. Considering the mandatory duty-cycle constraint (see Section 5.1.2), this results in overall limited downlink performance margins. Based on the assumption of Class A LoRaWAN nodes, Figure 5.4 depicts the average downlink throughput of a LoRaWAN network. As illustrated, the LoRaWAN downlink for Class A devices is based on two consecutive downlink receive windows. Following an uplink message, a Class A end device opens a first receive window (RX1) typically one second later. The first receive window is opened one second after termination and on the same frequency channel as the previous uplink message. The second receive window (RX2) is typically opened 2 s after uplink transmission and based on a dedicated downlink channel at 869.525 MHz with a mandatory duty cycle of 10 % and a limitation to data rate class DR0. This results in a low average downlink throughput of about 31.2 bit/s, which excludes a large number of downlink-intensive and safety-critical applications, such as update or upgrade functions.

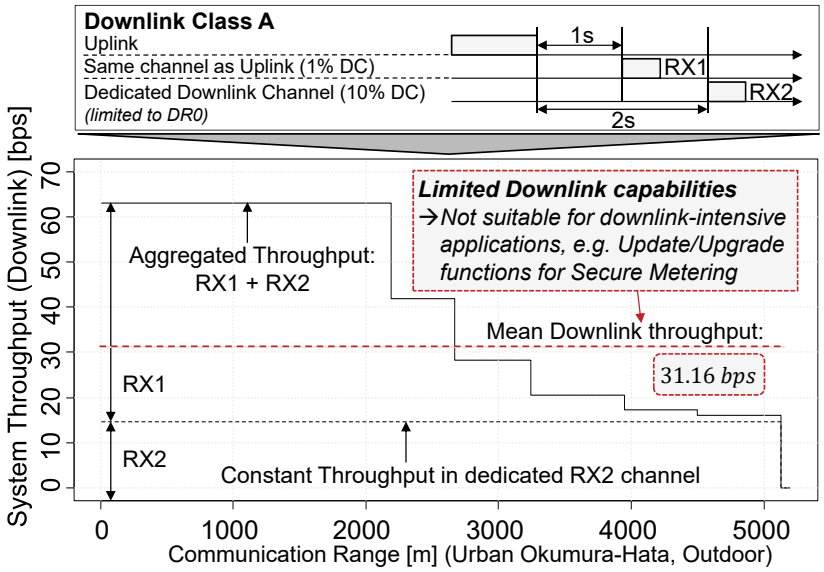

Fig. 5.4: System throughput (downlink) utilizing maximum duty-cycle capabilities ©[2019] IEEE. Reprinted, with permission, from [93].

In the uplink communication direction, the pure ALOHA channel-access scheme is implemented. In this context, the parallel uplink communications of different spreading factors are orthogonal to each other, permitting each data rate class to be modeled as an independent ALOHA process. The maximum system throughput of each data rate class is closely related to the number of devices and can be determined using the known ALOHA model equation $S = G \cdot e^{-2G}$, whereby S is the normalized channel throughput

and G the channel traffic [8]. The integration of the capture effect takes into account that two parallel transmissions with an Received Signal Strength Indicator (RSSI) delta greater than 6 dB for co-channel rejection [592] do not interfere with each other. In this case, only one packet is dropped, and the packet wirh higher received power is assumed to be decoded successfully, resulting in increased scalability. ALOHA equations can be adapted to $S = \frac{G \cdot e^{-2G}}{2} \cdot (1 + e^{G})$ [140]. In order to allow the derivation of guaranteed performance in addition to maximum uplink system throughput or scalability, the existing model is extended to incorporate the derivation of latency bounds, whereby latency is defined as the sum of Time-on-Air (ToA) and forced time off (T_{off}) by regulatory duty-cycle restrictions. Consequently, every interfered transmission increases the time on air $ToA+T_{off}$. Therefore, the mean latency $\overline{\tau_{DR}}$ can be defined based on the maximum packet collisions per transmission $\frac{G}{S}$ as depicted in the following equation:

$$\begin{aligned} \overline{\tau_{DR}} &= \frac{G}{S} \cdot (ToA + T_{off}) - T_{off} \\ &= (e^{2G} - 1) \cdot (ToA + T_{off}) + ToA \end{aligned} \tag{5.1}$$

Equation 5.1 is enhanced to consider the 99 %-quantile of defined 5G mMTC latency requirements $\tau_{99\%}$.

$$\tau_{99\%} = \log_{1-e^{-2G}}(0.01) \cdot (ToA + T_{off}) + ToA \tag{5.2}$$

The derived model can be modified and configured depending on desired parameter scenarios. Figure 5.5 illustrates the results for an exemplary configuration of a 32 B payload, 3 channels, and the maximum duty cycle. Without considering the capture effect, the maximum throughput of approximately 3.3 kbit/s can be achieved with a fleet size of about 900 subscribers. This can be further increased by about 50 % due to the additional consideration of the capture effect, resulting in a maximum throughput of about 4.7 kbit/s for a uniformly distributed number of subscribers of 1350. However, taking into account the limiting data-rate class DR0, this is simultaneously accompanied by an increased average latency of 400 s, which corresponds to an increase of about 12.5 %.

5.1.3.1 LoRaWAN Contribution to 5G

LoRaWAN technology is emerging as a very good solution in the unlicensed spectrum band to support 5G mMTC targets. Although LoRaWAN does not support the required 164 dB with a loss of 151 dB, it can cover the targeted area of one square kilometer, which is defined for the 5G mMTC connection density target. At any rate, it has a very good range and deep indoor availability even for urban areas [304]. Furthermore, Figure 5.6 illustrates the impact of various LoRaWAN parameter configurations on maximum scalability in view of 5G mMTC connection density and latency requirements (see Section 5.1.1). It can be shown that the 5G mMTC parameter set results in a significant contribution of 10 % for three 125 kHz frequency channels, which can be increased up

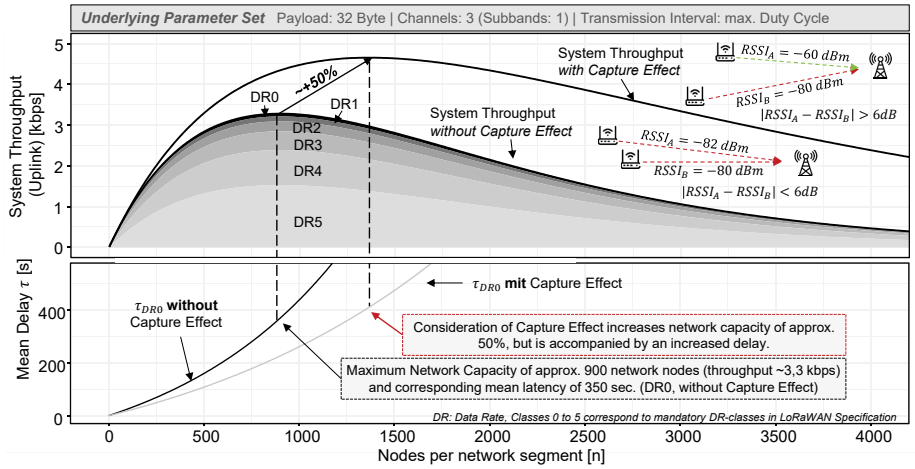

Fig. 5.5: System throughput (uplink) and latency utilizing maximum duty cycle capabilities ©[2019] IEEE. Reprinted, with permission, from [93].

to 25 % when the 8 channels are considered to cover 5G mMTC targets of one million devices per square kilometer. However, these results are obtained without considering the 5G mMTC latency requirement of 10 s. If the latency requirement is taken into account to ensure a certain quality of service, scalability is reduced by 25 % for a 50 %-quantile or up to 70 % under consideration of a 99 %-quantile.

When deviating from the 5G mMTC traffic pattern and considering other payload sizes, it can be seen that this factor has only a minor impact on scalability. By contrast, the variation of the transmission interval has a fundamental effect. In the case of a low transmission interval of only 12 hours, more than 50 % can be met of the overall connection density of one million devices per square kilometer without consideration of a latency requirement (99 %-quantile). By contrast, almost 20 % can be met when taking into account a latency requirement. So far, the results described have been focused on the assessment of unacknowledged traffic in the uplink direction. When ACK packets in the downlink direction are included, the downlink indicates a significantly constrained scalability of the LoRaWAN network. Even for a transmission interval of 12 hours, the scalability decreases to about 14 250 subscribers per square kilometer, which corresponds to a reduction of 97 % and is consistent with the limitations of the LoRaWAN downlink. Overall, depending on the application scenario and configuration, a very significant contribution of LoRaWAN technology to the 5G mMTC goals can be deduced, though the application field should be limited to non-time-critical sensor applications.

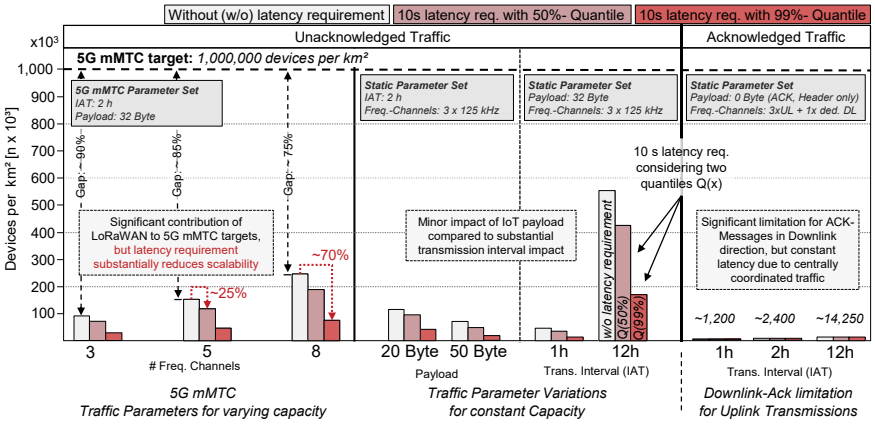

Fig. 5.6: Impact of various LoRaWAN parameter configurations on maximum scalability considering 5G mMTC connection density and latency requirements ©[2019] IEEE. Reprinted, with permission, from [93].

5.1.4 Data-Driven Capacity Improvements

In this section, optimization methods are presented that further enhance the performance of the previously discussed LPWAN technologies in unlicensed frequency bands.

5.1.4.1 Dynamic Spectrum Management to Improve Scalability of Time-Critical Sensor Applications

Due to their simple and cost-effective technical viability, a steadily increasing number of LPWAN are operated in unlicensed frequency bands. For each user, this leads to a large number of possible interference effects caused by a wide variety of technologies, each using different channel access methods. There is no central coordination as in licensed mobile radio frequencies. Despite mandatory interference mitigation techniques, quality of service in terms of availability, latency, etc. cannot be guaranteed due to uncoordinated channel access in unlicensed frequency ranges. In order to tackle these challenges, a data-driven spectrum management procedure is proposed (see Figure 5.7). This approach relies on SDR -based spectrum sensing to gather information on channel occupation, which is used to predict future spectral utilization.

The predicted activity profiles are used for spectrum management in order to intelligently schedule future transmissions. Three scheduling approaches, namely Restricted Access Window, Weighted Restricted Access Window, and Coordinated Restricted Access Window, have been developed and evaluated. An overview of these approaches is shown in Figure 5.8, a brief description of each approach is given below. For evaluation, an externally defined Key Performance Indicator (KPI) is required, which is represented by the expected latency in this work. As LoRaWAN is the technology of choice, an

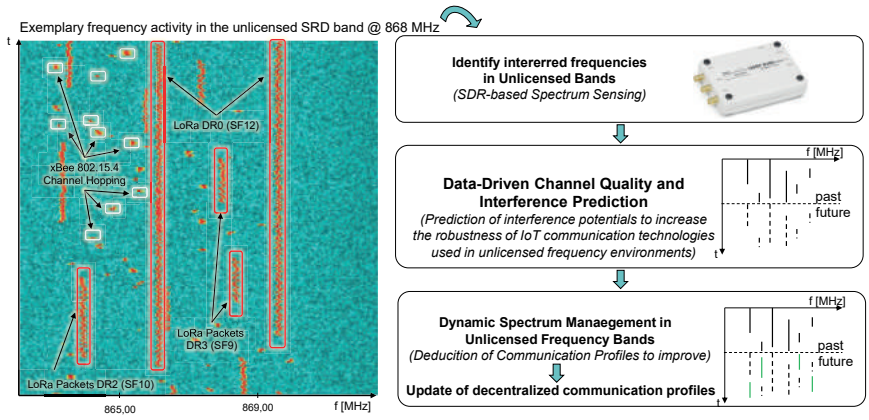

Fig. 5.7: Analyzing and reducing the impact of interference in unlicensed bands.

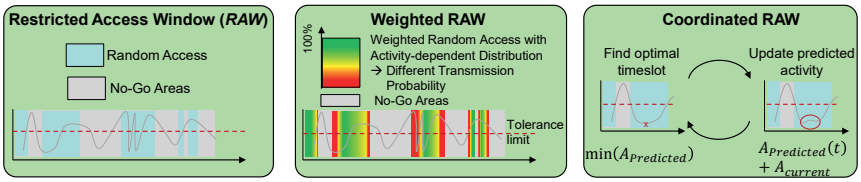

Fig. 5.8: Concept of the dynamic spectrum-management approach.

ALOHA channel access is used as the basis and a duty cycle has to be observed. The expected latency is modelled as described in Section 5.1.3

The three optimization methods Restricted Access Window, Weighted Restricted Access Window, and Coordinated Restricted Access Window are described below.

Restricted Access Window (RAW) The basic idea and the name are derived from IEEE 802.11ah technology. The time ranges are first divided into no-go areas and random access areas. In this case, the channel is accessed randomly in areas below the drawn-in threshold and access is avoided in the no-go areas. Restricted Access Window is the simplest developed method. This method is a coarse-grained method, which means that the gateway does not determine the transmission time for each subscriber; rather, the gateway only transmits the appropriate transmission time through beacons and the subscribers randomly select a transmission time.

Weighted Restricted Access Window In this method, the Restricted Access Window method is extended by a weighted access probability. This means the time ranges with an activity level below the tolerance limit have a higher chance to be selected.

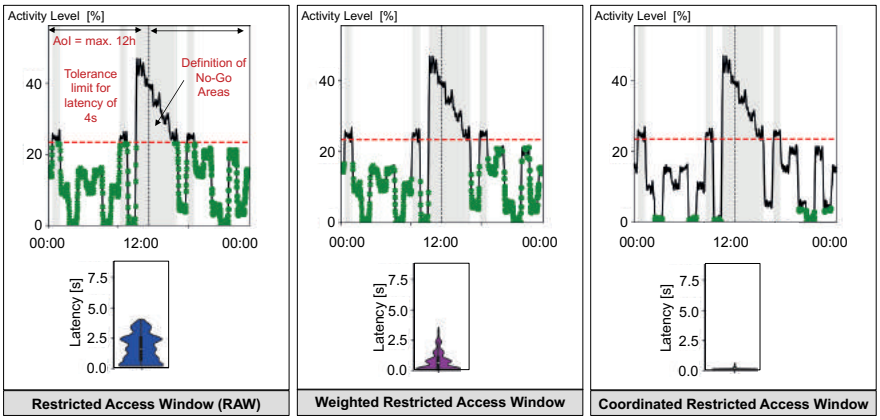

Fig. 5.9: Comparison of the three scheduling approaches with regard to potential latency reduction.

In this work, a quadratic functions is used as weighting and the weighting is calculated for each time point using Equation 5.3.

$$W_t = (A_{Tolerance} - A_t)^2 \tag{5.3}$$

This method is a coarse-grained method and the gateway transmits matching transmission time points with their weights by beacons. The participants finally choose a transmission time point by a weighted random.

Coordinated Restricted Access Window (CRAW) This procedure finds the optimal transmission time for each participant by assigning the time with the minimum activity to the first participant. Then the activity profiles are updated with the resulting activity by first participants. This procedure is repeated for the other participants.
The cyclic steps of the procedure are as follows:

1. finding the minimum activity in profile and assigning this time to the participant;
2. updating the profile in view of the activity change at the assigned time point; and
3. repeating the procedures for the next participant.

Compared with the previous methods, CRAW is a fine-grained scheduling method. This means that the gateway must communicate the transmission time to each subscriber. Therefore, this method requires many resources. In CRAW, the influence of each participant on the activity is taken into account.

Comparison of Scheduling Performance In this section, the three developed methods are evaluated using the presented scenarios. The predicted daily profiles from the 868 MHz study are used as a basis. In Figure 5.9, the results are shown and the used daily profiles are represented at the top.

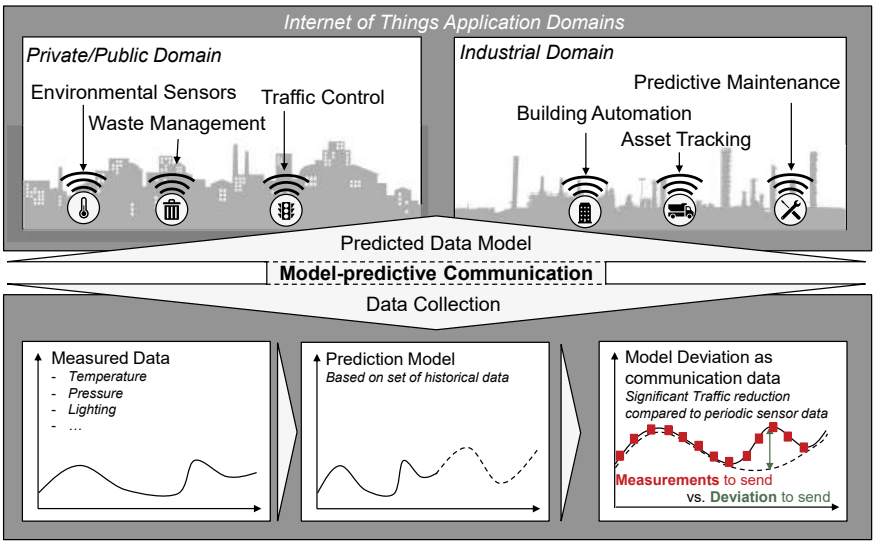

Fig. 5.10: Model-predictive communication in Internet of Things environments ©[2020] IEEE. Reprinted, with permission, from [30].

It can be seen that the expected latency can be significantly reduced by using more intelligent scheduling mechanisms, which avoids times with high channel activity. Simply avoiding regions with a mean latency over a given threshold with the RAW approach sets the baseline potential, which can be optimized using a weighting to reduce the mean latency below 1 s in this study. Providing more intelligence for the scheduling method with CRAW pushes all observed latencies below 1 s.

However, this approach generates an increased computational effort, which has to be taken into account as a trade off.

5.1.4.2 Data-Driven Model-Predictive Communication

In this section, we propose a data-driven approach to reduce communication efforts by leveraging knowledge about the underlying sensor data in IoT systems. In order to keep devices simple, the data transmission of IoT devices typically follows a regular pattern of equidistant time intervals. This leaves a high potential for optimizing the efficiency of spectral resource usage. Therefore, this study proposes a model-predictive communication framework that leverages knowledge about the underlying sensor data and allows IoT devices to rate the value of observations in order to decrease communication effort and free up spectral resources for other parties. This approach potentially increases the number of devices considered in the scheduling of resources in licensed frequency bands and reduces the likelihood of interference in unlicensed bands. Figure 5.10 depicts the concept of this work.

The underlying approach is not to transmit every measurement, only those that deviate significantly from a predetermined model. Therefore, two time series forecasting mechanisms are examined in this study to generate such a model of a temperature sensor setup. Because these methods simply rely on past data, the use case may be readily changed. This approach has been evaluated by using a dataset originating from an environmental indoor sensor located in a residential area in Dortmund, Germany. From the 1st of January 2019 until the 19th of November 2019, the system gathered temperature, humidity, and CO_2 concentration at a frequency of around 5 minutes. To minimize model complexity, the dataset was resampled to 30-minute time steps in this work. The raw data can be accessed via [29]. The forecasting algorithms used in this work are an autoregression based approach and a neural network approach, which are described in the following sections. Both models take advantage of a decomposition approach, which is discussed more below.

Seasonal and Trend Decomposition Using LOESS (STL) In this work, the Seasonal and Trend decomposition using LOESS (STL) method [540] is used to extract typical properties of the underlying data, such as a daily profile for temperature data. The algorithm consists of two loops: the inner loop uses LOcally wEighted Scatterplot Smoothing (LOESS) [141] to extract the seasonal and trend components and the outer loop is used to minimize the impact of outliers by computing robustness coefficients. An example decomposition of a daily temperature profile from the dataset is shown in Figure 5.11.

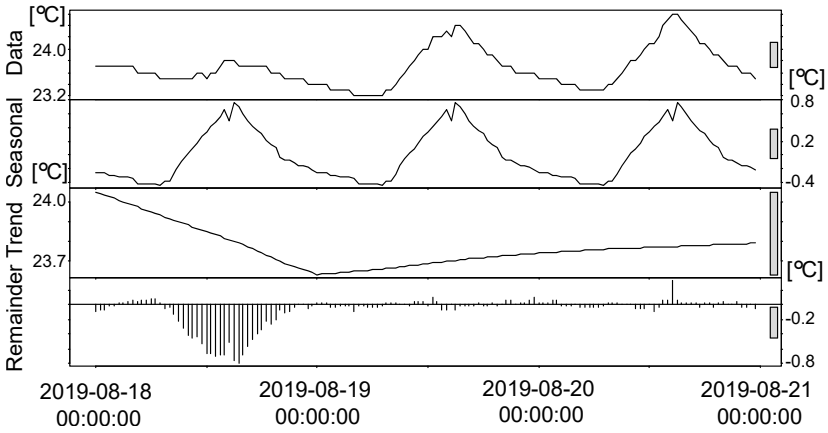

Fig. 5.11: Decomposition of measured temperature data from 18th of August until 21st of August using STL ©[2020] IEEE. Reprinted, with permission, from [30].

The example shows strong daily seasonality and a local weather trend. Additionally, a significant remainder is present that could not be related to either the trend or seasonality component. The decomposition can be used to aid forecast algorithms by subtracting the seasonal component before applying the prediction, and adding it back to the resulting time series afterwards. As the seasonal component typically changes slowly over time, it is possible to reduce the complexity of the analyzed time series in order to decrease prediction errors.

AutoRegressive Integrated Moving Average (ARIMA) The ARIMA algorithm is a state-of-the-art time series forecasting method and one of the most widely used. It is composed of three components: the autoregressive part $AR(p)$, entailing the past values of the original series; the integrated part $I(d)$, related to differencing in order to make the time series stationary; and the moving average component $MA(q)$ marking the model errors. The parameter set (p,d,q) defines the order of each component and therefore indicates the specific ARIMA model in use. In detail, p indicates the number of considered past values, d is the differencing degree and q specified the considered previous error terms. To simplify the application in this work, the forecast package for the statistical programming language R is used. The language contains an implementation of the ARIMA algorithm, which allows automatic parameter set selection for every model realization [294], where a unit test procedure checking for stationarity is used to specify d, while p and q are found by minimizing Akaike's Information Criterion (AIC) . A prediction for the 21st of August 2019 using ARIMA(0,1,0) based on a training period of three previous days together with the 95 % and 80 % prediction intervals as well as the actual measurement as test data is depicted in Figure 5.12

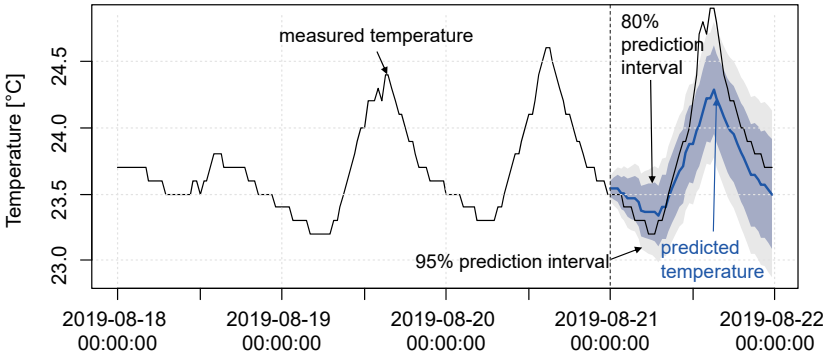

Fig. 5.12: Forecast for one sensor using STL + ARIMA(0,1,0) for the 21st of August 2019 with 95 % and 80 % prediction intervals ©[2020] IEEE. Reprinted, with permission, from [30].

It is evident that a good match is achieved between the predicted and the measured temperature, with most of the measured temperature values lying inside the 80 % prediction interval.

Long Short-Term Memory (LSTM) As a variant of Recurrent Neural Networks (RNN) developed by Hochreiter et al. [277], LSTM networks have an improved ability to learn long-term dependencies in a dataset, which makes them appropriate for time series prediction tasks. The main structure of LSTM networks consists of concatenated cells that are linked together by constant cell states and the input flow. An input gate is used to regulate the influence of the cells input, while a forget gate filters previous cell states.

The implementation used in this work is based on the well-known keras python framework with the Theano library as a backend. In order to keep the training duration managable and avoid decreasing the model accuracy by deeper networks as stated in [358], a single network layer is used. 50 LSTM cells were used within this layer, however the influence of the number of cells was seen to be minor.

Figure 5.13 compares LSTM and ARIMA performance for an exemplary prediction based on the data of the 21st of August 2019 with the three previous days as an input in terms of required communication events with different tolerance ranges of 0.5 °C, 1 °C, and 2 °C, respectively. The tolerance ranges depict the need for communication events, as only deviations higher than the tolerance should be transmitted.

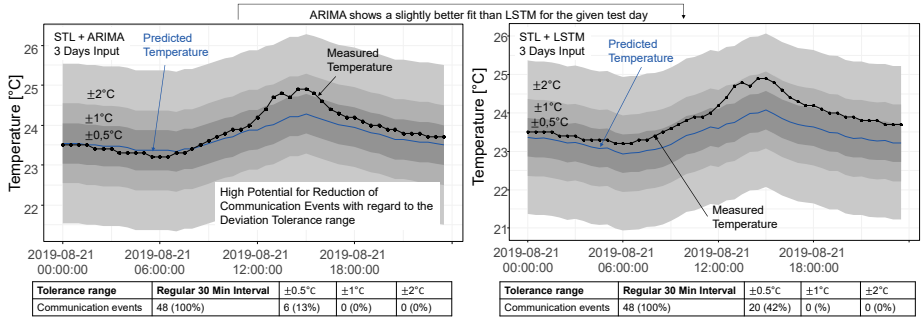

Tolerance range	Regular 30 Min Interval	±0.5°C	±1°C	±2°C
Communication events	48 (100%)	6 (13%)	0 (0%)	0 (0%)

Tolerance range	Regular 30 Min Interval	±0.5°C	±1°C	±2°C
Communication events	48 (100%)	20 (42%)	0 (%)	0 (0%)

Fig. 5.13: Forecast for one sensor using STL + ARIMA(0,1,0) and LSTM for the 21nd of August 2019 with tolerance ranges of ± 0.5 °C,± 1 °C and ± 2 °C. ARIMA shows a slightly higher potential in decreasing communication effort for ± 0.5 °C tolerance. ©[2020] IEEE. Reprinted, with permission, from [30].

ARIMA can achieve a potential reduction of communication events by 87 % for this day with a tolerance of 0.5 °C, where LSTM provides a lower potential of 58 %. However, for applications where a higher tolerance of 1 °C is sufficient, both algorithms provide enough accuracy to save up to 100 % of the communication effort.

Statistical Performance Evaluation Both implemented approaches are statistically evaluated using the Root Mean Square Error (RMSE) as a metric for the modeling errors and for the potential reduction of communication events. To make both modeling concepts comparable, a sliding window approach as depicted in Figure 5.14 is used for validation. Each sliding set from the total dataset has a configurable number of input

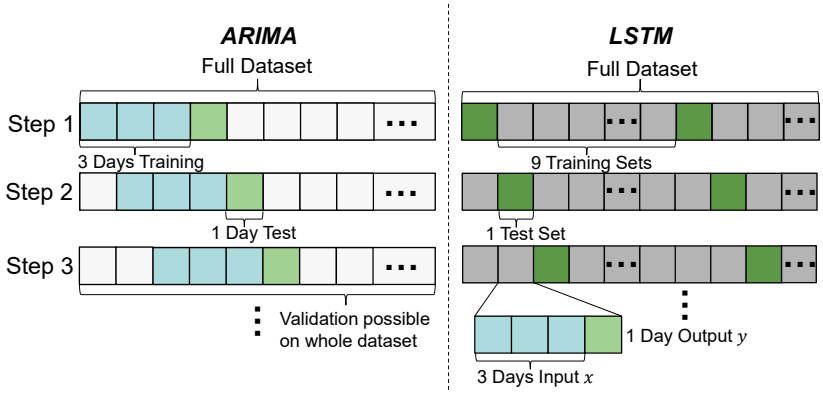

Fig. 5.14: Walk forward validation with sliding window (constant training period length) in each algorithm's variant. ARIMA takes t days of training data and produces a one day forecast that is compared to the test day, LSTM needs input/output pairs for training and is then tested on a held back test set. ©[2020] IEEE. Reprinted, with permission, from [30].

days followed by one output day. An advantage of ARIMA is the ability to sufficiently predict future data relying solely on a small data basis. LSTM, on the other hand, needs a large portion of input/output pairs as training data in order to learn essential features enabling the sensor data prediction. The impact of varying numbers of input days from 3 to 12 on the prediction of one forecast day has been addressed in this work. A 10-fold cross-validation is used to further validate the forecast results, with a 90-10 % split between training and test data. The resulting modeling error of both approaches for varying number of input days represented by the RMSE distribution is illustrated in Figure 5.15.

It can be seen that the impact of the input period length is small, but both models have a slightly lower error with smaller input periods. In general, LSTM provides a doubled mean error of around 0.3 °C when compared with ARIMA, producing a mean error of around 0.15 °C for all input period lengths. LSTM also shows a higher spread of the observed errors, except from a small amount of outliers experienced with ARIMA for longer input periods. The latter lead to the conclusion that data lying further in the past is less relevant for predicting the temperature values of the following day, increasing the amount of false forecasts observed. These results allow an estimation of the tolerance ranges in which the evaluated models can be used. A tolerance range of ±

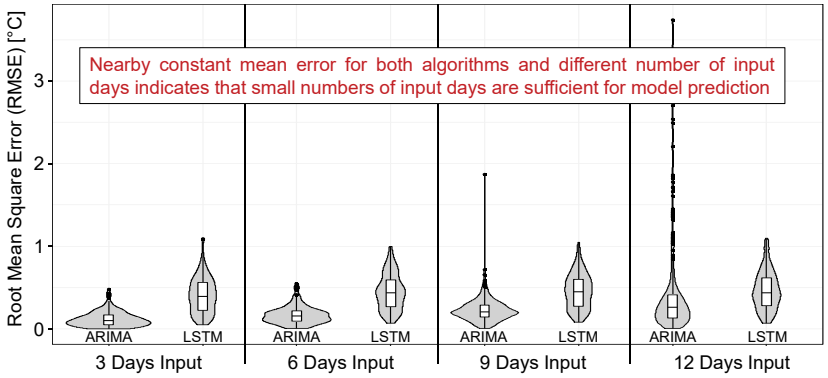

Fig. 5.15: RMSE for both approaches and varying training period. ARIMA shows a smaller RMSE than LSTM, however the mean error for both approaches remains nearly constant. ©[2020] IEEE. Reprinted, with permission, from [30].

0.5 °C was therefore chosen as the minimum tolerance for evaluating the potential of the model-predictive approach. Two supplementary tolerance ranges of ± 1 °C and ± 2 °C were evaluated to show the dependence of performance on the chosen tolerance for different applications. This analysis is carried out for the underlying temperature sensor system and depicted in Figure 5.16.

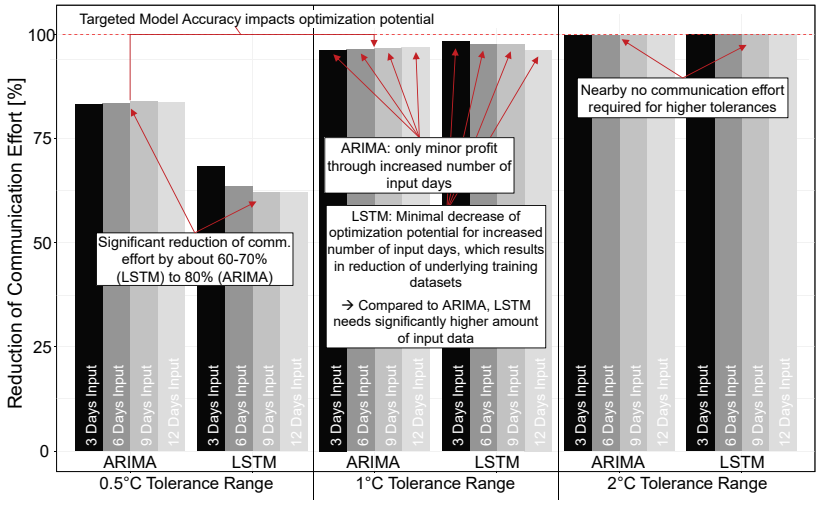

Fig. 5.16: Potential reduction in communication effort provided by using ARIMA and LSTM depending on varying input periods and tolerance ranges. ©[2020] IEEE. Reprinted, with permission, from [30].

A potential reduction of communication effort by more than 60 % with a tolerance of ± 0.5 °C can be observed with both models, with ARIMA reaching more than 80 % potential reduction. LSTM performance decreases with a larger input period, which can be explained by a reduced number of input/output pairs available for training. As expected from the error distribution in Figure 5.15, both models perform almost equally well for higher tolerance ranges with more than a 90 % reduction for ± 1 °C tolerance to nearly no communication effort for ± 2 °C tolerance.

Finally, both modeling approaches have the potential to cut sensor systems' communication effort significantly. Due to its superior overall results and a much higher efficiency in terms of input data needs and computational effort, the ARIMA algorithm is the favored method.

5.1.4.3 Outlook and Future Work

Even requirement profiles of current 5G mMTC applications are in a state of constant evolution. While classic network dimensioning is largely based on stochastic behavior and correlated traffic volumes, the share of event-driven machine communication will increase dramatically. The key challenge will be the realization of reliable critical alarm communication in the face of unpredictable behavior. In this context, a requirement migration towards mission-critical 6G MTC networks is inevitable [429]. To ensure that a guaranteed quality of service can be achieved for application classes that can barely be predicted with a reasonable amount of resources, the available resources must be allocated dynamically and as a function of defined costs (frequency utilization, delay times, energy consumption, …). For this purpose, the definition of new service classes is needed defining targeting latency, Block Error Rates (BLER), and its service characterization, and highlighting the need for future 6G systems to leverage application-domain information about the predictability of resource requirements and conditions. The new service classes are shown in Figure 5.17.

5.1.5 Acknowledgments

In addition to the CRC 876, part of the work has been supported by the federal state of North Rhine-Westphalia and the "European Regional Development Fund" (EFRE) 2014 - 2020 in the course of the "OPUS" project under grant number EFRE0800885, as well as the German Federal Ministry of Education and Research (BMBF) in the project A-DRZ (13N14857).

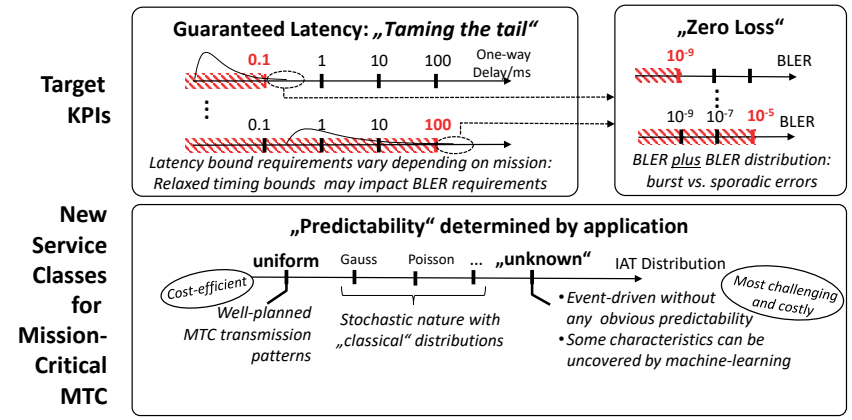

Fig. 5.17: Requirement migration towards mission-critical 6G MTC networks [429]. Used under CC BY 4.0 (https://creativecommons.org/licenses/by/4.0/).

5.2 Resource-Efficient Vehicle-to-Cloud Communications

Benjamin Sliwa

Abstract: Big vehicular data is anticipated to become the new fuel for catalyzing the further development of connected and autonomous driving. Vehicles themselves will act as mobile sensor nodes that actively sense their environment and gather meaningful data for novel crowdsensing-enabled services such as the distributed generation of high-definition maps, traffic monitoring, and predictive maintenance. However, the implied tremendous increase in massive Machine-Type Communication (mMTC) represents an enormous challenge for the coexistence of different resource-consuming applications and entities within the limited radio spectrum. A promising approach for achieving relief through a more resource-efficient usage of existing network resources is the utilization of client-based intelligence. Novel communications paradigms such as anticipatory mobile networking aim to improve decision processes within wireless communication systems by explicitly taking context information into account. In the context of vehicular crowdsensing, these methods exploit the delay-tolerant nature of the targeted applications for scheduling the data transfer with respect to the expected resource efficiency. If the current radio channel and network load conditions do not allow a resource-efficient transmission, the data transfer process is postponed and the acquired data is aggregated locally in favor of a better transmission opportunity in the near future along the expected vehicular trajectory.

In the following, the different evolution phases of the novel Channel-aware Transmission (CAT) scheme are presented. These are characterized by a sequential introduction of different machine learning methods. While the basis CAT approach applies a probabilistic channel-access mechanism based on measurements of the Signal-to-Noise-plus-Interference Ratio (SINR), Machine Learning CAT (ML-CAT) applies supervised learning for predicting the currently achievable data rate using features from the network context, the mobility context, and the application context domain.

This approach is then further extended by Reinforcement Learning CAT (RL-CAT) through the autonomous detection and exploitation of favorable transmission opportunities. Finally, Blackspot-Aware Contextual Bandit (BC-CB) integrates a priori knowledge about the geospatially-dependent uncertainties of the prediction model, which is uncovered by unsupervised machine learning.

It is shown that machine learning-aided opportunistic data transfer is not only able to increase the average data rate of the individual transmissions; it also achieves a massive reduction of the occupied network resources and the power consumption of the mobile device. The price to pay is an increase of the Age of Information (AoI) of the sensor measurements. In addition to the presentation of the novel opportunistic

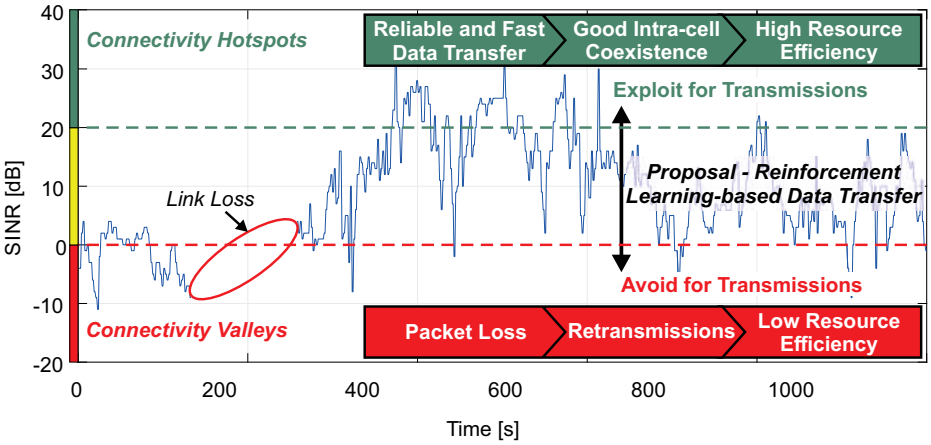

Fig. 5.18: Example for the temporal dynamics of the SINR in vehicular environments.

data-transfer approaches, new machine learning enabled methods for simulating these anticipatory mobile networks are presented, discussed, and validated.

5.2.1 Introduction

According to a recent white paper [6] published by the 5G Automotive Association (5GAA), predictive Quality of Service (QoS)—e.g. the prediction of the data rate along a vehicular trajectory—is expected to become a key enabling method for future connected and autonomous driving.

Although machine learning has already started to penetrate all areas of wireless communications [714], the current 5G standardization efforts focus on implementing intelligence on the network infrastructure side [1]. However, as discussed in initial visionary works [17], it is anticipated that not only the trend of replacing mathematical models with machine learning-based equivalents will continue, but also that pervasive intelligence will be a key driver for the further cellular network evolution. These developments are closely related to the arising anticipatory mobile networking paradigm [107, 630], which aims to improve decision processes within wireless communication systems through explicit consideration of context knowledge and machine learning-based data analysis.

Figure 5.18 shows a real world trace of the Signal-to-Interference-plus-Noise Ratio (SINR) acquired along a vehicular trajectory. It can be seen that vehicular communication channels are characterized by short-term and large-term fluctuations. This behavior is the result of a superposition of distance variations between sender and receiver, mobility-related factors, and obstacle-related signal variation due to shadowing, reflection, and refraction.

In order to guarantee a reliable data transfer, the mobile device reduces the achievable transmission efficiency in favor of better data integrity during challenging radio channel periods. As a result of the implied overhead, a large amount of network resources is unavailable for transmitting the actual payload data.

5.2.2 Related Work

Opportunistic data transfer implements the idea of postponing the data transfer to of delay-tolerant applications to situations where a higher resource efficiency can be achieved due to better radio channel conditions. Acquired data is stored in a local buffer until a favorable transmission opportunity is detected and the whole data buffer is transferred.

Channel-Aware Transmission (CAT) [296, 298], which represents the foundation for the further machine learning-based enhancements presented in this contribution, utilizes Signal-to-Interference-plus-Noise Ratio (SINR) measurements for client-based opportunistic data transmission based on the known significance of downlink quality indicators for assessing the uplink radio channel quality [297]. The probabilistic medium access is performed as

$$p(t) = \begin{cases} 0, & \Delta t \leq \Delta t_{\min} \\ \left(\frac{\text{SINR}}{\text{SINR}_{\max}}\right)^{\alpha}, & \Delta t_{\min} < \Delta t < \Delta t_{\max} \\ 1, & \Delta t \geq \Delta t_{\max} \end{cases} \tag{5.4}$$

whereas Δt is the elapsed time since the last transmission has been performed, $\Delta t_{\min}$ is a minimum inter-packet gap in order to avoid overly frequent medium access, and $\Delta t_{\max}$ is an application-specific deadline for the Age of Information (AoI) of the sensor data packets. Through configuration of α, it can be defined how much the transmission scheme prefers very high metric values within the transmission process.

5.2.3 Machine Learning-Enabled Opportunistic Vehicle-to-Cloud Communication

Although CAT has been demonstrated to achieve significant benefits in comparison with conventional data transmission approaches, recent analyses [638] have shown that physical layer indicators such as SINR have only a limited significance for estimating the achievable data rate. Since the latter is inversely proportional to the transmission duration, it is directly related to the resource occupation time. As a result, the maximization of the end-to-end data rate contributes to improving the intra-cell resource efficiency. For the exploitation of this property, the novel data transfer schemes build upon predictions of the achievable end-to-end data rate.

The methodological evolution of context-aware data transmission approaches is summarized in Figure 5.19. The different evolution stages are characterized through

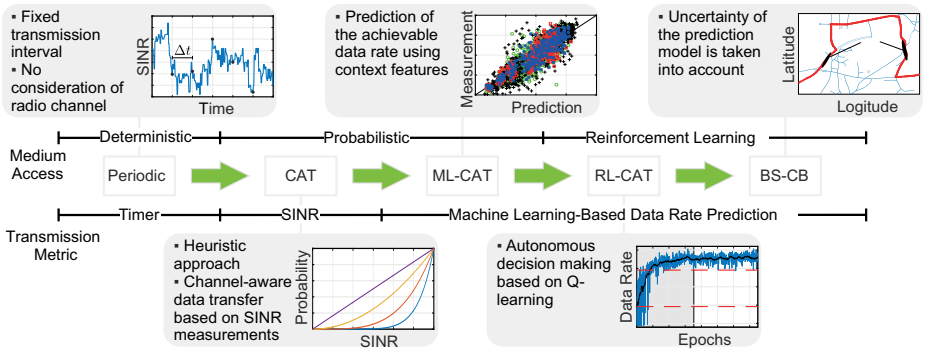

Fig. 5.19: Evolution of channel-sensitive solution approaches for resource-efficient vehicle-to-cloud communication.

a sequential introduction of solution approaches from different machine learning disciplines:

Channel-Aware Transmission (CAT) [296, 298] uses a probabilistic medium-access approach that takes SINR into account.

Machine Learning CAT (ML-CAT) [628, 629, 632] utilizes features from the network, mobility, and application domains for predicting the currently achievable end-to-end data rate, which is then used as the radio channel assessment metric.

Reinforcement Learning CAT (RL-CAT) [636] replaces the heuristic medium access approach with a Q-learning mechanism to autonomously detect and exploit favorable transmission opportunities.

Black Spot-Aware Contextual Bandit (BC-CB) [622, 625] incorporates a priori knowledge about the geospatially dependent uncertainties of the predictions model as a measurement of trust into the latter.

In the following, the enabling methods and novel data-transmission schemes are introduced. Additional details and analyses of various parameter variants are discussed in more detail in the referenced scientific publications.

5.2.3.1 End-to-End Data-Rate Prediction in Vehicular Networks

The considered dataset contains context traces in multiple vehicular evaluation scenarios (campus, urban, suburban, highway). Using the native `Android` Application Programming Interface (API), context indicators from different logical domains are acquired:

- **Network context features $\mathbf{x}_{\text{net}}$**: RSRP, RSRQ, SINR, CQI, TA, Carrier frequency
- **Mobility context features $\mathbf{x}_{\text{mob}}$**: Velocity, Cell Id
- **Application context features $\mathbf{x}_{\text{app}}$**: Payload size of the sensor data packet to be transmitted

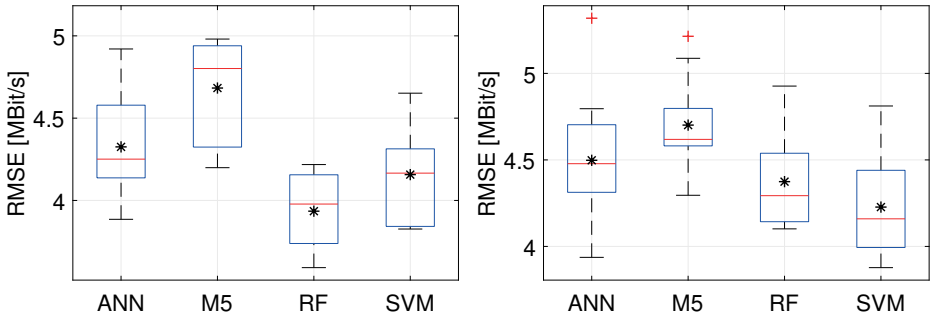

(a) Uplink **(b)** Downlink

Fig. 5.20: Comparison of the resulting data-rate prediction accuracy achieved by different machine learning models.

In addition to these passive indicators, the measured data rate of active transmissions in uplink and downlink direction with a random payload size ranging from 0.1 MB to 10 MB is determined every 10 s. An in-depth analysis of the statistical properties of the measurements is given in [638].

Using the resulting feature set $\mathbf{X}$ composed from the individual context vectors $\mathbf{x} = (\mathbf{x}_{\text{net}}, \mathbf{x}_{\text{mob}}, \mathbf{x}_{\text{app}})$, we trained a machine learning model f_{ML} on the corresponding data rate measurements $\mathbf{y}$ such that $f_{\text{ML}} : \mathbf{X} \rightarrow \mathbf{y}$. For this purpose, different regression models are considered:

Artificial Neural Network (ANN) with sigmoid action, two hidden layers, ten neurons per hidden layer, learning rate $\eta = 0.1$, momentum $\alpha = 0.001$, and 500 training epochs.

M5 Regression Tree (M5)

Random Forest (RF) with 100 trees and a maximum tree depth of 15.

Support Vector Machine (SVM) trained via Sequential Minimal Optimization (SMO) with Radial Basis Function (RBF) kernel, regularization parameter $C = 1.0$, and kernel coefficient $\gamma = 1.0$.

The training process is carried out using LIghtweight Machine learning for IoT Systems (LIMITS) [633], which provides high-level automation features for the well-known Waikato Environment for Knowledge Analysis (WEKA) framework and allows the export of `C/C++` implementations of the trained prediction models.

Figure 5.20 shows the Root Mean Square Error (RMSE) of the 10-fold cross-validation in both transmission directions. It can be seen that there are only minor differences for the more complex models if they are properly tuned. Even for the much simpler M5 model, a comparably high prediction accuracy is achieved. The RF model achieves the lowest prediction errors in the uplink direction. In contrast to ANNs and SVMs, another advantage of this approach is a significantly lower complexity for the hyperparameter tuning. As a consequence of these considerations, the further analysis focuses on

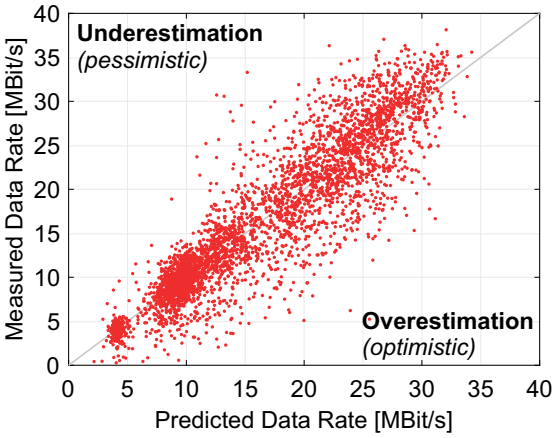

Fig. 5.21: Comparison of RF-based data-rate predictions and corresponding measurements. The diagonal line corresponds to a hypothetical perfect prediction model.

utilizing the RF model for performing the data-rate predictions. A scatterplot of the resulting uplink model is shown in Figure 5.21.

5.2.3.2 Machine Learning CAT (ML-CAT)

The basic idea of ML-CAT is to extend the CAT scheme with a machine learning-based metric for assessing the radio channel quality. While the latter is represented by the predicted data rate $\tilde{S}(t) = f_{\mathrm{ML}}(\mathbf{x}(t))$, the value range of the probabilistic transmission model is implicitly related to the value range of the SINR metric (0 dB to 40 dB according to [298]). Therefore, a normalization $\Theta(t)$ based on the value range $[\Phi_{\min}, \Phi_{\max}]$ of the transmission metric $\Phi(t)$ is defined as

$$\Theta(t) = \frac{\Phi(t) - \Phi_{\min}}{\Phi_{\max} - \Phi_{\min}} \tag{5.5}$$

The transmission probability $p(t)$ is then computed in analogy to Equation 5.4

$$p(t) = \begin{cases} 0, & \Delta t \leq \Delta t_{\min} \\ \Theta(t)^{\alpha}, & \Delta t_{\min} < \Delta t < \Delta t_{\max} \\ 1, & \Delta t \geq \Delta t_{\max} \end{cases} \tag{5.6}$$

5.2.3.3 Reinforcement Learning CAT (RL-CAT)

With RL-CAT, the previously probabilistic medium access is replaced by a reinforcement learning approach. A schematic illustration of the interactions between the different logical entities is shown in Figure 5.22. The model consists of three core components:

– The actual opportunistic data transfer is realized as an **agent** that learns to perform the possible actions—local buffering in expectation of future improvements or

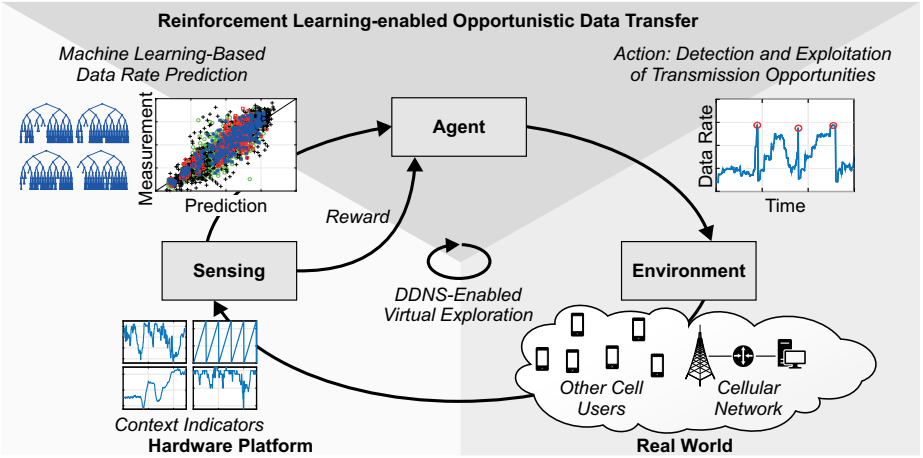

Fig. 5.22: Interaction between the logical entities for reinforcement learning-enabled opportunistic data transfer.

transmission of the whole sensor data buffer—through observation of the resulting *rewards*.

- The **environment** is represented by the cellular network. In contrast to conventional reinforcement learning, which assumes that the actions taken by the agent have a significant impact of the state of the latter within the environment, external impact factors have a dominant influence on the end-to-end behavior.
- **Sensing** is performed using the actual hardware platform based on the measurable context indicators. The raw measurements are brought together using an RF-based data-rate prediction model.

The reinforcement learning-based action selection process utilizes a decision table Q for assessing the expected rewards of the possible actions a_{IDLE} and a_{TX} based on a given state represented by the context tuple $\mathbf{c_t} = \left(\tilde{S}(t), \Delta t\right)$. Based on the available measurements, the action to be executed is determined as $a = \arg\max_a Q(\mathbf{c_t}, a)$. The classical Q-learning update process can be formulated as

$$Q(\mathbf{c_t}, a) = (1 - \alpha) \cdot Q(\mathbf{c_t}, a) + \alpha \left[r_a + \lambda \cdot \arg\max_a Q(\mathbf{c_{t+1}}, a) \right] \tag{5.7}$$

whereas α represents the learning rate, λ is the discount factor, and r_a is the reward of the taken action a. However, as pointed out earlier, the agents impact on its own state can be regarded as negligible: even if the agent was capable of performing "optimal" actions, the achievable end-to-end performance would be still impacted by non-controllable factors such as the network quality and the traffic load caused by other users. Therefore, a myopic approach that focuses on optimizing the immediate reward of the taken actions is implemented by setting $\lambda = 0$, which results in the

simplified formula

$$Q(\mathbf{c_t}, a) = (1 - \alpha) \cdot Q(\mathbf{c_t}, a) + \alpha \cdot r_a. \tag{5.8}$$

The action-specific reward functions are defined as

$$r_{\mathrm{TX}}(S, \Delta t) = \underbrace{w \cdot \frac{(S - S^*)}{S_{\max}}}_{\text{Data rate optimization}} + \underbrace{(1 - w) \cdot \frac{\Delta t}{\Delta t_{\max}}}_{\text{AoI optimization}} \tag{5.9}$$

and

$$r_{\mathrm{IDLE}}(\Delta t) = \begin{cases} \Omega & \Delta t \geq \Delta t_{\max} \\ 0 & \text{else} \end{cases}. \tag{5.10}$$

Hereby, the parameter w allows us to control the fundamental trade-off between data-rate optimization and AoI reduction, S^* is the target data rate, and $S_{\max}$ represents the upper data rate bound of the empirical measurements. Although there is no immediate reward if no data transfer is initiated, Ω serves as a punishment factor if the buffering time Δt exceeds the application-specific deadline $\Delta t_{\max}$.

Instead of performing a large number of real-world transmissions for training the reinforcement learning mechanisms, a Data-Driven Network Simulation (DDNS) setup is implemented according to [637]. In contrast to classical system-level network simulation, which requires a large number of assumptions and simplifications for setting up virtual representations of concrete real world scenarios, DDNS makes use of a combination of machine learning models and empirical context traces. This black spot approach does not require us to explicitly model communicating entities and achieves not only a close-to-reality representation of real-world behavior but also a massive computational efficiency.

5.2.3.4 Black Spot-Aware Contextual Bandit (BS-CB)

While BS-CB builds upon the reinforcement learning-based medium access approach of RL-CAT, it introduces additional mechanisms for accessing trust in the data-rate predictions. Moreover, it replaces the Q-learning component by a contextual bandit reinforcement model. A detailed description of BC-CB is given in [622].

In order to improve the data-rate prediction accuracy, the concept of black spot regions is introduced. Within those areas, the properties of the geographical environment lead to a significant increase in the location-specific prediction error (e.g., related to an increased handover probability). If knowledge about the presence about those black spots is available, transmissions can be postponed in order to avoid severe mismatches of predictions and measurements. For this purpose, BS-CB leverages a priori data about previous measurements in the targeted scenarios. Based on k-means-enabled unsupervised learning, the data is clustered into N_c clusters. For each cluster, the clusterwise RMSE is computed and compared with a given cluster threshold $\mathrm{RMSE}_{\max}$. All clusters that exceed the defined threshold are treated as black spot clusters and fitted to ellipses.

During the application phase of the transmission mechanism, the vehicle performs an ellipse test to check if it is currently within a black spots region. If the condition is fulfilled, the transmission process is postponed.

The actual reinforcement learning process is modeled as a contextual bandit that proposed an action a (either `IDLE` for further local buffering or `TX` for data transmission) using

$$a = \arg\max_{a \in \mathbf{A}} \left(\underbrace{\hat{\theta}_a^T \mathbf{c}}_{\text{Estimated reward}} + \underbrace{\alpha \sqrt{\mathbf{c}^T \mathbf{A}_a^{-1} \mathbf{c}}}_{\text{CB}} \right) \tag{5.11}$$

whereas $\hat{\theta}$ corresponds to the ridge regression coefficients of action a. $\mathbf{c} = \left(\tilde{S}(t), \Delta t \right)$ is the d-dimensional context tuple for the predicted data rate $\tilde{S}(t)$ and the current buffering delay Δt.
The degree of exploration is controlled using the greediness parameter δ

$$\alpha = 1 + \sqrt{\frac{\ln(2/\delta))}{2}}. \tag{5.12}$$

After either the `IDLE` or the `TX` has been performed, the regression coefficients are updated as

$$\hat{\theta}_a \leftarrow \mathbf{A}_a^{-1} \mathbf{b}_a \tag{5.13}$$

with

$$\mathbf{b}_a \leftarrow \mathbf{b}_a + r_a \cdot \mathbf{c}. \tag{5.14}$$

For determining the actual rewards of the chosen actions, the reward functions of RL-CAT are re-utilized according to Equation 5.9 and Equation 5.10.

5.2.4 Results of the Real-World Performance Comparison

For the performance evaluation of the novel machine learning-enabled methods, a 25 km long evaluation track with varying environmental characteristics, speed limitations, and building densities is considered. For each transmission scheme, ten real-world drive tests are performed. Hereby, a virtual sensor application generates 50 kB of sensor data per second, which is buffered locally until a transmission decision is taken for the whole buffer.

Figure 5.23 shows the achieved end-to-end data rate of the transmission schemes. While the basic channel-sensitive approach of CAT is already able to achieve a significant improvement of the data rate, the latter is highly increased through the introduction of machine learning-based channel assessment. Moreover, the reinforcement learning-based data transfer results in additional gains. In comparison to conventional fixed interval data transmission, BS-CB achieves performance improvements of 195 % in uplink and 223 % in downlink direction.

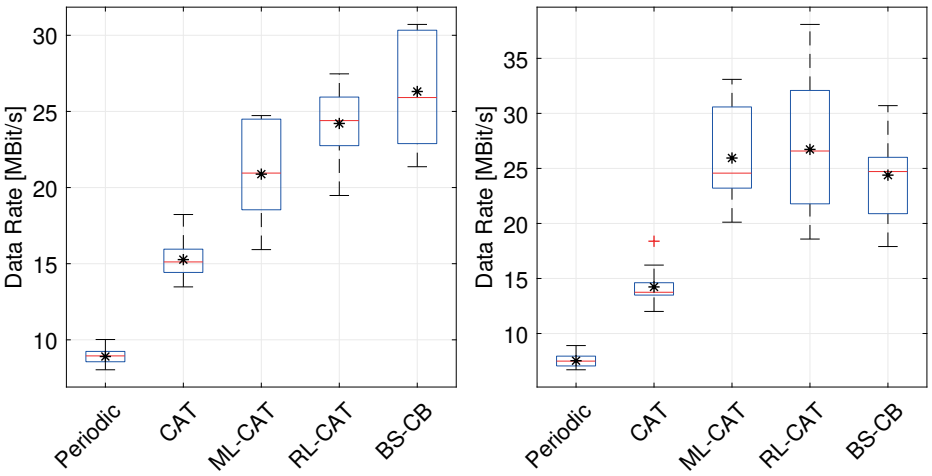

(a) Uplink **(b)** Downlink

Fig. 5.23: End-to-end data rate of the different transmission schemes.

As shown in Figure 5.24, the apparently selfish goal of data rate maximization contributes to improving the good of all: all opportunistic data transfer methods are able to achieve a significantly better resource efficiency than the conventional approach. Although the methodological evolution is also represented in the achieved results, there are only minor differences between the machine learning-enabled methods. Here, BS-CB reduces the number of occupied cell resources by around 85 % in both transmission directions.

5.2.5 Outlook and Future Work

Due to its enabling character for all presented transmission schemes, future work should focus on optimizing the accuracy of the data-rate prediction model. A major limitation of client-based prediction approaches is their limited insight into the current traffic load within the cell. Future networks could compensate this limitation through active announcement of network infrastructure knowledge about the traffic load, e.g., acquired through novel 5G mechanisms such as the Network Data Analytic Function (NWDAF). As shown by a recent feasibility study [626], the integration of available network knowledge reduces the RMSE by 25 % in the uplink and 30 % in the downlink direction.

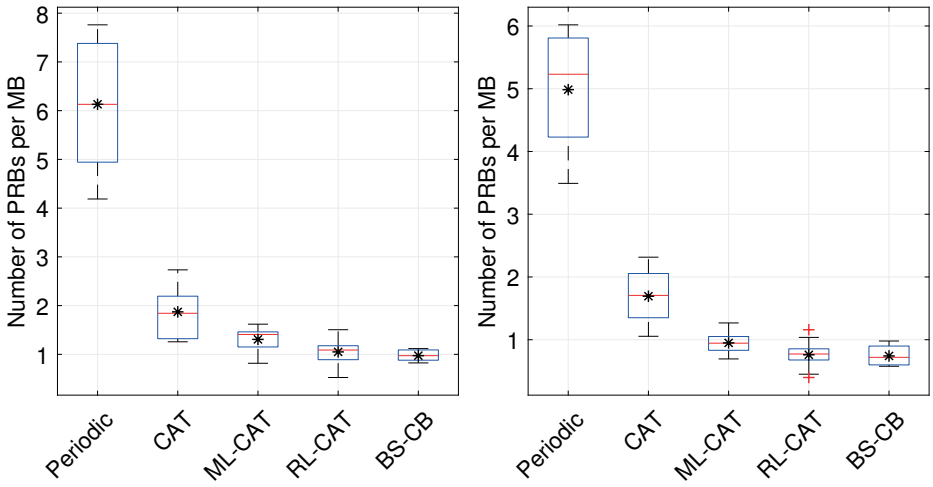

(a) Uplink **(b)** Downlink

Fig. 5.24: Resource efficiency of the different transmission schemes.

5.3 Mobile-Data Network Analytics Highly Reliable Networks

Robert Falkenberg
Karsten Heimann
Benjamin Sliwa

Abstract: The analysis of mobile network data is a fundamental requirement for the development and invention of novel networking approaches that fulfill the rapidly growing requirements and demands on those networks. This process requires the identification and a thorough investigation of shortcomings in existing field deployments, independent of the network operators and/or the network equipment vendors. Despite public standardization by the 3rd Generation Partnership Project (3GPP), cellular networks are developed and operated as closed systems that provide a predefined networking service to the subscriber while disclosing only a minimum of system-related information such as signal strength or quality in the User Equipment (UE).

However, researchers often require a deeper insight into network functionality, especially when it comes to considering network load and occasional congestions while still maintaining the privacy of the regular network users. With this knowledge, future devices may predict their achievable throughput passively under the current load and channel conditions without the need of triggering a transmission just for the sake of throughput measurements that in turn induces (unnecessary) network load. They may leverage this predicted information for e.g. load balancing, network selection, or service adaptation. Since cellular networks are centrally governed by the base stations, which assign the spectral resources by explicit and fine-grained signaling to each active device in the coverage area, information about the cell-wide resource utilization is already "in the air". For performance reasons and unlike the ciphered payload exchange between UE and the base station, the control messages that carry the resource assignments are not encrypted.

However, these messages are scrambled by a device- and session-specific Radio Network Temporary Identifier (RNTI), which is essential for the proper interpretation and validation of those messages and which is exchanged only once at the beginning of each session. This section describes the achievements of the CRC 876 in extracting these control messages of new and already active sessions efficiently and reliably over the air and without the need for expensive specialized hardware. The methodology of the underlying control channel analysis is embedded into a comprehensive open-source software framework Fast Analysis of LTE Control Channels (FALCON), which uses Software-Defined Radios (SDRs) to capture the base station's signal and accurately extracts the control messages in real time on a regular computer system.

In subsequent case studies, supervised learning is used to leverage the disclosed network-load information from short-term observations for the prediction of the expected data rate and ultimately uses this as a metric for dynamic network selection to achieve the highest throughput over the fastest network connection currently available.

5.3.1 Introduction

The steady increase in data traffic in mobile networks, triggered by the rapidly growing number of human and machine network subscribers, poses a challenge to both network operators and the services that depend on them in face of the limited radio spectrum for meeting the simultaneously growing demands on quality of service. Achievable data rates depend on the one hand on the cell bandwidth and signal quality, and on the other hand on the activity of other cell users competing simultaneously for the available radio resources. One of the possible strategies is to use higher frequency ranges, in which higher bandwidths and thus more spectral resources are available for transmission. However, due to the inherently higher signal attenuation, these frequency ranges are only suitable for covering smaller areas, so that region-wide coverage is only economical with a correspondingly high user density. In order to meet the growing demands in the remaining areas and to counteract bottlenecks, the usage efficiency of the available resources must therefore be further increased. For example, subscribers could switch to less busy networks or perform delay-tolerant data transmissions only when channel and load situations are favorable.

However, both the research and the application of such mitigation strategies require the accurate measurement of both signal quality and instantaneous network load in order to identify overload situations without creating unnecessary load themselves, e.g. in the form of test transmissions. Even though mobile devices measure signal strength and quality autonomously, present mobile networks allow users and external observers only a very limited insight into the momentary resource utilization of the cell. Although the total occupancy of radio resources can be determined by spectrum analysis (cf. top row in Figure 5.25), the actual degree of contention in the case of full occupancy remains concealed since the number of served subscribers cannot be identified in the spectrum (cf. last two columns in Figure 5.25).

In 4G and 5G networks, the distribution of spectral resources is governed by the base station, which explicitly allocates its resources to single active subscribers via special control channels. For efficiency reasons, these allocation messages are not encrypted, but reliable decoding requires knowledge of the addressee's RNTI, since the attached checksum is scrambled with it. As a result, UEs can only read the assignments that affect them. Inactive users without assigned RNTI and external observers can decode only the assignments for specially reserved RNTIs that concern general system information or paging.

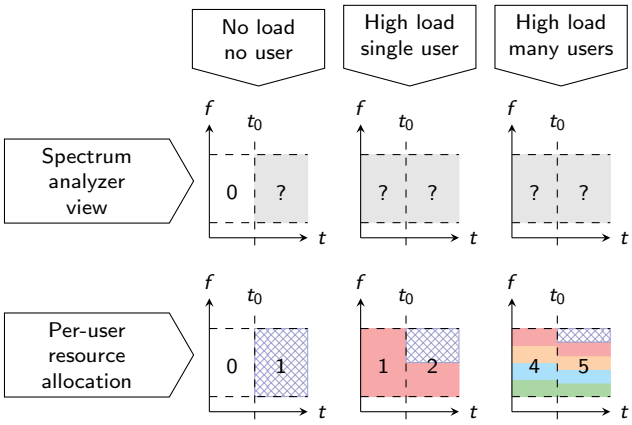

Fig. 5.25: If a new user starts a transmission at t_0, the number of allocated spectral resources (shaded in blue) depends on the number and activity of other participants (bottom row, other colors). In the spectrum analysis (upper row), the allocations of individual users are indistinguishable, and the number of participants remains unknown. ©[2016] IEEE. Reprinted, with permission, from [195].

In the course of the CRC 876 research, efficient control channel analysis methods for finding valid RNTIs have been developed and evaluated, enabling passive load sensing of the mobile network and thus providing valuable information for a client-side data-rate prediction. The presented approaches are directly applicable to public 4G cellular networks and enable real-time discovery of all resource allocations using off-the-shelf PCs and SDRs. A comprehensive reference implementation is provided in the form of the open-source framework FALCON. Using the collected data and derived features to characterize the network load, supervised learning is used to train prediction models that enable data-rate prediction whose accuracy significantly exceeds previous approaches based purely on signal strength. Applied simultaneously to multiple cellular networks, the prediction enables the UE to perform predictive network selection in order to transmit data over the network with the most promising data rate, especially during high-load periods. The prediction accuracy, achievable data rate gain, and impact on UE energy consumption are evaluated using case study data collected in public mobile networks.

The following sections are structured as follows: Section 5.3.2 presents related work in the area of control channel analysis. Subsequently, Section 5.3.3 discusses methods for analyzing control channels and presents some further implications that can be derived from observing the cell activity. Key findings are summarized and a conclusion is drawn in Section 5.3.5.

5.3.2 Related Work

The assessment of the current mobile network connection, especially in terms of link throughput, is commonly done by means of active probing [289]. This means that a data transmission is triggered to measure the throughput currently available for the device under test. However, this approach loads network and radio resources, which may be omitted by a purely passive measurement or prediction.

The device itself has only a limited view on the current network load, yet it provides performance indicators such as Reference Signal Received Power (RSRP) and Reference Signal Received Quality (RSRQ), which can be used only for a rough forecast of the achievable data rate of subsequent transmissions [297]. Authors in [392] additionally utilize details from lower protocol layers and the chipset.

More promising solution approaches need to consider further information that is usually outside the scope of the mobile device. For this reason, expensive commercial tools with special hardware requirements[695] as well as off-the-shelf SDRs and open-source protocol stacks allow tailored solutions based on deep insights into the signaling protocol behavior and related routines. In terms of the SDR-based approach, especially LTEye from [359] and Online Watcher for LTE (OWL) from [108] deal with the analysis of the control channel for resource allocations to infer the current resource utilization and concurrently active users. As will be detailed in the next section, LTEye suffers from numerous false-positive detections, while OWL constitutes a solid, real-time capable approach which only detects new devices though. In contrast to that, our approach FALCON [196] implements improved detection capabilities of the resource utilization in mobile networks and is even able to forecast or recommend the most performant network at a given time.

5.3.3 Control Channel Analysis

In current mobile networks, radio resources are divided by time and frequency in the manner of a two-dimensional resource grid. The resource grid spans the cell bandwidth in frequency domain and the time domain is divided into a nested and periodic structure of symbols, slots, subframes, and frames. The smallest resource unit in 4G and 5G networks is the Resource Element (RE), which corresponds to a single subcarrier of an Orthogonal Frequency Division Multiplexing (OFDM) symbol. According to a predefined pattern, some REs are used to broadcast synchronization sequences or provide reference levels for equalization. REs without a special purpose serve as resources for the transmission of any other data, including control and payload messages. These spare REs are grouped into equal-sized Resource Blocks (RBs), which are the smallest unit of resources that can be allocated to individual UEs. In common 4G Networks, a RE spans 12 subcarriers in frequency domain and 7 symbols in time, which corresponds to

a bandwidth of 180 kHz and a duration of 0.5 ms. However, resource allocations always apply for both slots (0.5 ms) of a subframe (1 ms).

The allocation of RBs is organized centrally by the base station and signaled to the UEs via a special control channel, namely Physical Downlink Control Channel (PDCCH), which is located in the first 1, 2, or 3 symbols of every subframe and spans the entire cell bandwidth. It includes assignments for transmissions in both directions, the downlink (i.e. from the base station to the UE) and the uplink (i.e. from the UE to the base station). These apply in the downlink (DL) direction for the current subframe or in the uplink (UL) direction 4 subframes later to give the UE enough time to prepare.

From a logical point of view, the PDCCH consists of a sequence of Control Channel Elements (CCEs), each comprising 36 REs, whose total number is calculated from the cell bandwidth and the number of occupied OFDM symbols. These CCEs carry the encoded Downlink Control Information (DCI) for single UEs, which contain the RB allocation, the Modulation and Coding Scheme (MCS), the power control commands, and further control information required for decoding or encoding the payload in the allocated resources. Base stations use rate 1/3 channel coding, interleaving, and rate matching for each emitted DCI data structure to provide FEC and to fit the encoded sequence into L, $L \in \{1, 2, 4, 8\}$ consecutive CCEs. The aggregation level L is selected by the base station according to the channel conditions of the addressee to ensure proper reception. Any additional or spare space within the L CCEs is filled by cyclic repetitions of the encoded sequence and interleaving ensures an even distribution of repeated bits.

Prior to the encoding, each DCI is appended with its 16-bit Cyclic Redundancy Check (CRC) checksum, which is additionally scrambled (via binary XOR) with the RNTI of the addressee. Conversely, receiving UEs only consider decoded DCI where the CRC matches their current RNTI.[1]

Since the PDCCH has no table of contents, only blind decoding of the CCEs can determine whether relevant information is present. This also includes all possible combinations resulting from different L. To reduce the number of decoding attempts for a UE, the standard defines a search space function that restricts the search space to a maximum of 22 evenly distributed locations according to RNTI, subframe, and L.

Furthermore, the standard defines numerous DCI formats for different transmission modes, which depend on the number of antennas used and the capabilities of the UE and the evolved NodeB (eNodeB) as 4G base station. Transmission modes are negotiated both when connections are established and dynamically depending on the channel conditions. The DCI formats differ in their length and consequently in the length of the encoded sequence. However, the same circular approach is used for rate matching, i.e. to populate the CCEs so that the initial format is no longer apparent in the transmitted

1 The UE also tracks special reserved RNTIs for system information and paging as required.

sequence. End devices therefore decode the sequences multiple times, assuming any DCI formats (or their lengths) that comply with the specified transmission mode.

Consequently, an external observer faces the following challenges when decoding the entire content of the PDCCH:

1. Numerous decoding attempts covering all locations of the PDCCH including all possible combinations of DCI format length and aggregation level are computationally expensive and result in many invalid DCI candidates.
2. Discard all DCI candidates with CRC that do not match a valid RNTI.
3. Find valid RNTIs within the set of decoded DCI candidates.

The chicken-or-egg situation resulting from the last two points can be resolved in several ways: LTEye [359] re-encodes each decoded DCI candidate, compares the encoded sequence with the received bits on the channel, and discards any candidates that deviate by a certain degree. But in the presence of noise or interference, we show that this approach is highly inaccurate and leads to numerous false decisions. In a more robust approach, [108] follows the initial connection establishment of joining UEs, which contains the RNTI assignment in plain text, and builds up a list of valid RNTIs. However, RNTIs of UEs that entered the cell before the monitoring remain undetected. Therefore, OWL follows the approach of [359] as a fallback. By contrast, the authors of UnCover Information in Mobile Access Networks (U-CIMAN) [771] propose to first accept any DCI candidate and to decode the potential payload in the allocated RBs. If this attempt fails due to an invalid CRC of the payload, the DCI is discarded. The approach involves a significant computational cost due to the larger amount of data and the more complex decoding procedure for the payload data.

In this area, CRC 876 has made substantial contributions to a resource-efficient yet reliable control channel analysis that is especially suitable for short-term monitoring in order to estimate the total cell load. In [195] we propose a histogram-based approach in conjunction with an inverse application of the search space function to identify valid RNTIs and decode the corresponding DCI candidates. First, DCI candidates decoded from all possible locations, formats, and aggregation levels are validated with respect to their permitted positions, as the eNodeB never places DCI outside their associated search space. This approach reliably filters 80-90 % of all candidates, including invalid DCI.

The following filter stage first collects the RNTIs of all DCI candidates in a history. According to an attached histogram, all candidates are discarded whose RNTI in the histogram does not exceed a threshold value k. It is based on the fact that active participants receive multiple assignments within a short period of time and that their RNTIs occur more frequently than the random RNTIs that result from decoding with incorrect parameters. An example is given in Figure 5.26.

Length of history and threshold value are optimization parameters that allow a trade-off between the probability of false positive detection, the minimum required activity of individual UEs, and the detection delay [194].

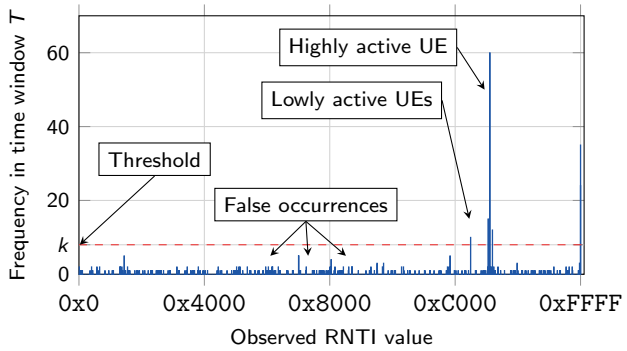

Fig. 5.26: Example for the histogram-based RNTI validation approach. RNTIs of active UEs appear with high frequency, while RNTIs of false DCI uniformly spread over the entire value range with low frequency. ©[2016] IEEE. Reprinted, with permission, from [195].

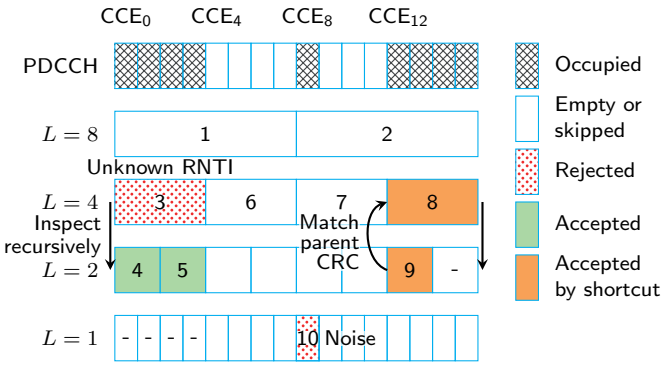

Fig. 5.27: Recursive PDCCH analysis with the *short-cut* decoding approach for immediate yet reliable discovery of unseen RNTIs. ©[2019] IEEE. Reprinted, with permission, from [196].

To improve the detection speed of unseen RNTIs, which is especially important for short-term observations, we proposed a novel *short-cut* decoding approach in [196]. The approach exploits the scheme of how the eNodeB populates the CCEs with the encoded DCI sequence. Although in most cases such sequence fits into a single CCE, operators configure the eNodeB to use higher aggregation levels in order to increase the robustness against distortions. Conversely, due to the circular repetition, a properly cropped sequence still allows for a correct decoding of the DCI. Therefore, if both decoding of the full and the shortened sequence result in the same DCI and CRC, the associated RNTI can be assumed as valid and the DCI shall be accepted. This approach can be implemented efficiently by combining a breadth-first search with a depth-first search for each location as shown in Figure 5.27. The top line shows the PDCCH as a sequence of consecutive CCEs, which are either occupied or empty according to the placement done by the eNodeB. Empty CCEs, recognizable by insufficient signal

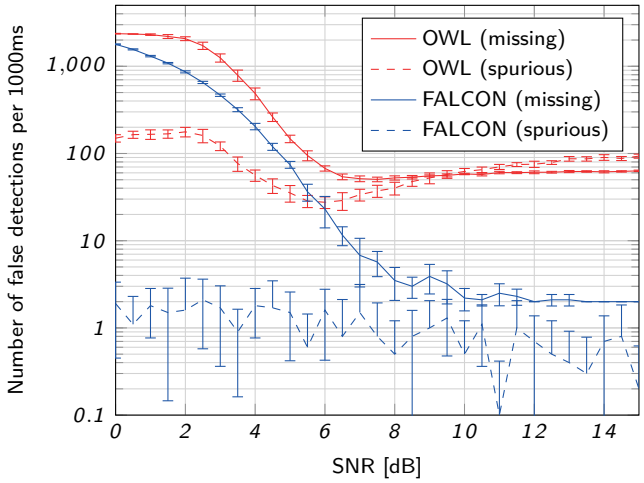

Fig. 5.28: Simulation results for decoding the PDCCH of 1 s Long Term Evolution (LTE) signal with FALCON and OWL for different Signal to Noise Ratio (SNR) values. Results are averaged over 10 repetitions; the standard deviation is shown as error bars.

power, can be skipped as they cannot contain any meaningful information. Conversely, occupied CCEs may also not have meaningful information, as the received signal power may originate from a neighbor cell with overlapping PDCCH in time and frequency. The breadth-first search component starts with aggregation level $L = 8$ and sequentially decodes all locations at this level (second line). In the given example, the two possible locations (1) and (2) are skipped, as each overlaps at least one empty CCE, and hence does not form a continuous sequence. The search continues with $L = 4$, inspecting the first location (3) with continuous CCE occupation by decoding the sequence for all DCI formats. If this inspection does not result in any DCI with a known RNTI, the depth-first component is activated and the location is inspected recursively using the next-smaller aggregation level. In the given example both locations (4) and (5) contain DCI with known RNTIs; overlapping locations (e.g. for $L = 1$) are marked as checked. As the recursion terminates, the breadth-first search continues with locations (6) and (7) both being skipped. Next, location (8) contains a valid DCI at $L = 4$ but the RNTI has not yet been seen. However, the recursive inspection of the shortened sequence, given by the first half at location (9), returns the same DCI and RNTI. As this only happens for valid DCI, the RNTI is immediately added to the list of known RNTIs, the DCI is accepted, and overlapping locations are marked as checked.

To enable detection even in the case of poor signal quality, where the bisected sequence can no longer be decoded correctly, histogram-based validation can be employed afterwards. If a recursive descent does not discover a known RNTI, all potential RNTIs along the descent path are added to the history and RNTIs exceeding a threshold are added to the active set as described above.

Robustness and reliability of the combined approach in FALCON is presented in Figure 5.28. It shows the number of missed and false DCI messages as solid and dashed lines, respectively, as a function of the SNR on the channel after analyzing a well-defined LTE signal for the duration of 1 s. Starting with 0 dB for poor radio conditions, the SNR is increased in steps of 0.5 dB to 15 dB, representing an excellent signal. For each step, the figure shows the average and the standard deviation over 10 repetitions. Furthermore, the figure also contains the results of OWL, which relies on the re-encoding approach for short-term observations.

Independently of the SNR, the amount of spurious DCI messages stays at a negligible level at FALCON, whereas OWL produces at least 30 and up to 100 false detections. In general, for all covered SNR, FALCON misses significantly fewer DCI messages than OWL. Especially for SNR values greater than 7 dB the number of missed DCI messages undershoots 10 for FALCON, while OWL remains on a level between 50 and 100. Thus, the comparison of both approaches reveals the robustness and reliability of FALCON.

Similar results are achieved in the field as an activity histogram of each RNTI over 5 s as shown in Figure 5.29. Blue circles represent the number of resource allocations detected by FALCON for each RNTI, and red crosses show the results of OWL. It is evident that the most active RNTI concentrate in a small value interval, indicating that the eNodeB assigns RNTIs consecutively to new UEs. The peak region moves over time towards larger RNTI values, as shown in the results 5 min and 10 min later. On the other hand, OWL reports numerous spurious DCI messages with random RNTIs, which are uniformly distributed over the entire value range with very low frequency.

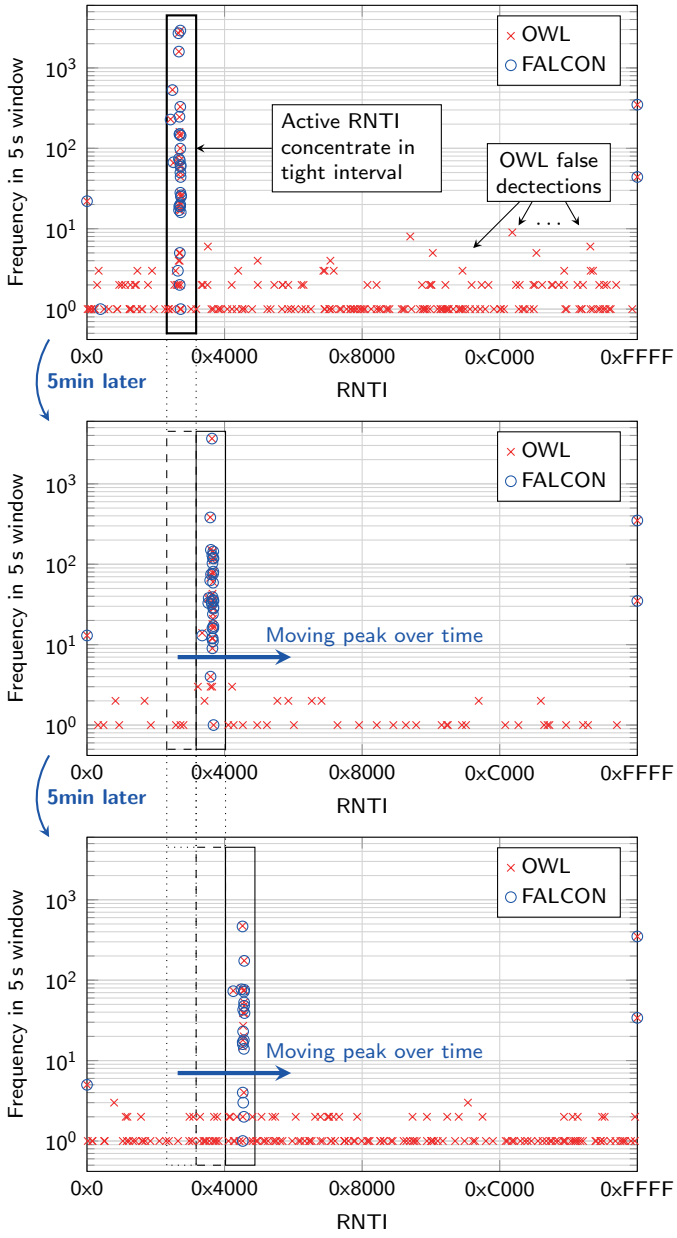

Fig. 5.29: Activity histogram of individual RNTI detections from field measurements. The highlighted peak moving over time indicates a RNTI assignment strategy of the base station. ©[2019] IEEE. Reprinted, with permission, from [196].

5.3.4 Cooperative Data Rate Prediction Leveraging FALCON

Client-based data-rate prediction is a key enabler for anticipatory mobile networking. By considering network context measurements to estimate the end-to-end transmission efficiency—represented by the predicted data rate—mobile clients can actively contribute to optimizing the intra-cell resource efficiency by scheduling data-intense transmission to resource-efficient connectivity hotspots [622]. However, the accuracy of client-based data-rate prediction methods is inherently limited since the UEs are only aware of the radio channel conditions but not of the network load.

In addition to the passive context measurements of purely client-based data-rate prediction according to [622], FALCON allows the derivation of additional features (number of Physical Resource Blocks (PRBs) and UEs, Transport Block Size (TBS), MCS) that are correlated to the current network load of the cell. For a proof-of-concept evaluation, the following feature sets are derived for the two transmission directions after an initial feature importance analysis:

- **Uplink feature set**: RSRP, RSRQ, velocity, payload size, number of PRBs, number of UEs, cell ID
- **Downlink feature set**: RSRQ, velocity, payload size, number of PRBs, number of UEs, TBS, MCS, cell ID

It can be seen that the uplink direction is more sensitive to the radio channel conditions while the downlink performance highly depends on the intra-cell traffic load.

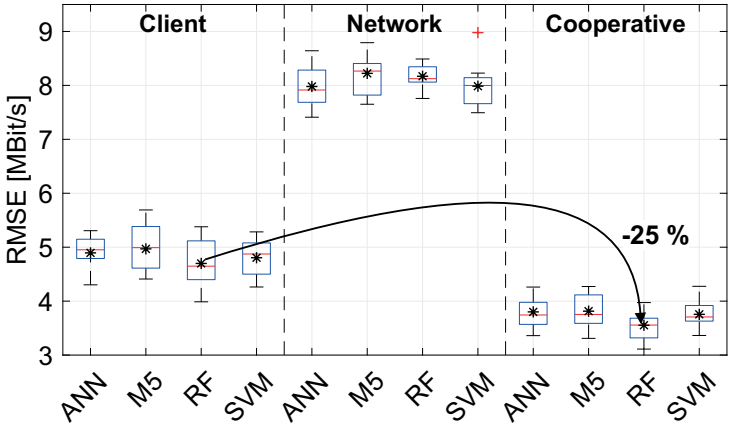

Fig. 5.30: End-to-end uplink data-rate prediction: performance comparison of different prediction approaches and machine learning models. ©[2020] IEEE. Reprinted, with permission, from [626].

The resulting prediction accuracy of different machine learning models (Artificial Neural Network (ANN), M5 regression tree, Random Forest (RF), Support Vector Ma-

chine (SVM)) is shown in Figure 5.30. Although the RF model achieves the highest overall accuracy, there are only minor differences between the machine learning models. However, significant differences can be observed for the three different data-rate prediction methods. While the purely client-based approach is mostly unaware of the traffic load and the purely network-based approach yields a high prediction error due to the absence of radio channel information, the cooperative prediction method reduces the average Root Mean Squared Error (RMSE) by 25 % in the uplink direction. As further analyzed in [626], similar improvements are also achieved in the downlink direction.

These initial results show that the context-awareness and the predictability of mobile communications can be significantly improved by combining client measurements with network-side information. Therefore, future networks such as 6G should actively provide network-load information to the clients in order to allow them to actively participate in network management functions.

5.3.5 Conclusion

FALCON is a novel open-source and SDR-based instrument for LTE control channel analysis that allows the reliable monitoring of the resource allocations of LTE cells in real-time. Through the application of shortcut-precoding, a fast DCI integrity check is achieved and the list of active RNTIs—which corresponds to an estimation of the number of active users—is derived. Aided by a histogram approach, the accuracy of FALCON is maintained even during low signal quality periods. The revealed network-side information is of particular value for intelligent networking methods that utilize end-to-end predictions for their decision making, such as resource-efficient vehicle-to-cloud communications that is discussed in Section 5.2. As purely client-based data-rate prediction approaches that rely on network context measurements are unaware of the current network load, their achievable prediction accuracy is inherently limited. As demonstrated in a first real-world proof-of-concept study [626], the incorporation of FALCON offers the potential to improve client-based data-rate prediction methods by up to 25 % in the uplink and 30 % in the downlink direction.

5.4 Machine Learning-Enabled 5G Network Slicing

Caner Bektas
Fabian Kurtz
Dennis Overbeck
Christian Wietfeld

Abstract: In this contribution, the different parts of the end-to-end network slicing concept are presented, including the Core Network (CN) and the Radio Access Network (RAN), while highlighting the differences and similarities of both domains.

Further, prototypical implementations and empirical evaluations of 5G network slicing are discussed, deepening the understanding of network slicing and identifying possible advantages and challenges. The predictability of user traffic in the respective network slices poses such a challenge, as resources in the RAN–in contrast to resources in the CN–are subject to fluctuations based on channel quality. Critical infrastructures typically require very low latencies in the single-digit milliseconds range and are thus considered ultra-Reliable Low Latency Communication (uRLLC) . To mitigate latency-intensive scheduling requests and grant operations, resources in the RAN have to be pre-allocated for uRLLC slices.

This operation, also known under the term Configured Grants (CGs), pre-allocates resources for, say, high-priority slices, so that User Equipments (UEs) are able to send data without asking for resources, which reduces the scheduling latency down to zero. The simplest method for calculating CGs is based on static allocations, which has one major drawback: unused resources are wasted, and thus, can not be used by remaining slices, effectively lowering spectral efficiency. Here, we present SAMUS (Slice-Aware Machine Learning-based Ultra-Reliable Scheduling), a data-driven method to predict resources in the future based on real data, e.g., solar activity in smart grid slices, to reduce latencies while maintaining high spectral efficiency.

5.4.1 Introduction to 5G End-to-End Network Slicing

Critical infrastructures, such as energy networks, logistics, or autonomous transportation, are becoming more and more automated to further increase efficiency. Automation is often achieved by the self-organization of processes and actors via mobile communication systems. As many different vertical industries are reliant on mobile communication, a highly diverse set of Key Performance Indicators (KPIs) need considering by the communication systems.

Until recently, dedicated communication networks were the go-to solution in order to meet these divergent criteria, as they can be designed specifically for the needs of the respective critical infrastructures. Consequentially, the fifth generation of mobile networks (5G) aims to unify these different and partly contradictory set of requirements into a single physical infrastructure. Employing Software-Defined Networking (SDN) and Network Function Virtualization (NFV) techniques, 5G *network slicing* is integrated into the 5G standard. By utilizing virtual dedicated networks called network slices on top of a single physical communication network, various vertical industries can be automated, as shown in Figure 5.31.

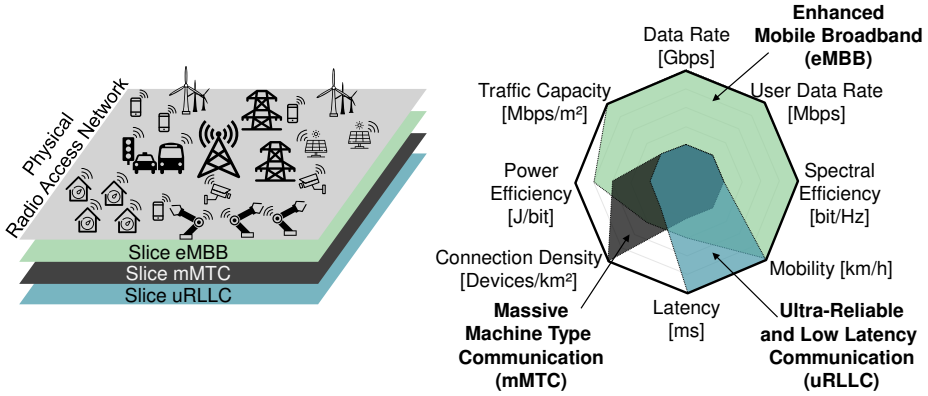

Fig. 5.31: Network slicing as a key enabler for fulfilling all specific service requirements simultaneously.

5.4.2 5G Core Network Slicing

5.4.2.1 Description and Methodology

The virtualization of network resources depicts a main pillar of 5G networking as illustrated in Figure 5.32.

The creation of multiple isolated network partitions known as slices can independently and efficiently manage different use cases with their respective demands on QoS or other guarantees. For 5G communication, three main categories are defined. The first category is enhanced Mobile Broadband (eMBB), which is used for data-rate intensive services (up to 20 Gbit/s). This category comprises ultra-high resolution video streaming as well as fixed wireless broadband and Augmented respectively Virtual Reality (AR/VR) . The second category is massive Machine Type Communication (mMTC) , designed for the emerging Internet of Things (IoT) and Industry 4.0 applications, both of which introduce a significant increase in inter-machine communication. But

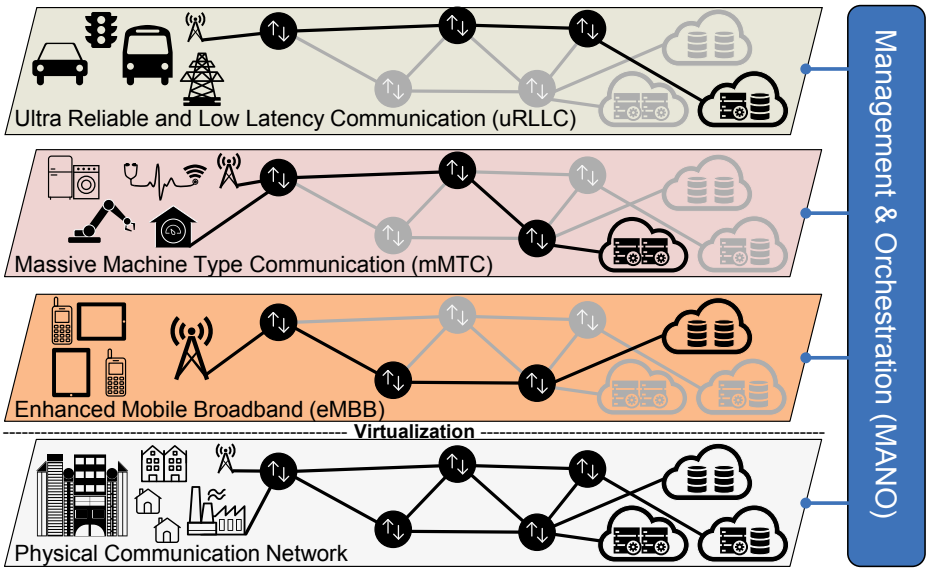

Fig. 5.32: Overview of a sliced 5G communication network.

a massive amount of devices comes with its own set of challenges in, say, electrical power grid application scenarios such as smart metering. The third major category is the ultra-Reliable and Low Latency Communication (uRLLC) service. This category comprises services such as Intelligent Transportation Systems (ITS) with Floating Car Data (FCD)-based Vehicle-to-X communication. Here, mission-critical and latency-sensitive applications are addressed. Our approach builds on NFV and SDN, which are closely related. With NFV, hard- and software is decoupled and functionalities are abstracted in order to achieve highly flexible communication infrastructures for enabling cloud computing. Virtual Network Functions (VNFs) are now able to run on Commercial-Off-The-Shelf (COTS) server platforms . By using the complementary approach of NFV and SDN, the controller can dynamically route traffic flows between the VNFs, while being deployed as a VNF itself. In addition to the utilization of SDN and NFV, our concept is based on queuing strategies utilizing the Hierarchical Token Bucket (HTB) . On the bare-metal and virtualized data-plane devices, the switching software Open vSwitch (OVS) is deployed. Furthermore, a Management and Network Orchestration (MANO) controller is implemented. This controller creates a main bridge in each switch, which includes the respective physical ports. By that, one bridge per slice is added or removed as needed. The concept is depicted in Figure 5.33.

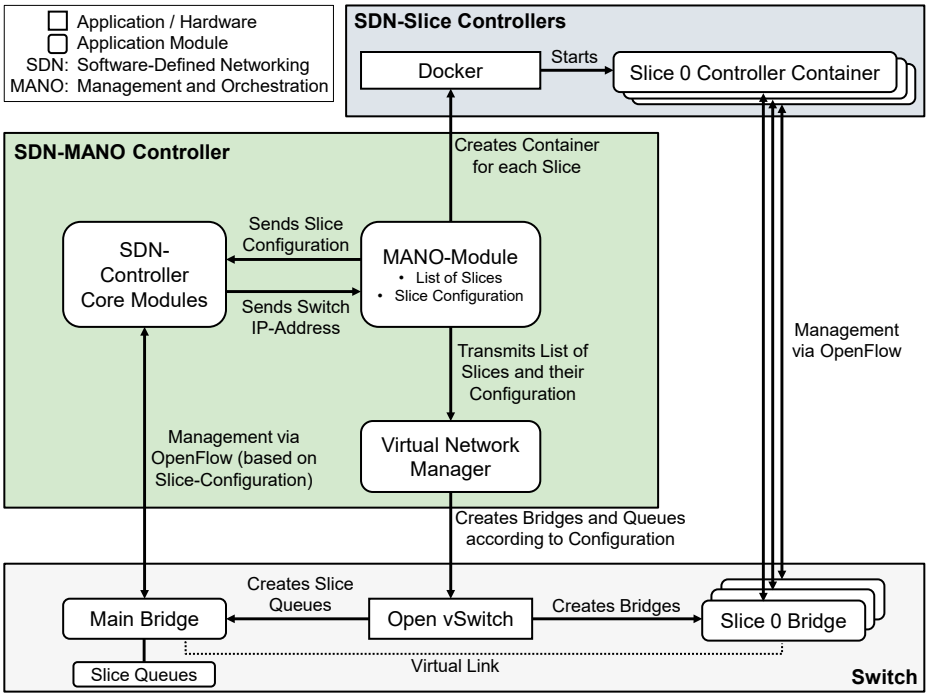

Fig. 5.33: The developed network slicing architecture.

The slice bridges comprise virtual ports, which are residing within the main bridge. The orchestration is done via the MANO controller, which dynamically instantiates slice controllers (e.g. via Docker), which in turn can be optimized for each application scenario. In the event of traffic entering the data plane, the MANO controller assigns packets to the respective bridge. There, the flow is mapped to the respective QoS queue and virtual destination port on the main bridge regarding the specific protocol or other criteria of the packet. This is done by the respective slice's controller. For each hop, the slice controller repeats this procedure of directing the flows to the main bridge. Unknown flows or not specified matches are handled on a best-effort basis. While this first part focuses on wired 5G communications, compatibility with the air interface slicing technologies presented in later sections.

5.4.2.2 Empirical Evaluation of 5G Core Network Slicing

Overview of the Testing Environment The testbed scenario is depicted in Figure 5.34. Six servers are assigned in pairs for each of the different use cases. These servers function as hosts to either send or receive data traffic over the sliced network. Furthermore, four machines are designated as SDN controllers, where three of them act as slice controllers running Floodlight and one is the MANO controller employing Ryu. For

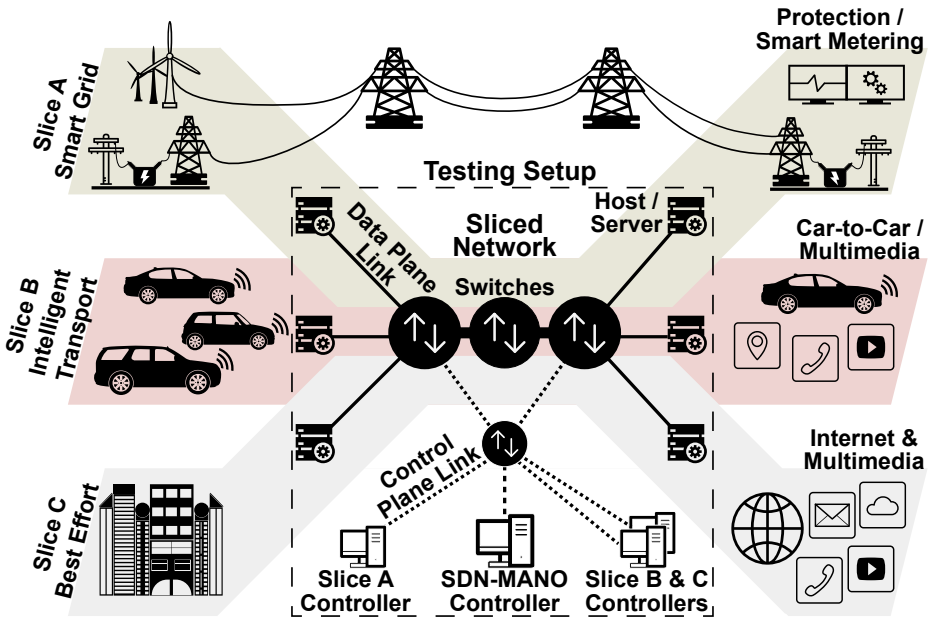

Fig. 5.34: Evaluation scenario in the testing setup.

measurement interference avoidance, three different networks, namely out-of-band control, maintenance, and sliced data-plane network, are in use. The Precision Time Protocol (PTP) is utilized to synchronize the controller clocks with a maximum deviation of 153 µs and a mean deviation of 16 µs. The underlying network load is generated via iPerf2 community edition and consists of User Datagram Protocol (UDP) packets. The maximum performance of the evaluated methods needs to be determined by considering the different layers of the ISO-OSI stack. The Ethernet frame size on the 2nd layer of the OSI model is 1512 B, which is used as a point of reference. Since performance evaluations are located on layer 4, the payload (i.e. goodput) results in 1470 B, which is 97.2 % of the layer 2 data rate. The following measurements were repeated at least 100 times with a minimum duration of 1 min per run.

Evaluation Scenarios Scenario A depicts a performance study, where key aspects from 5G and critical infrastructure communication are evaluated such as delay and data rate for varying network loads. Therefore, the overhead of our approach is determined to demonstrate the efficient use of resources. By using 100 Mbit/s Ethernet links, possible limitations can be avoided while simultaneously affording the option of CPU load monitoring during testing. Moreover, the independence of the slices from each other is verified, so any detrimental effects of errors or overload in one slice harming other slices can be precluded. Within the evaluation, the network load is increased in steps,

Tab. 5.1: Slices and traffic flows of the critical infrastructure communications scenario.

Slice (descending priority)	Use case	5G service class	Priority within slice	Hard min. data rate [Mbit/s]	Max. delay [ms]
Smart grids	Protection (IEC 61850)	uRLLC	Highest	50	1
	Smart metering	mMTC	High	200	20
Intelligent Transportation Systems	Floating car data	uRLLC	Highest	100	1
	Passenger Internet	eMBB	Low	450	10
Best-effort	Multimedia	None	Lowest	None	100

reaching beyond the maximum usable data rate/goodput, i.e. 97.2 % of the nominal layer 2 link capacity.

This approach represents cases in which end users try to use more resources than allocated for their respective slice, thus serving to demonstrate slice independence. The misconfiguration by operators of sliced communication networks is simulated as well. For this, two slices whose combined data rate exceeds the underlying physical network's data rate are configured. Scenario B represents a scalability analysis, where a viable approach for deployment in largescale, multitenant communication infrastructure is demonstrated. Since the number of slices should not influence the overall network performance, the delay performance for no, 2, 8, and 16 slices is analyzed. The available data rate is shared equally among the slices, with traffic streams utilizing 100 % of the respective slice's capacity. This ensures the exclusivity of side effects caused by slicing and not by network congestion or other factors. Furthermore, the validity of slice isolation and the stability of end-to-end delay is examined. For this, seven out of 21 slices are subjected to UDP-based traffic with data rates above the allocated limit. In contrast to scenario A, 1 Gbit/s Ethernet is used to stress test the concept. Finally, scenario C depicts a critical infrastructure communication including FCD of ITS and the IEC 61850 SG protocol.

Since both use cases are considered as uRLLC 5G services and therefore assigned the highest priority, the slices on which they are transmitted is allocated equal priority, including the dedicated SDN controllers. Smart metering (representing mMTC) and passenger internet (eMBB) are included in the related slices to demonstrate the ability of traffic distinction in our solution. Moreover, a best-effort slice is included for handling multimedia traffic and perpetually transmitting low-priority data at 950 Mbit/s, roughly consuming the maximum available layer 4 goodput of the 1 Gbit/s network. Therefore, if another slice needs a specific data rate, the network is overloaded and reallocated due to its differing priorities. The maximum tolerable delays for each priority data is given in Table 5.1. For the evaluation, the given data rates were achieved by bundling

multiple traffic flows. However, the bundling of constant data rates can be found in real-world use cases such as smart metering.

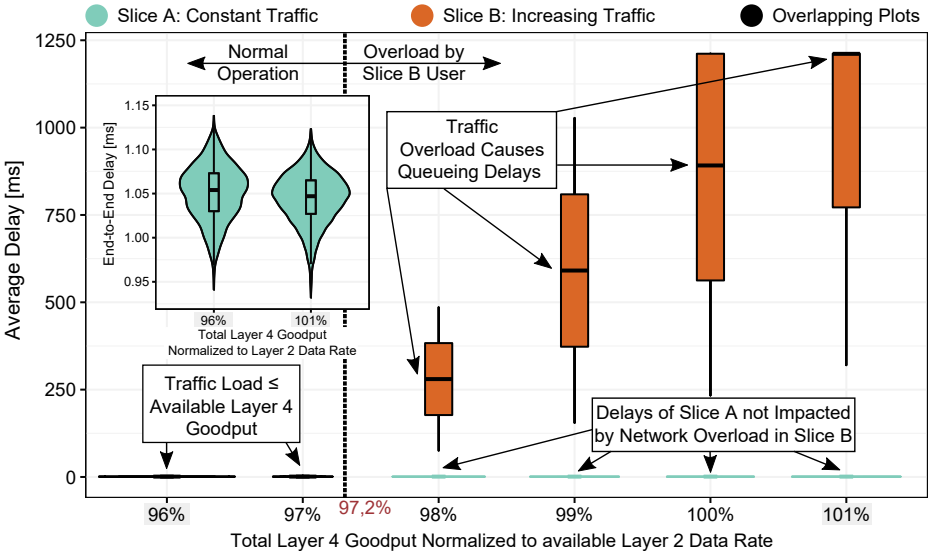

Fig. 5.35: End-to-end delays of two slices for varying traffic loads.

Evaluation Results In Figure 5.35 the end-to-end delays of two slices for varying traffic loads are depicted [360]. The physical 100 Mbit/s Ethernet network is shared fairly; slices A and B transmit UDP packets. Below the aforementioned limit of 97.2 %, the median end-to-end delay is located at 1.05 ms with a variance of approximately 0.05 ms. Nevertheless, when step-wise exceeding the limit at slice B, an overload situation is created resulting in increased delays. At 101 % load the median delays rise sharply up to 1.212 ms. However, the delays at slice A remain unaffected, even compared to no slicing as depicted with the enlarged violin plots. Hence, the isolation of the slices is shown. The misconfiguration by the operator is simulated and depicted in Figure 5.36.

Slice A receives a data rate of 40 Mbit/s (38.9 Mbit/s effectively on layer 4). The total sum of queue data rates (depicted on the x-axis) should not exceed the theoretical layer 2 limit of the 100 Mbit/s Ethernet link. This maximum is calculated as the ratio between frame sizes at layers 2 and 1, which amounts to 1512 B/1532 B = 98.7 %. In the event of misconfiguration, slice B tries to utilize resources, which do not exist. Therefore, slice B cannot maintain the layer 4 goodput. While slice A consumes the HTB tokens and remains stable, the data rate of slice B levels out to 56.9 Mbit/s, which is below the configured 60 Mbit/s queue data rate on layer 2. Overhead in terms of achieved throughput is not observed and therefore confirms expectations. Figure 5.37 depicts scenario B evaluation results.

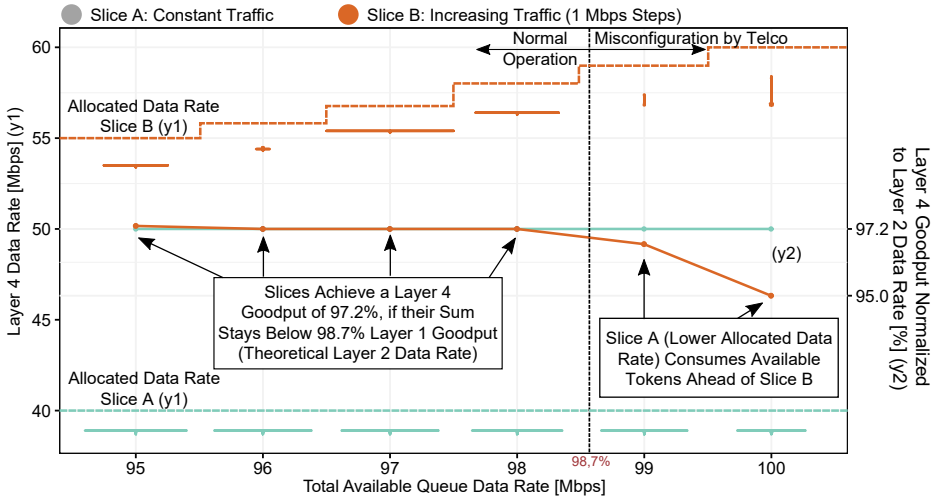

Fig. 5.36: Impact analysis of misconfiguration by the physical network operator.

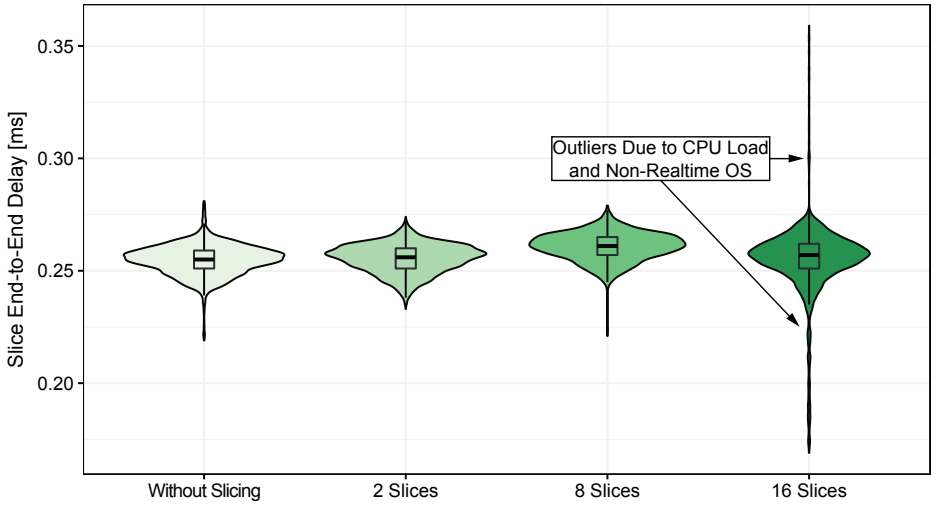

Fig. 5.37: Evaluation results of multiple slices on end-to-end delays.

The end-to-end delays of dedicated, sliced networks are given with an overall capacity of 1 Gbit/s. The data rate is fairly distributed between the slices with all of them in idle mode except one. With as many as sixteen slices, the delays remain stable. However, outliers of up to 0.36 ms may be a result of CPU context switches, which are required since the hardware provides a maximum of eight threads and queues. The outliers down to 0.17 ms are presumably caused by the non-realtime reduced timer/interrupt coalescing of the Network Interface Card (NIC) and Operating System (OS), which is triggered by the raise in computational load. Therefore, for highly sliced networks, performance optimizations of the developed source, real-time kernels and higher thread-count CPUs are to be pursued. The following stress test is given in Figure 5.38.

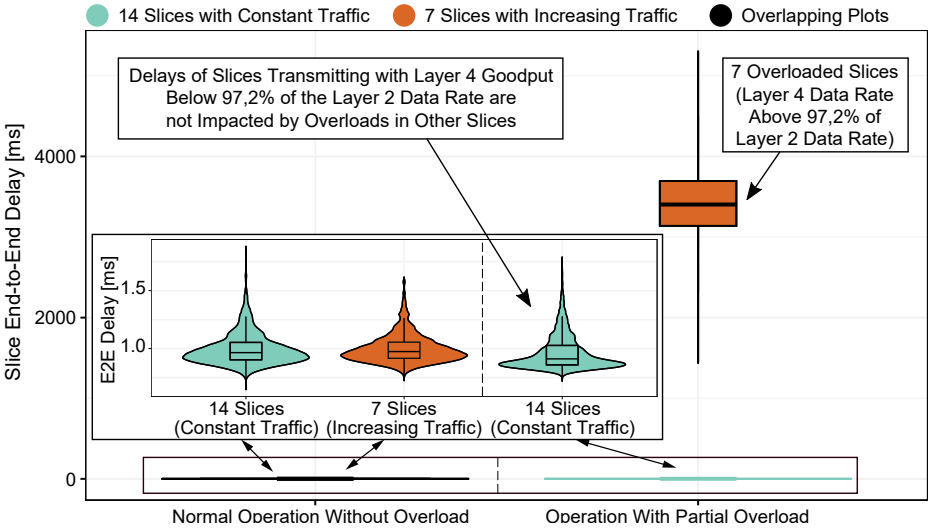

Fig. 5.38: Stress testing of scalability with partial overload.

Under normal operation (left-hand side), 21 coexisting slices can fully utilize their allocated data rates with a stable median delay of 1 ms. In comparison with previous tests, the delays are higher, because of more slices sharing a slower physical network of 100 Mbit/s. On the right-hand side, partial overload is simulated, resulting in seven slices trying to exceed their limits and accordingly causing increasing delays up to 3.5 s. However, the other slices stay unaffected. Therefore, the isolation of network slices remains equally robust even with high loads in several slices. Furthermore, the data rate remains stable across all realized scalability tests.

Finally, Figure 5.39 summarizes the measured data rate of the traffic flows given in Table 5.1. It starts with only one traffic flow of 950 Mbit/s, which fully utilizes the physical network on a best-effort basis. Therefore, even though the traffic continues throughout the test, network resources can be allocated to higher priority slices. Thus,

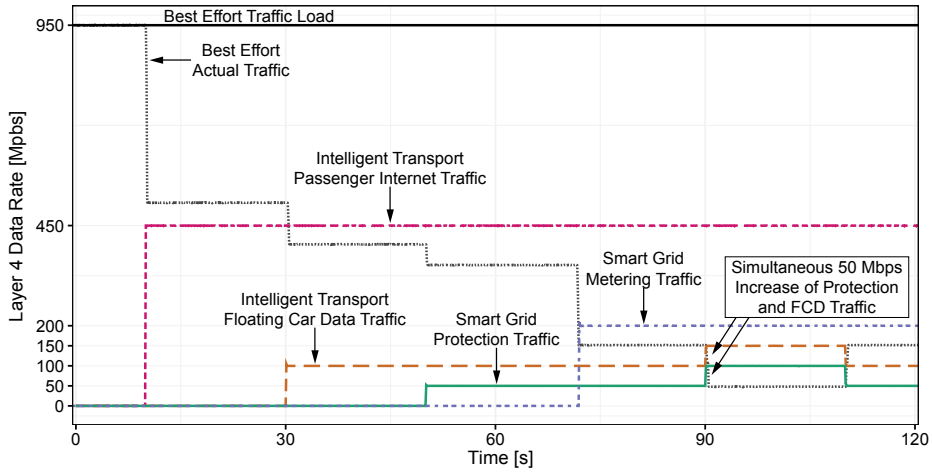

Fig. 5.39: Critical infrastructure communication scenario—data rate allocation and slice isolation.

when passenger Internet and FCD traffic of the ITS slice is generated, the best-effort throughput is reduced nearly instantaneously. The same happens when introducing protection and smart metering traffic on the SG slice. For an especially critical test case, the protection and FCD traffic is simultaneously increased to 50 Mbit/s at 90 s into the measurement. Figure 5.40 depicts the end-to-end delay of the slices.

As shown, hard service guarantees are provided during these transitions. Best-effort typically stays below the set boundary of 100 ms. Nevertheless, outliers of about 350 ms occur, which are induced by slice overloads. The delay for smart metering and passenger internet stays below 3ms with a median of approximately 1.3 ms and therefore satisfy service-level guarantees. The outliers result from the starting phase. During slice reconfiguration, the violins of protection and FCD traffic show slight delay variance, which does not affect the requirements since it stays mostly below 0.5 ms.

5.4.3 Data-Driven 5G Network Slicing in the Radio Access Network

In contrast to the previous sections, in which network slicing in the core network was discussed, the focus here is on the Radio Access Network (RAN). In this context, the data-driven aspect of network slicing becomes more important, as low-latency slices can be realized only via the prediction of emerging network traffic. This relation is described in the following subsections.

5.4.3.1 Introduction to Data-Driven Network Slicing

Previously, the three main service types eMBB, uRLLC, and mMTC were introduced. The balance between uRLLC slices and eMBB slices is particularly challenging to maintain

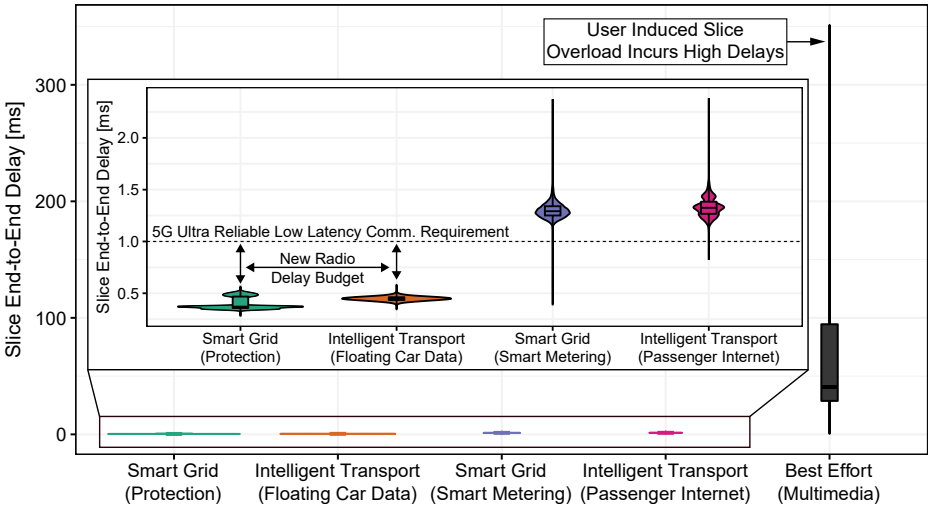

Fig. 5.40: Critical infrastructure communication scenario–end-to-end slice delays.

within Radio Resource Management (RRM), i.e. the network scheduler that is a crucial part of realizing network slicing within the RAN. To understand this relation, end-to-end latency components, which were derived from [493], are depicted in Figure 5.41.

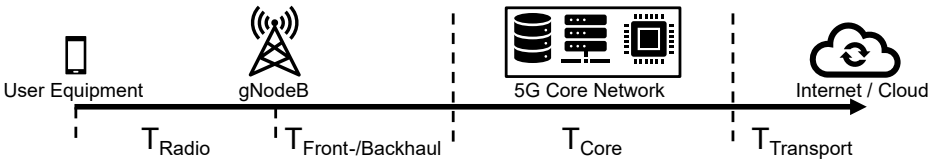

Fig. 5.41: Components in cellular networks which induce latency [493]. ©[2021] IEEE. Reprinted, with permission, from [54].

- $T_{Transport}$: Latency caused by the transmission of data through the transport network as when a web page is retrieved from the Internet.
- T_{Core}: Once data has been transmitted via the radio interface or from the transport network, it goes through the core network. This introduces additional latency, firstly because it is transmitted over an additional network, but also because the IP packets are unpacked and packed into different protocols required by the mobile network.
- $T_{Front-/Backhaul}$: The connection between gNodeB (5G base station) and core network introduces additional latency.
- T_{Radio}: The physical properties of the transmission channel are the main cause of radio latency, but the scheduler (T_{Sched}) also adds a significant delay.

The latency induced by the network scheduler (from here on T_{Sched}) is part of the T_{Radio} component, which is heavily reliant on the type of services or slices present in the 5G network. To further illustrate this, suppose there are two slices configured within the communication network, one uRLLC slice and one eMBB slice. From now on, the uplink direction of data transmission is focused (from UE to gNodeB). The UEs or the applications within the eMBB slice are data-rate intensive, which means that *as many Resource Blocks (RBs) as possible* are to be scheduled. By contrast, applications in uRLCC slices are not data rate intensive, but are to be scheduled *as fast as possible* to minimize the overall end-to-end latency (cf. Figure 5.41). This means that in order to minimize the scheduling latency, it is crucial to issue the scheduling grants before a request is even generated. For this, the so-called Configured Grant (CG) or proactive scheduling will be introduced in 5G [375]. As the name suggests the scheduling grants can be configured in advance to ensure Quality of Service (QoS) requirements . The major challenge, however, is that this requires a prediction of future data demands and channel qualities to allocate the required amounts of RBs for each network slice. The exact prediction of RBs for the uRLLC slice is crucial in this process because end-to-end latency will increase significantly if the predicted RBs are too low, which will induce retransmissions. If the predicted RBs are too high, the unused RBs will be wasted and thus not available for other network slices within the cell. This in order affects the aforementioned balance between uRLCC and eMBB slices, as the required prediction will induce prediction errors and thus waste resources for the data rate-intensive eMBB slices.

The remainder of this section will describe a data-driven CG-based scheduling and simulation framework called *SAMUS* [54], or *Slice-Aware Machine learning-based Ultra-reliable Scheduling*.

5.4.3.2 Description and Methodology of SAMUS: Slice-Aware Machine Learning-based Ultra-Reliable Scheduling - A Data-Driven Network Slicing Framework

5G-RGS (5G Resource Grid Simulation) Framework Figure 5.42 provides an overview of all modules, inputs, and outputs of the SAMUS system. As can be seen, the SAMUS system is not only comprised of the actual SAMUS scheduler prototype but additionally includes the 5G Resource Grid Simulation (5G-RGS) framework, which was specifically developed to evaluate the SAMUS scheduler prototype. There, channel conditions and data amounts of each User Equipment (UE) (or the external data used to predict the amounts) are provided as input to both modules. The 5G-RGS framework is then able to calculate resulting data rates and packet latencies ($\propto T_{Sched}$) based on the aforementioned channel conditions, the Transmission Time Interval (TTI) , as well as the allocated RBs. The last is a product of the data-driven SAMUS scheduler prototype, which generates CGs in the form of resource grid allocations based on external data, a process which will be described later in this section.

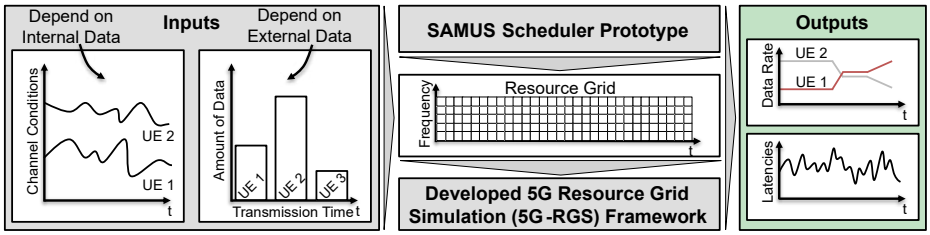

Fig. 5.42: SAMUS system overview including all modules, inputs, and outputs. ©[2021] IEEE. Reprinted, with permission, from [54].

In order to further detail the data rate calculation process of the 5G-RGS framework, Equation 5.15 is provided,

$$\textit{Data Rate (Mbit/s)} = 10^{-6} \cdot \sum_{n=1}^{N_{UE}} \left(\frac{TBS^{(n)}}{TTI} \cdot N_{TTI} \right) \tag{5.15}$$

where N_{UE} describes the UE amount of the slice, $TBS^{(n)}$ the Transport Block Size (TBS) available for the n-th UE in bit, and N_{TTI} describes how many TTIs are available in a second (the default here is the New Radio (NR) specification of 1 ms).
Moreover, packet latencies are calculated via Equation 5.16:

$$\textit{Latency (ms)} = (I_S - I_C) \cdot TTI \tag{5.16}$$

where I_S represents the scheduling interval of a final packet bit transmission and I_C the interval of packet creation. Note that latency components like retransmissions or other components of (T_{Radio}) are neglected. Based on scenarios in [53], the 5G-RGS framework was successfully validated.

SAMUS Scheduler Prototype As can be seen on Figure 5.43, the inputs for the SAMUS scheduler prototype are comprised of channel conditions or Channel Quality Indicators (CQIs) and data amounts of each UE or Buffer Status Reports (BSR), which are generated from historical data (hence, data-driven). Apart from the fact that data-driven (low-latency) CGs can be generated, traditional (latency-intensive) Scheduling Requests (SRs) can also be processed by the SAMUS scheduler. The CGs, if predicted correctly based on historical data, can reduce the scheduling latency T_{Sched} down to zero. For generating the RBs and the data-driven CGs, the ARIMA (Auto-Regressive Integrated Moving Average) method is utilized to predict the future traffic data demands and CQIs.

To ensure safe operation of mission-critical slices, resources of critical applications are allocated first, while granting the remaining RBs to best-effort (eMBB, no QoS) slices (based on the *Greedy Network Slicing Scheduler* in [53]). Traditional scheduling is used whenever a packet could not be transmitted due to prediction errors (effectively increasing scheduling latency). As a result, a resource grid in the form of a matrix is

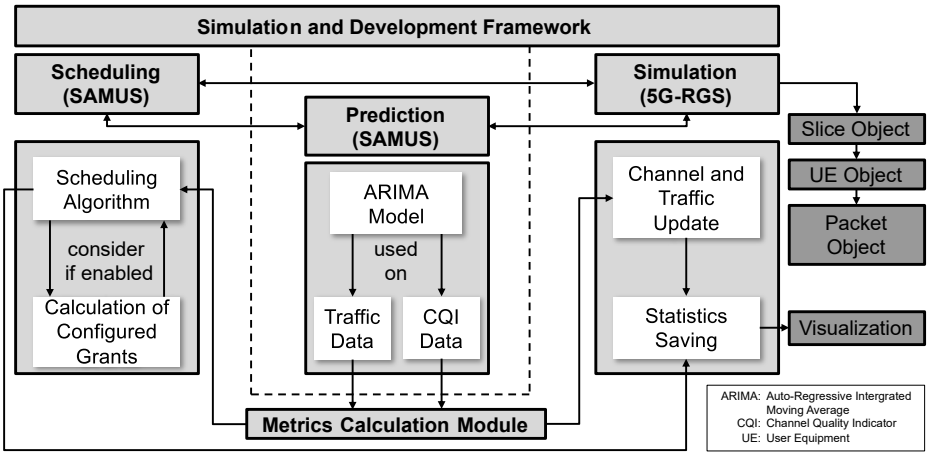

Fig. 5.43: Overall SAMUS framework including all modules and their interactions. ©[2021] IEEE. Reprinted, with permission, from [54].

passed on to the 5G-RGS, which calculates and protocols the Key Performance Indicators (KPIs), e.g., data rate and latencies. Finally, CQI and BSR values are updated and a new cycle is initiated.

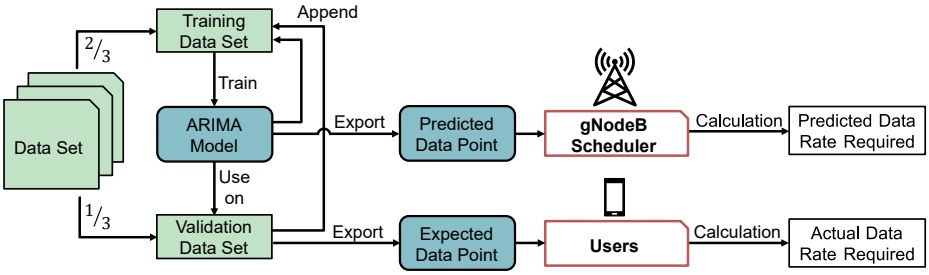

Fig. 5.44: Flow chart of SAMUS's prediction module utilizing ARIMA for training and operation. ©[2021] IEEE. Reprinted, with permission, from [54].

The details of the ARIMA-based prediction module are depicted in Figure 5.44 as a flow chart. There, the dataset associated with a slice is split up into training and validation datasets with a ratio of $\frac{2}{3}$ and $\frac{1}{3}$, respectively. Subsequently, the ARIMA model is trained in the course of offline learning based on the training dataset in order to predict future data, which in turn is utilized by the SAMUS scheduler to generate CGs. The data rate that corresponds with these CGs is the so-called data rate *predicted*. By contrast, the *actual* data rate required is calculated based on the validation dataset. The value for the actual data rate required results from the actually transmitted packets within the simulation, which are generated in order to test the prediction quality of the ARIMA

module. Also, online learning is facilitated to further optimize predictions based on the newly acquired data during the simulations.

These simulations, which were utilized to evaluate the SAMUS scheduler prototype are described in the next section.

5.4.3.3 Evaluation of the Data-Driven Scheduler Prototype SAMUS

Evaluation Scenario and Parameters As indicated in the previous sections, the SAMUS framework was evaluated based on a realistic network slicing scenario . In order to be able to compare the novel approach of the SAMUS scheduler to traditional methods as well as to present different trade-off strategies, so-called *modes* were designed and utilized.

The following modes were configured and evaluated:

- *Mode 1*: Traditional scheduling with request and grant method (without CGs)
- *Mode 2*: Fixed amount of RBs (fixed CGs)
  - *Mode 2.1*: **Average** historical data rate used as amount of fixed grants (fixed **optimistic** approach)
  - *Mode 2.2*: **Maximum** historical data rate used as amount of fixed grants (fixed **pessimistic** approach)
- *Mode 3*: **Data-driven** CGs (predicted based on ARIMA)
  - *Mode 3.1*: Predicted CGs as is (**No over-provisioning**)
  - *Mode 3.2*: *Over-provisioned* predicted CGs (**With 10 % over-provisioning**)

The mode configuration as well as other simulation parameters like configured and simulated network slices are listed in Figure 5.45.

The three realized slices in the evaluation scenario are also depicted in Figure 5.46, which are defined as follows:

- *Smart Grid (SG) slice (uRLCC - Highest priority)*: The *Smart Grid* slice is modeled after photovoltaic systems transmitting data to regulate energy generation. The National Renewable Energy Laboratory (NREL)[2] provides open data for solar activity, which is used to train the ARIMA model and generate data traffic proportional to the solar activity.
- *Electric Vehicle (EV) charging slice (uRLLC - High priority)*: EV charging-point occupancy data of the German city Bonn[3] was gathered and data traffic based on this dataset is generated for the EV charging slice.
- *Best-Effort (BE) slice (eMBB - Low priority)*: A constant rate of 18.96 Mbps is generated, which corresponds to the remaining capacity of the cell, to simulate devices with high data rate demands and to measure the remaining data rate within the

2 See https://www.nrel.gov/grid/solar-power-data.html.

3 See https://new-poi.chargecloud.de/bonn.

General Settings

Channel Bandwidth	20 MHz
5G Subcarrier Spacing	15 kHz
Channel Quality	Fixed Modulation and Coding Scheme (MCS) of 15
5G MCS Index Table	64QAM
SR Occasion	every 4 ms
Packet TTI	1 ms
Simulated Time	1 h

Slice-Specific Settings

Smart Grid (uRLLC)	Electric Vehicle Charging (uRLLC)	Best Effort (eMBB)
5 UEs	4 UEs	2 UEs
variable aggregated throughput	variable aggregated throughput	18.96 Mbps aggregated throughput

Mode-Specific Settings

Mode 1	Mode 2.1	Mode 2.2	Mode 3.1	Mode 3.2
No Configured Grants (CG)	Fixed CGs (optimistic)	Fixed CGs (pessimistic)	Predicted CGs	Predicted CGs
No Overprovisioning (OP)	No OP	No OP	No OP	10% OP

Fig. 5.45: Settings and parameters of the simulation framework and the different modes utilized in the evaluation of the SAMUS framework. ©[2021] IEEE. Reprinted, with permission, from [54].

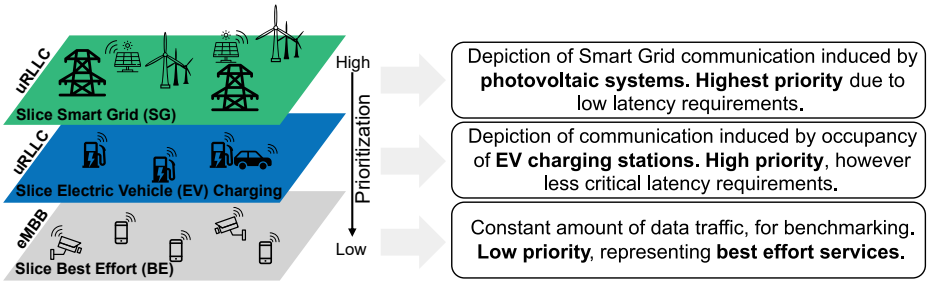

Fig. 5.46: Evaluation scenario comprising mission-critical and best-effort network slices to analyze different trade-off strategies within the SAMUS system. ©[2021] IEEE. Reprinted, with permission, from [54].

non-critical eMBB slices, after mission-critical slices are served by the SAMUS scheduler.

In the following section, the evaluation of the SAMUS framework is presented based on this scenario.

Evaluation Results The 5G-RGS framework described before was utilized to evaluate the SAMUS scheduler prototype based on the modes presented earlier, which represent different trade-off strategies between the balance of uRLLC latency and the eMBB data rate. For this, a 60 min interval was analyzed (cf. Figure 5.45), which represents a time frame of highly dynamic activity within the different slices such as the time of sunrise in the SG slice or the time of rush hour in the EV slice.

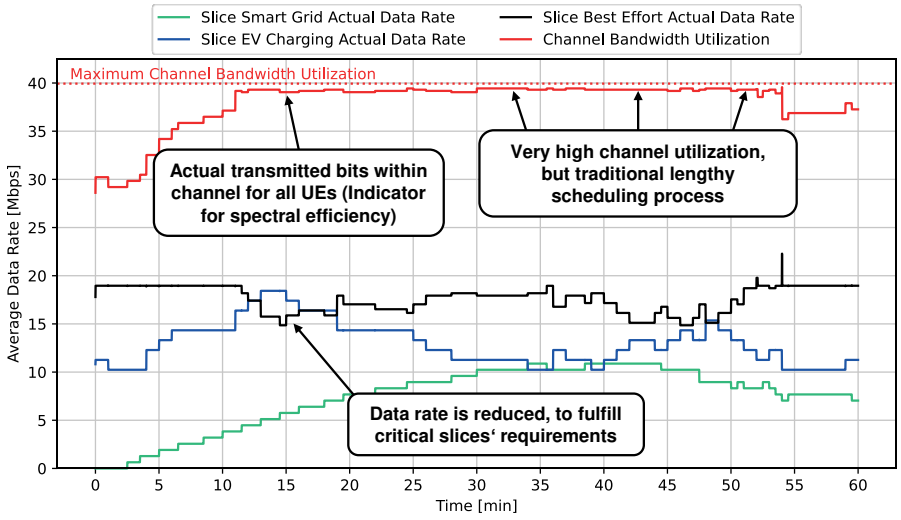

Fig. 5.47: Data-rate progressions for the different network slices in mode 1 (5G parameters only–traditional scheduling). ©[2021] IEEE. Reprinted, with permission, from [54].

In Figure 5.47, the results for mode 1 are depicted, where the average slice data rate in Mbit/s is plotted as a function of the simulation time in min. The dotted and solid red lines represent the maximum and the actual channel bandwidth utilization, respectively. This indicates the efficiency of resource usage, i.e., high channel utilization means low RB wastage. The green, black, and red solid lines represent the average uplink data rate transmitted for the SG, EV charging, and Best-Effort (BE) slices, respectively. Based on the *Greedy Network Slicing* method, it becomes clear that the available RBs are allocated to the higher priority slices at the expense of the BE data rate. This is the

desired behavior of the utilized traditional scheduling requests and grants, because RBs are distributed exactly as required and no resources are wasted. However, the main disadvantage of this approach is that it leads to very high scheduling latency (T_{Sched}). This connection becomes more clear when looking at Figure 5.48.

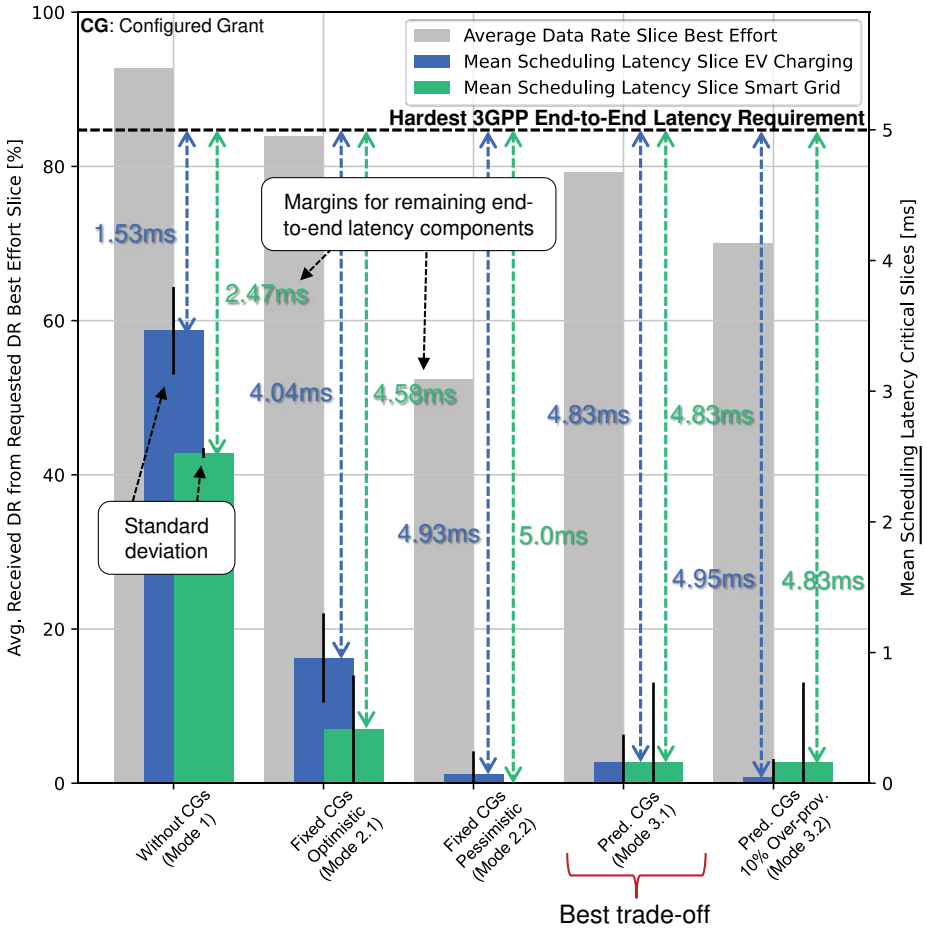

Fig. 5.48: Average BE data rates versus mean and standard deviation of high priority slice latencies (averaging window of 2 s and hardest 3GPP latency requirement according to 3GPP 23.501 [3]) for all modes. Margins for remaining latency components are indicated by the arrows in the respective colors of the slices. ©[2021] IEEE. Reprinted, with permission, from [54].

In this figure, two different y-axes are depicted describing the ratio of the average transmitted data rate from the actual data rate required by the BE slice on the left axis (gray bar plot) as well as the mean scheduling latency (T_{Sched}) of the high priority slices

on the right axis (green bar: SG slice; blue bar: EV charging slice; black lines: standard deviation), respectively. The different modes are listed on the *x*-axis. By looking at the results for mode 1, it can be seen that the high bandwidth utilization represented by the ratio of transmitted data to required data is very high at 92.74 %. At the same time however, only 1.53 ms and 2.47 ms margins for other latency components based on the hardest 3GPP end-to-end latency requirements [3] are left for the EV charging and SG slices, respectively. This results from the utilization of lengthy traditional scheduling mechanisms.

For comparison, the results of the modes in Figure 5.49 can be consulted. As for the modes 2.1 and 2.2, depicted in Figure 5.49a and Figure 5.49b, the channel bandwidth utilization drops for both approaches, especially for the pessimistic approach. This is the result of the fixed allocation of RBs to the mission-critical slices. However, the effect of this method on the latency becomes clear again with a look at Figure 5.48. There it can be seen that the margins for the end-to-end latency, especially for the pessimistic approach, increase to almost 5 ms, since the scheduling latency drops to almost zero due to the constant availability of resources. By contrast, the data rate efficiency of the BE slice drops down to 52.22 %. Thus, the fixed CGs represent a very latency-focused approach, whereas mode 1 maximizes channel utilization.

The data-driven ARIMA-based mode 3, which is the major contribution of the SAMUS framework, represents a good balance between these two extremes, as can be seen by looking at the data rates in Figure 5.49c, 5.49d and the latencies in Figure 5.48 for modes 3.1 and 3.2, respectively. Moreover, the channel bandwidth utilization is relatively high with an almost 80 % ratio of actual to requested data rate within the BE slice. Additionally, as data amounts of the mission-critical slices can be predicted very well, and thus, data can be instantly transmitted, margins for other latency components of 4.83 ms to 4.95 ms can be observed. The scheduling latency is zero most of the time.

5.4.4 Conclusion

In this section, we presented *5G network slicing* approaches for both the core network and the radio access network. While the same goal is pursued in both domains, the implementation is all the more differentiated. Especially in the RAN, machine learning-supported methods will be indispensable, since the prediction of upcoming data traffic is a prerequisite for implementing low-latency slices, while still maintaining high spectral efficiency. This relation was shown here in this section on the basis of our SAMUS approach, which is able to efficiently trade-off resources between different slice types.

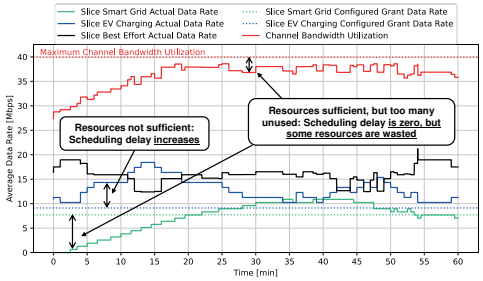

(a) Mode 2.1: Optimistic Fixed CGs

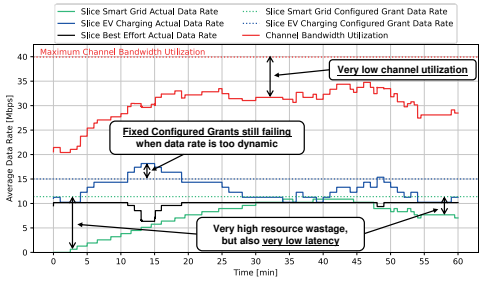

(b) Mode 2.2: Pessimistic Fixed CGs

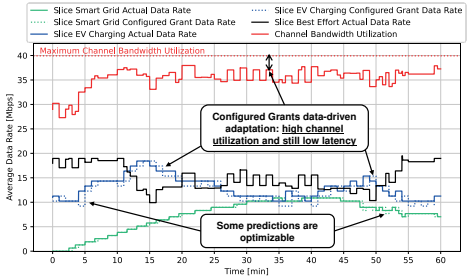

(c) Mode 3.1: Over-Provisioned Predicted CGs

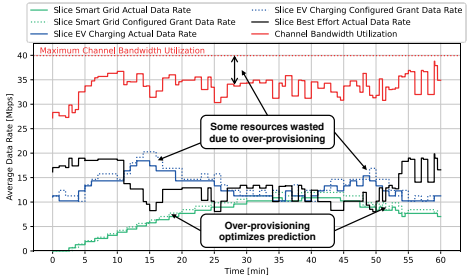

(d) Mode 3.2: 10 % Over-Provisioned Predicted CGs

Fig. 5.49: Comparison of Data Rates Within the Defined Slices and For All Modes Based on the Defined Evaluation Scenario. ©[2021] IEEE. Reprinted, with permission, from [54].

5.4.5 Acknowledgments

In addition to the CRC 876, part of his work has been supported by the Ministry of Economic Affairs, Innovation, Digitalisation and Energy of the State of North Rhine-Westphalia (MWIDE NRW) along with the Competence Center 5G.NRW under grant number 005-01903-0047 and by the Federal Ministry for Economic Affairs and Energy (BMWi) in the course of the project 5Gain under the funding reference 03EI6018C.

5.5 Potential of Millimeter Wave Communications

Karsten Heimann
Simon Häger
Christian Wietfeld

Abstract: For mobile communication networks, radio spectrum resources have always been a scarce commodity. With the cultivation of *millimeter Wave (mmWave)* wavelengths, a vast amount of spectrum at frequencies above 24.25 GHz has become available to serve the demands of *enhanced mobile broadband* services and applications of fifth-generation of mobile communications (5G). However, the higher carrier frequencies compared with the heretofore allotted spectrum comes with novel challenges for the operation of a cellular network: The more significant propagation losses require directional/beam antennas and their directivity needs to be adjusted permanently and individually per user. In addition, the poor obstacle penetration necessitates a careful beam alignment based on Line-Of-Sight (LOS) conditions. In case of obstructions, signal reflection paths need to be leveraged, which may be volatile and time-consuming to discover. By means of signal quality measurements, a self-contained beam tracking may maintain the LOS or virtual LOS via reflections to mobile devices. As a further feature, the directional knowledge of the base station antenna beams can even be exploited for a bearing-like localization approach allowing for an enhanced network positioning service compared with cell-level approaches. The sophisticated Software-Defined Radio (SDR)-based mmWave platform allows for the experimental evaluation of the mentioned features. The results prove the potential of mmWave communications for various vehicular and logistics use cases. The lessons learned will go into future research directions such as smart radio environments. The novel technology of Reconfigurable Intelligent Surface (RIS) is a promising strategy for improving the capabilities of the general environment to supply better radio conditions to a wireless channel in non-LOS conditions. For example, a RIS can purposefully redirect the base station's mmWave pencil beam to reach a device in an obstructed area and thus extend the network coverage. Future integrated, radar-like sensing capabilities of communication networks are expected to operate at mmWave frequencies due to large bandwidth, high directionality, and low multipath features promising high-quality measurements. We show that the channel information of current mmWave systems, beam orientation in particular, already enables novel sensing applications.

5.5.1 Introduction

Besides the growing number of Internet-of-Things (IoT) devices with low data rate but high coverage requirements and applications demanding reliable or low-latency data transfer, another main direction of impact of fifth-generation mobile networks (5G) focuses on the enhanced Mobile Broadband (eMBB) services. It is believed that future applications such as augmented, virtual, or extended reality necessitate a high-performance wireless network infrastructure.

While optimizing the utilization of the traditional *sub 6 GHz* radio spectrum in terms of spectral efficiency, this resource is already heavily used. However, with 5G the third Generation Partnership Project (3GPP) targets additional spectral resources in the mmWave domain (particularly from 24.25 GHz to 52.6 GHz in Frequency Range 2 (FR2)) [2, Table 5.1-1]. Frequencies in the THz domain will also be targeted in future mobile networks promising even larger bandwidths–and use thus a vast amount of resources. Although these resources ought to enable an enhanced throughput at the air interface, novel challenges arise due to the higher frequencies.

Unlike the popular misconception, the more severe path loss itself is not the main issue, because higher frequencies allow for an increased antenna gain within the same space constraints. With this, the path loss itself is more than compensated. Nevertheless, the increased antenna gain is achieved by a more distinct directivity, which demands a proper antenna alignment. Phased Array Antennas (PAAs) resolve that issue by interconnecting multiple antenna elements, so that a sophisticated superposition of the processed signals allows for an adjustable radiation characteristic known as *beamforming* or *spatial filtering*. A PAA applies phase shifts to the signals of the individual antenna elements. For example, at a Uniform Linear Array (ULA), N antenna elements are uniformly spaced with some distance d (mostly at half a wavelength, so $d = \frac{\lambda}{2}$). To create a beam directivity that points towards a direction θ, phase shifts of 0, ϕ, 2ϕ to $(N-1)\phi$ are applied to the respective antenna elements 0, 1, 2 to $N-1$ with $\phi = \frac{2\pi}{\lambda} d \cos\theta$ [39, Chapter 6]. Put simply, the number of elements N determines the beamwidth and antenna gain. In general, a larger N leads to a higher gain and a more focused beam.

This means that a steerable directivity is feasible and can be achieved electrically or by software. The transmitter antenna's *beam* can be dynamically aligned to a receiving antenna and vice versa facilitating radio propagation by high transmit-and-receive antenna gains. However, radio signaling as part of the control plane of the Radio Access Network (RAN) needs to carry out this alignment task in a timely manner, which could be challenging due to the volatile radio conditions and the users' mobility.

For example, the alignment could be performed by means of a potentially time-consuming discovery procedure such as *beam sweeping*. The coverable angular space is iteratively sampled by switching the beam through different pointing directions. In doing so, the beamwidth constitutes a trade-off between a higher gain and a reduced number of iterations required to sample the complete angular space. While a precise

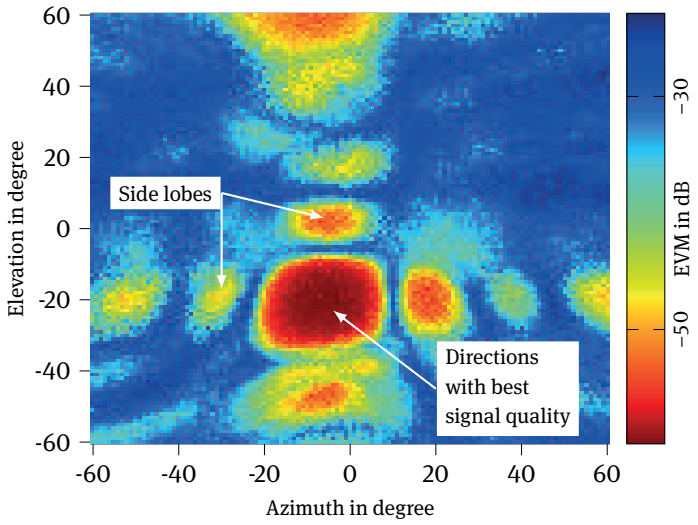

Fig. 5.50: Exemplary heatmap illustration of signal quality measurement during an exhaustive search *(beam sweep)*. The signal quality is given as Error Vector Magnitude (EVM) with lower values representing better signal qualities. The red area represents the beam-pointing directions with a suited signal quality. ©[2020] IEEE. Reprinted, with permission, from [268].

beam alignment is generally feasible in the analog domain, a number of quantized main lobe/beam-pointing directions as large as the number of antenna elements is often used to span an angular grid, where the selectable beams have the least possible overlap [148, Chapter 6].

Although the exhaustive sweep procedure can be accelerated by using multiple beams in parallel, each beam requires its separate RF-chain which are expensive with regard to their costs and energy demands. For this reason, it is believed, that analog or hybrid beamforming, where only a small number of parallel beams is available, is applicable for mmWave communications.

During a sweep, the measurements of the signal quality can be interpreted as a *heatmap*. Figure 5.50 depicts such a heatmap, with the most-suitable directions represented by the red spots. In a mobile network like 5G, the base station continuously transmits some reference signals at different beam directions in the downlink, so the User Equipment (UE) is able to select the strongest one, while performing a sweep with its receiving beam. Since such systems are defined for Time Division Duplex (TDD), channel reciprocity can be assumed and the UE can use the determined beam configuration for initially accessing the network and reporting back the suited beam direction pair. These directions can subsequently be used for further transmissions/receptions until the mobile device has moved or some obstruction occurs, which means the measured beam—dependent signal quality becomes outdated.

5.5.2 Beam Tracking for Interruption-Free High-Performance Communications to Mobile Devices

A proper beam alignment is as crucial to establishing a communication link as it is to maintaining it by tracking the mobile UE. The main drawback of the exhaustive search is its large search space and exploration time. Also, during this exploration, the beam points in various directions with weak signal quality, which may lead to a heavily reduced radio link performance or even a connection loss. For this reason, other procedures take into account a position or previous direction information and potentially the device mobility to facilitate an interruption-free utilization of the radio resources for purposeful data transmissions. This means, that once a proper alignment is initially discovered, *beam tracking* is preferably applied to follow the device movement. Only in case of a connection loss due to, say, sudden blockage, another comprehensive exploration might be required for radio link recovery.

In our works [269, 270], we analyze the applicability of beam tracking for supplying mobile users with mmWave radio links.

As a proof of concept, the position of a mid–flight drone/Unmanned Aerial Vehicle (UAV) is recorded by an optical reference system allowing for a geometry-based, precise calculation of the required beam-pointing direction. Figure 5.51 gives an overview of the experimental setup. The UAV movement describes an arc at a fixed distance of 1.8 m from the stationary active antenna/PAA. The central experiment logic controls this movement, processes the UAV position, sends corresponding beam-pointing commands and logs the measured performance indicators such as signal quality and data rate. In addition to the PAA's beam alignment, the passive horn antenna at the UAV can be aligned horizontally by means of the UAV's yaw rotation.

The evaluation results are condensed in the time-series graphs of Figure 5.52. When only the yaw rotation of the UAV is used to align the passive antenna at the UAV, the communication link is active only within a small range around the center direction, which is where the PAA is configured to point at in the static case. On the contrary, when only the PAA's pencil beam is continuously aligned towards the UAV, the misalignment of the horn leads to connection losses. Since the horn has a wider beamwidth, the tolerance for a misalignment is larger. Finally, when both transmitter and receiver antenna are continuously aligned to each other, a stable link is observed in terms of a constantly high data rate of about 2.8 Gbit/s. This proves the general applicability of mmWave communications utilizing PAAs for beam alignment in scenarios with mobile users.

Since external position knowledge might not always be available and could require an additional, beam-alignment independent control link (for example at a conventional sub-6 GHz band) for reliable reporting, a self-contained beam tracking approach based on signal quality measurements is evaluated in [270]. Besides keeping the beam at a direction with a still acceptable signal quality, better beam-pointing directions need to be explored during a temporary impairment of the signal quality. In general, there is a

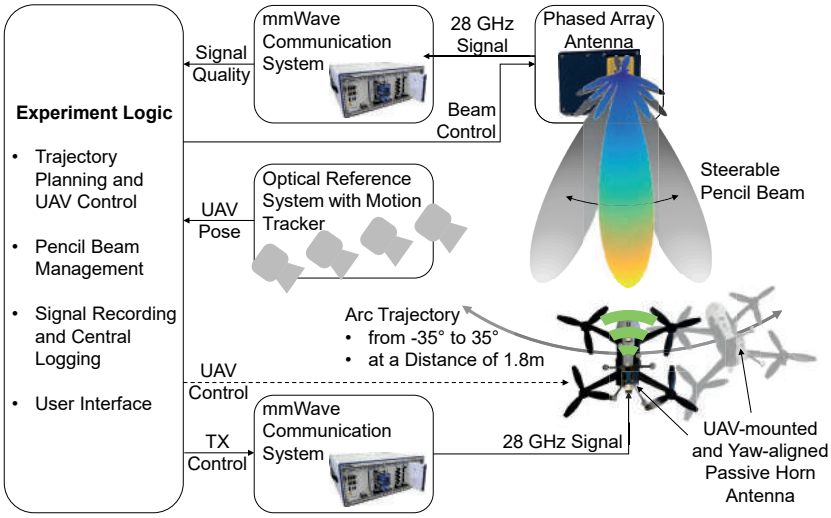

Fig. 5.51: Experimental setup for real-time mmWave beam alignment studies with a flying UAV based on [269]. The motion capture system provides position information of the UAV, which is processed to a beam-pointing direction and sent to the PAA as a control command.

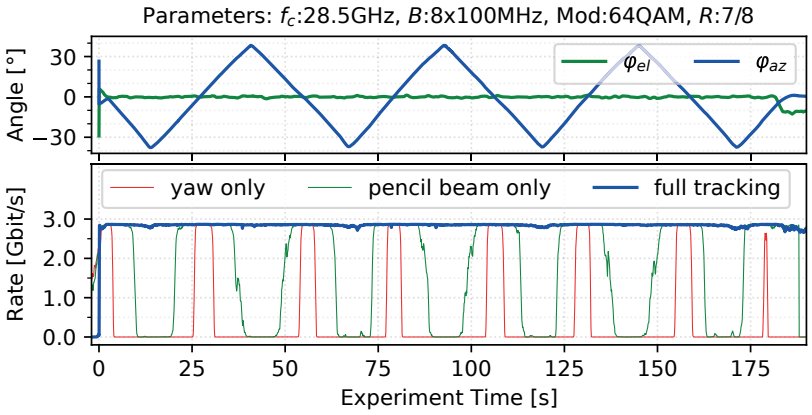

Fig. 5.52: Evaluation of mmWave beam alignment with a mid-flight UAV based on [269]. While connection losses occur without updating the directivity, a seamlessly high data rate can be achieved by tracking the movement of the mid-flight UAV with both transmitter and receiver antennas.

trade-off between the detail of exploration and the perceived signal quality by leveraging the aligned beam's gain, since the beam needs to be intentionally misaligned to explore the device's moving direction. Assuming that the device motion tracking under Line-Of-Sight (LOS) conditions requires only gradual changes in the beam-pointing direction, the search space can be substantially limited to the adjacent directions. For example in [270], a 3×3 grid in the azimuth and elevation plane of the angular space of beam directions is spanned centered at the last acceptable signal-quality direction. In doing so, every scanning cycle consists of as few as nine signal quality measurements and the subsequent search grid is centered on the direction with the highest signal quality. Although this appears to misalign beams in most cases, the small amount of measurements per cycle allow for low grid spacing below the beam width as long as the sample rate is significantly higher than the device's relative angular velocity. With this, the connection can still be maintained during the exploration. The minor reduction in antenna gain due to the slight misalignment can be compensated for by the communication system. For an experimental evaluation of this approach, the device motion is emulated in a reproducible fashion with a precise reference by using a rail system.

The statistical results of this empirical analysis is depicted as violin plots in Figure 5.53. While the signal quality is represented by the Error Vector Magnitude (EVM), where a lower value corresponds to a higher signal quality, the link performance is evaluated in terms of data rate. The emulated mobile device velocity is converted to the related maximal tracking dynamics from the antenna's perspective. A small exploration grid spacing of $\Delta = 1^\circ$ reduces the decline of the antenna gain due to a reduced misalignment only at low relative velocities, since this step size is not sufficient to keep track of the motion at higher dynamics. A larger spacing of $\Delta = 5^\circ$ deteriorates the link performance or may even lead to connection losses due to severe beam misalignments during exploration. Finally, the grid spacing needs to be fitted to both the device's velocity and the antenna's beamwidth. In the conducted test setup, a grid spacing of 3° empirically turned out to be a reasonable tradeoff. The results of laboratory evaluation thus prove, that a high–performance communication link can be maintained even for considerable device velocities.

With respect to the utilization of the novel radio resources at the mmWave domain, this beam tracking approach allows for efficient utilization of the spectrum by reducing the link outage due to lengthy exploration phases. In addition, the directional transmissions via *pencil beams* facilitate dense spatial reusability of these resources, since the interference within the mobile network is reduced.

As an outlook, future beam tracking techniques may incorporate reinforcement learning approaches to solve the exploration–exploitation tradeoff dilemma between a comprehensive exploration of beam-pointing directions with their associated signal qualities and a perfectly aligned beam with ideal signal quality conditions for data transmissions. In doing so, dynamic and reactive adjustments of the search space

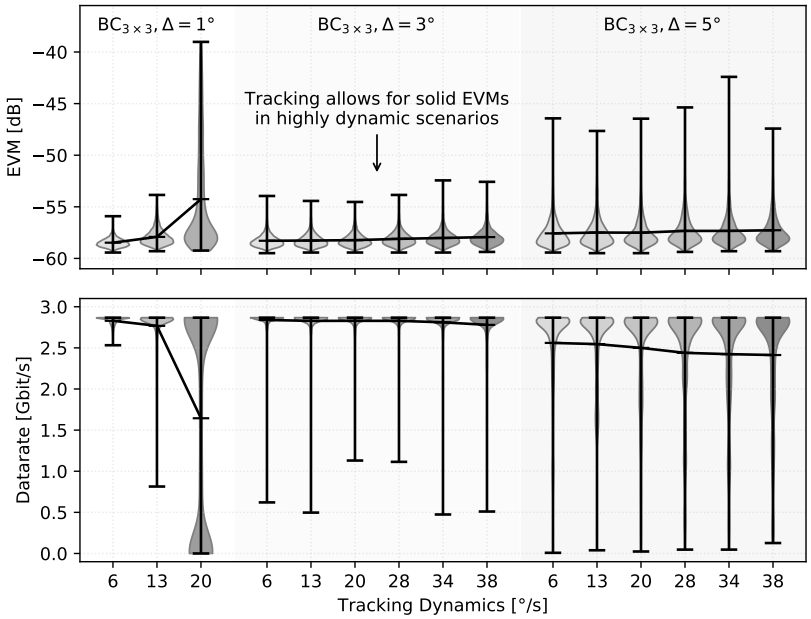

Fig. 5.53: Statistical beam-tracking evaluation. For exploring grid spacings Δ of 1°, 3° and 5°, the signal quality in terms of Error Vector Magnitude (EVM) and the link performance in terms of data rate is analyzed under different tracking dynamics. The tracking dynamics correspond to different mobile device velocities that are reproducibly emulated by a rail system. The constantly low EVM values at $\Delta = 3^\circ$ constitute a reasonable configuration with a stable link performance even at the highest tracking dynamics. ©[2019] IEEE. Reprinted, with permission, from [270].

(grid shape and spacing, for example) are conceivable according to the anticipated movement of the device.

5.5.3 Dual-use of Beam Alignment Information for Positioning of Mobile Devices

Although a proper beam alignment embodies a new challenge to mobile networks, once gathered the direction information could also be used for a bearing-based positioning service, as addressed in [268]. Conventionally, cellular network-based positioning utilizes signal-strength measurements in conjunction with propagation loss models, signal travel time, or propagation delay measurements (such as those used for the *timing advance* mechanism) for distance-based positioning or *lateration*. With the necessity for directional transmissions, angle-based methods utilizing direction information as a bearing are conceivable at the mmWave domain. By means of two or more intersecting bearings and known base station positions, a user-position estimate can be provided on

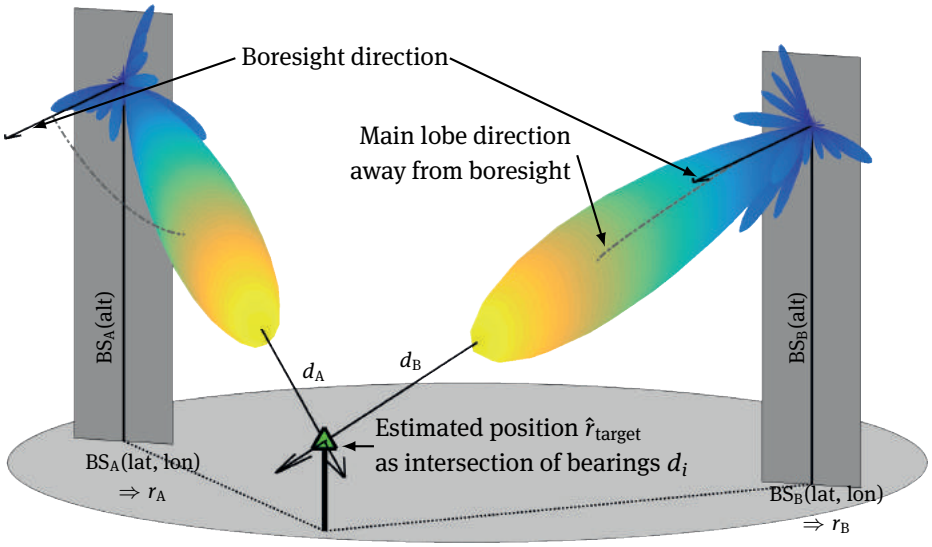

Fig. 5.54: A cross-bearing-based positioning utilizing mmWave beam alignment information. With known base station positions and beam directions as bearings, the mobile device position can be estimated as the intersection of the bearings. ©[2020] IEEE. Reprinted, with permission, from [268].

top of the ongoing wireless communication. In doing so, the accuracy strongly depends on the distance or constellation and the resolution of the direction finding.

Figure 5.54 illustrates the basic concept of this approach, which derives its origins from sea travel's cross-bearing. The estimated position $\hat{r}_{\text{target}}$ is defined as the position vector that minimizes the squared distance to the (two or more) lines spanned by the base station position vectors r_i and its beam-pointing direction d_i as direction vector. As a result, the position estimate is given as least squares approximation (with I as identity matrix):

$$\hat{r}_{\text{target}} = \left(\sum_i I - d_i d_i^T \right)^{-1} \left(\sum_i (I - d_i d_i^T) r_i \right) \quad .$$

From an exhaustive sweep, the direction vector d_i is estimated based on the beam-pointing direction with the highest signal quality. Since the area with reasonable signal quality turns out to be rather flat but noisy, this estimate is rather imprecise. For this reason, the centroid of the red region, which contains the highest signal quality, is taken as the improved direction estimate. In addition to the geometric location and orientation of the base station antennas, systematic deviations between the commanded and the actual pointing direction of the antenna beams are compensated. The experimental evaluation indicates the applicability of this approach and is depicted in Figure 5.55. In general, within the laboratory setup, a 3D Euclidean positioning error in the centimeter range is observed. The post-processing compensation for systematic deviations further

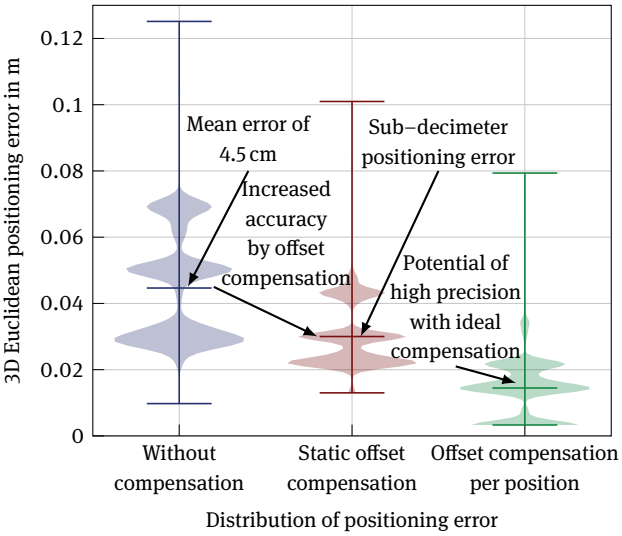

Fig. 5.55: Statistical evaluation of the positioning performance based on laboratory experiments. The 3D Euclidean positioning error lies in the centimeter range and can be further reduced by compensating systematic deviations in the direction estimates towards the UE. ©[2020] IEEE. Reprinted, with permission, from [268].

improves the estimated position and thus illustrates the potential of bearing-based mmWave positioning. Further details about the experiments can be found in [268].

The direction exploration as well as the post-processing compensation may be subject to machine learning-based optimization techniques introducing automated trade-off decisions between the resource utilization and positioning precision during runtime.

Nevertheless, hybrid procedures could combine distance measurements and bearings for a further enhanced positioning service of mobile networks. The large available bandwidth at the mmWave domain could be attractive for pseudoranging or Time Difference of Arrival (TDOA) considerations. The application of TDOA-based positioning utilizing the Ultra-Wideband (UWB) technology is analyzed in more detail in Section 3.5.

Due to the challenging propagation characteristics at the mmWave domain, a dense deployment of mmWave base stations is required and may lead to an enhanced system performance by utilizing approaches such as Coordinated MultiPoint (CoMP) or Dual Connectivity (DC), so the connection of one UE with multiple base stations at a time. Within a mobile network, the proposed positioning mechanism can be applied for both the downlink as well as the uplink direction. While the Angle of Departure (AoD) of the base station downlink beams could be signaled to the UE together with a map of base station positions to perform the positioning at the UE, the Angle of Arrival (AoA) of the base station uplink beams could be utilized to perform the task on the network side.

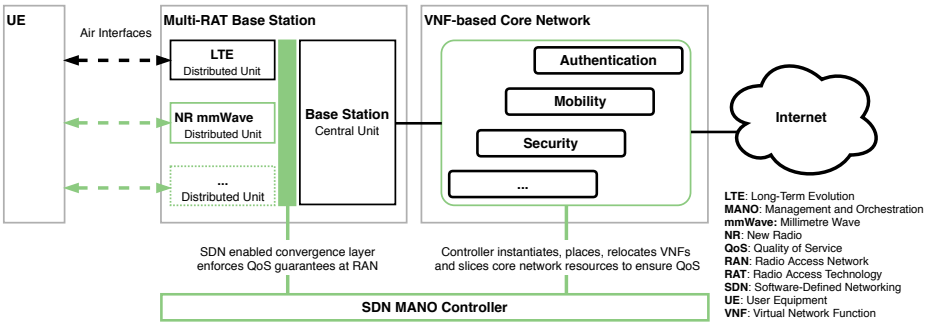

Fig. 5.56: End-to-end system architecture design for aggregated Radio Access Technologies (RATs) with prioritized services. The slices with different service requirements are orchestrated by the Software-Defined Networking (SDN) Management and Orchestration (MANO) controller. ©[2019] IEEE. Reprinted, with permission, from [266].

With the former providing reportedly sensitive information of base station positions to the UE, the latter alternative requires an active transmission of the UE, but no additional utilization of radio resources/signaling overhead leading to a resource-efficient positioning solution. In both cases, a preferably accurate direction estimate is required for positioning, which might be feasible only in case of a high-resolution sampling of the angular exploration space. Additionally, a Dilution of Precision (DOP) can be observed at acute angles between the intersecting bearings, so an elaborate placement of base station antennas might be advantageous.

5.5.4 Integration of High Priority mmWave Links into an End–to–End System Architecture

As part of 5G, the mmWave spectrum contributes to the available resources at the Radio Access Network (RAN). End-to-end applications between users (humans or machines) and services come with various requirements, which differ greatly from each other. Within an end-to-end system architecture, a mobile network needs to be agile and utilize the available resources at both, the core network and the RAN so that the application requirements can be fulfilled. As already elaborated in Section 5.4, network slicing is introduced to define virtual networks with certain configurations regarding throughput, latency, reliability, and others. Based on this, each application is dedicated to a specific *slice* that not only supplies the required performance, but also remains unaffected by traffic fluctuations or shortcomings of other slices of the same mobile network. To illustrate the potential of this slicing, our work [266] presents a system concept and an experimental evaluation of an unaffected and prioritized communication link among other *best-effort* traffic.

The overall system architecture design is depicted in Figure 5.56. To ensure the Quality of Service (QoS) requirements, the proposed Software–Defined Networking (SDN) Management and Network Orchestration (MANO) controller affects both RAN and core network. At the RAN, multi-RAT base stations are capable of directing the data traffic flow to and from UEs through different air interfaces according to the guidelines from the SDN MANO controller. For example, a conventional LTE link can be used in parallel with a 5G New Radio (NR) link at the mmWave domain at different Distributed Units (DUs) of the same base station. At the same time, the base station Central Unit (CU) is connected to the core network, where Virtual Network Functions (VNFs) dynamically allocate resources as required to operate the appropriate services. Finally, this design ensures end-to-end QoS within the whole mobile network. An experimental proof of concept study can be found in [266], where the Software-Defined Radio (SDR) and SDN-based components of the experimental setup allow for high flexibility and adaptability.

5.5.5 Intelligent Reflectors for Enhanced Propagation and Coverage under Non-Line-of-Sight Conditions

In addition to the discussed propagation loss and the need for directional transmissions, mmWave signals barely penetrate materials. As a consequence, the outdoor-to-indoor coverage is rather poor and obstructed areas need to rely on the presence of suited reflection paths. These reflection paths in turn are volatile and need to be explored by means of a potentially time-consuming discovery procedure such as the aforementioned *beam sweeping*. The beam management needs to provide routines to recover from link blockages and to switch between propagation paths, whenever the LOS condition varies. Taking the NLOS propagation into account, several challenges arise regarding the mobility support, which is doubtless a crucial feature of mobile radio networks. However, especially in dense urban scenarios, frequent LOS obstructions may demand sophisticated procedures to facilitate radio links via reflection paths.

With the novel concept of *smart radio environments* and the Reconfigurable Intelligent Surface (RIS) technology, the radio channel itself becomes modifiable to enhance the transmission performance. While much research concentrates on the optimization of transmitter and receiver techniques, the idea of this concept is to deploy elements (surfaces) with controllable reflection characteristics in the environment. Hence, these RISs act as dynamically controllable passive reflectors. In this way, they enable the purposeful utilization and adjustment of reflection paths allowing for an enhanced tracking capability of user devices with an obstructed line-of-sight to the base station. [46]

In our work [267], we highlight the potential of RISs for an enhanced mmWave network coverage at an urban campus scenario, as illustrated in Figure 5.57. The simulation model is based on our previous work [631] and extended to also account for RIS.

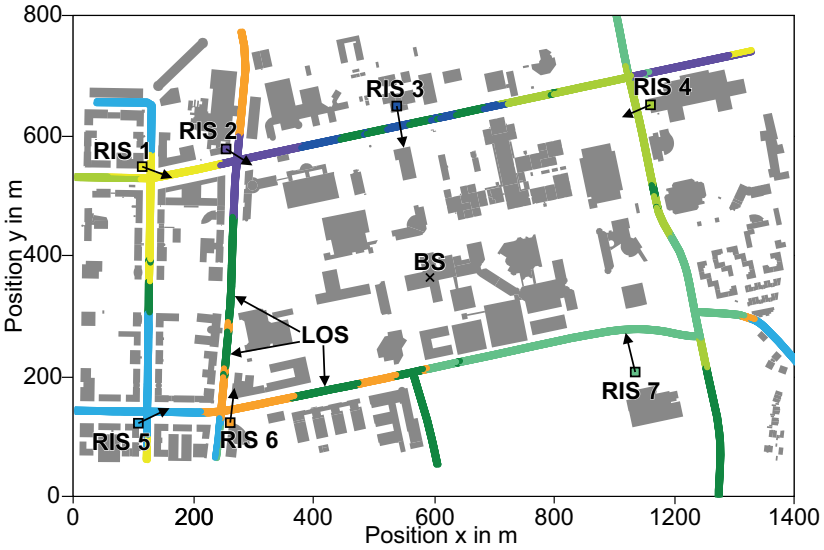

Fig. 5.57: Simulation scenario for a RIS-enhanced coverage study. With LOS coverage marked in dark green, colors according to the RISs are used for the road sections with NLOS conditions, which are covered by the respective RIS reflected paths. ©[2020] IEEE. Reprinted, with permission, from [267].

As depicted in Figure 5.58, the base station deployment only leads to a poor LOS coverage. The corresponding path loss lies only 23 % within the expected link budget of 142 dBm. However, the utilization of RIS reflection paths enhances the overall network coverage. For the alignment of the RIS reflection, distorted information about the target location is assumed. The true UE position is superimposed by a zero–mean normal distribution with standard deviation σ. The Empirical Cumulative Distribution Function (ECDF) of the path loss illustrates, that a coverage of 91 % is achievable in the case of $\sigma = 3$ m, for example. Even a comprehensive coverage is feasible due to the RIS placement in this evaluated scenario.

As a result, the deployment of RISs for smart radio environments may not only lead to an enhanced network coverage; it also allows for an improved efficiency in terms of energy and spectral resources, since the controlled reflections may reduce the exploration overhead of beam management algorithms as well as the required transmit power for sufficient signal strengths at the receiver. Nevertheless, an elaborate control of the RISs may require some radio resources for both measurements and signal suitable reflection paths. Finally, the RIS placement task may play a part in the network planning procedure and could be realized by means of machine learning approaches focusing on a cost–efficient way of providing a comprehensive network coverage.

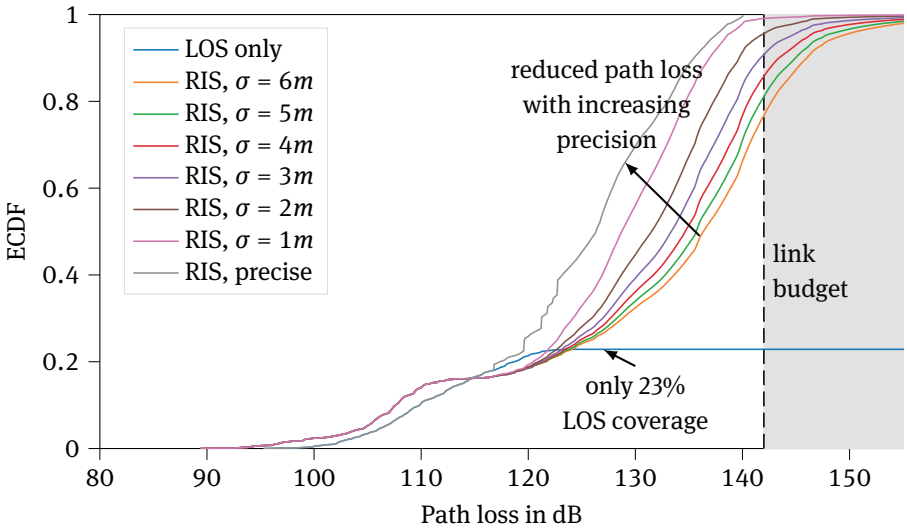

Fig. 5.58: Simulation results of RIS-enabled communication via reflection paths. Even with a slightly distorted target position information (zero-mean normal distribution with standard deviation σ), a comprehensive coverage of 91 % is achievable for $\sigma = 3$ m, while the pure LOS coverage amounts to only 23 % in this scenario. ©[2020] IEEE. Reprinted, with permission, from [267].

5.5.6 Towards Perceptive mmWave Networks by Channel Sensing

Over the past decade, it has been shown that radar and communication functionalities may be provided by Joint Communication and Radar/Radio Sensing (JCAS) systems, because the employed OFDM waveform of current 4G/5G networks and WLANs is also suitable for radar services [657]. Nowadays, a deep integration of radar-like sensing services is expected for 6G, thus allowing communication networks and their entities to become perceptive of the immediate surroundings. Such information may then be used to optimize the network performance by, say, supplying mmWave beam sweeping and tracking algorithms with user position and mobility information. Such information may also be used in the sub-6 GHz band, e.g. to assist handover decision making.

Radar systems are capable of detecting targets by analyzing the reflected waves of its own transmit signals over time. Through the use of large bandwidths and sweeping of highly directional antennas, it is possible to estimate distance, velocity, and angle information of the detected targets with high accuracy. Typical radar systems operate at very high frequencies, for example mmWaves or beyond, where multipath-based distortions are mitigated such that the high resolution due to bandwidth and directionality comes to fruition. Moreover, the imaging of the surroundings is also possible with radar technology. Considering the particular compatibility between radar requirements and mmWave communications, this is an opportunity for network operators to offer new

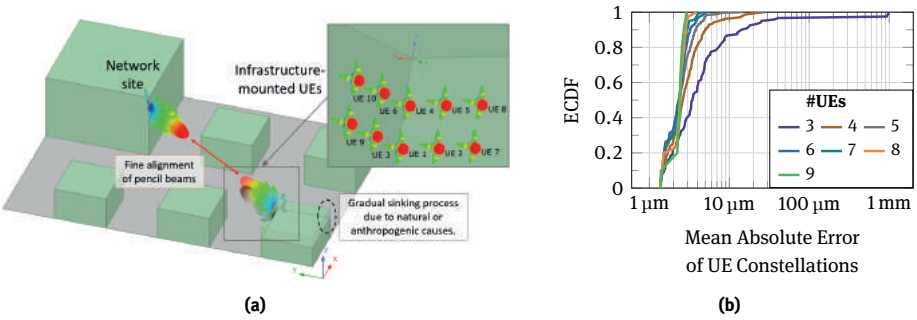

Fig. 5.59: *(a)* Scalability analysis of a subsidence process (gradual sinking of up to 5 mm) affecting a suburban house. *(b)* Distribution of mean incurred error throughout Δz-traversal for various UE mounting setups. ©[2021] IEEE. Reprinted, with permission, from [253].

services to the public, such as sensing-assisted traffic, but it also offers the prospect of process optimization in industrial facilities employing private network solutions.

However, the integration of radar functionality into mmWave communications still has a long way to go. For example, there is a need for hardware and signal processing enhancements. Nonetheless, radio-based sensing features such as user positioning have been available for more than two decades and steadily been enhanced ever since. Such sensing is enabled by analyzing the properties of one or more channels between the network and the user equipment entities. A large number of channel-based services, such as vehicle detection and classification (cf. Section 4.2), have already been proposed in literature [753]. These have been designed predominantly for sub-6 GHz WLANs, yet led to the recent launch of IEEE 802.11bf *Wi-Fi Sensing* standardization which even pertains to the mmWave domain.

With the inclusion of mmWave frequencies into the 5G standard, 5G positioning was successively adapted to allow the facilitation of mmWave beam information which, for example, enables angle-based positioning, thus enhancing the network's location services by new methods (cf. Section 5.5.3). Our work [253] followed a similar approach and considered the use of pencil beam orientation information to enhance traditional channel phase tracking-based measurements of relative motion and vibrations. By combining the movement information of several UEs along the LOS path beam orientations, we showed that millimeter range motions may be reconstructed in 3D space with less than 10 µm error. (See Figure 5.59 for the detailed results of a sample scenario.) Our full scalability analysis suggests that the usage of 4 to 5 distinct spatial link opportunities, which are expected in typical urban deployments between a single TX-RX pair, is a sensible choice. Therefore, high accuracy 3D motion tracking could in the future be conducted with a single-user device exploiting several distinct propagation paths to the network. Our ongoing work is evaluating the achievable beam orientation accuracy and the consequences of misalignment in the prior two contexts. Nonetheless, by showing

that future sensing features may already be realized with current mmWave technology, we point out the need for more research on this area because such techniques may also allow mobile networks to become more perceptive of its surroundings.

5.5.7 Concluding Remarks

In this section, the potential of mmWave communication has been elaborated with insights into several promising areas of research. The novel mmWave spectrum for mobile networks embodies a great opportunity for various future applications due to the vast amount of available spectral resources as well as the peculiar radio channel conditions. The demand for directional communication offers less interference and better spatial reuse of the spectral and time resources. According to the results of our presented works, the challenge of a proper beam alignment appears manageable and beam-based positioning can be provided as an additional feature. Also, there are concepts for integrating the novel spectrum into the overall mobile network capacity in terms of network slicing as shown in terms of a systems perspective. Last, the field of application will be further enhanced by the introduction of the novel concept of smart radio environments, where RISs support the propagation of mmWave beams and thus enhance the comprehensive network coverage in obstructed areas. Via an outlook on future perceptive networks, we showed that current networks could already partially enable novel sensing applications as expected for 6G. Therefore, research should further investigate and test sensing techniques that employ mmWave beam orientation information alongside the ongoing development of a 6G JCAS framework.

5.5.8 Acknowledgments

In addition to the CRC 876, part of this work has been partially supported by the Ministry of Economic Affairs, Innovation, Digitalization and Energy of the State of North Rhine-Westphalia (MWIDE NRW) along with the Competence Center 5G.NRW under grant number 005-01903-0047.

6 Privacy

6.1 Keynote: Construction of Inference-Proof Agent Interactions

Joachim Biskup

Abstract: To comply with the social issue of preserving privacy or pursuing other confidentiality requirements, we outline a broad range of conceptual solutions to a task of computing engineering: configuring the formal interactions of an individual's information system agent with the client agent of a communication partner in an inference-proof manner. Here inference-proofness means the following. A security mechanism shielding the system agent under the individual's control is reducing the information content of the messages sent to the client agent such that the partner would not be able to learn any information to be kept confidential under the individual's confidentiality concerns. This goal has to be provably guaranteed even if the communication partner as a rational reasoner will exploit not only a priori knowledge about the application underlying the communication acts but also additional background knowledge comprising both a complete specification of the interaction semantics and the full awareness of the security mechanism.

6.1.1 Foreword: Intended Audience

This contribution gathers, unifies, clarifies, and explains in depth the concepts and insights of a dedicated line of research and development within one of the basis subfields of IT-security, namely user-centric, self-determinative, and computer-supported enforcement of confidentiality interests, including the preservation of privacy at the discretion of the individuals involved. The own contributions started around two decades ago (see the brief bibliographic notes in Section 6.1.7), at the beginning inspecting the evaluation of sequences of closed queries by a database management system, such as a relational database system, the abstract semantics of which are based on a fragment of first-order logic. Understanding IT-security as a comprehensive problem of both organizational and computational issues, over the time, it becomes more and more demanding to expand to further operations like, e.g., transaction management for mixed query and update operations by more expressive information processing systems. At the end, dealing with procedural program execution as a service of any kind of knowledge- and belief-management system would be the ultimate goal.

The broader the range of particular operations by specific computing systems, in each case treated by appropriate highly sophisticated means, the more urgent the

https://doi.org/10.1515/9783110785982-006

need to identify useful abstractions for the computational issues and to reconsider the supportive organizational issues becomes. Regarding the first point, we abstract the object of protection to be the epistemic state of an intelligent computing agent participating in a multi-agent computing configuration. Regarding the second point, we explicitly expand all the intuitive assumptions underlying the various computational protection efforts into a framework of eight basic features, to which then the formal notion of the kind of protection we want to mathematically verify refers.

Experts in the fundamentals of IT-security might benefit from this article by learning the carefully elaborated essence of a large number of highly specialized publications. Computer scientists with a broader expertise in the field of confidentiality enforcement might be encouraged to generate a similar retrospective of their own line of research and, maybe, to fill some of the many gaps in the list of operations already treated. Researchers working on machine learning and embedded systems with a strong interest on security issues might also be triggered to fill those gaps for which they have appropriate expertise. Other researchers working on machine learning and embedded systems might gain detailed exemplary insights into the subtleties of integrating purely functional aspects with concise security considerations. They might further reflect on the related notions of (syntactic) data on the one hand and inferred (semantic) knowledge and belief on the other hand underlying their own work, and they might consider the design of an overall system architecture of their interest where the security measurements are appropriately located. Finally, admitting that this article deals only with a possibilistic version of confidentiality, all kinds of readers might think of and contribute to generalize the entire approach to probabilistic considerations.

6.1.2 Confidentiality-Preservation and Inference-Proofness

Since time immemorial, among many other activities, and in a closely intertwined manner, people have reasoned as individuals by acquiring, structuring, keeping, and exploiting *information* to make up their respective minds and behaved as social creatures by *communicating* with others. With the advent of computing technologies, individually dealing with information and socially communicating have been partly delegated to *computing agents*. On the one hand, the delegation is meant to facilitate routine tasks or even enhance human capabilities.

On the other, depending on the context, as delegators, individuals at their discretion or groups of them according to some socially accepted norm aim to still control the computing agents executing protocols as their delegatees, or at least the human delegators should appropriately configure the computing delegatees.

Being aware of the resulting reduction, and somehow simplifying, we can map concepts of human reasoning and communication to the inference protocols and interaction protocols of their computing agents and, correspondingly, actually performed human activities to protocol-complying computing process executions. Under such a

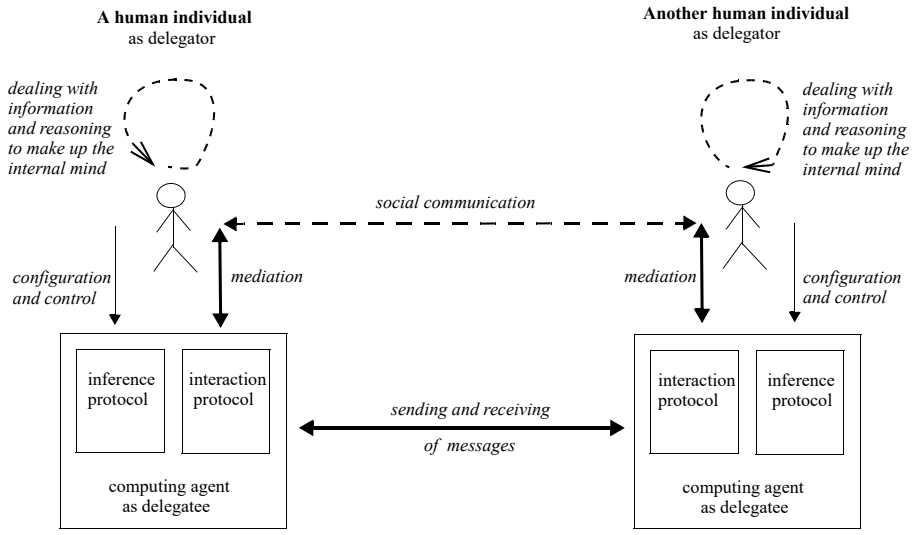

Fig. 6.1: Two reasoning and socially communicating human individuals and their protocol-based interacting computing agents as part of a larger community and the corresponding multi-agent computing configuration.

reduction, and even more simplifying, a group of human individuals is modeled to be complemented by a *multi-agent* computing configuration. In this model, each *human individual* controls a dedicated computing agent that, at least partly and by means of *protocol executions*, both deals with the information owned by that individual, in particular by internally deriving an epistemic state from a chosen information representation, and mediates the communications of that individual, in particular by sending and receiving *messages* according to one or more agreed interaction protocols. Figure 6.1 illustrates the sketched scenario.

Though, in principle, each individual can act in diverse roles and, correspondingly, each controlled computing agent can execute diverse protocols, we further specialize the model sketched above in focusing on only two individuals together with the respective computing agents. One individual is seen as an *information owner* controlling an *information system agent*, and the other individual is treated as a cooperating *communication partner* employing a *client agent*. Moreover, to enable *cooperation*, in principle the information owner is willing to *share information* with the communication partner. However, complying with *privacy issues* or pursuing other *confidentiality requirements*, as an exception from sharing, the information owner might want to hide some specific *pieces of information*.

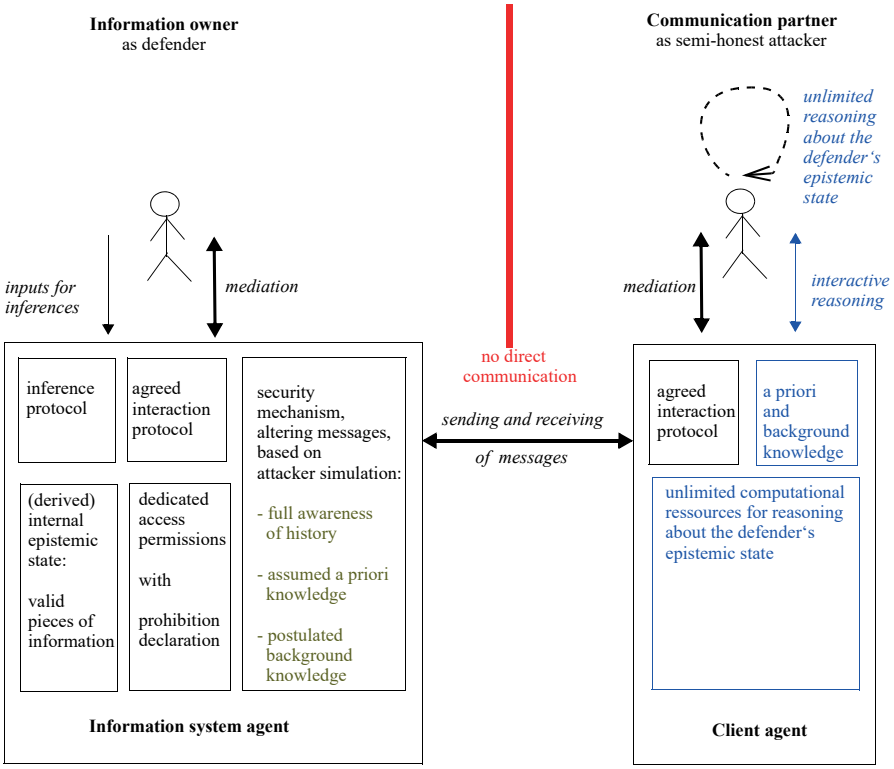

Fig. 6.2: The framework of a defending information owner with his information system agent and an attacking communication partner with his client agent (showing suspected but inaccessible parts in blue). See Figure 1 on page 82 of [73], © IFIP/Springer 2020.

Slightly more concisely, and visualized in Figure 6.2, we assume the following *framework* with eight *features*.

1. [Epistemic state of the information system agent as a single object of protection.] The human information owner does not deal with information processing and reasoning by himself but only provides the inputs to the information system agent under his *control*. At each point in time, that agent is internally deriving a formally defined *epistemic state*.
2. [Mediation of human communications by interacting computing agents.] Once having agreed on cooperation, the human information owner and his human communication partner do not communicate directly with each other, but only mediated by the computing agents under their respective control.
3. [Dedicated access permissions for information sharing.] Independently of the actual epistemic state, the information owner has granted dedicated *access permissions* to his communication partner. These permissions

declare that over the time the client agent of the partner may *interact* with the information system agent of the owner following some explicitly chosen *interaction protocols* that exclusively refer to the internal epistemic state of the information system agent.

4. [Exceptions by explicit prohibitions designating pieces of information.]
   Also independently of the actual epistemic state, the information owner has explicitly *declared* exceptions from the dedicated access permissions in the form of prohibitions. Each *prohibition* specifies a piece of information that the communication partner should not be able to learn. More precisely, with each prohibition expressed in terms of the information system agent and thus in reference to possible epistemic states, the communication partner should never be able to become sure about the actual *validity* in the epistemic state of the information system agent. In other words, from the partner's point of view it should always appear to be *possible* that the prohibited piece of information is *not valid* in the epistemic state of the information system.
5. [Partner suspected to reason about the validity of prohibitions.]
   Though the client agent is restricted to follow the interaction protocols of the access permissions exactly, the human communication partner can choose any sequence of permitted commands. Moreover, the partner is assumed to have unlimited computational resources when rationally reasoning about the validity or non-validity of a prohibited piece of information.
6. [Security mechanism implanted in the owner's information system agent.]
   To actually enforce the confidentiality requirements of the information owner, the information system agent is enhanced by some implanted *security mechanism* that should shield the underlying information processing from direct contact with the client agent. That security mechanism first inspects each message to be sent by the information system agent to the client agent according to the pertinent interaction protocol whether a *violation* of the information owner's confidentiality requirements would be enabled on the side of the communication partner. If this is the case, the security mechanism then *alters* the message such that the message is still as informative as possible but all options for a violation are blocked.
7. [Reasoning supported by a priori knowledge and background knowledge.]
   First of all, the communication partner's rational reasoning about the internal epistemic state of the information system agent is based on the messages exchanged by the respective computing agents. These messages are completely known to both agents. Additionally, the partner's rational reasoning is presumed to be supported by some *a priori knowledge* about the application dealt with in the cooperation between the two individuals involved and additional *background knowledge* comprising both a complete specification of the *interaction semantics* and the full *awareness of the security mechanism* (possibly even including the prohibition declaration) and, most notably, nothing else.

8. [Principle inaccessibility of the partner.]
   The internals of both the human communication partner and his client agent are considered to be principally inaccessible for the information owner and his system agent. This implies that the latter ones can only rely on *assumptions* about the details of the a priori knowledge and a *postulation* about the background knowledge available to the former ones.

These features can still be formally instantiated in various ways. In all cases, we follow a martial-sounding but common terminology, which ignores that in many scenarios an individual involved as communication partner will primarily be treated as cooperating in a friendly manner, rather than as a "total enemy". At least partially trusted for consciously sharing information in principle and correctly executing the agreed interaction protocols, the *communication partner*—together with the *client agent* controlled by him—is denoted as a *semi-honest attacker*, suspected to potentially aiming to maliciously infer the actual validity of pieces of information that the information owner has declared to be kept confidential. Accordingly, the *information owner*—together with the *information system agent* controlled by him—is denoted as the *defender*.

Now, the security mechanism has to invariantly enforce a suitable version of the following *security policy* of (possibilistic) *inference-proofness*, which also specifies the *attacker model*: For each prohibited piece of information ψ, the information content of messages sent to the attacking client agent—which is possibly enhanced by reasoning capabilities supplied by the human communication partner—during executions of agreed interaction protocols will never enable the attacking receiver to rationally infer that ψ is valid in the epistemic state, even when

- inspecting the complete *history* of preceding *interactions*,
- considering some *a priori knowledge* about the possible *epistemic states*,
- applying the *semantics* of the agreed *interaction protocols*, and
- being *aware* of the *functionality* of the *security mechanism*.

The concept of *rationality* on the side of the attacker is then captured by the following rephrasing of the still to be suitably versioned security policy of *inference-proofness* in terms of *indistinguishability* (as roughly visualized by Figure 6.3):

For each prohibited piece of information ψ,
for each epistemic state d satisfying the a priori knowledge,
for each sequence of messages $mes_1, \ldots, mes_k$
  exchanged during an interaction history
  complying with the agreed interaction protocols
  but potentially altered by the security mechanism
there exists an "*alternative*" epistemic state d' such that

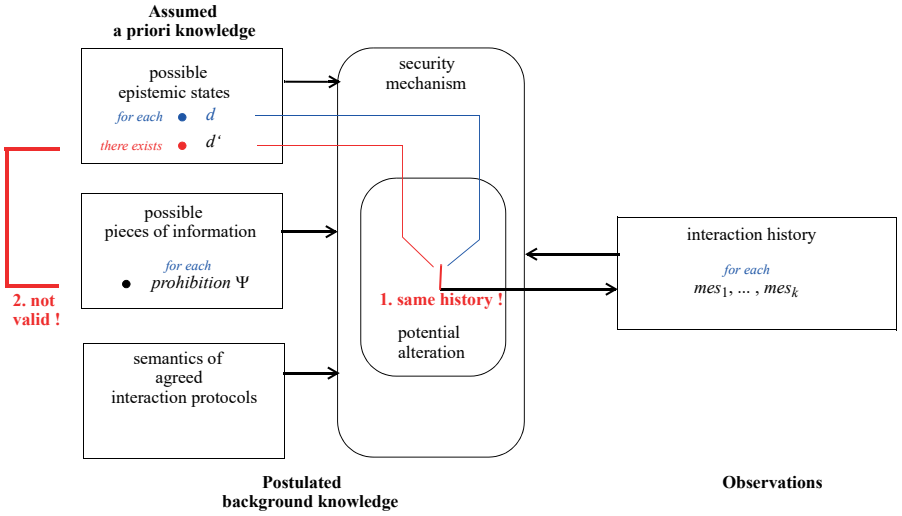

Fig. 6.3: A rough visualization of the requirement of inference-proofness: "for each ψ, for each d, for each $mes_1, \ldots, mes_k$, there exists d' such that ...".

1. the same sequence of messages would be generated, in compliance with the agreed interaction protocols and subjected to the alterations by the security mechanism, but
2. ψ is not valid in d'.

For this rephrasing, the epistemic state d is thought as actually derived by the information system agent—for short, "stored"—and might satisfy the prohibited piece of information ψ or not. The former case implies that the alternative state d' required to exist is different from d; in the latter case, the actually stored state d and the alternative state d' might be the same. Accordingly, declaring ψ as a prohibition does not intend to block any option that enables the attacker to infer the non-validity of ψ.

Confidentiality as inference-proofness could be trivially achieved by granting no access permissions at all or altering the information content of all messages sent to the attacker to nothing, violating any conflicting *availability requirements* and shutting down any communication mediated by the respective computing agents and, thus, making the whole thing useless. Accordingly, confidentiality requirements and availability requirements always have to be suitably balanced.

All our work focuses on the following three-level *conflict resolution* strategy:

1. As a general rule, some dedicated access permissions are granted for the sake of *availability*, to be freely enjoyed by the client agent, insofar as they do not conflict with level 2 of the strategy.

2. As exceptions from the general rule, specific prohibitions are declared for the sake of *confidentiality* that never must be violated and, thus, these prohibitions have to be enforced by alterations made by the security mechanism, but to comply with level 3 of the strategy only insofar as definitely necessary.
3. As a limitation for the effect of exceptions, the alterations made have to be minimal, again for the sake of *availability*.

Given the access permissions on the first level, the second and the third level lead to a combination of a *constraint solving problem* and an *optimization problem*.

Our simplified defender-attacker agent model still allows many instantiations and versions, respectively. For our concrete ongoing research, and accordingly for this article, we grossly distinguish the *structures* of three fundamental *data types* for an *information system*, namely *abstract* data sources, *propositional* knowledge or belief bases, and *first-order relational* databases.

For each of these data types—considering suitable refinements—we deal with the pertinent *operations*, which in our case are the *interactions* with a client agent, comprising in all cases at least

– *closed-query /yes–no-query evaluation* with *response* preparation,
  performed repeatedly with queries that in general are different;

and, depending on the refinement, additionally

– *open-query evaluation* with *response set* preparation,
  performed repeatedly with in general different queries;
– *view generation*,
  performed only once, since the attacker can freely employ a received view at his discretion instead of contacting the defender again;
– *view updating*,
  possibly performed from time to time, if manageable at all;
– *knowledge update transaction*,
  performed repeatedly usually intertwined with queries;
– *belief revision*,
  performed repeatedly usually intertwined with queries evaluated under *non-monotonic reasoning*;
– procedural *program execution*,
  performed repeatedly with in general different input parameters;
– *data outsourcing*,
  performed only once.

Tailored to the respective refinement, we propose and study *alterations* to the available interactions to ensure inference-proofness. There are two basic approaches to alterations, namely *weakening* the pertinent information about the actual epistemic state

and *lying* about the pertinent information about the actual epistemic state, allowing or even requiring appropriate refinements and *combinations* of weakening and lying.

6.1.3 A Generic Construction Methodology for Alterations

We can design and verify basic kinds of alterations according to the following generic construction methodology, which refers to the eight features of the framework described in Section 6.1.2. According to Feature 1, at each point in time, the defending information system agent is privately deriving its *actual* epistemic state. According to Feature 3, this state is then taken as the basis for the data contained in the messages to be sent to the attacking client agent during an interaction complying with some agreed protocol, for which some dedicated access permissions are granted. According to Feature 5, the attacker is suspected to aim at gaining as much information as possible about the defender's actual epistemic state, in particular whether a prohibited piece of information is valid in that state. In general, however, the attacker will face some uncertainty about that state, since (i) by Feature 2, there are no direct communication acts between the human individuals involved and, (ii) by Feature 6, the attacker has been separated from the defender's underlying information processing by the shield of the implanted security mechanism and the interactions are restricted to the exchange of messages. According to Feature 4, the attacker's uncertainty should always include that any prohibited piece of information might be *not* valid in the defender's actual epistemic state.

Conceptually, the attacker's uncertainty can be captured by the set of those epistemic states that appear to be *possible* to him. According to Feature 7, an epistemic state qualifies to be possible if it is compatible with both the potentially altered messages observed so far and the already initially available a priori knowledge and background knowledge. All the qualifying epistemic states together form the *least uncertainty* left to the attacker, i.e., the *best* achievement to satisfy his suspected curiosity:

- exactly one of the qualifying epistemic states is the actual state;
- all other qualifying epistemic states could possibly be the actual state as well;
- all non-qualifying epistemic states can definitely be excluded from being the actual one.

At the point of time t, we call the set $bestcv_t$ of the then-qualifying epistemic states the attacker's *best current view* (on the defender's actual epistemic state). In doing so we do not care whether or not the attacker really achieves exactly this optimal result. However, on the one hand, the kind of protection wanted by us is strongly based on the *presumed rationality* of the attacker: he definitely will never miss to identify an epistemic state as qualifying. On the other hand, he might be too lazy to exactly identify all actually non-qualifying epistemic states. In other words, the attacker might work with either the best current view $bestcv_t$ or any superset of it. Under this condition, the security

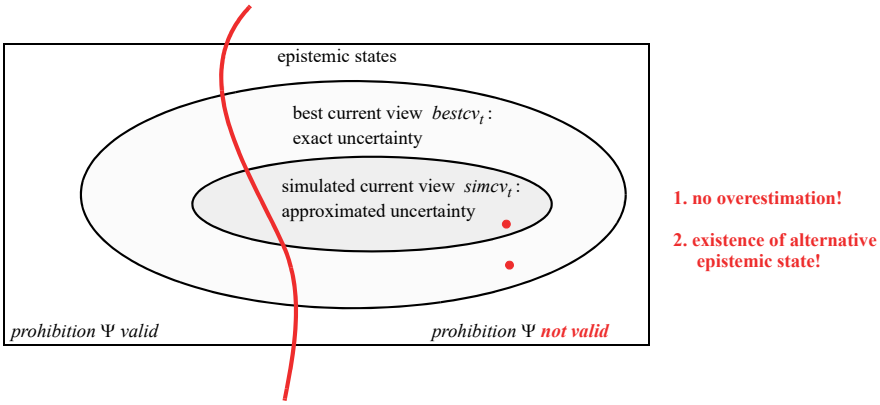

Fig. 6.4: A rough visualization of a pertinent security invariant.

policy of inference-proofness essentially requires that the best current view $bestcv_t$ still reflects *sufficient uncertainty*, namely that for all prohibited pieces of information ψ there exists an epistemic state $d' \in bestcv_t$ such that ψ is not valid in d'.
Fortunately, the defender's security mechanism does not need to determine the attacker's best current view $bestcv_t$, and the mechanism would not be able to carry out such a determination, due to the *inaccessibility* of the partner's internals according to Feature 8. Instead, it suffices to maintain an appropriate *simulated current view* $simcv_t$ that approximates the inaccessible behavior of the attacker and to enforce the following *security invariant* for all points in time t, or sometimes even a stronger one (as roughly visualized by Figure 6.4):

- $simcv_t \subseteq bestcv_t$, i.e., the defender approximates the attacker's uncertainty from below, potentially underestimating but never overestimating the uncertainty;
- for all prohibited pieces of information ψ there still exists an epistemic state $d' \in simcv_t$ such that ψ is not valid in d'.

Moreover, in general the security mechanism will work with some concise (i.e., algorithmically treatable) *representation* $rep(simcv)$ and algorithmically check the security invariant in terms of the pertinent representation.

Now, such a simulation would then be *initialized* at the point in time $t = 0$ by setting $rep(simcv_0)$ according to a suitably concise representation of the (assumed) a priori knowledge. This requires the natural *security precondition* that the a priori knowledge does not violate any prohibition. In fact we cannot prevent the attacker from "learning" what he is already sure of.

Inductively, at the point in time $t + 1$, the simulation should contain a suitably concise representation $rep(simcv_t)$ of an appropriate set $simcv_t$ of epistemic states. Then, for the functionally correct messages to be sent to the client agent—and in some cases

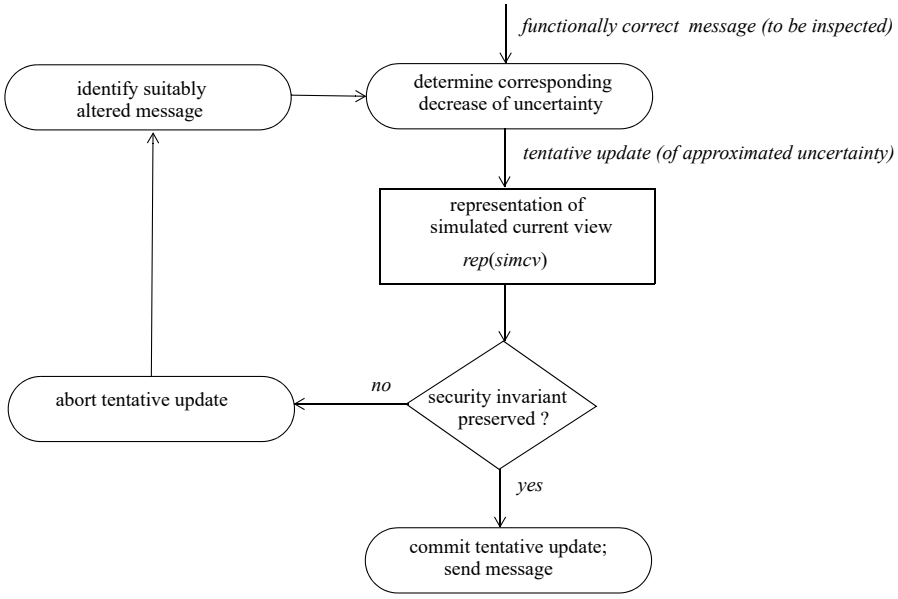

Fig. 6.5: A rough visualization of the basic steps of inductively enforcing a security invariant.

for some suitably determined modifications as well—the security mechanism first has to inspect the consequences of *tentatively* updating $rep(simcv_t)$ accordingly, i.e., whether or not the security invariant would be violated. In case of a violation, the security mechanism then has to identify suitable alterations that definitely avoid the violation. Finally, the possibly altered messages are actually sent out, and the simulation is actually updated accordingly, now ensuring the security invariant. Figure 6.5 visualizes the basic steps of these actions though the form is not universally applicable.
So we are left with the most crucial points of our construction methodology:

1. *maintaining a convenient data structure* for a "suitably concise representation $rep(simcv_t)$ of an appropriate set $simcv_t$ of epistemic states";
2. *checking tentative updates* for violations of the security invariant, preferably aiming at the approximately best computational complexity; and
3. as far as required, efficiently *identifying suitable alterations*, preferably without the need for a further explicit violation check.

Moreover, each specific situation dealt with might require some variations, in particular regarding the points in time to be considered. Let us first consider situations for which the defender's *epistemic state* is kept *fixed* over the time. If the sequence of interactions to be treated consists of only instantiations of the most basic interaction—*closed-query/yes–no-query evaluation* with response preparation—then the successive

points in time are essentially determined by the defender receiving the next query, which asks for the validity of a single piece of information in the defender's fixed epistemic state.
If the sequence of interactions also contains instantiations of the more advanced interaction of *open-query evaluation* with response-set preparation, then the response set can be formed by iteratively inspecting an internally generated sequence of closed queries, each of which is either a ground substitution of the open query (i.e., obtained by replacing free variables by constant symbols) or a specifically defined completeness sentence (dealing with "negative information"). Accordingly, for the point in time of receiving the open query, basically shared with the attacker, the defender privately determines a sequence of subpoints in time. Similarly, if the sequence of interactions also contains an instantiation of *view generation* for the fixed state, a response can be formed by internally inspecting a somehow exhaustive sequence of closed queries, also leading to a private sequence of subpoints in time. Alternatively, a view can be generated by inspecting suitably defined open queries, as explained above, leading to the respective private sequences of subpoints in time.

In general, however, the *epistemic state* of the defender might *change* over the time, raising additional issues, especially when it comes to finding an inference-proof way for

4. *resolving conflicts* between confidentiality and integrity, and
5. ensuring *backwards confidentiality*.

Basically, a confidentiality–integrity conflict might occur if the request for a change of the defender's epistemic state originates from the attacker. In general, with regard to pure functionality, the information system agent has to perform a *transaction* that first tentatively changes the state as requested and then checks whether or not the new state complies with all *semantic constraints* that are declared to be maintained as an *integrity invariant* according to the underlying application and, thus, assumed to be part of the *a priori knowledge*; for a relational database such a declaration might be part of the database schema, but in all cases such a declaration might also be expressed externally. In case of compliance, the transaction is *committed* making the tentative change persistent; otherwise, in case of non-compliance, the transaction is *aborted*, recovering the previous state. In both cases the requester is notified accordingly.

However, with regard to a security policy of inference-proofness, the respective functionally correct notification can be seen as a response set to one or more appropriately constructed closed queries. This response set might enable the attacker to learn the validity of some piece of information contained in the prohibition declaration, i.e., incorporating the response set to the (concise representation of the) *simulated current view* would violate the required *security invariant*. Consequently, the resulting conflict has to be resolved by either suitably modifying the transaction functionality or altering the response set leading to the actually returned notification, as already outlined for queries, or combining both activities.

In doing so, we also have to account for an *implicit information flow* that is caused by the overall control flow structure of a transaction as a *guarded command* in the form of an *if-then-else branching*: knowing the transaction semantics and being fully aware of the security mechanism, if the attacker can infer which branch has been selected, then he can also figure out whether and how the epistemic state has actually been changed and which alterations have actually been made to the functionally correct notification.

More generally, the defender's side of an interaction might have a more or less sophisticated overall *control flow structure* stemming from guarded commands like *if-then-else branching*, *repeat-repetitions*, *while-repetitions*, and similar procedural commands that can cause implicit information flows. Then the security mechanism typically inherits this potentially critical control flow structure. Moreover, the (code of the) security mechanism itself might have a critical control flow structure. Accordingly, the security mechanism has to appropriately treat an attacker's potential observation of a control path that has actually been chosen during a (hidden) execution on the defender's side and the resulting implicit information flow, similarly as the responses to explicit queries. As outlined above for transactions, the treatment might include to interpret such an observation as an implicit query.

The issue of *backwards confidentiality* results from the following observations about the consequences of an update of the epistemic state. First, previously released information about the validity of a prohibited piece of *information* might become *outdated* and, at least in general and whenever possible in an inference-proof way, the defender should suitably inform the attacking partner about the occurrence of an update and send him a pertinent *refreshment* of outdated information. Second, it can be shown that such a notification together with the refreshment and further information contained in messages sent at subsequent points in time might enable the attacker to infer the validity of a prohibited piece of information in the past at some preceding point in time. Accordingly, we have to strengthen the *security policy* of inference-proofness by requiring *continuous inference-proofness* to be enforced for the full range of all points in time that so far have happened, rather than just for the respective last one. This goal can actually be achieved by checking tentative updates of the representation of a simulated current view for *stronger versions* of the *security invariant*.

Basically, the general construction methodology for alterations proceeds iteratively, whether the points in time considered are externally determined by observably receiving/sending a message or only privately generated. However, for an interaction expected to be performed only once, in particular for *view generation* and *data outsourcing*, there could be only one point of time of interest, and thus one might look for a security mechanism that is working "more globally". Indeed, for view generation we have successfully designed and verified such a mechanism, in addition to the iteratively working ones inspecting sequences of appropriately formed queries, as mentioned before. Furthermore, the security mechanism for outsourcing data does not rely on such sequences. Nevertheless, these cases are also inspired by the approach to conceptually

set up some simulated counterpart to the best current view, possibly represented in some more manageable way.

6.1.4 Specific Constructions for Alterations

6.1.4.1 Weakening Including Refusing

The *weakening* approach to alterations of a harmful message comprises the special case of *refusing* to provide any *explicit information flow* as its extreme form.

In terms of a complete logic a refusal can *literally be represented* as the tautology $\chi \vee \neg\chi$, known by a rational attacker right from the beginning without having performed any interaction with the defender. However, literally replacing χ by the tautology $\chi \vee \neg\chi$ simply only if the actual validity of χ is harmful, without additionally caring about the potential harmfulness of its non-validity, equivalently by completeness, the fictitious validity of $\neg\chi$, would trigger an *implicit information flow* by means of the following kind of *meta-reasoning*, which exploits the postulated background knowledge about the security mechanism:

> For a (definitely flawed) security mechanism that, while inspecting a valid sentence χ, refuses on the harmfulness of χ but not on the harmfulness of the fictitious validity of $\neg\chi$, we would have the following equivalence: refusing occurs if and only if χ is valid and harmful. Thus, observing a refusal, the attacker could infer the validity of χ.

Notably, this kind of reasoning would just be caused by the careless handling of the critical *control flow structure* in the form of an if-then-else branching.

Accordingly, to ensure the required inference-proofness, the security mechanism has to make the two possible cases of validity and non-validity *indistinguishable* for the attacker, in the simplest way by refusing if and only if at least one of the cases is harmful.

In some situations weakening can also be achieved by using more general disjunctions. For example, let both ψ_1 and ψ_2 be prohibited pieces of information in isolation but not the disjunction $\psi_1 \vee \psi_2$, i.e., knowing the validity of $\psi_1 \vee \psi_2$ in the actual epistemic state is considered to be harmless, but figuring out which of the two disjuncts leads to the validity is harmful. Appropriately taking care of options for meta-reasoning similarly as for pure refusing and, additionally, of the potential entailments among several such disjunctions and of such disjunctions with other pieces of information, we might ensure inference-proofness by replacing a valid information ψ_i by the disjunction $\psi_1 \vee \psi_2$.

For all such kinds of weakening, the *literal representation* of weakened information, as conveyed in messages to the attacker,

- *correctly* reflects the actual epistemic state of the defender, and
- *notifies* the attacker about the fact of an actually performed weakening.

In terms of a logic, the correctness property allows us to solve the first three crucial points for the construction methodology roughly as follows:

1. A *simulated current view* is directly *represented* by the set of those sentences that so far are known to be valid in the defender's epistemic state, according to the a priori knowledge and the literal contents of messages.
2. A tentative update is independently performed with both the inspected sentence and its negation, respectively, and both versions are *checked for violations* of the security invariant by solving implication/entailment problems of the form "[current set of sentences together with tentatively added *sentence or its negation*, respectively] entails [sentence designating a *prohibition*]" by a pertinent *theorem prover*.
3. If any of the checks is positive, refusing is straightforwardly *identified* as the *suitable alteration* (or, possibly, a less easily defined but more informative disjunction that is stronger than a tautology but still harmless), and then the alteration can be notified in the corresponding message, leaving the current representing set of sentences untouched (or updating that set by adding the identified disjunction).

6.1.4.2 Lying

In contrast to weakening, the *lying* approach to alterations requires a sharp distinction between the *literal representation* of responses and the attacker's *rational sophisticated conclusions* about what he is literally observing.

To start with literal representations, consider for example a directly prohibited piece of information ψ in terms of a complete logic. The security mechanism would always have to pretend *literally* the non-validity of ψ. Hence, if the attacker is *aware of the security mechanism* including the prohibition declaration, he would not need to query any prohibition ψ: he will *always* receive the literal response that $\neg\psi$ is valid.

This feature implies that the defender cannot expect the enforcement of a declaration treating both ψ and $\neg\psi$ as prohibitions. Less obvious is a further consequence: even if the defender declares several pieces of information $\psi_1, \ldots, \psi_k$ as individual prohibitions, nevertheless under lying the disjunction $\psi_1 \vee \cdots \vee \psi_k$ has to be protected literally as well. For otherwise the attacker could perform the following *inconsistency-reasoning*:

> For a (definitely flawed) security mechanism that only lies literally on the (explicitly declared harmful) validity of $\psi_1, \ldots, \psi_k$ but not on the validity of $\psi_1 \vee \cdots \vee \psi_k$, we could get a literal representation of responses that contains the following inconsistent set of sentences: $\{\neg\psi_1, \ldots, \neg\psi_k, \psi_1 \vee \cdots \vee \psi_k\}$. Thus, observing such an inconsistency, the attacker could identify the occurrence of some lying in the literal representation and, furthermore, could be tempted to infer the validity of any sentence.

Accordingly, to ensure the required inference-proofness, the security mechanism better has to avoid running into such a somehow *"hopeless" situation*.

Moreover, regarding rational sophisticated conclusions about a *possibly lied* assertion about the non-validity of a prohibition ψ and the literal representation of this

assertion as $\neg\psi$, the *best current view bestcv* would contain *pairs* of epistemic states that differ (at least) in one state making the sentence ψ valid and the other state making the negated sentence $\neg\psi$ valid. So, evidently, a concise representation *rep*(*simcv*) of a *simulated current view simcv* containing only a sentence $\neg\psi$ that is possibly but not necessarily a *literal lie* would represent both alternatives (as already said above, in contrast to the properties of weakening).

The two issues have a conceptually simple and provably effective solution. First, we *replace a prohibition declaration* consisting exactly of the single pieces of information $\psi_1, \ldots, \psi_k$ by the singleton containing one disjunction $\psi_1 \vee \cdots \vee \psi_k$. Such a replacement considerably strengthens the security policy of inference-proofness: for each epistemic state d under consideration, the existence of an "alternative" state d' is required such that all of the ψ_i are *simultaneously not* valid in d'. Second, we form a concise *representation* of the *simulated current view* by the literally provided responses, being aware that the same literal representations for weakening and lying, respectively, essentially differ in their semantics.

Of course, similarly as for weakening but somewhat more subtly, lying is also due for any *harmful sentence* χ that could lead the attacker to believe in the validity of the disjunction of all ψ contained in the prohibition declaration.
In summary, for all kinds of lying, the *literal representation* of lied information, as conveyed in messages to the attacker,

- might *not correctly* reflect the actual epistemic state of the defender and, thus, might seriously *mislead* a naive receiver, and,
- naturally, *does not notify* the attacking receiver about the fact of an actually performed lying and, thus, lays the burden on the receiver of finding out whether or not some lying has *potentially* occurred.

Moreover, in terms of a logic, the first three crucial points for the construction methodology are solved roughly as follows;

1. A *simulated current view* is indirectly *represented* by the set of those sentences that so far have been *pretended* literally to be valid in the defender's epistemic state, according to the a priori knowledge and the content of messages.
2. A tentative update is *checked for violations* of the security invariant by solving one implication/entailment problem in the form "[current set of sentences together with tentatively added sentence] entails [sentence designating the *disjunction of all prohibitions*]" by a pertinent *theorem prover*.
3. If that check is positive, a lie on the tentatively added sentence is straightforwardly *identified* as the *suitable alteration* and then sent (without notification, of course) in the corresponding message, and that lie is inserted into the current representing set of sentences.

6.1.4.3 Combined Approaches

As we have seen, alterations by refusing on the one hand and by lying on the other are precipitated by different causes. Actually refusing is needed if *the validity or the non-validity* of the sentence inspected would be harmful regarding a *single* prohibition. Actually lying is needed if the *validity* of the inspected sentence would be harmful regarding the *disjunction* over the prohibition declaration. We might wonder whether and how we can do better by avoiding to always check the impact of both the validity and the non-validity of the sentence inspected and by never considering the strengthened security policy to protect the disjunction. This goal can be achieved by a suitable *combination of refusing and lying*:

- if the validity of the sentence inspected is not harmful regarding any prohibition, then return a message without any alteration;
- if the validity of the sentence inspected is harmful regarding some prohibition *and* also the non-validity of the sentence inspected is harmful regarding some prohibition, then return a message suitably indicating a refusal;
- if only the non-validity of the sentence inspected is harmless, then return a message suitably altered by the lie that literally pretends the non-validity.

Moreover, we can still *represent* a *simulated current view* by the set of those sentences that so far have been *pretended* literally to be valid.

6.1.4.4 Weakening versus Lying

At first glance, alterations by *lying* appear to be rather problematic, both from an ethical point of view and regarding the desired functionality. However, while lying is *ethically banned* in general, we all know of widely *accepted exceptions*, e.g., a white lie in a most critical situation, a small insincerity to avoid some larger offense, or an untruthful answer to an illegal request.

Moreover, in many cases *literally lying* might also functionally *disturb* the agreed cooperation between the communication partners, insofar as the receiver is behaving naively and "believing the lies" without further own reasoning. Even then, however, if an actually occurring alteration by literally lying does not affect an implicit or explicit *availability declaration*, we might argue that the respective interaction is at least beyond the agreement or even misusing it.

Finally, handling an interaction that is intended to change the defender's epistemic state by means of a transaction, we might face a confidentiality–integrity conflict for which an application of literal lying appears to be most natural, at least if the ethical and functional concerns are appropriately dealt with.

Above, we already distinguished between "lying" and "literal lying", and a deeper inspection of the issues indicates that such a distinction is crucial for a more information-theoretical discussion. To start with, for any kind of alteration, a sophisticated attacker can always distinguish between the conceptual notion of the *best current view*

and the pragmatically introduced notion of the *simulated current view* and its concise representation.

The *best current view bestcv* can be seen as a kind of an *inverse image*, as derived from the observed message history $mes_1, \ldots, mes_k$ and denoted by $bestcv = con\_mess^{-1}(mes_1, \ldots, mes_k)$. Here, we think of the message history being composed of k' many request messages $mes_{a_1}, \ldots, mes_{a_{k'}}$ received from the attacker, and k'' many reaction messages $mes_{d_1}, \ldots, mes_{d_{k''}}$ returned to the attacker, with $k = k' + k''$, where the latter were produced for the actual epistemic state by the defender's *message-generating* function possibly applying alterations, called *con_mess*.

Exploiting the postulated background knowledge, the attacker can thus determine *bestcv* in mathematical terms, and insofar as *con_mess* is an effectively computable function, i.e., its graph

$$\{\, (es, hist) \mid es \text{ is epistemic state,} \quad hist \text{ composed of received } hist_a \text{ and returned } hist_d \text{ is message history,} \quad con\_mess(hist_a, es) = hist_d \,\}$$

is recursively enumerable, the graph of the inverse function con_mess^{-1} is recursively enumerable as well and, thus, for each message history *hist* the inverse image $con\_mess^{-1}(hist)$ is also a recursively enumerable set. Notably, under this perspective, there is no conceptual difference between refusing and lying, or between any other approach to alterations.

The concise *representation rep*(*simcv*) of the *simulated current view simcv* is a technical means employed by the defender's security mechanism to effectively—and hopefully also efficiently—enforce the pertinent security invariant. However, exploiting the postulated background knowledge, the attacker can determine *rep*(*simcv*) as well. In fact, as outlined above, for weakening and lying that representation is just formed by the literal messages sent to the attacker.

Hence, in information-theoretical terms, the main difference between weakening (with the special case of refusing) on the one hand and any method that at least sometimes literally lies on the other hand can be described as follows:

- under *weakening* the representation of the simulated current view also represents the best current view in a *straightforward* way,
- whereas under *lying* the burden of the attacker to determine the best current view might be *much harder*.

In other words, for a sophisticated and rationally reasoning attacker, (i) there are no "real lies" but only literally ones, and (ii) lying causes an essential difference between what he can literally observe and what he can conceptually conclude, and in general imposes a high computational complexity on algorithmically determining the best current view as the least uncertainty.

A more detailed comparison of refusing and lying for a sequence of queries considers the longest prefix for which the responses have been correct (for short, called "longest honeymoon"). In turns out that refusing and lying are in general incomparable regarding this notion.

Somehow surprisingly, however, if we require to protect the disjunction over the prohibition declaration also for refusing, the information contents supplied by the two approaches are exactly the same, i.e., the conceptual best current views are always equal or, in other words, two epistemic states are *indistinguishable* for the attacker in case of refusing if and only if they are so in case of lying. Moreover, we can show that an actual refusal occurs if and only if a *potential lie* occurs. This result shows again that the information-theoretical difference between refusing and lying consists in the computational burden of the attacker to find out the grade of reliability of a message received: for refusing the attacker is explicitly notified by the defender; for lying the attacker has to find out by himself by means of rational reasoning (or just by simulating a defender that applies refusing).

6.1.5 Managing Computational Complexity

The overall computational *complexity of inference-proof* interactions is basically determined by the normal functionality of the information system under protection on the one hand and the *overhead* caused by the security mechanism on the other. Regarding the impact of the former, as far as applicable, the pertinent *logic underlying* the query evaluation appears to be most crucial. Regarding the impact of the latter, both the consideration of the *interaction history* (including the assumed *a priori knowledge*) by maintaining a concise *representation* of the *simulated current view* and the *number* and the *kinds* of *checking tentative updates* for satisfying or violating, respectively, the *security invariant* are most important. As far as applicable, the pertinent logic once again determines the costs.

In general we can expect a rather high or even practically infeasible level of complexity, and in some cases also beyond effective computability. Moreover, besides first of all treating the *constraint-solving problem* to achieve inference-proofness for the sake of confidentiality, additionally we always aim at still providing good *availability* and thus are facing the *optimization problem* to actually perform an alteration only if strictly needed. Each of the problems alone is known to be computationally hard in general, so will be their combination.

Whether or not a high computational complexity can be afforded might depend on the additional *timing constraints* of the desired communication acts. In particular, if the (attacking) communication partner expects to be served by the (defending) information owner online in real time, only a minor delay would be acceptable. By contrast, insofar as the interaction of *view generation* is initialized by the (defending) information owner, all the computations can be done offline and thus might last as

long as several hours. Accordingly, as a general heuristic we favor the shift of suitable parts of the overall computational burden to offline precomputations. The duration does not matter the communication partner at all or is bounded by the time normally spent by an interactively communicating partner for some other activity.

Given that at the core of any security mechanism we have to *check the tentative updates* of the representation of the simulated current view in terms of an underlying logic for complying with a security invariant, we can attempt to decrease the number and to lessen the complexity of such checks by restricting to special cases of the sentences used for formally expressing a priori knowledge, queries, and prohibitions. The best case would be that it suffices to relate (the validity of) a sentence inspected to the prohibitions in a straightforward way without the need to consider a simulation and thus the interaction history at all. Intuitively, this case could arise if both queries and prohibitions refer to elementary and mutually independent pieces of information.

More generally, the following guidelines for identifying computationally efficient cases have been successful: diminish the potential mutual dependence of the considered pieces of information about the defender's epistemic state and, thus, the *redundancy* contained in that state; and syntactically restrict the sentences expressing such pieces of information such that the pertinent *logical entailment* problems are easily solvable. In fact, as a first example and referring to the best case regarding *closed-query evaluation*, checking tentative updates can be done without considering the interaction history, and the entailment problems can be reduced to simple text comparisons under the following conditions: epistemic states are represented by *relational instances* of a *relational schema* with *functional dependencies* in Object Normal Form, i.e., they are in Boyce-Codd Normal Form and satisfy the Unique Key Property, and the epistemic states contain only atomic sentences, i.e., logical representations of single tuples. Further examples dedicatedly relax these requirements such that checking tentative updates can efficiently be implemented by means of SQL, even for restricted cases of *open-query evaluation*.

We can also employ a wide range of *approximation heuristics* to decrease the computational complexity, first of all by relaxing the *availability requirements* in order to facilitate the resulting *optimization problem* to minimize alterations.

6.1.6 Conclusions: Naive Illusion or Promising Hope?

Our main motivation has been to design and mathematically verify technical solutions to the social issue of preserving privacy, or of any other justified confidentiality concerns. Clearly, in general these goals require the consideration of a large range of psychological, social, institutional, legal, information-theoretical, and mathematical features. In public discussions about privacy preservation in the Information Age, some voices even claim that achieving privacy has become illusionary. Without discussing

arguments in detail, we just state that we do not share this view, as do, we believe, the legal bodies issuing the pertinent legislation.

The framework underlying our achievements focuses on a *very narrow* aspect of enforcing *confidentiality*: how to support a human information owner in hiding the validity of dedicated pieces of information, referred to as declared prohibitions, contained in the internal epistemic state of an information system agent under the owner's control, while that agent interacts with the client agent run by a semi-honest and rationally reasoning communication partner. Thus, the overall *target of protection* is *only* the internal state of a technical device whose well-defined interface to the outside world is supposed to be configurable and mastered at the discretion of the controlling human individual.

In the extreme case, that interface can just be totally disabled such that—under reasonable assumptions—no human can observe the internal state of the information system agent at all, thus preventing the *availability* of any information. So, conceptually starting with a disabled interface, the real problem is to gradually allow a flow of information from the internals of an otherwise completely shielded computing agent, while still guaranteeing the wanted kind of confidentiality of the declared prohibitions. The other way round, conceptually starting with a totally open interface, the real problem is to minimally confine such a flow of information, until the wanted kind of confidentiality according to the declared prohibitions is achieved, thus still preserving a maximal availability.

We want to emphasize that we are not aiming at anything more. Neither do we want to confine the information owner in chatting about what he has in his human mind, nor do we want to hinder the communication partner in observing real-world facts, nor to prevent him from exploiting any further information source. We care only about the conceptual information flow from the internal state of a technical device to an interacting computing agent based on protocol-complying exchanges of messages. Clearly, the occurrences of such flows might depend on additional circumstances, as captured by our assumptions and postulates about the communication partner regarding his a priori knowledge and his background knowledge, respectively.

So, our achievements are as promising as these assumptions and postulates are realistic and, furthermore, all the other features left aside by us can also be suitably dealt with.

6.1.7 Selected Bibliographic Notes

The study of inference-proof interactions of a logic-oriented information system started with two contributions about the interaction of closed-query evaluation with response preparation. Sicherman, Jonge, and Riet [611] suggested the refusing approach early, and Bonatti, Kraus, and Subrahmanian [96] later introduced the lying approach. Following

a first attempt to compare the two approaches, Biskup and Bonatti [78] set up a unifying framework for further developments.

Biskup and Bonatti [75, 76] introduced the combination of refusing and lying, for closed queries. Furthermore, Biskup and Bonatti [77] treated open-query evaluation with response-set preparation for a decidable relational submodel; considerably much later, Biskup, Bring, and Bulinski [79] reported on a partial prototype implementation with some optimizations. Among other attempts to restrict the expressibility of relational a priori knowledge, queries and prohibitions to enable inference-proofness in the spirit of access control by means of SQL only, Biskup, Embley, and Lochner [71] identified the impact of relational database schema normalization. Biskup and Weibert [86] extended all three approaches to alterations of responses to closed queries to an underlying incomplete propositional information system, which offers a third option (*don't know*) in addition to *yes* and *no*. Biskup, Gogolin, Seiler, and Weibert [81] added knowledge update transaction as a further interaction under lying, later also treated for refusing. Moreover, Biskup and Tadros [84] investigated the impact of non-monotonic reasoning for the interaction of belief revision. Biskup and Wiese [87] studied a concept of view generation as yet another interaction under lying—essentially, a as a combination of a restricted first-order logic satisfiability problem and a minimization problem. Later Biskup, Dahn, Diekmann, Menzel, Schalge, and Wiese [80] presented a prototype implementation exploring several heuristic optimizations and approximations. Biskup and Preuß [83] invented another method for view generation, later also extended to view updating, essentially based on weakening by means of disjunctions of prohibitions. Biskup, Tadros, and Zarouali [85] explored how to handle interactions expressed as procedural program executions in an inference-proof way, in particular exploiting methods of language-based security aiming at the security policy of non-interference. Biskup and Preuß [82] analyzed the fragmentation approach to secure data outsourcing. Finally, Biskup [73] studied the interaction of closed-query evaluation and view generation in the framework of abstract data sources, already presenting the framework reused in Section 6.1.2.

Some of these developments are discussed by Biskup [74]. Our narrower topic is embedded in some more general streams of research and related to many other specific topics, as outlined in Biskup [72]. Guarnieri, Marinovic, and Basin [243] and Halpern and O'Neill [255] are examples of taking a wider perspective, that includes probabilistic considerations.

Bibliography

[1] 3GPP. *3GPP TS 29.520 - 5G System; Network Data Analytics Services;Stage 3*. Tech. rep. 29.520. Version 15.3.0. 3rd Generation Partnership Project (3GPP), Mar. 2019 (cit. on p. 332).

[2] 3GPP. *NR; User Equipment (UE) radio transmission and reception; Part 1: Range 1 Standalone*. Tech. rep. 3rd Generation Partnership Project (3GPP), TS 38.101-1 V16.6.0, 2021 (cit. on p. 376).

[3] 3GPP. *System Architecture for the 5G System*. Tech. rep. 3rd Generation Partnership Project (3GPP), TS 23.501, Release 15.11.0, Sept. 2020 (cit. on pp. 371, 372).

[4] 3GPP TSG GERAN. *TR 36.888 V12.0.0: Study on provision of low-cost Machine-Type Communications (MTC) USer Equipment(UEs) based on LTE*. Tech. rep. 3GPP, 2013 (cit. on p. 264).

[5] 3GPP TSG GERAN. *TR 45.880 V13.1.0: Cellular system support for ultra-low complexity and low throughput Internet of Things (CIoT) (Release 13)*. Tech. rep. 3GPP, 2015 (cit. on p. 263).

[6] 5GAA. *White paper: Making 5G proactive and predictive for the automotive industry*. Tech. rep. 5G Automotive Association, Jan. 2020 (cit. on p. 332).

[7] C. Abrams. "Google's effort to prevent blindness shows AI challenges". In: *The Wall Street Journal* 1 (2019), p. 26 (cit. on p. 17).

[8] N. Abramson. "THE ALOHA SYSTEM: another alternative for computer communications". In: *Procs. of the Joint Computer Conference 1970*. ACM Press, 1970, p. 281. URL: http://portal.acm.org/citation.cfm?doid=1478462.1478502 (cit. on p. 318).

[9] F. Adib and D. Katabi. "See through walls with WiFi!" In: *SIGCOMM Computer Communication Review* 43.4 (Aug. 2013), pp. 75–86 (cit. on p. 246).

[10] C. C. Aggarwal and S. Y. Philip. "A condensation approach to privacy preserving data mining". In: *Procs. of the Int. Conference on Extending Database Technology 2004*. Springer. 2004, pp. 183–199 (cit. on p. 15).

[11] C. C. Aggarwal, C. Procopiuc, J. L. Wolf, P. S. Yu, and J. S. Park. "Fast Algorithms for Projected Clustering". In: *Procs. of the ACM SIGMOD Int. Conference on Management of Data 1999*. SIGMOD '99. New York, NY, USA: ACM, 1999, pp. 61–72 (cit. on p. 148).

[12] R. Agrawal and R. Srikant. "Privacy-preserving data mining". In: *Procs. of the ACM SIGMOD Int. Conference on Management of Data 2000*. 2000, pp. 439–450 (cit. on p. 15).

[13] D. Aha. "Tolerating noisy, irrelevant, and novel attributes in instance-based learning algorithms". In: *International Journal of Man-Machine Studies* 36.2 (1992), pp. 267–287 (cit. on pp. 139, 152).

[14] M. A. Akhloufi, S. Arola, and A. Bonnet. "Drones chasing drones: Reinforcement learning and deep search area proposal". In: *Drones* 3.3 (2019), p. 58 (cit. on p. 222).

[15] L. D. et al. "Incentive mechanisms for smartphone collaboration in data acquisition and distributed computing". In: *Procs. IEEE INFOCOM 2012*. 2012 (cit. on p. 298).

[16] Y. A. A. S. Aldeen, M. Salleh, and M. A. Razzaque. "A comprehensive review on privacy preserving data mining". In: *SpringerPlus* 4.1 (2015), pp. 1–36 (cit. on p. 15).

[17] S. Ali et al. *6G White Paper on Machine Learning in Wireless Communication Networks*. Apr. 2020. URL: https://arxiv.org/abs/2004.13875 (cit. on pp. 274, 332).

[18] S. Aljawarneh, A. Anguera, J. W. Atwood, J. A. Lara, and D. Lizcano. "Particularities of data mining in medicine: lessons learned from patient medical time series data analysis". In: *EURASIP Journal on Wireless Communications and Networking* 2019.1 (2019), pp. 1–29 (cit. on p. 15).

[19] M. Allhoff, A. Schönhuth, M. Martin, I. G. Costa, S. Rahmann, and T. Marschall. "Discovering motifs that induce sequencing errors". In: *BMC Bioinformatics* 14 Suppl 5 (2013), S1 (cit. on p. 55). **SFB876-B1**

https://doi.org/10.1515/9783110785982-007

[20] Y. Altintas, P. Kersting, D. Biermann, E. Budak, B. Denkena, and I. Lazoglu. “Virtual process systems for part machining operations”. In: *CIRP Annals* 63.2 (2014), pp. 585–605 (cit. on p. 157).

[21] B. Alzahrani, O. S. Oubbati, A. Barnawi, M. Atiquzzaman, and D. Alghazzawi. “UAV assistance paradigm: State-of-the-art in applications and challenges”. In: *Journal of Network and Computer Applications* 166 (2020), p. 102706 (cit. on p. 213).

[22] S. Amal, F. Felix, B. Jens, M. Katharina, W. Petra, and R. Wolfgang. “Simulation and Sensor Data Fusion for Machine Learning Application”. In: *Advanced Engineering Informatics* (2022) (cit. on p. 166).

[23] X. An, Q. Yang, and P. M. Bentler. “A latent factor linear mixed model for high-dimensional longitudinal data analysis”. In: *Statistics in medicine* 32.24 (2013), pp. 4229–4239 (cit. on p. 9).

[24] M. W. Anderson and I. Schrijver. “Next Generation DNA Sequencing and the Future of Genomic Medicine”. In: *Genes* 1.1 (June 2010). Number: 1Publisher: Molecular Diversity Preservation Int., pp. 38–69. URL: https://www.mdpi.com/2073-4425/1/1/38 (cit. on p. 55).

[25] D. W. Apley and J. Zhu. “Visualizing the effects of predictor variables in black box supervised learning models”. In: *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 82.4 (2020), pp. 1059–1086 (cit. on p. 17).

[26] S. Arabi, E. Sabir, H. Elbiaze, and M. Sadik. “Data gathering and energy transfer dilemma in UAV-assisted flying access network for IoT”. In: *Sensors* 18.5 (2018), p. 1519 (cit. on p. 213).

[27] L. Archibald, L. Phillips, D. Monnet, J. E. McGowan Jr, F. Tenover, and R. Gaynes. “Antimicrobial resistance in isolates from inpatients and outpatients in the United States: increasing importance of the intensive care unit”. In: *Clinical Infectious Diseases* 24.2 (1997), pp. 211–215 (cit. on p. 7).

[28] D. Ardila et al. “End-to-end lung cancer screening with three-dimensional deep learning on low-dose chest computed tomography”. In: *Nature medicine* 25.6 (2019), pp. 954–961 (cit. on pp. 5, 8).

[29] C. Arendt. “Raw environmental indoor sensor data”. In: (Nov. 2019). URL: http://doi.org/10.5281/zenodo.3555406 (cit. on p. 324).

[30] C. Arendt, S. Böcker, and C. Wietfeld. “Data-Driven Model-Predictive Communication for Resource-Efficient IoT Networks”. In: *Procs. of the IEEE World Forum on Internet of Things 2020*. Apr. 2020 (cit. on pp. 270, 323–328). **SFB876-A4**

[31] V. Arvind, J. Köbler, G. Rattan, and O. Verbitsky. “On the Power of Color Refinement”. In: *Procs. of the Int. Symposium on Fundamentals of Computation Theory 2015*. 2015, pp. 339–350 (cit. on p. 91).

[32] A. Asuncion and D. J. Newman. *UCI Machine Learning Repository*. 2007 (cit. on pp. 140, 152).

[33] J.-Y. Audibert and R. Munos. “Introduction to bandits: Algorithms and theory”. In: *ICML Tutorial on bandits* (2011) (cit. on p. 236).

[34] L. Babai and L. Kucera. “Canonical labelling of graphs in linear average time”. In: *Procs. of the Symposium on Foundations of Computer Science 1979*. IEEE, 1979, pp. 39–46 (cit. on p. 91).

[35] D. Baek, M. Hwang, H. Kim, and D.-S. Kwon. “Path planning for automation of surgery robot based on probabilistic roadmap and reinforcement learning”. In: *Procs. of the Int. Conference on Ubiquitous Robots 2018*. IEEE. 2018, pp. 342–347 (cit. on p. 6).

[36] B. M. Bai, B. Nalini, and J. Majumdar. “Analysis and detection of diabetes using data mining techniques – a big data application in health care”. In: *Emerging Research in Computing, Information, Communication and Applications*. Springer, 2019, pp. 443–455 (cit. on p. 15).

[37] L. Bai, L. Rossi, Z. Zhang, and E. R. Hancock. “An Aligned Subtree Kernel for Weighted Graphs”. In: *Procs. of the Int. Conference on Machine Learning 2015*. 2015, pp. 30–39 (cit. on p. 89).

[38] S. Bai, J. Z. Kolter, and V. Koltun. “An Empirical Evaluation of Generic Convolutional and Recurrent Networks for Sequence Modeling”. In: *arXiv: Computing Research Repository* (2018). DOI: arXiv:1803.01271 (cit. on p. 11).

[39] C. A. Balanis. *Antenna Theory: Analysis and Design*. 4th ed. Wiley, Feb. 2016 (cit. on p. 376).

[40] W. Balid, H. Tafish, and H. H. Refai. “Intelligent Vehicle Counting and Classification Sensor for Real-Time Traffic Surveillance”. In: *IEEE Transactions on Intelligent Transportation Systems* 19.6 (2018), pp. 1784–1794 (cit. on p. 245).

[41] A. Banks et al. “MQTT Version 3.1.1”. In: *OASIS standard* (2014) (cit. on p. 265).

[42] P. Bansal and K. M. Kockelman. “Forecasting Americans’ long-term adoption of connected and autonomous vehicle technologies”. In: *Transportation Research Part A: Policy and Practice* 95.8 (2017), pp. 49–63 (cit. on p. 293).

[43] M. K. Baowaly, C.-C. Lin, C.-L. Liu, and K.-T. Chen. “Synthesizing electronic health records using improved generative adversarial networks”. In: *Journal of the American Medical Informatics Association* 26.3 (2019), pp. 228–241 (cit. on p. 16).

[44] A. S. Barkousaraie, O. Ogunmolu, S. Jiang, and D. Nguyen. “A Fast Deep Learning Approach for Beam Orientation Optimization for Prostate Cancer IMRT Treatments”. In: *arXiv: Physics* (2019). DOI: arXiv:1905.00523 (cit. on p. 8).

[45] J. A. Barta, C. A. Powell, and J. P. Wisnivesky. “Global epidemiology of lung cancer”. In: *Annals of global health* 85.1 (2019) (cit. on p. 7).

[46] E. Basar, M. Di Renzo, J. De Rosny, M. Debbah, M.-S. Alouini, and R. Zhang. “Wireless Communications Through Reconfigurable Intelligent Surfaces”. In: *IEEE Access* 7 (2019), pp. 116753–116773 (cit. on p. 385).

[47] H. Bast, J. Sternisko, and S. Storandt. “Delay-robustness of transfer patterns in public transportation route planning”. In: *Procs. of the Workshop on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems 2013*. Vol. 33. Schloss Dagstuhl Leibniz-Zentrum fuer Informatik. 2013, pp. 42–54 (cit. on pp. 234, 235).

[48] H. Bast et al. “Fast routing in very large public transportation networks using transfer patterns”. In: *Procs. of the European Symposium on Algorithms 2010*. Springer. 2010, pp. 290–301 (cit. on pp. 234, 235).

[49] H. Bast et al. *Route Planning in Transportation Networks*. Tech. rep. MSR-TR-2014-4. Jan. 2014 (cit. on p. 233).

[50] G. E. Batista and M. C. Monard. “An analysis of four missing data treatment methods for supervised learning”. In: *Applied Artificial Intelligence* 17.5-6 (2003), pp. 519–533 (cit. on p. 13).

[51] S. Batra, Z. Huang, A. Petrenko, T. Kumar, A. Molchanov, and G. S. Sukhatme. “Decentralized control of quadrotor swarms with end-to-end deep reinforcement learning”. In: *Procs. of the Conference on Robot Learning 2022*. PMLR. 2022, pp. 576–586 (cit. on p. 222).

[52] É. Beke, A. Bódi, T. G. Katalin, T. Kovács, D. Maros, and L. Gáspár. “The role of drones in linking industry 4.0 and ITS Ecosystems”. In: *Procs. of the IEEE Int. Symposium on Computational Intelligence and Informatics 2018*. IEEE. 2018, pp. 000191–000198 (cit. on p. 212).

[53] C. Bektas, S. Böcker, F. Kurtz, and C. Wietfeld. “Reliable Software-Defined RAN Network Slicing for Mission-Critical 5G Communication Networks”. In: *Procs. of the IEEE Globecom Workshops 2019*. Waikoloa, Hawaii, USA, Dec. 2019. URL: https://ieeexplore.ieee.org/document/9024677 (cit. on p. 366).

[54] C. Bektas, D. Overbeck, and C. Wietfeld. “SAMUS: Slice-Aware Machine Learning-based Ultra-Reliable Scheduling”. In: *Procs. of the IEEE Int. Conference on Communications 2021*. Montreal, Canada, June 2021 (cit. on pp. 364–367, 369–371, 373). **SFB876-B4**

[55] A. Ben-Hur, D. Horn, H. Siegelmann, and V. Vapnik. “Support Vector Clustering”. In: *Journal of Machine Learning Research* 2 (Mar. 2002), pp. 125–137 (cit. on p. 148).

[56] J. L. Bentley and J. B. Saxe. “Decomposable Searching Problems I: Static-to-Dynamic Transformation”. In: *The Journal of Algorithms* 1.4 (1980), pp. 301–358 (cit. on p. 66).

[57] J. Bergstra and Y. Bengio. “Random search for hyper-parameter optimization.” In: *Journal of Machine Learning Research* 13.2 (2012) (cit. on p. 176).

[58] M. Bernas, B. Płaczek, and W. Korski. “Wireless network with Bluetooth low energy beacons for vehicle detection and classification”. In: *Procs. of the Int. Conference on Computer Networks 2018*. Ed. by P. Gaj, M. Sawicki, G. Suchacka, and A. Kwiecień. Cham: Springer Int. Publishing, 2018, pp. 429–444 (cit. on p. 246).

[59] H. M. Bette, L. Habel, T. Emig, and M. Schreckenberg. “Mechanisms of jamming in the Nagel-Schreckenberg model for traffic flow”. In: *Physical Review E* 95.012311 (Jan. 2017). URL: http://dx.doi.org/10.1103/PhysRevE.95.012311 (cit. on p. 287). **SFB876-B4**

[60] A. Bharti, R. Adeogun, and T. Pedersen. “Learning Parameters of Stochastic Radio Channel Models From Summaries”. In: *IEEE Open Journal of Antennas and Propagation* 1 (2020), pp. 175–188 (cit. on p. 274).

[61] S. S. Bhowmick and B. Seah. “Clustering and Summarizing Protein-Protein Interaction Networks: A Survey”. In: *IEEE Transactions on Knowledge Data Engineering* 28.3 (2016), pp. 638–658 (cit. on p. 86).

[62] G. Bhuvaneswari and R. Annamalai. “Development of a Solar Cell Model in MATLAB for PV Based Generation System”. In: *Procs. of the IEEE India Conference 2011*. 2011, pp. 1–5 (cit. on p. 197).

[63] S. Bickel, J. Bogojeska, T. Lengauer, and T. Scheffer. “Multi-Task Learning for HIV Therapy Screening”. In: *Procs. of the Int. Conference on Machine Learning 2008*. ICML ’08. New York, NY, USA: ACM, 2008, pp. 56–63 (cit. on p. 78).

[64] F. Biessmann, D. Salinas, S. Schelter, P. Schmidt, and D. Lange. “’Deep’ Learning for Missing Value Imputation in Tables with Non-Numerical Data”. In: *Procs. of the ACM Int. Conference on Information and Knowledge Management 2018*. 2018, pp. 2017–2025 (cit. on p. 13).

[65] M. Binder, F. Pfisterer, M. Lang, L. Schneider, L. Kotthoff, and B. Bischl. “mlr3pipelines - Flexible Machine Learning Pipelines in R”. In: *Journal of Machine Learning Research* 22.184 (2021), pp. 1–7. URL: http://jmlr.org/papers/v22/21-0281.html (cit. on p. 83).

[66] M. Birgersson and G. Hansson. “Data integration using machine learning: Automation of data mapping using machine learning techniques”. MA thesis. 2016 (cit. on p. 12).

[67] B. Bischl, M. Lang, O. Mersmann, J. Rahnenführer, and C. Weihs. “BatchJobs and BatchExperiments: Abstraction Mechanisms for Using R in Batch Environments”. In: *Journal of Statistical Software* 64.11 (2015), pp. 1–235. URL: http://www.jstatsoft.org/v64/i11 (cit. on p. 83). **SFB876-A3**

[68] B. Bischl et al. “Hyperparameter Optimization: Foundations, Algorithms, Best Practices and Open Challenges”. In: *arXiv: Statistics* (2021). DOI: arXiv:2107.05847v2 (cit. on p. 74). **SFB876-A3**

[69] B. Bischl et al. “mlr: Machine Learning in R”. In: *Journal of Machine Learning Research* 17.170 (2016), pp. 1–5. URL: http://jmlr.org/papers/v17/15-066.html (cit. on p. 82). **SFB876-A3**

[70] C. M. Bishop. *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Secaucus, NJ, USA: Springer-Verlag New York, Inc., 2006 (cit. on p. 221).

[71] J. Biskup, D. Embley, and J.-H. Lochner. “Reducing inference control to access control for normalized database schemas”. In: *Information Processing Letters* 106 (2008), pp. 8–12 (cit. on p. 412).

[72] J. Biskup. “Inference Control”. In: *Encyclopedia of Cryptography, Security and Privacy*. Ed. by S. Jajodia, P. Samarati, and M. Yung. Berlin, Heidelberg: Springer Berlin Heidelberg, 2021, pp. 1–7. URL: https://doi.org/10.1007/978-3-642-27739-9_753-2 (cit. on p. 412).

[73] J. Biskup. “Inference-Proof Monotonic Query Evaluation and View Generation Reconsidered”. In: *Procs. of Data and Applications Security and Privacy 2020*. Ed. by A. Singhal and J. Vaidya. Vol. 12122. LNCS. Springer, 2020, pp. 79–99 (cit. on pp. 394, 412).

[74] J. Biskup. “Selected Results and Related Issues of Confidentiality-Preserving Controlled Interaction Execution”. In: *Procs. of Foundations of Information and Knowledge Systems 2016*. Ed. by M. Gyssens and G. R. Simari. Vol. 9616. LNCS. Springer, 2016, pp. 211–234 (cit. on p. 412). **SFB876-A5**

[75] J. Biskup and P. A. Bonatti. “Controlled query evaluation for enforcing confidentiality in complete information systems”. In: *The International Journal of Information Security* 3.1 (2004), pp. 14–27 (cit. on p. 412).

[76] J. Biskup and P. A. Bonatti. “Controlled Query Evaluation for Known Policies by Combining Lying and Refusal”. In: *Annals of Mathematics and Artificial Intelligence* 40.1-2 (2004), pp. 37–62 (cit. on p. 412).

[77] J. Biskup and P. A. Bonatti. “Controlled query evaluation with open queries for a decidable relational submodel”. In: *Annals of Mathematics and Artificial Intelligence* 50.1-2 (2007), pp. 39–77 (cit. on p. 412).

[78] J. Biskup and P. A. Bonatti. “Lying versus refusal for known potential secrets”. In: *Data & Knowledge Engineering* 38.2 (2001), pp. 199–222 (cit. on p. 412).

[79] J. Biskup, M. Bring, and M. Bulinski. “Inference control of open relational queries under closed-world semantics based on theorem proving”. In: *The Information Systems Journal* 70 (2017), pp. 32–47 (cit. on p. 412).

[80] J. Biskup, C. Dahn, K. Diekmann, R. Menzel, D. Schalge, and L. Wiese. “Publishing inference-proof relational data: An implementation and experiments”. In: *Data & Knowledge Engineering* 120 (2019), pp. 1–44 (cit. on p. 412).

[81] J. Biskup, C. Gogolin, J. Seiler, and T. Weibert. “Inference-proof view update transactions with forwarded refreshments”. In: *Journal of Computer Security* 19 (2011), pp. 487–529 (cit. on p. 412).

[82] J. Biskup and M. Preuß. “Checking Inference-Proofness of Attribute-Disjoint and Duplicate-Preserving Fragmentations”. In: *Annals of Mathematics and Artificial Intelligence* (2019) (cit. on p. 412).

[83] J. Biskup and M. Preuß. “Information Control by Policy-Based Relational Weakening Templates”. In: *Procs. of the European Symposium on Research in Computer Security 2016*. Ed. by I. G. Askoxylakis, S. Ioannidis, S. K. Katsikas, and C. A. Meadows. Vol. 9879. LNCS. Springer, 2016, pp. 361–381 (cit. on p. 412).

[84] J. Biskup and C. Tadros. “Preserving confidentiality while reacting on iterated queries and belief revisions”. In: *Annals of Mathematics and Artificial Intelligence* 73.1-2 (2015), pp. 75–123 (cit. on p. 412). **SFB876-A5**

[85] J. Biskup, C. Tadros, and J. Zarouali. “Confidentiality enforcement by hybrid control of information flows”. In: *arXiv: Computing Research Repository* (2017). DOI: arXiv:1707.08482 (cit. on p. 412).

[86] J. Biskup and T. Weibert. “Keeping secrets in incomplete databases”. In: *International Journal of Information Security* 7.3 (2008), pp. 199–217 (cit. on p. 412).

[87] J. Biskup and L. Wiese. “A sound and complete model-generation procedure for consistent and confidentiality-preserving databases”. In: *Theoretical Computer Science* 412.31 (2011), pp. 4044–4072 (cit. on p. 412).

[88] A. Blagojević et al. “Artificial intelligence approach towards assessment of condition of COVID-19 patients-Identification of predictive biomarkers associated with severity of clinical condition and disease progression”. In: *Computers in Biology and Medicine* 138 (2021), p. 104869 (cit. on p. 11).

[89] A. Bloch, D. Berger, and J. Takala. “Understanding circulatory failure in sepsis”. In: *Intensive Care Medicine* 42.12 (2016), pp. 2077–2079 (cit. on p. 5).

[90] D. B. Blumenthal, N. Boria, J. Gamper, S. Bougleux, and L. Brun. “Comparing heuristics for graph edit distance computation”. In: *VLDB – The International Journal on Very Large Data Bases* 29.1 (2020), pp. 419–458. URL: https://doi.org/10.1007/s00778-019-00544-1 (cit. on pp. 88, 94).

[91] C. Bock, T. Gumbsch, M. Moor, B. Rieck, D. Roqueiro, and K. Borgwardt. “Association mapping in biomedical time series via statistically significant shapelet mining”. In: *Bioinformatics* 34.13 (2018), pp. i438–i446 (cit. on pp. 9, 15).

[92] C. Bock, M. Moor, C. R. Jutzeler, and K. Borgwardt. “Machine learning for biomedical time series classification: from shapelets to deep learning”. In: *Artificial Neural Networks*. Springer, 2021, pp. 33–71 (cit. on p. 9).

[93] S. Böcker, C. Arendt, P. Jörke, and C. Wietfeld. “LPWAN in the Context of 5G: Capability of LoRaWAN to Contribute to mMTC”. In: *Procs. of the IEEE World Forum on Internet of Things 2019*. IEEE, 2019 (cit. on pp. 262, 316, 317, 319, 320).

[94] A. Bommert and M. Lang. “stabm: Stability Measures for Feature Selection”. In: *Journal of Open Source Software* 6.59 (2021), p. 3010. URL: https://doi.org/10.21105/joss.03010 (cit. on p. 83). **SFB876-A3**

[95] A. Bommert, X. Sun, B. Bischl, J. Rahnenführer, and M. Lang. “Benchmark for Filter Methods for Feature Selection in High-Dimensional Classification Data”. In: *Computational Statistics & Data Analysis* 143 (Mar. 2020), p. 106839 (cit. on pp. 72, 81, 82). **SFB876-A3**

[96] P. A. Bonatti, S. Kraus, and V. S. Subrahmanian. “Foundations of Secure Deductive Databases”. In: *IEEE Transactions on Knowledge and Data Engineering* 7.3 (1995), pp. 406–422 (cit. on p. 411).

[97] M. Borowski, M. Görges, R. Fried, O. Such, C. Wrede, and M. Imhoff. “Medical device alarms”. In: (2011) (cit. on p. 5).

[98] G. E. P. Box, G. Jenkins, and G. C. Reinsel. *Time Series Analysis. Forecasting and Control*. 3rd. Englewood Cliffs: Prentice Hall, 1994 (cit. on p. 119).

[99] J. Boyan and M. Littman. “Packet routing in dynamically changing networks: A reinforcement learning approach”. In: *Advances in Neural Information Processing Systems 6: Procs. of the 1993 Conference*. 1994, pp. 671–678 (cit. on pp. 274, 275).

[100] G. Bradski. “The OpenCV Library”. In: *Dr. Dobb’s Journal of Software Tools* (2000) (cit. on p. 30).

[101] L. Breiman. “Bagging Predictors”. In: *Machine Learning* 24.2 (Aug. 1996), pp. 123–140. URL: https://doi.org/10.1023/A:1018054314350 (cit. on pp. 125, 169, 171).

[102] L. Breiman. “Random Forests”. In: *Machine Learning* 45.1 (Oct. 2001), pp. 5–32 (cit. on pp. 17, 152, 245).

[103] H. Breiteneder et al. “Biomarkers for diagnosis and prediction of therapy responses in allergic diseases and asthma”. In: *Allergy* 75.12 (2020), pp. 3039–3068 (cit. on p. 9).

[104] M. M. Breunig, H.-P. Kriegel, R. T. Ng, and J. Sander. “LOF: Identifying Density-Based Local Outliers”. In: *Procs. of the ACM SIGMOD Int. Conference on Management of Data 2000*.

SIGMOD '00. New York, NY, USA: Association for Computing Machinery, 2000, pp. 93–104. URL: https://doi.org/10.1145/342009.335388 (cit. on p. 14).

[105] A. Z. Broder. "On the resemblance and containment of documents". In: *Procs. of the Compression and Complexity of SEQUENCES 1997*. June 1997, pp. 21–29 (cit. on pp. 89, 93).

[106] G. Brown, J. L. Wyatt, and P. Tino. "Managing Diversity in Regression Ensembles". In: *Journal of Machine Learning Research* (2005), pp. 1621–1650 (cit. on p. 169).

[107] N. Bui, M. Cesana, S. A. Hosseini, Q. Liao, I. Malanchini, and J. Widmer. "A survey of anticipatory mobile networking: Context-based classification, prediction methodologies, and optimization techniques". In: *IEEE Communications Surveys & Tutorials* 19.3 (2017), pp. 1790–1821 (cit. on pp. 275, 332).

[108] N. Bui and J. Widmer. "OWL: A Reliable Online Watcher for LTE Control Channel Measurements". In: *Procs. of the Workshop on All Things Cellular: Operations, Applications and Challenges 2016*. ATC '16. New York, NY, USA: ACM, 2016, pp. 25–30 (cit. on pp. 345, 347).

[109] H. Bunke. "On a relation between graph edit distance and maximum common subgraph". In: *Pattern Recognition Letters* 18.8 (1997), pp. 689–694 (cit. on p. 87).

[110] J. Burridge. "Empirical Bayes analysis of survival time data". In: *Journal of the Royal Statistical Society. Series B* 43.1 (1981), pp. 65–75 (cit. on p. 63).

[111] S. Buschjäger, T. Liebig, and K. Morik. "Gaussian Model Trees for Traffic Imputation". In: *Procs. of the Int. Conference on Pattern Recognition Applications and Methods 2019*. SciTePress. 2019, pp. 243–254. URL: https://www.scitepress.org/PublicationsDetail.aspx?ID=g+tVIY+KNts=%5C&t=1 (cit. on p. 229). **SFB876-A1, SFB876-B4**

[112] M. D. Butti, H. Chanfreau, D. Martinez, D. García, E. Lacunza, and M. C. Abba. "BioPlat: a software for human cancer biomarker discovery". In: *Bioinformatics* 30.12 (2014), pp. 1782–1784 (cit. on p. 9).

[113] R. H. Byrd, P. Lu, J. Nocedal, and C. Zhu. "A Limited Memory Algorithm for Bound Constrained Optimization". In: *SIAM Journal on Scientific Computing (SISC)* 16.5 (Sept. 1995), pp. 1190–1208 (cit. on p. 168).

[114] D. L. Cameron et al. "GRIDSS: sensitive and specific genomic rearrangement detection using positional de Bruijn graph assembly". In: *Genome Research* 27.12 (Dec. 2017), pp. 2050–2060. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5741059/ (cit. on p. 57).

[115] K. S. Candan, R. Rossini, M. L. Sapino, and X. Wang. "sDTW: computing DTW distances using locally relevant constraints based on salient feature alignments". In: *arXiv: Computing Research Repository* (2012). DOI: arXiv:1208.0275 (cit. on p. 122).

[116] B.-T. Cao, S. Freitag, and G. Meschke. "A Hybrid RNN-GPOD Surrogate Model for Real-time Settlement Predictions in Mechanised Tunnelling". In: *Advanced Modeling and Simulation in Engineering Sciences* 3.1 (Mar. 2016), p. 5. URL: https://link.springer.com/article/10.1186/s40323-016-0057-9 (cit. on pp. 160, 161).

[117] H. Carr, J. Snoeyink, and U. Axen. "Computing contour trees in all dimensions". In: *Computational Geometry* 24.2 (2003), pp. 75–94 (cit. on p. 16).

[118] A. Cassini et al. "Attributable deaths and disability-adjusted life-years caused by infections with antibiotic-resistant bacteria in the EU and the European Economic Area in 2015: a population-level modelling analysis". In: *The Lancet Infectious Diseases* 19.1 (2019), pp. 56–66 (cit. on p. 7).

[119] E. R. Cavalcanti, J. A. R. de Souza, M. A. Spohn, R. C. d. M. Gomes, and A. F. B. F. d. Costa. "VANETs' research over the past decade: Overview, credibility, and trends". In: *SIGCOMM Computer Communication Review* 48.2 (May 2018), pp. 31–39 (cit. on pp. 274, 280).

[120] O. Chapelle, B. Schölkopf, and A. Zien. *Semi-Supervised Learning*. Cambridge, MA: MIT Press, 2006 (cit. on pp. 137, 144).

[121] P. Chapman et al. *CRISP–DM 1.0*. Tech. rep. The CRISP–DM Consortium, Aug. 2000 (cit. on p. 115).

[122] S. Chatterjee and A. S. Hadi. "Sensitivity analysis in linear regression". In: (1988) (cit. on p. 60).

[123] A. Chaudhary and S. S. Singh. "Lung cancer detection on CT images by using image processing". In: *Procs. of the Int. Conference on Computing Sciences 2012*. IEEE. 2012, pp. 142–146 (cit. on p. 5).

[124] H. Chaudhary. "Application of the Theory of a Single First Order Equation to Traffic Flow". In: *Journal of the Institute of Engineering* 9.1 (2014), pp. 175–181 (cit. on p. 237).

[125] N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer. "SMOTE: Synthetic Minority Over-Sampling Technique". In: *Journal of Artificial Intelligence Research* 16.1 (June 2002), pp. 321–357 (cit. on p. 125).

[126] N. V. Chawla, A. Lazarevic, L. O. Hall, and K. W. Bowyer. "SMOTEBoost: Improving Prediction of the Minority Class in Boosting". In: *Procs. of the ACM SIGKDD Conference on Knowledge Discovery and Data Mining 2003*. 2003 (cit. on p. 125).

[127] M. K. Cheezum, W. F. Walker, and W. H. Guilford. "Quantitative Comparison of Algorithms for Tracking Single Fluorescent Particles". In: *Biophysical Journal* 81.4 (2001), pp. 2378–2388. URL: https://www.sciencedirect.com/science/article/pii/S0006349501758845 (cit. on p. 28).

[128] J. Chen and B. Huang. "An In-Process Neural Network-Based Surface Roughness Prediction (INN-SRP) System Using a Dynamometer in End Milling Operations". In: *The Int. Journal of Advanced Manufacturing Technology* 21.5 (2003), pp. 339–347 (cit. on pp. 158, 160).

[129] J. Chen, S. Li, B. Dai, and G. Zhou. "Active Learning for Age Regression in Social Media". In: *Procs. of the Int. Symposium on Chinese Computational Linguistics and Natural Language Processing Based on Naturally Annotated Big Data 2016*. Ed. by M. Sun, X. Huang, H. Lin, Z. Liu, and Y. Liu. Cham: Springer Int. Publishing, 2016, pp. 351–362 (cit. on pp. 162, 163).

[130] K. Chen and L. Liu. "Privacy preserving data classification with rotation perturbation". In: *Procs. of the IEEE Int. Conference on Data Mining 2005*. IEEE. 2005, 4–pp (cit. on p. 15).

[131] S. Chen, B. Liu, M. Qian, and C. Zhang. "Kernel k-Means Based Framework For Aggregate Outputs Classification". In: *Procs. of the Int. Conference on Data Mining Workshops 2009*. 2009, pp. 356–361 (cit. on pp. 141, 151).

[132] X. Chen, H. Huo, J. Huan, and J. S. Vitter. "An efficient algorithm for graph edit distance computation". In: *Knowledge-Based Systems* 163 (2019), pp. 762–775 (cit. on pp. 88, 93, 94).

[133] Z. Chen, W. Chen, and Y. Shi. "Ensemble learning with label proportions for bankruptcy prediction". In: *Expert Systems with Applications* 146 (2020), pp. 113–155 (cit. on p. 139).

[134] C.-C. Chiu and C.-Y. Tsai. "A k-anonymity clustering method for effective data privacy preservation". In: *Procs. of the Int. Conference on Advanced Data Mining and Applications 2007*. Springer. 2007, pp. 89–99 (cit. on p. 15).

[135] K. Cho et al. "Learning phrase representations using RNN encoder-decoder for statistical machine translation". In: *Procs. of the 2014 Conference on Empirical Methods in Natural Language Processing*. Doha, Qatar: Association for Computational Linguistics, Oct. 2014, pp. 1724–1734. DOI: arXiv:1406.1078 (cit. on p. 11).

[136] R. Chrobok, O. Kaumann, J. Wahle, and M. Schreckenberg. "Three Categories of Traffic Data: Historical, Current, and Predictive". In: *IFAC Symposium on Control in Transportation Systems 2000*. Vol. 33. 9. 2000, pp. 221–226. URL: https://www.sciencedirect.com/science/article/pii/S1474667017381508 (cit. on p. 289).

[137] H. M. Chung and P. Gray. "Data mining". In: *Journal of Management Information Systems* 16.1 (1999), pp. 11–16 (cit. on p. 15).

[138] Y. Chung, S. Krishnan, and T. Kraska. “A data quality metric (DQM): how to estimate the number of undetected errors in data sets”. In: *Procs. of the VLDB Endowment* 10.10 (2016), pp. 1094–1105 (cit. on p. 12).

[139] I. Cima et al. “Cancer genetics-guided discovery of serum biomarker signatures for diagnosis and prognosis of prostate cancer”. In: *Procs. of the National Academy of Sciences* 108.8 (2011), pp. 3342–3347 (cit. on pp. 9, 10).

[140] R. Clark Robertson and Tri T Ha. “Effect of capture on throughput of variable length packet Aloha systems”. In: *Computer Communications* 17.12 (1994), pp. 836–842 (cit. on p. 318).

[141] W. S. Cleveland and S. J. Devlin. “Locally Weighted Regression : An Approach to Regression Analysis by Local Fitting Locally Weighted Regression : An Approach to Regression Analysis by Local Fiiting”. In: *Journal of the American Statistical Association* 83.403 (1988), pp. 37–41 (cit. on p. 324).

[142] H. Climente-González, C.-A. Azencott, S. Kaski, and M. Yamada. “Block HSIC Lasso: model-free biomarker detection for ultra-high dimensional data”. In: *Bioinformatics* 35.14 (2019), pp. i427–i435 (cit. on p. 9).

[143] L. Codeca, R. Frank, and T. Engel. “Luxembourg SUMO traffic (LuST) scenario: 24 hours of mobility for vehicular networking research”. In: *Procs. of the Vehicular Networking Conference 2015*. IEEE. 2015, pp. 1–8 (cit. on p. 239).

[144] D. Cohen-Steiner, H. Edelsbrunner, and D. Morozov. “Vines and vineyards by updating persistence in linear time”. In: *Procs. of the Symposium on Computational Geometry 2006*. 2006, pp. 119–126 (cit. on p. 16).

[145] B. Coifman and S. Kim. “Speed estimation and length based vehicle classification from freeway single-loop detectors”. In: *Transportation Research Part C: Emerging Technologies* 17.4 (2009), pp. 349–364. URL: https://www.sciencedirect.com/science/article/pii/S0968090X09000072 (cit. on p. 245).

[146] I. Corporation. *SLMD960H12L: IXOLAR High Efficiency SolarMD*. 2010 (cit. on p. 201).

[147] C. Cortes and V. N. Vapnik. “Support–Vector Networks”. In: *Machine Learning Journal* 20 (Sept. 1995), pp. 273–297 (cit. on p. 245).

[148] C. Cox. *An Introduction to 5G*. Wiley, Feb. 2021 (cit. on p. 377).

[149] D. R. Cox. “Regression Models and Life-Tables”. In: *Journal of the Royal Statistical Society. Series B* 34.2 (1972), pp. 187–220 (cit. on pp. 62, 63, 72, 79).

[150] R. Crane. “Prediction of Attenuation by Rain”. In: *IEEE Transactions on Communications* 28.9 (1980), pp. 1717–1733 (cit. on p. 246).

[151] D. Cristiani, F. Bottonelli, A. Trotta, and M. Di Felice. “Inventory Management through Mini-Drones: Architecture and Proof-of-Concept Implementation”. In: *Procs. of the IEEE Int. Symposium on A World of Wireless, Mobile and Multimedia Networks*. 2020, pp. 317–322 (cit. on p. 213).

[152] K. Crombecq, L. De Tommasi, D. Gorissen, and T. Dhaene. “A novel sequential design strategy for global surrogate modeling”. In: *Procs. of the Winter Simulation Conference 2009*. IEEE. 2009, pp. 731–742 (cit. on p. 17).

[153] B. Crowther. “Flocking of autonomous unmanned air vehicles”. In: *The Aeronautical Journal* 107.1068 (2003), pp. 99–109 (cit. on pp. 214, 215).

[154] L. E. Cue la Rosa and D. A. Borges Oliveira. “Learning from Label Proportions with Prototypical Contrastive Clustering”. In: *Procs. of the AAAI Conference on Artificial Intelligence 2022*. AAAI. 2022 (cit. on p. 139).

[155] F. Cus, M. Milfelner, and J. Balic. “An intelligent system for monitoring and optimization of ball-end milling process”. In: *Journal of Materials Processing Technology* 175.1-3 (2006), pp. 90–97 (cit. on p. 115).

[156] M. Dalla Mura, S. Prasad, F. Pacifici, P. Gamba, J. Chanussot, and J. A. Benediktsson. “Challenges and opportunities of multimodality and data fusion in remote sensing”. In: *Procs. of the IEEE* 103.9 (2015), pp. 1585–1601 (cit. on p. 115).

[157] C. Daniel and L. Mary. “Fusion of audio visual cues for vehicle classification”. In: *Procs. of the Int. Conference on Next Generation Intelligent Systems 2016*. 2016, pp. 1–4 (cit. on p. 245).

[158] C. A. L. Dantec et al. “A tale of two publics: Democratizing design at the margins”. In: *Procs. of the Conference on Designing Interactive Systems 2010*. 2010 (cit. on p. 298).

[159] R. Dara, S. Kremer, and D. Stacey. “Clustering unlabeled data with SOMs improves classification of labeled real-world data”. In: *Procs. of the Int. Joint Conference on Neural Networks 2002*. Vol. 3. 2002, pp. 2237–2242 (cit. on p. 139).

[160] G. Das, D. Gunopulos, and H. Mannila. “Finding Similar Time Series”. In: *Procs. of the Symposium on Principles of Data Mining and Knowledge Discovery 1997*. 1997, pp. 88–100 (cit. on p. 121).

[161] T. Davenport and R. Kalakota. “The potential for artificial intelligence in healthcare”. In: *Future Healthcare Journal* 6.2 (2019), p. 94 (cit. on p. 4).

[162] S. De, Y. Hu, V. V. Meduri, Y. Chen, and S. Kambhampati. “Bayeswipe: A scalable probabilistic framework for improving data quality”. In: *Journal of Data and Information Quality* 8.1 (2016), pp. 1–30 (cit. on p. 12).

[163] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan. “A fast and elitist multiobjective genetic algorithm: NSGA-II”. In: *IEEE Transactions on Evolutionary Computation* 6.2 (Apr. 2002), pp. 182–197 (cit. on p. 35).

[164] W. Delfmann, M. ten Hompel, W. Kersten, T. Schmidt, and W. Stölzle. “Logistics as a Science – Central Research Questions in the Era of the Fourth Industrial Revolution”. In: 11.1 (). URL: https://doi.org/10.23773/2018_9 (cit. on p. 110).

[165] D. Delling, T. Pajor, and R. Werneck. “Round-Based Public Transit Routing”. In: *Procs. of the Meeting on Algorithm Engineering and Experiments 2012*. Society for Industrial and Applied Mathematics, Jan. 2012 (cit. on p. 235).

[166] M. Demirbas, M. A. Bayir, C. G. Akcora, Y. S. Yilmaz, and H. Ferhatosmanoglu. “Crowdsourced sensing and collaboration using twitter”. In: *Procs. of the IEEE Int. Symposium on A World of Wireless, Mobile and Multimedia Networks 2010*. 2010 (cit. on p. 298).

[167] A. Demiriz, K. Bennett, K. P. Bennett, and M. J. Embrechts. “Semi-supervised clustering using genetic algorithms”. In: *Procs. of Artificial Neural Networks in Engineering 1999*. ASME Press, 1999, pp. 809–814 (cit. on p. 139).

[168] A. P. Dempster, N. M. Laird, and D. B. Rubin. “Maximum likelihood from incomplete data via the EM algorithm”. In: *Journal of the Royal Statistical Society. Series B* (1977), pp. 1–38 (cit. on pp. 141, 148).

[169] J. Demšar. “Statistical Comparisons of Classifiers over Multiple Data Sets”. In: *The Journal of Machine Learning Research* 7 (2006), pp. 1–30 (cit. on p. 153).

[170] B. Denkena, M.-A. Dittrich, and F. Uhlich. “Augmenting Milling Process Data for Shape Error Prediction”. In: *Procs. of the CIRP Conference on Manufacturing Systems 2016*. Vol. 57. 2016, pp. 487–491 (cit. on p. 158).

[171] M. A. DePristo et al. “A framework for variation discovery and genotyping using next-generation DNA sequencing data”. In: *Nature Genetics* 43.5 (May 2011). Number: 5Publisher: Nature Publishing Group, pp. 491–498. URL: https://www.nature.com/articles/ng.806 (cit. on pp. 48, 51, 55).

[172] A. Derrow-Pinion et al. “ETA Prediction with Graph Neural Networks in Google Maps”. In: *Procs. of the ACM Int. Conference on Information and Knowledge Management 2021*. 2021 (cit. on p. 240).

[173] G. Deshpande, L. Libero, K. R. Sreenivasan, H. Deshpande, and R. K. Kana. “Identification of neural connectivity signatures of autism using machine learning”. In: *Frontiers in Human Neuroscience* 7 (2013), p. 670 (cit. on p. 10).

[174] I. S. Dhillon, Y. Guan, and B. Kulis. “Kernel k-means: spectral clustering and normalized cuts”. In: *Procs. of the ACM SIGKDD Conference on Knowledge Discovery and Data Mining 2004*. Ed. by W. Kim, R. Kohavi, J. Gehrke, and W. DuMouchel. SIGKDD ’04. New York, USA: ACM Press, 2004, pp. 551–556 (cit. on pp. 141, 148).

[175] J. Dibbelt, T. Pajor, B. Strasser, and D. Wagner. “Intriguingly simple and fast transit routing”. In: *Procs. of the Int. Symposium on Experimental Algorithms 2013*. Springer. 2013, pp. 43–54 (cit. on p. 234).

[176] E. W. Dijkstra. “A note on two problems in connexion with graphs”. In: *Numerische Mathematik* 1.1 (1959), pp. 269–271 (cit. on p. 233).

[177] A. R. T. Donders, G. J. Van Der Heijden, T. Stijnen, and K. G. Moons. “A gentle introduction to imputation of missing values”. In: *Journal of Clinical Epidemiology* 59.10 (2006), pp. 1087–1091 (cit. on p. 13).

[178] Z. Dong, Y. Wu, M. Pei, and Y. Jia. “Vehicle Type Classification Using a Semisupervised Convolutional Neural Network”. In: *IEEE Transactions on Intelligent Transportation Systems* 16.4 (2015), pp. 2247–2256 (cit. on p. 245).

[179] F. Doshi-Velez and B. Kim. “Towards a Rigorous Science of Interpretable Machine Learning”. In: *arXiv: Computing Research Repository* (2017). DOI: arXiv:1702.08608 (cit. on p. 16).

[180] K. Drew et al. “Integration of over 9,000 mass spectrometry experiments builds a global map of human protein complexes”. In: *Molecular Systems Biology* 13.6 (2017), p. 932 (cit. on p. 86).

[181] P. Drineas, M. Magdon-Ismail, M. W. Mahoney, and D. P. Woodruff. “Fast approximation of matrix coherence and statistical leverage”. In: *The Journal of Machine Learning Research* 13.1 (2012), pp. 3475–3506 (cit. on p. 60).

[182] J. Du, C. Jiang, E. Gelenbe, L. Xu, J. Li, and Y. Ren. “Distributed data privacy preservation in IoT applications”. In: *IEEE Wireless Communications* 25.6 (2018), pp. 68–76 (cit. on p. 15).

[183] W. Du and Z. Zhan. “Using randomized response techniques for privacy-preserving data mining”. In: *Procs. of the ACM SIGKDD Conference on Knowledge Discovery and Data Mining 2003*. 2003, pp. 505–510 (cit. on p. 15).

[184] G. Duke, J. Briedis, and J. Green. “Survival of critically ill medical patients is time-critical”. In: *Critical Care and Resuscitation* 6.4 (2004) (cit. on p. 5).

[185] C. Dwork and A. Roth. “The Algorithmic Foundations of Differential Privacy”. In: *Foundations and Trends in Theoretical Computer Science* 9.3-4 (2014), pp. 211–407 (cit. on p. 232).

[186] J. F. Easton, C. R. Stephens, and M. Angelova. “Risk factors and prediction of very short term versus short/intermediate term post-stroke mortality: a data mining approach”. In: *Computers in Biology and Medicine* 54 (2014), pp. 199–210 (cit. on p. 15).

[187] C. Elkan. “Using the Triangle Inequality to Accelerate k-Means”. In: *Procs. of the Int. Conference on Machine Learning 2003*. 2003 (cit. on p. 151).

[188] T. Emmanuel, T. Maupong, D. Mpoeleng, T. Semong, M. Banyatsang, and O. Tabona. “A Survey On Missing Data in Machine Learning”. In: (2021) (cit. on p. 13).

[189] V. Emmerich, M. Döbele, T. Bauernhansl, D. Paulus-Rohmer, A. Schatz, and M. Weskamp. *Geschäftsmodell-Innovation durch Industrie 4.0-Chancen und Risiken für den Maschinen- und Anlagenbau*. 2015 (cit. on p. 104).

[190] H. J. Escalante et al. *Explainable and interpretable models in computer vision and machine learning*. Springer, 2018 (cit. on p. 17).

[191] M. Ester, H.-P. Kriegel, J. Sander, and X. Xu. “A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise”. In: *Procs. of the Int. Conference on Knowledge Discovery and Data Mining 1996*. AAAI Press, 1996, pp. 226–231 (cit. on p. 148).

[192] ETSI. *Short Range Devices (SRD) operating in the frequency range 25 MHz to 1 000 MHz; Part 2: Harmonised Std. covering the essential requirements of article 3.2 of Directive 2014/53/EU for non specific radio equipment*. 2016 (cit. on p. 315).

[193] *EU Science Hub Transport sector economic analysis*. https://ec.europa.eu/jrc/en/research-topic/transport-sector-economic-analysis. Accessed: 2021-02-01. 2021 (cit. on p. 213).

[194] R. Falkenberg, K. Heimann, and C. Wietfeld. “Discover Your Competition in LTE: Client-Based Passive Data Rate Prediction by Machine Learning”. In: *Procs. of the IEEE Global Communications Conference 2017*. Singapore, Dec. 2017, pp. 1–7 (cit. on p. 347). **SFB876-A4**

[195] R. Falkenberg, C. Ide, and C. Wietfeld. “Client-Based Control Channel Analysis for Connectivity Estimation in LTE Networks”. In: *Procs. of the IEEE Vehicular Technology Conference 2016*. Accepted for presentation. IEEE. Montréal, Canada, Sept. 2016 (cit. on pp. 344, 347, 348). **SFB876-A4**

[196] R. Falkenberg and C. Wietfeld. “FALCON: An accurate real-time monitor for client-based mobile network data analytics”. In: *Procs. of the IEEE Global Communications Conference 2019*. Waikoloa, Hawaii, USA: IEEE, Dec. 2019, pp. 1–7. URL: http://dx.doi.org/10.1109/GLOBECOM38437.2019.9014096 (cit. on pp. 345, 348, 351). **SFB876-A4**

[197] R. Falkenberg et al. “PhyNetLab: An IoT-based warehouse testbed”. In: *Procs. of the Federated Conference on Computer Science and Information Systems 2017*. Sept. 2017 (cit. on p. 200). **SFB876-A4**

[198] K. Fan, H. Zhang, S. Yan, L. Wang, W. Zhang, and J. Feng. “Learning a Generative Classifier from Label Proportions”. In: *Neurocomputing* 139 (Sept. 2014), pp. 47–55 (cit. on p. 142).

[199] L. Fan and G. Taylor. “Alter-CNN: An Approach to Learning from Label Proportions withApplication to Ice-Water Classification”. In: *Procs. of the Workshop on Learning and Privacy with Incomplete Data and Weak Supervision at the Conference on Neural Information Processing Systems 2015*. 2015 (cit. on p. 142).

[200] H. Fanaee-T and J. Gama. “Event labeling combining ensemble detectors and background knowledge”. In: *Progress in Artificial Intelligence* 2.2 (June 2014), pp. 113–127 (cit. on p. 164).

[201] U. M. Fayyad, G. Piatetsky-Shapiro, and P. Smyth. “From Data Mining to Knowledge Discovery: An Overview”. In: *Advances in Knowledge Discovery and Data Mining*. Ed. by U. M. Fayyad, G. Piatetsky-Shapiro, P. Smyth, and R. Uthurusamy. AAAI/MIT Press, 1996. Chap. 1, pp. 1–34 (cit. on p. 115).

[202] U. M. Fayyad and P. Smyth. “Towards Intelligent Trainable Tools for the Automated Analysis, Cataloging, and Searching of Digital Image Libraries: A Machine Learning Approach”. In: (1994) (cit. on p. 12).

[203] L. Feltrin, C. Buratti, E. Vinciarelli, R. D. Bonis, and R. Verdone. “LoRaWAN: Evaluation of Link- and System-Level Performance”. In: *IEEE Internet of Things Journal* 5.3 (June 2018), pp. 2249–2258 (cit. on p. 315).

[204] N. Femia, G. Petrone, G. Spagnuolo, and M. Vitelli. *Power Electronics and Control Techniques for Maximum Energy Harvesting in Photovoltaic Systems*. CRC Press, 2012 (cit. on p. 198).

[205] F. T. Fernandes, T. A. de Oliveira, C. E. Teixeira, A. F. de Moraes Batista, G. Dalla Costa, and A. D. P. Chiavegatto Filho. “A multipurpose machine learning approach to predict COVID-19 negative prognosis in São Paulo, Brazil”. In: *Scientific Reports* 11.1 (2021), pp. 1–7 (cit. on p. 11).

[206] F. Finkeldey, A. Saadallah, P. Wiederkehr, and K. Morik. “Real-time prediction of process forces in milling operations using synchronized data fusion of simulation and sensor data”.

In: *Engineering Applications of Artificial Intelligence* Volume 94 (Sept. 2020). URL: https://www.sciencedirect.com/science/article/pii/S0952197620301627?dgcid=coauthor%5C#fig6 (cit. on pp. 159, 162, 166, 167, 171). **SFB876-B3**

[207] F. Finkeldey, A. Wirtz, T. Merhofe, and P. Wiederkehr. "Learning-Based Prediction of Pose-Dependent Dynamics". In: *Journal of Manufacturing and Materials Processing* 4.3 (2020). URL: https://www.mdpi.com/2504-4494/4/3/85 (cit. on pp. 159, 172–175). **SFB876-B3**

[208] C. Fischer. "Runtime and Accuracy Issues in Three-dimensional Finite Element Simulation of Machining". In: *International Journal of Machining and Machinability of Materials* 6.1 (2009), p. 35 (cit. on p. 158).

[209] B. Fish and L. Reyzin. "On the Complexity of Learning a Class Ratio from Unlabeled Data". In: *Journal of Artificial Intelligence Research* (Dec. 2020), pp. 1333–1349 (cit. on p. 139).

[210] J. Foley, A. van Dam, and S. Feiner. *Computer Graphics: Principles and Practice*. 2-nd edition in C. The systems programming series. Reading, Mass., Menlo Park, Calif., and New York: Addison-Wesley Publishing, 1997 (cit. on p. 159).

[211] V. Fortino et al. "Machine-learning–driven biomarker discovery for the discrimination between allergic and irritant contact dermatitis". In: *Procs. of the National Academy of Sciences* 117.52 (2020), pp. 33474–33485 (cit. on pp. 9, 10).

[212] S. Freitag, B. Cao, J. Ninic, and G. Meschke. "Surrogate modeling for mechanized tunneling simulations with uncertain data". In: *Procs. of the Int. Workshop on Reliable Engineering Computing 2014*. 2014, pp. 44–63 (cit. on p. 160).

[213] J. H. Friedman. "Greedy function approximation: a gradient boosting machine". In: *Annals of Statistics* (2001), pp. 1189–1232 (cit. on p. 17).

[214] J. H. Friedman and B. E. Popescu. "Predictive learning via rule ensembles". In: *The Annals of Applied Statistics* 2.3 (2008), pp. 916–954 (cit. on p. 17).

[215] B. Friedrich. "The Effect of Autonomous Vehicles on Traffic". In: *Autonomous driving*. Berlin: Springer Open, 2016, pp. 317–334 (cit. on p. 295).

[216] J. Futoma, S. Hariharan, and K. Heller. "Learning to detect sepsis with a multitask Gaussian process RNN classifier". In: *Procs. of the Int. Conference on Machine Learning 2017*. PMLR. 2017, pp. 1174–1182 (cit. on p. 11).

[217] A. Gal, A. Mandelbaum, F. Schnitzler, A. Senderovich, and M. Weidlich. "Traveling time prediction in scheduled transportation with journey segments". In: *Information Systems* (2015) (cit. on p. 235).

[218] J. Gama and M. M. Gaber. *Learning from data streams: processing techniques in sensor networks*. Springer, 2007 (cit. on p. 126).

[219] X. Gao, Z. Chen, S. Tang, Y. Zhang, and J. Li. "Adaptive weighted imbalance learning with application to abnormal activity recognition". In: *Neurocomputing* 173 (2016), pp. 1927–1935 (cit. on p. 125).

[220] M. R. Garey and D. S. Johnson. *Computers and Intractability – A Guide to NP-Completeness*. Int. Computer Science Series. New York: Freeman, 1979 (cit. on p. 88).

[221] V. N. Garla and C. Brandt. "Ontology-guided feature engineering for clinical text classification". In: *Journal of Biomedical Informatics* 45.5 (2012), pp. 992–998 (cit. on p. 12).

[222] E. Garrison and G. Marth. "Haplotype-based variant detection from short-read sequencing". In: *arXiv: Quantitative Biology* (July 2012). DOI: arXiv:1207.3907 (cit. on p. 55).

[223] R. Geisberger, P. Sanders, D. Schultes, and D. Delling. "Contraction hierarchies: Faster and simpler hierarchical routing in road networks". In: *Procs. of the Int. Workshop on Experimental and Efficient Algorithms 2008*. Springer. 2008, pp. 319–333 (cit. on p. 234).

[224] E. I. George and R. E. McCulloch. "Variable selection via Gibbs sampling". In: *Journal of the American Statistical Association* 88.423 (1993), pp. 881–889 (cit. on pp. 63, 64).

[225] J. George, L. Mary, and R. K. S. “Vehicle detection and classification from acoustic signal using ANN and KNN”. In: *Procs. of the Int. Conference on Control Communication and Computing 2013*. 2013, pp. 436–439 (cit. on p. 245).

[226] L. N. Geppert, K. Ickstadt, A. Munteanu, J. Quedenfeld, and C. Sohler. “Random projections for Bayesian regression”. In: *Statistics and Computing* 27.1 (2017), pp. 79–101. URL: http://doi.org/10.1007/s11222-015-9608-z (cit. on pp. 59, 60, 69). **SFB876-C4**

[227] L. N. Geppert, K. Ickstadt, A. Munteanu, and C. Sohler. “Streaming statistical models via Merge & Reduce”. In: *International Journal of Data Science and Analytics* 10.4 (2020), pp. 331–347. URL: https://doi.org/10.1007/s41060-020-00226-0 (cit. on pp. 59, 66, 68–70). **SFB876-C4**

[228] S. Gezici and H. V. Poor. “Position Estimation via Ultra-Wide-Band Signals”. In: *Procs. of the IEEE* 97.2 (Feb. 2009), pp. 386–403 (cit. on p. 181).

[229] Z. Gheid and Y. Challal. “Efficient and privacy-preserving k-means clustering for big data mining”. In: *Procs. of the IEEE International Conference on Trust, Security and Privacy in Computing and Communications 2016*. IEEE. 2016, pp. 791–798 (cit. on p. 15).

[230] N. Ghosh et al. “Estimation of tool wear during CNC milling using neural network-based sensor fusion”. In: *Mechanical Systems and Signal Processing* 21.1 (2007), pp. 466–479 (cit. on p. 158).

[231] M. Goddard. “The EU General Data Protection Regulation (GDPR): European regulation that has a global impact”. In: *International Journal of Market Research* 59.6 (2017), pp. 703–705 (cit. on p. 231).

[232] M. Goerigk, M. Knoth, M. Müller-Hannemann, M. Schmidt, and A. Schöbel. “The Price of Robustness in Timetable Information”. In: *Procs. of the Workshop on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems 2011*. Ed. by A. Caprara and S. Kontogiannis. Vol. 20. OpenAccess Series in Informatics (OASIcs). Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2011, pp. 76–87. DOI: 10.4230/OASIcs.ATMOS.2011.76. URL: http://drops.dagstuhl.de/opus/volltexte/2011/3268 (cit. on p. 234).

[233] A. Goldstein, A. Kapelner, J. Bleich, and E. Pitkin. “Peeking inside the black box: Visualizing statistical learning with plots of individual conditional expectation”. In: *Journal of Computational and Graphical Statistics* 24.1 (2015), pp. 44–65 (cit. on p. 17).

[234] J. R. González, L. Armengol, X. Solé, E. Guinó, and V. Moreno. *SNPassoc: SNPs-Based Whole Genome Association Studies*. R package version 1.9-2. 2014. URL: https://CRAN.R-project.org/package=SNPassoc (cit. on p. 61).

[235] I. Goodfellow. “NIPS 2016 Tutorial: Generative Adversarial Networks”. In: *arXiv: Computing Research Repository* (2016). DOI: arXiv:1701.00160 (cit. on p. 16).

[236] I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. MIT Press, 2016. URL: http://www.deeplearningbook.org (cit. on pp. 8, 245).

[237] I. Goodfellow et al. “Generative Adversarial Nets”. In: *Advances in Neural Information Processing Systems 27: Procs. of the 2014 Conference*. 2014, pp. 2672–2680. URL: http://papers.nips.cc/paper/5423-generative-adversarial-nets.pdf (cit. on p. 38).

[238] N. Görnitz, M. Kloft, and U. Brefeld. “Active and Semi-Subpervised Data Domain Description”. In: *Procs. of the European Conference on Machine Learning and Knowledge Discovery in Databases 2009*. Ed. by W. Buntine, M. Grobelnik, D. Mladenic, and J. Shawe-Taylor. Berlin, Heidelberg: Springer, 2009, pp. 407–422 (cit. on p. 125).

[239] K. Gouda and M. Hassaan. “CSI_GED: An efficient approach for graph edit similarity computation”. In: *Procs. of the IEEE Int. Conference on Data Engineering 2016*. 2016, pp. 265–276 (cit. on p. 88).

[240] M. Goyal, T. Knackstedt, S. Yan, and S. Hassanpour. “Artificial intelligence-based image classification for diagnosis of skin cancer: Challenges and opportunities”. In: *Computers in Biology and Medicine* (2020), p. 104065 (cit. on p. 17).

[241] E. Graf, C. Schmoor, W. Sauerbrei, and T. Schumacher. “Assessment and comparison of prognostic classification schemes for survival data”. In: *Statistics in Medicine* 18.17-18 (1999), pp. 2529–2545 (cit. on p. 65).

[242] B. Grüning et al. “Bioconda: sustainable and comprehensive software distribution for the life sciences”. In: *Nature Methods* 15.7 (July 2018), pp. 475–476 (cit. on p. 44).

[243] M. Guarnieri, S. Marinovic, and D. A. Basin. “Securing Databases from Probabilistic Inference”. In: *Procs. of the Computer Security Foundations Workshop 2017*. 2017, pp. 343–359 (cit. on p. 412).

[244] T. Gumbsch, C. Bock, M. Moor, B. Rieck, and K. Borgwardt. “Enhancing statistical power in temporal biomarker discovery through representative shapelet mining”. In: *Bioinformatics* 36.Supplement 2 (2020), pp. i840–i848 (cit. on p. 15).

[245] E. L. Gurevich, V. Temchura, K. Überla, and A. Zybin. “Analytical Features of Particle Counting Sensor based on Plasmon assisted Microscopy of Nano Objects”. In: *Sensors and Actuators B: Chemical* B 160.1 (Sept. 2011), pp. 1210–1215. URL: http://www.sciencedirect.com/science/article/pii/S0925400511008525 (cit. on pp. 24, 25). **SFB876-B2**

[246] L. Habel, C. Ide, C. Wietfeld, and M. Schreckenberg. “Improving Vehicular Traffic Simulations Using Real-Time Information on Environmental Conditions”. In: *Procs. of the IEEE Vehicular Technology Conference 2016*. IEEE. Montreal, Canada: IEEE, Sept. 2016. URL: http://dx.doi.org/10.1109/VTCFall.2016.7881044 (cit. on p. 288). **SFB876-B4**

[247] L. Habel, A. Molina, T. Zaksek, K. Kersting, and M. Schreckenberg. “Traffic simulations with empirical data – How to replace missing traffic flows?” In: *Traffic and Granular Flow ’15*. Ed. by V. L. Knoop and W. Daamen. Springer, May 2016, pp. 491–498. URL: http://dx.doi.org/10.1007/978-3-319-33482-0_62 (cit. on pp. 230, 289). **SFB876-B4**

[248] L. Habel and M. Schreckenberg. “Asymmetric Lane Change Rules for a Microscopic Highway Traffic Model”. In: *Procs. of Int. Conference on Cellular Automata for Research and Industry 2014*. Vol. 8751. LNCS. Krakow, Poland: Springer, Sept. 2014, pp. 620–629. URL: http://dx.doi.org/10.1007/978-3-319-11520-7_66 (cit. on p. 287). **SFB876-B4**

[249] F. Hadiji, A. Molina, S. Natarajan, and K. Kersting. “Poisson Dependency Networks: Gradient Boosted Models for Multivariate Count Data”. In: *Machine Learning Journal* 100.2 (2015), pp. 477–507. URL: http://dx.doi.org/10.1007/s10994-015-5506-z (cit. on pp. 230, 289). **SFB876-B4**

[250] S. Haendeler, A. Lewandowski, and C. Wietfeld. “Passive Detection of Wrong Way Drivers on Motorways Based on Low Power Wireless Communications”. In: *Procs. of the Vehicular Technology Conference 2014*. 2014, pp. 1–5 (cit. on pp. 242, 246, 247).

[251] M. Haferkamp, B. Sliwa, and C. Wietfeld. “A Low Cost Modular Radio Tomography System for Bicycle and Vehicle Detection and Classification”. In: *Procs. of the IEEE Int. Systems Conference 2021*. Vancouver, Canada, Apr. 2021 (cit. on pp. 243, 246). **SFB876-A4, SFB876-B4**

[252] M. Haferkamp et al. “Radio-based Traffic Flow Detection and Vehicle Classification for Future Smart Cities”. In: *Procs. of the IEEE Vehicular Technology Conference 2017 Workshop on Connecting All Things for Enabling Smart Cities (CONTEST)*. Sydney, June 2017 (cit. on pp. 243, 246). **SFB876-B4**

[253] S. Häger, S. Böcker, S. Jamali, T. Reinsch, and C. Wietfeld. “A Novel System Architecture for Small-Scale Motion Sensing Exploiting 5G mmWave Channels”. In: *Procs. of the IEEE Globecom Workshops 2021*. Madrid, Spain, Dec. 2021. DOI: 10.1109/GCWkshps52748.2021.9682166 (cit. on p. 388). **SFB876-A4**

[254] R. N. Hall. “Electron-Hole Recombination in Germanium”. In: 87.2 (1952), pp. 387–387. URL: https://link.aps.org/doi/10.1103/PhysRev.87.387 (cit. on p. 197).

[255] J. Y. Halpern and K. R. O’Neill. “Secrecy in Multiagent Systems”. In: *ACM Transactions on Information and System Security* 12.1 (2008) (cit. on p. 412).

[256] M. Hamer and R. D’Andrea. “Self-Calibrating Ultra-Wideband Network Supporting Multi-Robot Localization”. In: *IEEE Access* 6 (2018), pp. 22292–22304 (cit. on p. 190).

[257] F. Harary and R. Norman. “Some properties of line digraphs”. In: *Rendiconti del Circolo Matematico di Palermo* 9.2 (May 1960), pp. 161–168. URL: http://dx.doi.org/10.1007/BF02854581 (cit. on p. 237).

[258] P. E. Hart, N. J. Nilsson, and B. Raphael. “A Formal Basis for the Heuristic Determination of Minimum Cost Paths”. In: *Autonomous Mobile Robots*. Ed. by S. S. Iyengar and A. Elfes. (also to appear in IEEE Transactions on Systems Science, and Cybernetics, 1968). Los Alamitos, Calif. [u.a.]: IEEE Computer Soc. Press, 1991, pp. 375–382 (cit. on pp. 233, 235).

[259] T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning: Data Mining, Inference and Prediction*. 2nd. Statistics. Springer, 2009 (cit. on p. 146).

[260] T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer series in statistics. New York, USA: Springer, 2001 (cit. on p. 124).

[261] Q. A. Hathaway et al. “Machine-learning to stratify diabetic patients using novel cardiac biomarkers and integrative genomics”. In: *Cardiovascular Diabetology* 18.1 (2019), pp. 1–16 (cit. on p. 10).

[262] M. Hayes. “Enables: Research Infrastructure Position Paper”. In: *Position Paper* 3.3 (2019), p. 58 (cit. on p. 262).

[263] K. He, G. Gkioxari, P. Dollár, and R. B. Girshick. “Mask R-CNN”. In: *arXiv: Computing Research Repository* (2017). DOI: arXiv:1703.06870 (cit. on p. 32).

[264] R. Hecht-Nielsen. “Theory of the Backpropagation Neural Network”. In: *Neural Networks for Perception (Vol. 2): Computation, Learning, Architectures*. USA: Harcourt Brace & Co., 1992, pp. 65–93 (cit. on p. 31).

[265] D. Heckerman, D. M. Chickering, C. Meek, R. Rounthwaite, and C. M. Kadie. “Dependency Networks for Collaborative Filtering and Data Visualization”. In: *Procs. of the Conference on Uncertainty in Artificial Intelligence 2000*. 2000, pp. 264–273. URL: https://dslpitt.org/uai/displayArticleDetails.jsp?mmnu=1%5C&smnu=2%5C&article_id=31%5C&proceeding_id=16 (cit. on p. 230).

[266] K. Heimann, P. Gorczak, C. Bektas, F. Girke, and C. Wietfeld. “Software-Defined End-to-End Evaluation Platform for Quality of Service in Non-Standalone 5G Systems”. In: *Procs. of the IEEE Int. Systems Conference 2019*. IEEE, 2019 (cit. on pp. 384, 385). **SFB876-A4, SFB876-B4**

[267] K. Heimann, A. Marsch, B. Sliwa, and C. Wietfeld. “Reflecting Surfaces for Beyond Line-Of-Sight Coverage in Millimeter Wave Vehicular Networks”. In: *Procs. of the IEEE Vehicular Networking Conference 2020*. Dec. 2020. URL: https://ieeexplore.ieee.org/document/9318411 (cit. on pp. 385–387). **SFB876-A4, SFB876-B4**

[268] K. Heimann, J. Tiemann, S. Böcker, and C. Wietfeld. “Cross-Bearing based Positioning as a Feature of 5G Millimeter Wave Beam Alignment”. In: *Procs. of the IEEE Vehicular Technology Conference 2020*. Antwerp, Belgium, May 2020. URL: https://ieeexplore.ieee.org/document/9129182 (cit. on pp. 377, 381–383). **SFB876-B4**

[269] K. Heimann, J. Tiemann, S. Böcker, and C. Wietfeld. “On the Potential of 5G mmWave Pencil Beam Antennas for UAV Communications: An Experimental Evaluation”. In: *Procs. of the Int. ITG Workshop on Smart Antennas 2018*. Mar. 2018 (cit. on pp. 378, 379). **SFB876-A4, SFB876-B4**

[270] K. Heimann, J. Tiemann, D. Yolchyan, and C. Wietfeld. "Experimental 5G mmWave Beam Tracking Testbed for Evaluation of Vehicular Communications". In: *IEEE 5G World Forum* (Sept. 2019). URL: https://ieeexplore.ieee.org/document/8911692 (cit. on pp. 378, 380, 381). **SFB876-A4, SFB876-B4**

[271] D. Helbing. *Verkehrsdynamik: neue physikalische Modellierungskonzepte*. Springer-Verlag, 2013 (cit. on p. 226).

[272] C. Hernandez, C. Mella, G. Navarro, A. Olivera-Nappa, and J. Araya. "Protein complex prediction via dense subgraphs and false positive analysis". In: *PLoS One* 12.9 (2017), e0183460 (cit. on p. 86).

[273] J. Hernández-González, I. Inza, I. Granado, O. C. Basurko, J. A. Fernandes, and J. A. Lozano. "Aggregated outputs by linear models: An application on marine litter beaching prediction". In: *Information Sciences* 481 (2019), pp. 381–393 (cit. on p. 139).

[274] J. Hernández-González, I. Inza, and J. A. Lozano. "Learning Bayesian network classifiers from label proportions". In: *Pattern Recognition* 46.12 (2013), pp. 3425–3440. URL: http://www.sciencedirect.com/science/article/pii/S0031320313002069 (cit. on p. 142).

[275] S. Herranz et al. "The milling of airframe components with low rigidity: A general approach to avoid static and dynamic problems". In: *Procs. of the I MECH E Part B Journal of Engineering Manufacture* 219 (Nov. 2005), pp. 789–802 (cit. on p. 157).

[276] A. Higgins and E. Kozan. "Modeling train delays in urban networks". In: *Transportation Science* 32.4 (1998), pp. 346–357 (cit. on p. 234).

[277] S. Hochreiter and J. Schmidhuber. "Long Short-Term Memory". In: *Neural Computation* (1997), pp. 1735–1780 (cit. on pp. 5, 11, 326).

[278] F. Holsboer. "How can we realize the promise of personalized antidepressant medicines?" In: *Nature Reviews Neuroscience* 9.8 (2008), pp. 638–646 (cit. on p. 3).

[279] M. ten Hompel, B. Vogel-Heuser, and T. Bauernhansl. *Handbuch Industrie 4.0*. Springer, 2020 (cit. on pp. 103, 110).

[280] S. Hoogendoorn, P. Bovy, and W. Daamen. *Pedestrian and evacuation dynamics*. 2002 (cit. on p. 227).

[281] G. Hooker. "Discovering additive structure in black box functions". In: *Procs. of the ACM SIGKDD Conference on Knowledge Discovery and Data Mining 2004*. 2004, pp. 575–580 (cit. on p. 17).

[282] R. P. Horgan and L. C. Kenny. "'Omic' technologies: genomics, transcriptomics, proteomics and metabolomics". In: *The Obstetrician & Gynaecologist* 13.3 (2011), pp. 189–195 (cit. on p. 9).

[283] D. W. Hosmer Jr, S. Lemeshow, and R. X. Sturdivant. *Applied logistic regression*. Vol. 398. John Wiley & Sons, 2013 (cit. on p. 5).

[284] Q. Hou et al. "RankProd combined with genetic algorithm optimized artificial neural network establishes a diagnostic and prognostic prediction model that revealed C1QTNF3 as a biomarker for prostate cancer". In: *EBioMedicine* 32 (2018), pp. 234–244 (cit. on p. 10).

[285] R. Houari, A. Bounceur, A. K. Tari, and M. T. Kecha. "Handling missing data problems with sampling methods". In: *Procs. of the Int. Conference on Advanced Networking Distributed Systems and Applications 2014*. IEEE. 2014, pp. 99–104 (cit. on p. 13).

[286] C.-T. Hsueh, D. Liu, and H. Wang. "Novel biomarkers for diagnosis, prognosis, targeted therapy and clinical trials". In: *Biomarker Research* 1.1 (2013), pp. 1–2 (cit. on p. 9).

[287] M. Hu. "Cambridge Analytica's black box". In: *Big Data & Society* 7.2 (2020), p. 2053951720938091 (cit. on p. 14).

[288] D. Huang, T. T. Allen, W. I. Notz, and N. Zeng. "Global Optimization of Stochastic Black-Box Systems via Sequential Kriging Meta-Models". In: *Journal of Global Optimization* 34.3 (Mar. 2006), pp. 441–466 (cit. on p. 80).

[289] J. Huang et al. “An In-Depth Study of LTE: Effect of Network Protocol and Application Behavior on Performance”. In: *Procs. of the ACM Conference on SIGCOMM 2013*. SIGCOMM ’13. New York, NY, USA: ACM, 2013, pp. 363–374 (cit. on p. 345).

[290] N. Hubballi, B. K. Patra, and S. Nandi. “NDoT: nearest neighbor distance based outlier detection technique”. In: *Procs. of the Int. Conference on Pattern Recognition and Machine Intelligence 2011*. Springer. 2011, pp. 36–42 (cit. on p. 14).

[291] T. Huck, A. Westenberger, M. Fritzsche, T. Schwarz, and K. Dietmayer. “Precise timestamping and temporal synchronization in multi-sensor fusion”. In: *Procs. of the IEEE Intelligent Vehicles Symposium 2011*. IEEE. 2011, pp. 242–247 (cit. on p. 120).

[292] U. Hunkeler, H. L. Truong, and A. Stanford-Clark. “MQTT-S – A publish/subscribe protocol for Wireless Sensor Networks”. In: *Procs. of the Int. Conference on Communication Systems Software and Middleware and Workshops 2008*. IEEE. 2008, pp. 791–798 (cit. on p. 220).

[293] S. L. Hyland et al. “Early prediction of circulatory failure in the intensive care unit using machine learning”. In: *Nature Medicine* 26.3 (2020), pp. 364–373 (cit. on pp. 5, 6, 11).

[294] R. Hyndman and Y. Khandakar. “Automatic Time Series Forecasting: The forecast Package for R”. In: *Journal of Statistical Software, Articles* 27.3 (2008), pp. 1–22. URL: https://www.jstatsoft.org/v027/i03 (cit. on p. 325).

[295] K. Ickstadt, M. Schäfer, and M. Zucknick. “Toward Integrative Bayesian Analysis in Molecular Biology”. In: *Annual Review of Statistics and Its Application* 5.1 (Mar. 2018), pp. 141–167. URL: https://doi.org/10.1146/annurev-statistics-031017-100438 (cit. on p. 63). **SFB876-C4**

[296] C. Ide. “Resource-efficient LTE machine-type communication in vehicular environments”. PhD thesis. TU Dortmund University, Mar. 2016 (cit. on pp. 333, 334). **SFB876-B4**

[297] C. Ide, R. Falkenberg, D. Kaulbars, and C. Wietfeld. “Empirical Analysis of the Impact of LTE Downlink Channel Indicators on the Uplink Connectivity”. In: *Procs. of the IEEE Vehicular Technology Conference 2016*. Accepted for presentation. IEEE. Nanjing, China, May 2016 (cit. on pp. 333, 345). **SFB876-A4, SFB876-B4**

[298] C. Ide et al. “LTE Connectivity and Vehicular Traffic Prediction based on Machine Learning Approaches”. In: *Procs. of the IEEE Vehicular Technology Conference 2015*. Boston, USA, Sept. 2015. URL: http://ieeexplore.ieee.org/xpl/login.jsp?tp=%5C&arnumber=7391019%5C&url=http%5C%3A%5C%2F%5C%2Fieeexplore.ieee.org%5C%2Fxpls%5C%2Fabs_all.jsp%5C%3Farnumber%5C%3D7391019 (cit. on pp. 333, 334, 336). **SFB876-B4**

[299] L. G. Jaimes, I. J. Vergara-Laurens, and A. Raij. “A Survey of Incentive Techniques for Mobile Crowd Sensing”. In: *IEEE Internet of Things Journal* 2.5 (2015) (cit. on p. 298).

[300] A. Jain, M. Murty, and P. Flynn. “Data clustering: a review. ACM Computing Surveys”. In: 3.31 (1999), pp. 264–323 (cit. on p. 124).

[301] G. H. John and P. Langley. “Estimating Continuous Distributions in Bayesian Classifiers”. In: *Procs. of the Conference on Uncertainty in Artificial Intelligence 1995*. San Francisco: Morgan Kaufmann, 1995, pp. 338–345 (cit. on pp. 151, 152).

[302] J. W. Johnson. “Adapting Mask-RCNN for Automatic Nucleus Segmentation”. In: *arXiv: Computing Research Repository* (2018). DOI: arXiv:1805.00500 (cit. on p. 32).

[303] D. R. Jones, M. Schonlau, and W. J. Welch. “Efficient Global Optimization of Expensive Black-Box Functions”. In: *Journal of Global Optimization* 13.4 (1998), 455–492. URL: http://link.springer.com/article/10.1023/A:1008306431147 (cit. on p. 71).

[304] P. Jörke, S. Böcker, F. Liedmann, and C. Wietfeld. “Urban channel models for smart city IoT-networks based on empirical measurements of LoRa-lmks at 433 and 868 MHz”. In: *Procs. of the IEEE Int. Symposium on Personal, Indoor and Mobile Radio Communications 2017*. Vol. 2017-Octob. 2017 (cit. on pp. 315, 318). **SFB876-A4**

[305] P. Jörke, R. Falkenberg, and C. Wietfeld. “Power Consumption Analysis of NB-IoT and eMTC in Challenging Smart City Environments”. In: *Procs. of the IEEE Global Communications*

Conference Workshops, Workshop on Green and Sustainable 5G Wireless Networks 2018. IEEE, Dec. 2018 (cit. on pp. 264, 267, 268).

[306] P. Jörke, J. Güldenring, S. Böcker, and C. Wietfeld. "Coverage and Link Quality Improvement of Cellular IoT Networks with Multi-Operator and Multi-Link Strategies". In: *Procs. of the IEEE Vehicular Technology Conference 2019*. IEEE, Apr. 2019 (cit. on pp. 269, 270). **SFB876-A4**

[307] A. Juliani et al. "Unity: A General Platform for Intelligent Agents". In: *arXiv: Computing Research Repository* (2018). DOI: arXiv:1809.02627 (cit. on pp. 220–222).

[308] J. Jumper et al. "Highly accurate protein structure prediction with AlphaFold". In: *Nature* 596.7873 (2021), pp. 583–589 (cit. on p. 19).

[309] N. Kaempchen and K. Dietmayer. "Data synchronization strategies for multi-sensor fusion". In: *Procs. of the IEEE Conference on Intelligent Transportation Systems 2003*. Vol. 85. 1. 2003, pp. 1–9 (cit. on p. 120).

[310] H. Kagermann, W. Wahlster, and J. Helbig. *Recommendations for Implementing the Strategic Initiative INDUSTRIE 4.0*. 2013 (cit. on p. 104).

[311] V. Kann. "On the Approximability of the Maximum Common Subgraph Problem". In: *Procs. of the Annual Symposium on Theoretical Aspects of Computer Science 1992*. Vol. 577. LNCS. Springer, 1992, pp. 377–388 (cit. on p. 88).

[312] T. Karras, M. Aittala, J. Hellsten, S. Laine, J. Lehtinen, and T. Aila. "Training Generative Adversarial Networks With Limited Data". In: *arXiv: Computing Research Repository* (2020). DOI: arXiv:2006.06676 (cit. on p. 39).

[313] M. W. Kattan and T. A. Gerds. "The index of prediction accuracy: an intuitive measure useful for evaluating risk prediction models". In: *Diagnostic and Prognostic Research* 2.1 (May 2018). URL: https://doi.org/10.1186/s41512-018-0029-2 (cit. on p. 73).

[314] G. Katz, E. C. R. Shin, and D. Song. "ExploreKit: Automatic Feature Generation and Selection". In: *Procs. of the IEEE Int. Conference on Data Mining 2016*. 2016 (cit. on p. 122).

[315] G. Ke et al. "LightGBM: A Highly Efficient Gradient Boosting Decision Tree". In: *Advances in Neural Information Processing Systems 30: Procs. of the 2017 Conference*. Ed. by I. Guyon et al. Vol. 30. Curran Associates, Inc., 2017. URL: https://proceedings.neurips.cc/paper/2017/file/6449f44a102fde848669bdd9eb6b76fa-Paper.pdf (cit. on pp. 5, 9).

[316] R. M. Keenan and L. .-. Tran. "Fall detection using Wi-Fi signals and threshold-based activity segmentation". In: *Procs. of the IEEE Int. Symposium on Personal, Indoor and Mobile Radio Communications 2020*. 2020, pp. 1–6 (cit. on p. 246).

[317] B. S. Kerner, S. L. Klenov, and M. Schreckenberg. "Simple cellular automaton model for traffic breakdown, highway capacity, and synchronized flow". In: *Physical Review E* 84.4 (2011), p. 46110 (cit. on p. 294).

[318] A. A. Khalil, A. J. Byrne, M. A. Rahman, and M. H. Manshaei. "REPlanner: Efficient UAV Trajectory-Planning using Economic Reinforcement Learning". In: *Procs. of the IEEE Int. Conference on Smart Computing 2021*. IEEE. 2021, pp. 153–160 (cit. on p. 222).

[319] A. Khan, Y. Wu, C. C. Aggarwal, and X. Yan. "NeMa: Fast Graph Search with Label Similarity". In: *Procs. of the Very Large Data Bases Endowment 2013* 6.3 (2013), pp. 181–192. URL: http://www.vldb.org/pvldb/vol6/p181-khan.pdf (cit. on p. 89).

[320] P. Kidger, J. Morrill, and T. Lyons. "Generalised Interpretable Shapelets for Irregular Time Series". In: *arXiv: Computing Research Repository* (2020). DOI: arXiv:2005.13948 (cit. on p. 9).

[321] O. Kienzle. "Die Bestimmung von Kräften und Leistungen an spanenden Werkzeugen und Werkzeugmaschinen". In: *VDI-Z* 94.11 (1952), pp. 299–305 (cit. on p. 160).

[322] H.-J. Kim, N.-O. Jo, and K.-S. Shin. "Optimization of cluster-based evolutionary undersampling for the artificial neural networks in corporate bankruptcy prediction". In: *Expert Systems with Applications* 59 (2016), pp. 226–234 (cit. on p. 125).

[323] J. Kim, D. Choi, and C. Li. “Inves: Incremental Partitioning-Based Verification for Graph Similarity Search”. In: *Advances in Database Technology - Procs. of the Int. Conference on Extending Database Technology 2019*. Ed. by M. Herschel, H. Galhardas, B. Reinwald, I. Fundulaki, C. Binnig, and Z. Kaoudi. OpenProceedings.org, 2019, pp. 229–240 (cit. on pp. 93, 94).

[324] D. Kinane et al. “Intelligent Synthesis and Real-time Response using Massive Streaming of Heterogeneous Data (INSIGHT) and its anticipated effect on Intelligent Transport Systems (ITS) in Dublin City, Ireland”. In: *Procs. of the ITS European Congress 2014*. 2014 (cit. on p. 238).

[325] A. K. Kiss, D. Hajdu, D. Bachrathy, and G. Stepan. “Operational stability prediction in milling based on impact tests”. In: *Mechanical Systems and Signal Processing* 103 (2018), pp. 327–339 (cit. on p. 173).

[326] J. P. Klein and M. L. Moeschberger. *Survival Analysis: Techniques for Censored and Truncated Data*. 2nd ed. Statistics for Biology and Health. Springer, 2003 (cit. on p. 79).

[327] L. A. Klein, M. K. Mills, and D. P. Gibson. “Traffic detector handbook: Third edition - volume I”. In: 2006 (cit. on p. 244).

[328] W. Knoll et al. “Surface Plasmon Fluorescence Techniques for Bioaffinity Studies.” In: *Handbook of Surface Plasmon Resonance*. Ed. by R. B. M. Schasfoort and A. J. Tudos. RSC Publishing, 2008, pp. 275–310 (cit. on p. 24).

[329] P. G. Knoops et al. “A machine learning framework for automated diagnosis and computer-assisted planning in plastic and reconstructive surgery”. In: *Scientific Reports* 9.1 (2019), pp. 1–12 (cit. on p. 6).

[330] E. M. Knorr, R. T. Ng, and V. Tucakov. “Distance-based outliers: algorithms and applications”. In: *The VLDB Journal – The International Journal on Very Large Data Bases* 8.3-4 (Feb. 2000), pp. 237–253 (cit. on p. 14).

[331] M. H. Ko, G. West, S. Venkatesh, and M. Kumar. “Using dynamic time warping for online temporal fusion in multisensor systems”. In: *Information Fusion* 9.3 (2008), pp. 370–388 (cit. on p. 166).

[332] M. F. Al-Kobaisi. “The Jawetz, Melnick & Adelberg’s Medical Microbiology 24th Edition”. In: *Sultan Qaboos University Medical Journal* 7.3 (Dec. 2007), pp. 273–275. URL: https://journals.squ.edu.om/index.php/squmj/article/view/1338 (cit. on p. 23).

[333] R. Kobayashi, Y. Mutuka, and T. Harada. “Risk Consistent Multi-Class Learning from Label Proportions”. In: *arXiv: Computing Research Repository* (Mar. 2022). DOI: arXiv:2203.12836 (cit. on p. 139).

[334] R. Kohavi, G. John, R. Long, and D. Manley. “MLC++: A Machine Learning Library in C++”. In: *Tools with Artificial Intelligence* (1994), pp. 740–743 (cit. on p. 169).

[335] G. Köksal, I. Batmaz, and M. C. Testik. “A review of data mining applications for quality improvement in manufacturing industry”. In: *Expert Systems with Applications* 38.10 (2011), pp. 13448–13467 (cit. on p. 115).

[336] R. I. Kondor and J. D. Lafferty. “Diffusion kernels on graphs and other discrete input spaces”. In: *Procs. of the Int. Conference on Machine Learning 2002*. 2002, pp. 315–322. URL: citeseer.ist.psu.edu/kondor02diffusion.html (cit. on p. 229).

[337] D. Kong, J. Zhu, C. Duan, L. Lu, and D. Chen. “Bayesian linear regression for surface roughness prediction”. In: *Mechanical Systems and Signal Processing* 142 (2020), p. 106770 (cit. on p. 158).

[338] H. J. de Koning et al. “Reduced lung-cancer mortality with volume CT screening in a randomized trial”. In: *New England Journal of Medicine* 382.6 (2020), pp. 503–513 (cit. on p. 7).

[339] B. Konrad, D. Lieber, and J. Deuse. “Striving for Zero Defect Production: Intelligent Manufacturing Control through Data Mining in Continuous Rolling Mill Processes”. In: *Robust*

Manufacturing Control: Procs. of the CIRP Sponsored Conference RoMaC 2012. Ed. by K. Windt. Lecture Notes in Production Engineering. CIRP. Berlin, Heidelberg: Springer-Verlag, 2013, pp. 215–229. URL: http://link.springer.com/chapter/10.1007%5C%2F978-3-642-30749-2_16 (cit. on pp. 114, 116–118, 132). **SFB876-B3**

[340] R. P. H. Kooyman et al. *Handbook of Surface Plasmon Resonance*. Ed. by R. B. M. Schasfoort and A. J. Tudos. The Royal Society of Chemistry, 2008, P001–403. URL: http://dx.doi.org/10.1039/9781847558220 (cit. on p. 21).

[341] Y. Koren, R. Bell, and C. Volinsky. "Matrix factorization techniques for recommender systems". In: *Computer* 42.8 (2009), pp. 30–37 (cit. on p. 13).

[342] J. Köster, L. J. Dijkstra, T. Marschall, and A. Schönhuth. "Varlociraptor: enhancing sensitivity and controlling false discovery rate in somatic indel discovery". In: *Genome Biology* 21.1 (Apr. 2020), p. 98. URL: https://doi.org/10.1186/s13059-020-01993-6 (cit. on pp. 48–50, 53).

[343] J. Köster, F. Mölder, C. Schröder, T. Hartmann, J. Forster, and D. Laehnemann. *snakemake-workflows/dna-seq-varlociraptor: Version 2.0.5*. July 2021. URL: https://doi.org/10.5281/zenodo.5112918 (cit. on p. 54).

[344] J. Köster and S. Rahmann. "Massively parallel read mapping on GPUs with the q-group index and PEANUT". In: *PeerJ* 2 (2014), e606 (cit. on pp. 45, 57). **SFB876-C1**

[345] J. Köster and S. Rahmann. "Snakemake: a scalable bioinformatics workflow engine". In: *Bioinformatics* 28.19 (2012). PMID: 22908215, pp. 2520–2522. URL: http://dx.doi.org/10.1093/bioinformatics/bts480 (cit. on pp. 44, 95).

[346] H. Kotthaus et al. "RAMBO: Resource-Aware Model-Based Optimization with Scheduling for Heterogeneous Runtimes and a Comparison with Asynchronous Model-Based Optimization". In: *Procs. of the Int. Conference on Learning and Intelligent Optimization 2017*. 2017, pp. 180–195. URL: https://www.springerprofessional.de/en/rambo-resource-aware-model-based-optimization-with-scheduling-fo/15164982 (cit. on p. 77). **SFB876-A3**

[347] I. Koutsopoulos. "Optimal incentive-driven design of participatory sensing systems". In: *Procs. IEEE INFOCOM 2013*. 2013 (cit. on p. 298).

[348] D. Kozakov et al. "The ClusPro web server for protein-protein docking". In: *Nature Protocols* 12.2 (2017), pp. 255–278 (cit. on p. 86).

[349] D. Krajzewicz, J. Erdmann, M. Behrisch, and L. Bieker. "Recent Development and Applications of SUMO - Simulation of Urban MObility". In: *International Journal On Advances in Systems and Measurements* 5.3&4 (Dec. 2012), pp. 128–138 (cit. on p. 239).

[350] E. Kretschmann. "Determination of Optical Constants of Metals by Excitation of Surface Plasmons". In: *Zeitschrift für Physik* 241.4 (1971), pp. 313–324. URL: %3CGo%20to%20ISI%3E://WOS:A1971I477700003 (cit. on p. 24).

[351] N. Kriege, P. Giscard, and R. C. Wilson. "On Valid Optimal Assignment Kernels and Applications to Graph Classification". In: *arXiv: Computing Research Repository* (2016). DOI: arXiv:1606.01141 (cit. on p. 89). **SFB876-A6**

[352] N. M. Kriege, P.-L. Giscard, F. Bause, and R. C. Wilson. "Computing Optimal Assignments in Linear Time for Graph Matching". In: *arXiv: Computing Research Repository* arXiv:1901.10356 (2019) (cit. on p. 88). **SFB876-A6**

[353] N. M. Kriege, F. D. Johansson, and C. Morris. "A Survey on Graph Kernels". In: *Applied Network Science* 5.1 (2020), p. 6. URL: https://doi.org/10.1007/s41109-019-0195-3 (cit. on p. 87). **SFB876-A6**

[354] W. Kritzinger, M. Karner, G. Traar, J. Henjes, and W. Sihn. "Digital Twin in manufacturing: A categorical literature review and classification". In: *IFAC-PapersOnLine* 51.11 (2018), pp. 1016–1022 (cit. on p. 108).

[355] A. Krogh and J. Vedelsby. “Neural Network Ensembles, Cross Validation, and Active Learning”. In: *Advances in Neural Information Processing Systems 7: Procs. of the 1994 Conference*. 1995 (cit. on p. 170).

[356] I. Krontiris and A. Albers. “Monetary incentives in participatory sensing using multi-attributive auctions”. In: *International Journal of Parallel, Emergent and Distributed Systems* 27.4 (2012) (cit. on pp. 298, 299).

[357] H. Kück and N. de Freitas. “Learning to Classify Individuals Based on Group Statistics”. In: *Procs. of the Conference in Uncertainty in Artificial Intelligence 2005*. 2005, pp. 332–339 (cit. on pp. 139, 140).

[358] A. Kulkarni, A. Seetharam, A. Ramesh, and J. Herath. “DeepChannel: Wireless Channel Quality Prediction Using Deep Learning”. In: *IEEE Transactions on Vehicular Technology* (2019), pp. 1–1 (cit. on p. 326).

[359] S. Kumar, E. Hamed, D. Katabi, and L. Erran Li. “LTE Radio Analytics Made Easy and Accessible”. In: *Procs. of the ACM Conference on SIGCOMM 2014*. New York, NY, USA: ACM, 2014, pp. 211–222 (cit. on pp. 345, 347).

[360] F. Kurtz, C. Bektas, N. Dorsch, and C. Wietfeld. “Network Slicing for Critical Communications in Shared 5G Infrastructures - An Empirical Evaluation”. In: *Procs. of the IEEE Int. Conference on Network Softwarization 2018*. IEEE. Canada, June 2018. URL: https://ieeexplore.ieee.org/document/8460110 (cit. on p. 360).

[361] A. Kuzmichev, J. Skolnik, A. Zybin, and R. Hergenroder. “Absolute analysis of nanoparticle suspension with surface plasmon microscopy”. In: *Analytical Chemistry* 90 (2018), pp. 10732–10737 (cit. on p. 26). **SFB876-B2**

[362] V. Lagani, G. Kortas, and I. Tsamardinos. “Biomarker signature identification in ’omics’ data with multi-class outcome”. In: *Computational and Structural Biotechnology Journal* 6.7 (2013), e201303004 (cit. on p. 9).

[363] P. Lakhani and B. Sundaram. “Deep learning at chest radiography: automated classification of pulmonary tuberculosis by using convolutional neural networks”. In: *Radiology* 284.2 (2017), pp. 574–582 (cit. on p. 10).

[364] C. Lam, D. Yi, M. Guo, and T. Lindsey. “Automated detection of diabetic retinopathy using deep learning”. In: *AMIA Summits on Translational Science Proceedings* 2018 (2018), p. 147 (cit. on p. 6).

[365] J. Lamas, P.-M. Castro-Castro, A. Dapena, and F. Vazquez-Araujo. “Vehicle classification using the discrete fourier transform with traffic inductive sensors”. In: *Sensors* 15 (Oct. 2015) (cit. on p. 245).

[366] M. Lamraoui, M. Barakat, M. Thomas, and M. E. Badaoui. “Chatter detection in milling machines by neural network classification and feature selection”. In: *Journal of Vibration and Control* 21.7 (2015), pp. 1251–1266 (cit. on p. 160).

[367] M. Lang. “Automatische Modellselektion in der Überlebenszeitanalyse”. Dissertation. Dortmund: TU Dortmund University, 2015. URL: http://dx.doi.org/10.17877/DE290R-7245 (cit. on pp. 71, 75–77, 83).

[368] M. Lang, B. Bischl, and D. Surmann. “batchtools: Tools for R to Work on Batch Systems”. In: *The Journal of Open Source Software* 2.10 (2017) (cit. on p. 83).

[369] M. Lang, H. Kotthaus, P. Marwedel, C. Weihs, J. Rahnenführer, and B. Bischl. “Automatic Model Selection for High-Dimensional Survival Analysis”. In: *Journal of Statistical Computation and Simulation* 85.1 (2014), pp. 62–76. URL: http://www.tandfonline.com/eprint/fwxNXMbwDakhMqRrdMsJ/full (cit. on pp. 72, 75, 83). **SFB876-A3**

[370] M. Lang et al. “Mlr3: A Modern Object-Oriented Machine Learning Framework in R”. In: *Journal of Open Source Software* 4.44 (2019). 00000, p. 1903. URL: https://joss.theoj.org/papers/10.21105/joss.01903 (cit. on p. 83). **SFB876-A3**

[371] B. Langmead and S. L. Salzberg. “Fast gapped-read alignment with Bowtie 2”. In: *Nature Methods* 9.4 (2012), pp. 357–359. URL: http://www.nature.com/nmeth/journal/v9/n4/full/nmeth.1923.html (cit. on p. 45).

[372] G. Lanza et al. *Acatech Studie: Wandlungsfähige, menschzentrierte Strukturen in Fabriken und Netzwerken der Industrie 4.0*. 2018 (cit. on p. 104).

[373] K. Lasers. *Clubmax FB4 series Model Specific Operation Manual V271020*. Kvant Lasers. 2020 (cit. on p. 220).

[374] R. A. Laskowski, F. Gerick, and J. M. Thornton. “The structural basis of allosteric regulation in proteins”. In: *FEBS Letters* 583.11 (2009), pp. 1692–1698 (cit. on p. 95).

[375] F. Laue, P. Karunakaran, and R. Schober. “Detection Schemes and Model Mismatch Analysis for 5G Configured-Grant Access for URLLC”. In: *Procs. of the IEEE Globecom Workshops 2019*. Dec. 2019 (cit. on p. 365).

[376] S. Laur, H. Lipmaa, and T. Mielikäinen. “Cryptographically private support vector machines”. In: *Procs. of the ACM SIGKDD Conference on Knowledge Discovery and Data Mining 2006*. 2006, pp. 618–624 (cit. on p. 15).

[377] N. D. Lawrence. “Data Science and Digital Systems: The 3Ds of Machine Learning Systems Design”. In: *arXiv: Computing Research Repository* (2019). DOI: arXiv:1903.11241 (cit. on p. 107).

[378] A. Ledergerber, M. Hamer, and R. D’Andrea. “A robot self-localization system using one-way ultra-wideband communication”. In: *Procs. of the IEEE/RSJ Int. Conference on Intelligent Robots and Systems 2015*. 2015, pp. 3131–3137 (cit. on p. 190).

[379] H. K. Lee, R. Barlovic, M. Schreckenberg, and D. Kim. “Mechanical restriction versus human overreaction triggering congested traffic states”. In: *Physical Review Letters* 92.23 (2004), p. 238702 (cit. on pp. 286, 287, 290, 293, 294, 296).

[380] J.-S. Lee and B. Hoh. “Dynamic pricing incentive for participatory sensing”. In: *Pervasive and Mobile Computing* 6.6 (2010) (cit. on p. 298).

[381] J.-S. Lee and B. Hoh. “Sell your experiences: A market mechanism based incentive for participatory sensing”. In: *Procs. of IEEE PerCom 2010*. 2010 (cit. on pp. 298, 299).

[382] K. H. Lee, S. Chakraborty, and J. Sun. “Bayesian variable selection in semiparametric proportional hazards model for high dimensional survival data”. In: *International Journal of Biostatistics* 7.1 (2011) (cit. on pp. 62–64).

[383] S. Lee et al. “Robust Selection of Cancer Survival Signatures from High-Throughput Genomic Data Using Two-Fold Subsampling”. In: *PLoS ONE* 9 (2014), e108818 (cit. on pp. 72, 81). **SFB876-C1, SFB876-A3**

[384] B. Lela, D. Bajić, and S. Jozić. “Regression analysis, support vector machines, and Bayesian neural network approaches to modeling surface roughness in face milling”. In: *The International Journal of Advanced Manufacturing Technology* 42.11-12 (2009), pp. 1082–1088 (cit. on p. 158).

[385] J. E. Lenssen, V. Shpacovitch, D. Siedhoff, P. Libuschewski, R. Hergenröder, and F. Weichert. “A Review of Nano-Particle Analysis with the PAMONO-Sensor”. In: *Biosensors: Advances and Reviews* (2017), pp. 81–100 (cit. on pp. 21, 29–32, 35). **SFB876-B2**

[386] J. E. Lenssen, V. Shpacovitch, and F. Weichert. “Real-Time Virus Size Classification Using Surface Plasmon PAMONO Resonance and Convolutional Neural Networks”. In: *Bildverarbeitung für die Medizin 2017*. 2017, pp. 98–103. URL: http://dx.doi.org/10.1007/978-3-662-54345-0_26 (cit. on p. 37). **SFB876-B2**

[387] J. E. Lenssen et al. “Real-Time Low SNR Signal Processing for Nanoparticle Analysis with Deep Neural Networks”. In: *Procs. of the Int. Conference on Bio-inspired Systems and Signal Processing 2018*. (Best paper award). SciTePress, 2018. URL: https://doi.org/10.5220/0006596400360047 (cit. on pp. 30, 31). **SFB876-B2**

[388] J. Lerouge, Z. Abu-Aisheh, R. Raveaux, P. Héroux, and S. Adam. “New binary linear programming formulation to compute the graph edit distance”. In: *Pattern Recognition* 72 (2017), pp. 254–265 (cit. on pp. 88, 94, 95).

[389] M. Lesser. “Charge-coupled device (CCD) image sensors”. In: *High Performance Silicon Imaging*. Ed. by D. Durini. Second Edition. Woodhead Publishing Series in Electronic and Optical Materials. Woodhead Publishing, 2020 (cit. on p. 21).

[390] H. Li. “Minimap2: versatile pairwise alignment for nucleotide sequences”. In: *arXiv: Quantitative Biology* (Aug. 2017). DOI: arXiv:1708.01492 (cit. on p. 45).

[391] H. Li and R. Durbin. “Fast and accurate long-read alignment with Burrows-Wheeler transform”. In: *Bioinformatics* 26.5 (2010), pp. 589–595. URL: http : // bioinformatics . oxfordjournals.org/content/26/5/589 (cit. on pp. 45, 49).

[392] P. Li, L. He, X. Hu, Y. Zhang, L. Li, and X. Wu. “Concept based Short Text Stream Classification with Topic Drifting Detection”. In: *Procs. of the IEEE Int. Conference on Data Mining 2016*. 2016 (cit. on p. 345).

[393] S. C.-X. Li and B. Marlin. “A scalable end-to-end gaussian process adapter for irregularly sampled time series classification”. In: *Advances in Neural Information Processing Systems 29: Procs. of the 2016 Conference*. Vol. 29. 2016 (cit. on p. 13).

[394] T. Li, A. K. Sahu, A. Talwalkar, and V. Smith. “Federated learning: Challenges, methods, and future directions”. In: *IEEE Signal Processing Magazine* 37.3 (2020), pp. 50–60 (cit. on p. 14).

[395] X. Li et al. “Automatic IMRT planning via static field fluence prediction (AIP-SFFP): a deep learning algorithm for real-time prostate treatment planning”. In: *Physics in Medicine & Biology* 65.17 (2020), p. 175014 (cit. on p. 8).

[396] Y. Li et al. “Multi-platform omics analysis reveals molecular signature for COVID-19 pathogenesis, prognosis and drug target discovery”. In: *Signal Transduction and Targeted Therapy* 6.1 (2021), pp. 1–11 (cit. on p. 11).

[397] Y. Liang and P. Zhao. “Similarity Search in Graph Databases: A Multi-Layered Indexing Approach”. In: *Procs. of the IEEE Int. Conference on Data Engineering 2017*. IEEE Computer Society, 2017, pp. 783–794 (cit. on pp. 93, 94).

[398] O. Liberg, M. Sundberg, E. Wang, J. Bergman, and J. Sachs. *Cellular Internet of Things: Technologies, Standards, and Performance*. Elsevier Science, 2017 (cit. on pp. 263, 264).

[399] P. Libuschewski. “Exploration of Cyber-Physical Systems for GPGPU Computer Vision-Based Detection of Biological Viruses”. PhD thesis. Dortmund, Germany: TU Dortmund University, 2017. URL: http://dx.doi.org/10.17877/DE290R-17952 (cit. on pp. 29, 30, 38). **SFB876-B2**

[400] P. Libuschewski, P. Marwedel, D. Siedhoff, and H. Müller. “Multi-Objective Energy-Aware GPGPU Design Space Exploration for Medical or Industrial Applications”. In: *Procs. of the Int. Conference on Signal-Image Technology and Internet-Based Systems 2014*. Ed. by I. C. Society. IEEE Computer Society, 2014, pp. 637–644. URL: http://dx.doi.org/10.1109/SITIS.2014.11 (cit. on p. 30). **SFB876-B2**

[401] P. Libuschewski, D. Siedhoff, C. Timm, A. Gelenberg, and F. Weichert. “Fuzzy-enhanced, Real-time capable Detection of Biological Viruses Using a Portable Biosensor”. In: *Procs. of the Int. Joint Conference on Biomedical Engineering Systems and Technologies 2013*. 2013, pp. 169–174 (cit. on pp. 29, 30). **SFB876-B2**

[402] P. Libuschewski, D. Siedhoff, C. Timm, and F. Weichert. “Mobile Detektion viraler Pathogene durch echtzeitfähige GPGPU-Fuzzy-Segmentierung”. In: *Bildverarbeitung für die Medizin 2013* (2013), pp. 326–331. URL: http://dx.doi.org/10.1007/978-3-642-36480-8_57 (cit. on p. 36). **SFB876-B2**

[403] P. Libuschewski et al. “Multi-Objective Computation Offloading for Mobile Biosensors via LTE”. In: *Procs. of the Int. Conference on Wireless Mobile Communication and Healthcare*

2014. 2014, pp. 226–229. URL: dx.doi.org/10.4108/icst.mobihealth.2014.257374 (cit. on pp. 35, 36). **SFB876-B2, SFB876-A4**

[404] D. Lieber. "Data Mining in der Qualitätslenkung am Beispiel der Stabstahlproduktion". PhD thesis. TU Dortmund University, 2018 (cit. on pp. 132, 133). **SFB876-B3**

[405] D. Lieber, B. Konrad, J. Deuse, M. Stolpe, and K. Morik. "Sustainable Interlinked Manufacturing Processes through Real-Time Quality Prediction". In: *Leveraging Technology for a Sustainable World: Procs. of the CIRP Conference on Life Cycle Engineering 2012*. Ed. by D. A. Dornfeld and B. S. Linke. CIRP. Berlin, Heidelberg: Springer-Verlag, 2012, pp. 393–398. URL: http://link.springer.com/chapter/10.1007/978-3-642-29069-5_67 (cit. on pp. 116, 118, 119, 133). **SFB876-B3**

[406] T. Liebig, C. Körner, and M. May. "Fast Visual Trajectory Analysis Using Spatial Bayesian Networks". In: *Procs. of the IEEE International Conference on Data Mining Workshops 2009*. IEEE Computer Society, 2009, pp. 668–673. URL: http://doi.ieeecomputersociety.org/10.1109/ICDMW.2009.44 (cit. on p. 228).

[407] T. Liebig, C. Körner, and M. May. "Scalable Sparse Bayesian Network Learning for Spatial Applications". In: *Procs. of the IEEE International Conference on Data Mining Workshops 2008*. IEEE Computer Society, 2008, pp. 420–425. URL: http://dx.doi.org/10.1109/ICDMW.2008.124 (cit. on p. 228).

[408] T. Liebig, Z. Xu, and M. May. "Incorporating Mobility Patterns in Pedestrian Quantity Estimation and Sensor Placement". In: *Procs. of the Int. Workshop on Citizen Sensor Networks CitiSens 2012*. Ed. by J. Nin and D. Villatoro. Springer, 2013, pp. 67–80. URL: http://link.springer.com/content/pdf/10.1007%5C%2F978-3-642-36074-9_7 (cit. on p. 229).

[409] T. Liebig, Z. Xu, M. May, and S. Wrobel. "Pedestrian Quantity Estimation with Trajectory Patterns". In: *Procs. of the European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases 2012*. Springer, 2012, pp. 629–643. URL: http://link.springer.com/chapter/10.1007%5C%2F978-3-642-33486-3_40 (cit. on p. 229).

[410] T. Liebig. "Privacy Preserving Centralized Counting of Moving Objects". In: *AGILE 2015*. Ed. by F. Bacao, M. Y. Santos, and M. Painho. Lecture Notes in Geoinformation and Cartography. Springer Int. Publishing, 2015, pp. 91–103. URL: http://dx.doi.org/10.1007/978-3-319-16787-9_6 (cit. on p. 239).

[411] T. Liebig, S. Peter, M. Grzenda, and K. Junosza-Szaniawski. "Dynamic Transfer Patterns for Fast Multi-modal Route Planning". In: *Societal Geo-innovation: Selected papers of the 20th AGILE Conference on Geographic Information Science*. Ed. by A. Bregt, T. Sarjakoski, R. van Lammeren, and F. Rip. Cham: Springer Int. Publishing, 2017, pp. 223–236. URL: http://dx.doi.org/10.1007/978-3-319-56759-4_13 (cit. on pp. 233, 235).

[412] T. Liebig and M. Sotzny. "On Avoiding Traffic Jams with Dynamic Self-Organizing Trip Planning". In: *Procs. of the Int. Conference on Spatial Information Theory 2017*. Ed. by E. Clementini, M. Donnelly, M. Yuan, C. Kray, P. Fogliaroni, and A. Ballatore. Vol. 86. Leibniz Int. Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2017, 17:1–17:12. URL: http://drops.dagstuhl.de/opus/volltexte/2017/7761 (cit. on pp. 235, 240).

[413] J. Lin, E. Keogh, L. Wei, and S. Lonardi. "Experiencing SAX: a novel symbolic representation of time series". In: *Data Mining and Knowledge Discovery* 15.2 (2007), pp. 107–144 (cit. on p. 123).

[414] T. Lin, P. Dollár, R. B. Girshick, K. He, B. Hariharan, and S. J. Belongie. "Feature Pyramid Networks for Object Detection". In: *arXiv: Computing Research Repository* (2016). DOI: arXiv:1612.03144 (cit. on p. 32).

[415] T. Lin et al. "Microsoft COCO: Common Objects in Context". In: *arXiv: Computing Research Repository* (2014). DOI: arXiv:1405.0312 (cit. on pp. 8, 32).

[416] C. Lippert, J. Listgarten, Y. Liu, C. M. Kadie, R. I. Davidson, and D. Heckerman. “FaST linear mixed models for genome-wide association studies”. In: *Nature Methods* 8.10 (2011), pp. 833–835 (cit. on p. 9).

[417] B. Liu and H. Zhu. “Energy-effective data gathering for UAV-aided wireless sensor networks”. In: *Sensors* 19.11 (2019), p. 2506 (cit. on p. 213).

[418] S. Liu, X. Fang, Y. Guo, C. Liu, and J. Yang. “Theory and Realization of Secondary Task Assignment for Multi-UAV Pickup Based on Green Scheduling”. In: *Complexity* 2021 (2021) (cit. on p. 213).

[419] F. Llinares-López, L. Papaxanthos, D. Bodenham, D. Roqueiro, C. Investigators, and K. Borgwardt. “Genome-wide genetic heterogeneity discovery with categorical covariates”. In: *Bioinformatics* 33.12 (2017), pp. 1820–1828 (cit. on p. 9).

[420] J. Long, E. Shelhamer, and T. Darrell. “Fully convolutional networks for semantic segmentation”. In: *Procs. of the IEEE Conference on Computer Vision and Pattern Recognition 2015*. 2015, pp. 3431–3440 (cit. on p. 176).

[421] M. López-Ibáñez, J. Dubois-Lacoste, L. P. Cáceres, M. Birattari, and T. Stützle. “The irace package: Iterated racing for automatic algorithm configuration”. In: *Operations Research Perspectives* 3 (2016), pp. 43–58. URL: https://doi.org/10.1016/j.orp.2016.09.002 (cit. on p. 75).

[422] D. Lymberopoulos and J. Liu. “The Microsoft Indoor Localization Competition: Experiences and Lessons Learned”. In: *IEEE Signal Processing Magazine* 34.5 (Sept. 2017), pp. 125–140 (cit. on p. 193).

[423] D. Lymberopoulos, J. Liu, et al. *Microsoft Indoor Localization 2018*. Apr. 2018. URL: https://www.microsoft.com/en-us/research/event/microsoft-indoor-localization-competition-ipsn-2018/ (cit. on p. 193).

[424] X. Ma and L. Gao. “Discovering protein complexes in protein interaction networks via exploring the weak ties effect”. In: *BMC Systems Biology* 6 Suppl 1 (2012), S6 (cit. on p. 86).

[425] A. Machanavajjhala, D. Kifer, J. Gehrke, and M. Venkitasubramaniam. “L-diversity: Privacy beyond k-anonymity”. In: *The ACM Transactions on Knowledge Discovery from Data* 1.1 (2007) (cit. on p. 15).

[426] J. MacQueen. “Some Methods for Classification and Analysis of Multivariate Observations”. In: *Procs. of the Berkeley Symposium on Mathematical Statistics and Probability 1967*. Ed. by L. M. L. Cam and J. Neyman. Vol. 1. University of California Press, 1967, pp. 281–297 (cit. on pp. 124, 148).

[427] H. Maeda, Y. Sekimoto, T. Seto, T. Kashiyama, and H. Omata. “Road damage detection using deep neural networks with images captured through a smartphone”. In: *Computer Aided Civil and Infrastructure Engineering* 33.2 (2018), pp. 1–15 (cit. on p. 309).

[428] M. Mahdavi et al. “A machine learning based exploration of COVID-19 mortality risk”. In: *Plos one* 16.7 (2021), e0252384 (cit. on pp. 10, 11).

[429] N. H. Mahmood et al. “Machine type communications: key drivers and enablers towards the 6G era”. In: *EURASIP Journal on Wireless Communications and Networking* 2021.1 (June 2021), p. 134. URL: https://doi.org/10.1186/s13638-021-02010-5 (cit. on pp. 329, 330).
SFB876-A4

[430] P.-F. Marteau and S. Gibet. “On recursive edit distance kernels with application to time series classification”. In: *IEEE transactions on neural networks and learning systems* 26.6 (2014), pp. 1121–1133 (cit. on pp. 9, 11).

[431] M. Martin. “Cutadapt removes adapter sequences from high-throughput sequencing reads”. In: *EMBnet.journal* 17.1 (2011) (cit. on p. 55).

[432] M. Masoudinejad, J. Emmerich, D. Kossmann, A. Riesner, M. Roidl, and M. ten Hompel. “Development of a measurement platform for indoor photovoltaic energy harvesting in materials

handling applications". In: *Procs. of the Int. Renewable Energy Congress 2015*. Mar. 2015, pp. 1–6. URL: http://www.sciencedirect.com/science/article/pii/S2210670715300329 (cit. on p. 203). **SFB876-A4**

[433] M. Masoudinejad. "Data-Sets for Indoor Photovoltaic Behavior in Low Lighting Conditions". In: 5.2 (2020), p. 32. URL: https://www.mdpi.com/2306-5729/5/2/32 (cit. on pp. 203, 204). **SFB876-A4**

[434] M. Masoudinejad. "Modeling Energy Supply Unit of Ultra-Low Power Devices with Indoor Photovoltaic Harvesting". PhD Thesis. TU Dortmund University, 2020. URL: https://eldorado.tu-dortmund.de/handle/2003/39765 (cit. on pp. 195, 196, 200). **SFB876-A4**

[435] M. Masoudinejad, J. Emmerich, D. Kossmann, A. Riesner, M. Roidl, and M. ten Hompel. "A Measurement Platform for Photovoltaic Performance Analysis in Environments with Ultra-Low Energy Harvesting Potential". In: *Sustainable Cities and Society* 25 (2016), pp. 74–81 (cit. on p. 203). **SFB876-A4**

[436] F.-I. für Materialfluss und Logistik IML Dortmund. "Logistik entdecken, Magazin des Fraunhofer-Instituts für Materialfluss und Logistik IML Dortmund". In: *Logistik entdecken* 2019.19 (2019) (cit. on p. 220).

[437] F. Mattern and C. Floerkemeier. "From the Internet of Computers to the Internet of Things". In: *From Active Data Management to Event-Based Systems and More*. Springer, 2010, pp. 242–259 (cit. on p. 105).

[438] M. May. *Eight ways machine learning is assisting medicine*. 2021 (cit. on p. 4).

[439] J. D. Mazimpaka and S. Timpf. "A visual and computational analysis approach for exploring significant locations and time periods along a bus route". In: *Procs. of the ACM SIGSPATIAL Int. Workshop on Computational Transportation Science 2016*. ACM. 2016, pp. 43–48 (cit. on pp. 233, 234).

[440] W. McLaren et al. "The Ensembl Variant Effect Predictor". In: *Genome Biology* 17.1 (June 2016), p. 122. URL: https://doi.org/10.1186/s13059-016-0974-4 (cit. on p. 56).

[441] M. J. Mehlman. "Medical practice guidelines as malpractice safe harbors: illusion or deceit?" In: *Journal of Law, Medicine & Ethics* 40.2 (2012), pp. 286–300 (cit. on p. 19).

[442] P. Melville. "Creating Diverse Ensemble Classifiers to Reduce Supervision". PhD thesis. Austin, TX, USA, 2005 (cit. on p. 163).

[443] G. Mendel. "Versuche über Pflanzen-Hybriden". In: 4 (1866). Place: BrünnSeperatabdruck aus dem IV. Bande der Verhandlungen des naturforschenden Vereines; mit handschriftlichen Korrekturen Mendels., pp. 3–47. URL: http://vlp.mpiwg-berlin.mpg.de/references?id=lit26745 (cit. on pp. 48, 51).

[444] I. Mierswa. "Non-Convex and Multi-Objective Optimization in Data Mining". PhD thesis. TU Dortmund University, 2008 (cit. on p. 147).

[445] I. Mierswa and K. Morik. "Automatic Feature Extraction for Classifying Audio Data". In: *Machine Learning Journal* 58 (2005), pp. 127–149 (cit. on pp. 119, 122, 123).

[446] Mikhaylov, Konstantin et al. "Analysis of Capacity and Scalability of the LoRa Low Power Wide Area Network Technology". In: *European Wireless 2016* (2016), pp. 119–124 (cit. on p. 315).

[447] T. M. Mitchell. *Machine Learning*. New York: McGraw Hill, 1997 (cit. on pp. 139, 142).

[448] J. Mitra et al. "Statistical machine learning to identify traumatic brain injury (TBI) from structural disconnections of white matter networks". In: *NeuroImage* 129 (2016), pp. 247–259 (cit. on p. 10).

[449] H.-C. Möhring, P. Wiederkehr, K. Erkorkmaz, and Y. Kakinuma. "Self-optimizing machining systems". In: *CIRP Annals* 69.2 (2020), pp. 740–763 (cit. on p. 158).

[450] F. Mölder et al. "Sustainable data analysis with Snakemake". In: *F1000Research* 10 (Jan. 2021), p. 33. URL: https://f1000research.com/articles/10-33/v1 (cit. on p. 44). **SFB876-C1**

[451] A. Molina, A. Munteanu, and K. Kersting. “Core Dependency Networks”. In: *Procs. of the AAAI Conference on Artificial Intelligence 2018*. 2018. URL: https://www.aaai.org/ocs/index.php/AAAI/AAAI18/paper/view/16847 (cit. on p. 230). **SFB876-B4, SFB876-C4**

[452] A. Molina, S. Natarajan, and K. Kersting. “Poisson Sum-Product Networks: A Deep Architecture for Tractable Multivariate Poisson Distributions”. In: *Procs. of the AAAI Conference on Artificial Intelligence 2017*. 2017, pp. 2357–2363 (cit. on p. 230). **SFB876-B4, SFB876-A6**

[453] C. Molnar, G. Casalicchio, and B. Bischl. “Interpretable machine learning–a brief history, state-of-the-art and challenges”. In: *Procs. of the Joint European Conference on Machine Learning and Knowledge Discovery in Databases 2020*. Springer. 2020, pp. 417–431 (cit. on p. 17).

[454] M. Moor, M. Horn, B. Rieck, D. Roqueiro, and K. Borgwardt. “Early recognition of sepsis with Gaussian process temporal convolutional networks and dynamic time warping”. In: *Procs. of the Machine Learning for Healthcare Conference 2019*. PMLR. 2019, pp. 2–26 (cit. on pp. 10, 11).

[455] M. Moor, B. Rieck, M. Horn, C. R. Jutzeler, and K. Borgwardt. “Early prediction of sepsis in the ICU using machine learning: a systematic review”. In: *Frontiers in medicine* 8 (2021), p. 348 (cit. on p. 11).

[456] E. Moradi, A. Pepe, C. Gaser, H. Huttunen, J. Tohka, A. D. N. Initiative, et al. “Machine learning framework for early MRI-based Alzheimer’s conversion prediction in MCI subjects”. In: *Neuroimage* 104 (2015), pp. 398–412 (cit. on pp. 9, 10).

[457] M. Morenza-Cinos, V. Casamayor-Pujol, J. Soler-Busquets, J. L. Sanz, R. Guzmán, and R. Pous. “Development of an RFID inventory robot (AdvanRobot)”. In: *Robot Operating System (ROS)*. Springer, 2017, pp. 387–417 (cit. on p. 217).

[458] K. Morik and S. Wessel. “Incremental Signal to Symbol Processing”. In: *Making Robots Smarter – Combining Sensing and Action through Robot Learning*. Ed. by K. Morik, V. Klingspor, and M. Kaiser. Kluwer Academic Publ., 1999. Chap. 11, pp. 185–198 (cit. on p. 123).

[459] D. Mottin, M. Lissandrini, Y. Velegrakis, and T. Palpanas. “Exemplar queries: a new way of searching”. In: *VLDB – The International Journal on Very Large Data Bases* 25.6 (2016), pp. 741–765. URL: https://doi.org/10.1007/s00778-016-0429-2 (cit. on p. 89).

[460] M. Moya, M. Koch, and L. Hostetler. “One-class classifier networks for target recognition applications.” In: *Procs. of the World Congress on Neural Networks 1993*. Int. Neural Network Society, 1993, pp. 797–801 (cit. on p. 125).

[461] M. Müller. “Dynamic time warping”. In: *Information Retrieval for Music and Motion* (2007), pp. 69–84 (cit. on p. 121).

[462] M. Müller-Hannemann and M. Schnee. “Efficient timetable information in the presence of delays”. In: *Robust and Online Large-Scale Optimization*. Springer, 2009, pp. 249–272 (cit. on p. 234).

[463] A. Munteanu, C. Schwiegelshohn, C. Sohler, and D. P. Woodruff. “On Coresets for Logistic Regression”. In: *Advances in Neural Information Processing Systems 31: Procs. of the 2018 COnference*. 2018. URL: http://papers.nips.cc/paper/7891-on-coresets-for-logistic-regression (cit. on p. 59). **SFB876-A2, SFB876-C4**

[464] Y. Murakami, T. Magome, K. Matsumoto, T. Sato, Y. Yoshioka, and M. Oguchi. “Fully automated dose prediction using generative adversarial networks in prostate cancer patients”. In: *PloS one* 15.5 (2020), e0232697 (cit. on p. 6).

[465] A. Murrenhoff, M. Roidl, and M. ten Hompel. “Control concept for ü r virtualized and learnable ä capable material flow systems”. In: *Logistics Journal: Proceedings* 2019.12 (2019) (cit. on pp. 219, 220).

[466] D. R. Musicant, J. M. Christensen, and J. F. Olson. “Supervised Learning by Training on Aggregate Outputs”. In: *Procs. of the IEEE Int. Conference on Data Mining 2007*. Washington, DC, USA: IEEE Computer Society, Oct. 2007, pp. 252–261 (cit. on pp. 137, 139, 141).

[467] K. Nagel and M. Schreckenberg. “A cellular automaton model for freeway traffic”. In: *Journal de Physique 1* 2 (1992), pp. 2221–2229 (cit. on pp. 285, 287).

[468] F. Naumann, A. Bilke, J. Bleiholder, and M. Weis. “Data fusion in three steps: Resolving inconsistencies at schema-, tuple-, and value-level”. In: *IEEE Data Engineering Bulletin* 29.2 (2006), pp. 21–31 (cit. on pp. 166, 167).

[469] A. Nazabal, P. M. Olmos, Z. Ghahramani, and I. Valera. “Handling incomplete heterogeneous data using vaes”. In: *Pattern Recognition* 107 (2020), p. 107501 (cit. on p. 13).

[470] R. A. Nazib and S. Moh. “Routing protocols for unmanned aerial vehicle-aided vehicular ad hoc networks: A survey”. In: *IEEE Access* 8 (2020), pp. 77535–77560 (cit. on p. 274).

[471] T. R. Network. “The Cancer Genome Atlas Program”. In: (2021). URL: http://cancergenome.nih.gov/ (cit. on p. 64).

[472] A. Neumann, C. Aichele, M. Lindner, and S. Wunderlich. “Better approach to mobile ad-hoc networking (BATMAN)”. In: *IETF draft* (2008), pp. 1–24 (cit. on p. 273).

[473] J. Neumann, A. W. Burks, et al. *Theory of self-reproducing automata*. Vol. 1102024. University of Illinois press Urbana, 1966 (cit. on p. 227).

[474] K. Ni et al. “Sensor network data fault types”. In: *ACM Transactions on Sensor Networks (TOSN)* 5.3 (2009), pp. 1–29 (cit. on p. 119).

[475] T. L. and Nico Piatkowski and Christian Bockermann and Katharina Morik. “Predictive Trip Planning - Smart Routing in Smart Cities”. In: *Procs. of the Joint EDBT/ICDT Workshops 2014*. Vol. 1133. CEUR-WS.org, 2014, pp. 331–338 (cit. on p. 230). **SFB876-A1**

[476] B. Niehoefer, F. Schweikowski, S. Lehnhausen, and C. Wietfeld. “Cloud-aided SDR Solution for Lane-Specific Vehicle Positioning via Local Interference Compensation”. In: *Procs. of IEEE Aerospace Conference 2014*. Big Sky, USA, Mar. 2014 (cit. on p. 242).

[477] B. Niehöfer, F. Schweikowski, and C. Wietfeld. “LOcal interferenCe compensATion (LOCATe) for GNSS-based Lane-Specific Positioning of Vehicles”. In: *Procs. of the IEEE Vehicular Technology Conference 2016*. Nanjing, China, May 2016 (cit. on p. 242).

[478] F. Niemann et al. “Lara: Creating a dataset for human activity recognition in logistics using semantic attributes”. In: *Sensors* 20.15 (2020), p. 4083 (cit. on p. 218).

[479] NREL. *Reference Air Mass 1.5 Spectra*. 2012 (cit. on p. 201).

[480] K. O’Shea and R. Nash. “An introduction to convolutional neural networks”. In: *arXiv: Computing Research Repository* (2015). DOI: arXiv:1511.08458 (cit. on p. 8).

[481] G. Oddi, D. Macone, A. Pietrabissa, and F. Liberati. “A proactive link-failure resilient routing protocol for MANETs based on reinforcement learning”. In: *Procs. of the Mediterranean Conference on Control Automation 2012*. 2012, pp. 1259–1264 (cit. on pp. 274, 275).

[482] I. Oleaga, C. Pardo, J. J. Zulaika, and A. Bustillo. “A machine-learning based solution for chatter prediction in heavy-duty milling machines”. In: *Measurement* 128 (2018), pp. 34–44 (cit. on p. 158).

[483] R. Olfati-Saber. “Flocking for multi-agent dynamic systems: Algorithms and theory”. In: *IEEE Transactions on automatic control* 51.3 (2006), pp. 401–420 (cit. on p. 215).

[484] D. L. Olson and D. Delen. *Advanced Data Mining Techniques*. Springer-Verlag Berlin Heidelberg, 2008 (cit. on p. 115).

[485] S. Oota. “Somatic mutations – Evolution within the individual”. In: *Methods*. RNA-Seq: Methods and Applications 176 (Apr. 2020), pp. 91–98. URL: https://www.sciencedirect.com/science/article/pii/S1046202318303827 (cit. on p. 53).

[486] O. S. Oubbati, M. Atiquzzaman, P. Lorenz, M. H. Tareque, and M. S. Hossain. “Routing in flying ad hoc networks: Survey, constraints, and future challenge perspectives”. In: *IEEE Access* 7 (2019), pp. 81057–81105 (cit. on p. 274).

[487] T. Ozturk, M. Talo, E. A. Yildirim, U. B. Baloglu, O. Yildirim, and U. R. Acharya. “Automated detection of COVID-19 cases using deep neural networks with X-ray images”. In: *Computers in biology and medicine* 121 (2020), p. 103792 (cit. on p. 6).

[488] S. L. Pardau. “The California consumer privacy act: towards a European-style privacy regime in the United States”. In: *Journal of Technology Law & Policy* 23 (2018), p. 68 (cit. on p. 14).

[489] R. B. Parikh, S. Teeple, and A. S. Navathe. “Addressing bias in artificial intelligence in health care”. In: *The Journal of the American Medical Association (JAMA)* 322.24 (2019), pp. 2377–2378 (cit. on p. 17).

[490] H. Park, H. Lee, and C. Seok. “High-resolution protein-protein docking by global optimization: recent advances and future challenges”. In: *Current Opinion in Structural Biology (COSB)* 35 (2015), pp. 24–31 (cit. on p. 86).

[491] J. Park, K. Law, R. Bhinge, M. Chen, D. Dornfeld, and S. Rachuri. “Realtime energy prediction for a milling machine tool using sparse Gaussian process regression”. In: Oct. 2015, pp. 1451–1460 (cit. on p. 158).

[492] K. Parry, L. N. Geppert, A. Munteanu, and K. Ickstadt. “Cross-Leverage Scores for Selecting Subsets of Explanatory Variables”. In: *arXiv: Statistics* (2021). DOI: arXiv:2109.08399 (cit. on pp. 60, 61). **SFB876-C4**

[493] I. Parvez, A. Rahmati, I. Guvenc, A. I. Sarwat, and H. Dai. “A Survey on Low Latency Towards 5G: RAN, Core Network and Caching Solutions”. In: *IEEE Communications Surveys Tutorials* 20.4 (May 2018), pp. 3098–3130 (cit. on p. 364).

[494] G. Patrini, R. Nock, T. Caetano, and P. Rivera. “(Almost) No Label No Cry”. In: *NIPS 27*. Curran Associates, Inc., 2014, pp. 190–198 (cit. on pp. 140, 141).

[495] D. Paul, E. Bair, T. Hastie, and R. Tibshirani. “Preconditioning for feature selection and regression in high-dimensional problems”. In: *The Annals of Statistics* 36.4 (Aug. 2008), pp. 1595–1618 (cit. on p. 81).

[496] W. R. Pearson. “Selecting the Right Similarity-Scoring Matrix”. In: *Current Protocols in Bioinformatics* 43 (2013), pp. 1–9 (cit. on p. 86).

[497] M. Pellegrini, M. Baglioni, and F. Geraci. “Protein complex prediction for large protein protein interaction networks with the Core&Peel method”. In: *BMC Bioinformatics* 17.Suppl 12 (2016), p. 372 (cit. on p. 86).

[498] B. Peng, T. Bergs, D. Schraknepper, F. Klocke, and B. Döbbeler. “A hybrid approach using machine learning to predict the cutting forces under consideration of the tool wear”. In: *Procedia CIRP* 82 (2019), pp. 302–307 (cit. on p. 158).

[499] N. Piatkowski, S. Lee, and K. Morik. “Spatio-temporal random fields: compressible representation and distributed estimation”. In: *Machine Learning* 93.1 (2013), pp. 115–139. URL: http://link.springer.com/article/10.1007%5C%2Fs10994-013-5399-7 (cit. on p. 230). **SFB876-A1**

[500] *Picard toolkit*. http://broadinstitute.github.io/picard/. 2019 (cit. on p. 55).

[501] B. Pinkas. “Cryptographic techniques for privacy-preserving data mining”. In: *ACM SIGKDD Explorations Newsletter* 4.2 (2002), pp. 12–19 (cit. on p. 15).

[502] I. Pires, N. Garcia, N. Pombo, and F. Flórez-Revuelta. “From data acquisition to data fusion: a comprehensive review and a roadmap for the identification of activities of daily living using mobile devices”. In: *Sensors* 16.2 (2016), p. 184 (cit. on p. 166).

[503] D. Plakhotnik et al. “Integration of Process Monitoring Data into CAM Simulation.” In: (2018). IVth Int. Conference on High Speed Machining, Donostia/San Sebastian - Spain, 17-18 April 2018 (cit. on p. 158).

[504] A. van der Pol E. andBoing, P. Harrison, A. Sturk, and R. Nieuwland. “Classification, Functions, and Clinical Relevance of Extracellular Vesicles”. In: *Pharmacological Reviews* 64 (2012), pp. 676–705 (cit. on p. 23).

[505] E. van der Pol, A. N. Boing, E. L. Gool, and R. Nieuwland. “Recent developments in the nomenclature, presence, isolation, detection and clinical impact of extracellular vesicles”. In: *Journal of Thrombosis and Haemostasis* 14.1 (2016), pp. 48–56. URL: %3CGo%20to%20ISI%3E://WOS:000370661100005 (cit. on p. 23).

[506] H. Poon and P. Domingos. “Sum-product networks: A new deep architecture”. In: *Procs. of the IEEE Int. ConferenceComputer Vision Workshops 2011*. IEEE. 2011, pp. 689–690 (cit. on p. 230).

[507] F. Potorti et al. “Comparing the Performance of Indoor Localization Systems through the EvAAL Framework”. In: *Sensors* 17.10 (2017) (cit. on p. 192).

[508] A. Pottmeier, C. Thiemann, A. Schadschneider, and M. Schreckenberg. “Mechanical Restriction Versus Human Overreaction: Accident Avoidance and Two-Lane Traffic Simulations”. In: *Traffic And Granular Flow*. Ed. by Springer. Springer, 2005, pp. 503–508 (cit. on p. 293).

[509] J. A. Preiss, W. Honig, G. S. Sukhatme, and N. Ayanian. “Crazyswarm: A large nano-quadcopter swarm”. In: *Procs. of the IEEE Int. Conference on Robotics and Automation 2017*. IEEE. 2017, pp. 3299–3304 (cit. on pp. 216, 217, 221).

[510] W. N. Price, S. Gerke, and I. G. Cohen. “Potential liability for physicians using artificial intelligence”. In: *JAMA* 322.18 (2019), pp. 1765–1766 (cit. on p. 19).

[511] N. C. Purandare and V. Rangarajan. “Imaging of lung cancer: implications on staging and management”. In: *The Indian journal of radiology & imaging* 25.2 (2015), p. 109 (cit. on p. 8).

[512] T. G. Purdie, R. E. Dinniwell, D. Letourneau, C. Hill, and M. B. Sharpe. “Automated planning of tangential breast intensity-modulated radiotherapy using heuristic optimization”. In: *International Journal of Radiation Oncology* Biology* Physics* 81.2 (2011), pp. 575–583 (cit. on p. 8).

[513] Y. Qian, Q. Tong, and B. Wang. “Multi-Class Learning from Label Proportions for Bank Customer Classification”. In: *Procedia Computer Science* 162 (2019). Procs. of the Int. Conference on Information Technology and Quantitative Management 2019, pp. 421–428 (cit. on p. 139).

[514] N. Quadrianto, A. J. Smola, T. S. Caetano, and Q. V. Le. “Estimating Labels from Label Proportions”. In: *Journal of Machine Learning Research* 10 (Dec. 2009), pp. 2349–2374 (cit. on pp. 139, 151).

[515] J. Quedenfeld and S. Rahmann. “Variant tolerant read mapping using min-hashing”. In: *arXiv: Quantitative Biology* (Feb. 2017). DOI: arXiv:1702.01703 (cit. on p. 47).

[516] J. R. Quinlan. “Induction of Decision Trees”. In: *Machine Learning* 1.1 (Mar. 1986), pp. 81–106. URL: http://portal.acm.org/citation.cfm?id=637962.637969 (cit. on p. 152).

[517] S. Rahmann, M. Martin, J. H. Schulte, J. Köster, T. Marschall, and A. Schramm. “Identifying Transcriptional miRNA Biomarkers by Integrating High-Throughput Sequencing and Real-Time PCR Data”. In: *Methods* 59.1 (Jan. 2013), pp. 154–163 (cit. on p. 45). **SFB876-C1, SFB876-B1**

[518] G. Raja, S. Anbalagan, V. S. Narayanan, S. Jayaram, and A. Ganapathisubramaniyan. “Inter-UAV collision avoidance using Deep-Q-learning in flocking environment”. In: *Procs. of the IEEE Ubiquitous Computing, Electronics & Mobile Communication Conference 2019*. IEEE. 2019, pp. 1089–1095 (cit. on p. 222).

[519] S. Rajab, M. O. Al Kalaa, and H. Refai. “Classification and speed estimation of vehicles via tire detection using single-element piezoelectric sensor”. In: *Journal of Advanced Transportation* 50.7 (2016), pp. 1366–1385 (cit. on p. 245).

[520] A. K. Ramachandran Venkatapathy, H. Bayhan, F. Zeidler, and M. ten Hompel. “Human Machine Synergies in Intra-Logistics: Creating a Hybrid Network for Research and Technologies”. In: *Procs. of the Workshop on Information Technologies for Logistics 2017* (2017) (cit. on p. 218). **SFB876-A4**

[521] B. Rao, A. G. Gopi, and R. Maione. “The societal impact of commercial drones”. In: *Technology in Society* 45 (2016), pp. 83–90 (cit. on pp. 212, 213).

[522] T. S. Rappaport et al. “Wireless Communications and Applications Above 100 GHz: Opportunities and Challenges for 6G and Beyond”. In: *IEEE Access* 7 (2019), pp. 78729–78757 (cit. on p. 246).

[523] T. Rausch, T. Zichner, A. Schlattl, A. M. Stütz, V. Benes, and J. O. Korbel. “DELLY: structural variant discovery by integrated paired-end and split-read analysis”. In: *Bioinformatics* 28.18 (Sept. 2012), pp. i333–i339. URL: https://doi.org/10.1093/bioinformatics/bts378 (cit. on p. 55).

[524] S. Reddy, D. Estrin, M. Hansen, and M. Srivastava. “Examining micro-payments for participatory sensing data collections”. In: *Procs. of the ACM Int. Conference on Ubiquitous Computing 2010*. 2010 (cit. on p. 304).

[525] A. Rejeb, K. Rejeb, S. J. Simske, and H. Treiblmaier. “Drones for supply chain management and logistics: a review and research agenda”. In: *International Journal of Logistics Research and Applications* (2021), pp. 1–24 (cit. on p. 213).

[526] C. W. Reynolds. “Flocks, herds and schools: A distributed behavioral model”. In: *Procs. of the Conference on Computer Graphics and Interactive Techniques 1987*. 1987, pp. 25–34 (cit. on p. 214).

[527] M. T. Ribeiro, S. Singh, and C. Guestrin. “"Why Should I Trust You?": Explaining the Predictions of Any Classifier”. In: *Procs. of the ACM SIGKDD Conference on Knowledge Discovery and Data Mining 2016*. KDD ’16. New York, NY, USA: ACM, 2016, pp. 1135–1144. URL: http://doi.acm.org/10.1145/2939672.2939778 (cit. on p. 17).

[528] M. T. Ribeiro, S. Singh, and C. Guestrin. “Anchors: High-precision model-agnostic explanations”. In: *Procs. of the AAAI Conference on Artificial Intelligence 2018*. Vol. 32. 1. 2018 (cit. on p. 17).

[529] M. T. Ribeiro, S. Singh, and C. Guestrin. “Model-agnostic Interpretability of Machine Learning”. In: *arXiv: Statistics* (2016). DOI: arXiv:1606.05386 (cit. on p. 17).

[530] J. Richter, H. Kotthaus, B. Bischl, P. Marwedel, J. Rahnenführer, and M. Lang. “Faster Model-Based Optimization through Resource-Aware Scheduling Strategies”. In: *Procs. of the Int. Conference on Learning and Intelligent Optimization 2010*. Vol. 10079. Lecture Notes in Computer Science (LNCS). Springer Int. Publishing, 2016, pp. 267–273. URL: http://link.springer.com/chapter/10.1007/978-3-319-50349-3_22 (cit. on p. 77). **SFB876-A3**

[531] J. Richter, K. Madjar, and J. Rahnenführer. “Model-based optimization of subgroup weights for survival analysis”. In: *Bioinformatics* 35.14 (July 2019), pp. i484–i491. URL: https://doi.org/10.1093/bioinformatics/btz361 (cit. on pp. 72, 78, 81, 83). **SFB876-A3**

[532] M. S. Ridgely and M. D. Greenberg. “Too many alerts, too much liability: sorting through the malpractice implications of drug-drug interaction clinical decision support”. In: *Saint Louis University Journal of Health Law & Policy* 5 (2011), p. 257 (cit. on p. 19).

[533] N. Rieke et al. “The future of digital health with federated learning”. In: *NPJ digital medicine* 3.1 (2020), pp. 1–7 (cit. on p. 14).

[534] K. Riesen, M. Ferrer, and H. Bunke. “Approximate Graph Edit Distance in Quadratic Time”. In: *IEEE/ACM Transactions on Computational Biology and Bioinformatics* (2015). Epub ahead of print (cit. on pp. 88, 102).

[535] K. Riesen and H. Bunke. “Approximate graph edit distance computation by means of bipartite graph matching”. In: *Image and Vision Computing* 27.7 (2009), pp. 950–959 (cit. on pp. 88, 94).

[536] A. P. Rifai, H. Aoyama, N. H. Tho, S. Z. M. Dawal, and N. A. Masruroh. “Evaluation of turned and milled surfaces roughness using convolutional neural network”. In: *Measurement* 161 (2020), p. 107860 (cit. on p. 158).

[537] E. Rivers et al. “Early goal-directed therapy in the treatment of severe sepsis and septic shock”. In: *New England Journal of Medicine* 345.19 (2001), pp. 1368–1377 (cit. on p. 5).

[538] A. Rizk-Jackson et al. “Evaluating imaging biomarkers for neurodegeneration in pre-symptomatic Huntington’s disease using machine learning techniques”. In: *Neuroimage* 56.2 (2011), pp. 788–796 (cit. on p. 10).

[539] S. J. Rizvi and J. R. Haritsa. “Maintaining data privacy in association rule mining”. In: *Procs. of the Int. Conference on Very Large Data Bases 2002*. Elsevier. 2002, pp. 682–693 (cit. on p. 15).

[540] Robert B Cleveland, William S. Cleveland, Jean E. McRae, and Irma Terpenning. *STL: A Seasonal-Trend decomposition Procedure Based on Loess*. 1990 (cit. on p. 324).

[541] E. Roberts, B. A. Bassett, and M. Lochner. “Bayesian Anomaly Detection and Classification”. In: *arXiv: Statistics* (2019). DOI: arXiv:1902.08627 (cit. on p. 14).

[542] B. Røder, S. Nielsen, P. Magnussen, A. Engquist, and N. Frimodt-Møller. “Antibiotic usage in an intensive care unit in a Danish university hospital.” In: *The Journal of antimicrobial chemotherapy* 32.4 (1993), pp. 633–642 (cit. on p. 7).

[543] C. Römer, J. Hiry, C. Kittl, T. Liebig, and C. Rehtanz. “Charging control of electric vehicles using contextual bandits considering the electrical distribution grid”. In: *Procs. of the Int. Workshop on Knowledge Discovery from Mobility and Transportation Systems, co-located with the European Conference on Machine Learning and Principles and Practice of Knowledge Discovery 2018*. 2018, (accepted) (cit. on p. 240). **SFB876-B4**

[544] O. Ronneberger, P.Fischer, and T. Brox. “U-Net: Convolutional Networks for Biomedical Image Segmentation”. In: *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. Vol. 9351. LNCS. Springer, 2015, pp. 234–241 (cit. on p. 39).

[545] R. N. Rooks, E. M. Simonsick, L. M. Klesges, A. B. Newman, H. N. Ayonayon, and T. B. Harris. “Racial disparities in health care access and cardiovascular disease indicators in Black and White older adults in the Health ABC Study”. In: *Journal of aging and health* 20.6 (2008), pp. 599–614 (cit. on p. 18).

[546] A. Roth and T. Liebig. “Transforming PageRank into an Infinite-Depth Graph Neural Network”. In: *Procs. of the European Conference on Machine Learning and Knowledge Discovery in Databases 2022*. Springer. 2022, (accepted) (cit. on p. 240).

[547] A. Roth, K. Wüstefeld, and F. Weichert. “A Data-Centric Augmentation Approach for Disturbed Sensor Image Segmentation”. In: *Journal of Imaging* 7.10 (Oct. 2021). DOI: 10.3390/jimaging7100206. URL: https://www.mdpi.com/2313-433X/7/10/206 (cit. on pp. 38–40). **SFB876-B2, SFB876-B4**

[548] A. Roth. *Einführung und Umsetzung von Industrie 4.0: Grundlagen, Vorgehensmodell und Use Cases aus der Praxis*. Springer-Verlag, 2016 (cit. on pp. 103, 104).

[549] M. C. P. Rubio et al. “A Realistic Evaluation of Indoor Robot Position Tracking Systems: The IPIN 2016 Competition Experience”. In: *Measurement* 135 (Mar. 2019), pp. 151–162. URL: http://www.sciencedirect.com/science/article/pii/S0263224118310728 (cit. on p. 192).

[550] I. Ruczinski, C. Kooperberg, and M. LeBlanc. “Logic Regression”. In: *Journal of Computational and Graphical Statistics* 12.3 (2003), pp. 475–511 (cit. on p. 60).

[551] A. Ruepp et al. “CORUM: the comprehensive resource of mammalian protein complexes – 2009”. In: *Nucleic acids research* 38.suppl 1 (2010), pp. D497–D501 (cit. on pp. 87, 95).

[552] S. Rüping. “SVM Classifier Estimation From Group Probabilities”. In: *Procs. of the Int. Conference on Machine Learning 2010*. 2010, pp. 911–918 (cit. on pp. 140, 141, 151, 152).

[553] H. Ruppel, L. De Vaux, D. Cooper, S. Kunz, B. Duller, and M. Funk. “Testing physiologic monitor alarm customization software to reduce alarm rates and improve nurses’ experience of alarms in a medical intensive care unit”. In: *PloS one* 13.10 (2018), e0205901 (cit. on p. 5).

[554] O. Russakovsky et al. “Imagenet large scale visual recognition challenge”. In: *International journal of computer vision* 115.3 (2015), pp. 211–252 (cit. on p. 8).

[555] J. Russell. *Frequently Asked Questions. The American Road and Transportation Builders Association (ARTBA)*. URL: https://www.artba.org/about/ (cit. on p. 297).

[556] A. Saadallah, L. Moreira-Matias, R. Sousa, J. Khiari, and J. Jenelius E.and Gama. “BRIGHT - Drift-Aware Demand Predictions for Taxi Networks”. In: *Procs. of the IEEE Int. Conference on Data Engineering 2018*. IEEE Transactions on Knowledge and Data Engineering, 2018 (cit. on p. 169). **SFB876-B3**

[557] A. Saadallah, O. Abdulaaty, J. Büscher, T. Panusch, K. Morik, and J. Deuse. “Early Quality Prediction using Deep Learning on Time Series Sensor Data”. In: *Procs. of the CIRP Conference on Manufacturing Systems 2022*. Elsevier, 2022 (cit. on p. 135). **SFB876-B3**

[558] A. Saadallah, J. Büscher, O. Abdulaaty, T. Panusch, J. Deuse, and K. Morik. “Explainable Predictive Quality Inspection using Deep Learning inElectronics Manufacturing”. In: *Procs. of the CIRP Conference on Manufacturing Systems 2022*. Elsevier, 2022 (cit. on p. 135). **SFB876-B3**

[559] A. Saadallah, A. Egorov, B.-T. Cao, S. Freitag, K. Morik, and G. Meschke. “Active Learning for Accurate Settlement Prediction Using Numerical Simulations in Mechanized Tunneling”. In: *CIRP Manufacturing Systems Conference 2019* (2019). URL: https://www.cirp-cms2019.org/ (cit. on pp. 160, 161). **SFB876-B3**

[560] A. Saadallah, F. Finkeldey, K. Morik, and P. Wiederkehr. “Stability prediction in milling processes using a simulation-based machine learning approach”. In: *Procs. of the CIRP Conference on Manufacturing Systems 2018*. Elsevier, 2018 (cit. on pp. 158–160, 162–165, 171). **SFB876-B3**

[561] A. Saadallah, M. Jakobs, and K. Morik. “Explainable Online Deep Neural Network Selection using Adaptive Saliency Maps for Time Series Forecasting”. In: *Procs. of the Joint European Conference on Machine Learning and Knowledge Discovery in Databases 2021*. To appear. 2021 (cit. on p. 127). **SFB876-B3**

[562] A. Saadallah and M. Katharina. “Online Ensemble Aggregation using Deep Reinforcement Learning for Time Series Forecasting”. In: 2021 (cit. on p. 127). **SFB876-B3**

[563] A. Saadallah, N. Piatkowski, F. Finkeldey, P. Wiederkehr, and K. Morik. “Learning Ensembles in the Presence of Imbalanced Classes”. In: *ICPRAM: International Conference on Pattern Recognition Applications and Methods 2019* (2019). URL: http://www.icpram.org/Home.aspx (cit. on p. 125).

[564] A. Saadallah, F. Priebe, and M. Katharina. “A Drift-based Dynamic Ensemble Members Selection using Clustering for Time Series Forecasting”. In: *Procs. of the European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases 2019*. Sept. 2019 (cit. on pp. 120, 127). **SFB876-B3**

[565] A. Saadallah, M. Tavakol, and M. Katharina. “An Actor-Critic Ensemble Aggregation Model for Time-Series Forecasting”. In: 2021 (cit. on p. 127). **SFB876-B3**

[566] A. Saadi, D. U. Himmelstein, S. Woolhandler, and N. I. Mejia. “Racial disparities in neurologic health care access and utilization in the United States”. In: *Neurology* 88.24 (2017), pp. 2268–2275 (cit. on p. 18).

[567] A. Sachan, D. Roy, and P. Arun. “An analysis of privacy preservation techniques in data mining”. In: *Advances in computing and information technology*. Springer, 2013, pp. 119–128 (cit. on p. 15).

[568] T. Sachweh, D. Boiar, and T. Liebig. “Differentially Private Learning from Label Proportions”. In: *Procs. of the ECML Workshop on Parallel, Distributed, and Federated Learning 2021*. 2021, accepted (cit. on p. 233). **SFB876-B4**

[569] A. Sadeghnejad Barkousaraie, O. Ogunmolu, S. Jiang, and D. Nguyen. “A fast deep learning approach for beam orientation optimization for prostate cancer treated with intensity-modulated radiation therapy”. In: *Medical physics* 47.3 (2020), pp. 880–897 (cit. on p. 6).

[570] A. Sadeghnejad-Barkousaraie, G. Bohara, S. Jiang, and D. Nguyen. “A reinforcement learning application of a guided Monte Carlo Tree Search algorithm for beam orientation selection in radiation therapy”. In: *Machine Learning: Science and Technology* 2.3 (2021), p. 035013 (cit. on p. 8).

[571] R. Saket, A. Raghuveer, and B. Ravindran. “On Combining Bags to Better Learn from Label Proportions”. In: *Procs. of the Int. Conference on AI and Statistics 2022*. PMLR, 2022 (cit. on p. 139).

[572] J. Saleema, P. D. Shenoy, K. Venugopal, and L. Patnaik. “Cancer prognosis prediction model using data mining techniques”. In: *Data Mining and Knowledge Engineering* 6.1 (2014), pp. 21–29 (cit. on p. 15).

[573] C. Sánchez Claros and A. Tramontano. “Detecting mutually exclusive interactions in protein-protein interaction maps”. In: *PLoS One* 7.6 (2012), e38765 (cit. on p. 95).

[574] A. Sanfeliu and K. S. Fu. “A Distance measure between attributed relational graphs for pattern recognition.” In: *IEEE Transactions on Systems, Man, and Cybernetics* 13.3 (1983), pp. 353–362 (cit. on p. 87).

[575] R. Sardar, A. Sharma, and D. Gupta. “Machine learning assisted prediction of prognostic biomarkers associated with COVID-19, using clinical and proteomics data”. In: *Frontiers in genetics* 12 (2021) (cit. on p. 11).

[576] R. E. Schapir. “A brief introduction to boosting”. In: *Procs. of the Int. Joint Conference on Artificial Intelligence 1999*. 1999 (cit. on p. 125).

[577] R. B. M. Schasfoort. *Handbook of Surface Plasmon Resonance*. RSC Publishing, 2008 (cit. on p. 24).

[578] M. Schirmer, R. D’Amore, U. Z. Ijaz, N. Hall, and C. Quince. “Illumina error profiles: resolving fine-scale variation in metagenomic sequencing data”. In: *BMC Bioinformatics* 17.1 (Mar. 2016), p. 125. URL: https://doi.org/10.1186/s12859-016-0976-y (cit. on p. 49).

[579] J. Schmitt, J. Boenig, T. Borggraefe, G. Beitinger, and J. Deuse. “Predictive model-based quality inspection in electronics manufacturing using Machine Learning and Edge Cloud Computing”. In: *Advanced Engineering Informatics (ADVEI)* (2020) (cit. on pp. 116, 128, 130). **SFB876-B3**

[580] J. Schmitt and J. Deuse. “Modellbasierte Prüfprozesse. Einsatz von Data-Mining-Verfahren zur industriellen Qualitätssicherung”. In: *Zeitschrift für wirtschaftlichen Fabrikbetrieb (ZWF)* 114.4 (Apr. 2019), pp. 191–193 (cit. on pp. 128, 130, 131). **SFB876-B3**

[581] B. Schölkopf, J. C. Platt, J. C. Shawe-Taylor, A. J. Smola, and R. C. Williamson. “Estimating the Support of a High-Dimensional Distribution”. In: *Neural Computation* 13.7 (July 2001), pp. 1443–1471 (cit. on p. 125).

[582] B. Schowe and K. Morik. “Fast-Ensembles of Minimum Redundancy Feature Selection”. In: *Ensembles in Machine Learning Applications*. Ed. by O. Okun, G. Valentini, and M. Re. Studies in Computational Intelligence. Springer, 2011, pp. 75–95 (cit. on p. 122).

[583] A. Schramm et al. “Mutational dynamics between primary and relapse neuroblastomas”. In: *Nature Genetics* 47.8 (Aug. 2015), pp. 872–877. URL: http://www.nature.com/ng/journal/vaop/ncurrent/full/ng.3349.html (cit. on pp. 44, 53, 54, 56). **SFB876-C1**

[584] G. Schuh, R. Anderl, J. Gausemeier, M. ten Hompel, and W. Wahlster. *Acatech Studie: Industrie 4.0 Maturity Index – Die Digitale Transformation von Unternehmen Gestalten*. 2017 (cit. on pp. 105, 106).

[585] C. Schüler, M. Patchou, B. Sliwa, and C. Wietfeld. “Robust machine learning-enabled routing for highly mobile vehicular networks with PARRoT in ns-3”. In: *Procs. of the Workshop on Ns-3*. Virtual, June 2021 (cit. on p. 276). **SFB876-A4, SFB876-B4**

[586] C. Schüler, B. Sliwa, and C. Wietfeld. “Towards Machine Learning-Enabled Context Adaption for Reliable Aerial Mesh Routing”. In: *Procs. of the IEEE Vehicular Technology Conference 2021*. Virtual Event, Sept. 2021 (cit. on pp. 274, 277, 278, 282, 283). **SFB876-A4, SFB876-B4**

[587] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov. “Proximal Policy Optimization Algorithms”. In: *arXiv: Computing Research Repository* (2017). DOI: arXiv:1707.06347 (cit. on p. 223).

[588] M. Schulte, J. Köster, S. Rahmann, and A. Schramm. “Cancer evolution, mutations, and clonal selection in relapse neuroblastoma”. In: *Cell Tissue Research* 372.2 (May 2018), pp. 263–268 (cit. on p. 44). **SFB876-C1**

[589] M. Schumacher, H. Binder, and T. Gerds. “Assessment of survival prediction models based on microarray data”. In: *Bioinformatics* 23.14 (2007), pp. 1768–1774 (cit. on p. 65).

[590] P. Schwab et al. “Real-time prediction of COVID-19 related mortality using electronic health records”. In: *Nature communications* 12.1 (2021), pp. 1–16 (cit. on pp. 10, 11).

[591] N. Semiconductor. *nRF51822 Product Specification*. Tech. rep. Technical report, Nordic Semi, 2014 (cit. on p. 216).

[592] Semtech. *SX1272/73 -860 MHz to 1020 MHz Low Power Long Range Transceiver*. Tech. rep. Rev. 3.1. 2017, p. 129 (cit. on p. 318).

[593] B. Settles. *Active Learning*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool Publishers, 2012 (cit. on pp. 161, 163).

[594] A. Severyn and A. Moschitti. “Automatic feature engineering for answer selection and extraction”. In: *Procs. of the Conference on Empirical Methods in Natural Language Processing 2013*. 2013, pp. 458–467 (cit. on p. 122).

[595] A. Shahrjooihaghighi, H. Frigui, X. Zhang, X. Wei, B. Shi, and A. Trabelsi. “An ensemble feature selection method for biomarker discovery”. In: *Procs. of the IEEE Int. Symposium on Signal Processing and Information Technology 2017*. IEEE. 2017, pp. 416–421 (cit. on p. 9).

[596] C. Shang, A. Palmer, J. Sun, K.-S. Chen, J. Lu, and J. Bi. “VIGAN: Missing view imputation with generative adversarial networks”. In: *Procs. of the IEEE Int. Conference on Big Data 2017*. IEEE. 2017, pp. 766–775 (cit. on p. 13).

[597] X. Shao et al. “Conditional sum-product networks: Imposing structure on deep probabilistic architectures”. In: *Procs. of the Int. Conference on Probabilistic Graphical Models 2020*. PMLR. 2020, pp. 401–412 (cit. on pp. 231, 240). **SFB876-B4**

[598] X. Shao et al. “Conditional Sum-Product Networks: Modular Probabilistic Circuits via Gate Functions”. In: *International Journal of Approximate Reasoning* (2022), pp. 298–313 (cit. on p. 231). **SFB876-B4**

[599] S. Sharma, H. Mohammadmoradi, M. Heydariaan, and O. Gnawali. “Device-free activity recognition using ultra-wideband radios”. In: *Procs. of the Int. Conference on Computing, Networking and Communications 2019*. 2019, pp. 1029–1033 (cit. on p. 246).

[600] Z. Shelby et al. “The Constrained Application Protocol”. In: *Internet Engineering Task Force (IETF)* (2014) (cit. on p. 266).

[601] Z. Shen. “Cancer biomarkers and targeted therapies”. In: *Cell & bioscience* 3.1 (2013), pp. 1–2 (cit. on p. 9).

[602] N. Shervashidze, P. Schweitzer, E. van Leeuwen, K. Mehlhorn, and K. Borgwardt. “Weisfeiler–Lehman Graph Kernels”. In: *Journal of Machine Learning Research* 12 (2011), pp. 2539–2561 (cit. on pp. 89–91).

[603] S. Shiraishi and K. L. Moore. “Knowledge-based prediction of three-dimensional dose distributions for external beam radiotherapy”. In: *Medical physics* 43.1 (2016), pp. 378–387 (cit. on p. 8).

[604] W. Shockley and W. T. Read. “Statistics of the Recombinations of Holes and Electrons”. In: 87.5 (1952), pp. 835–842. URL: https://link.aps.org/doi/10.1103/PhysRev.87.835 (cit. on p. 197).

[605] C. Shorten and T. M. Khoshgoftaar. “A survey on image data augmentation for deep learning”. In: *Journal of Big Data* 6.1 (2019), pp. 1–48 (cit. on p. 177).

[606] V. Shpacovitch and R. Hergenröder. “Optical and surface plasmonic approaches to characterize extracellular vesicles. A review”. In: *Analytica Chimica Acta* 1005 (2017), pp. 1–15. URL: https://doi.org/10.1016/j.aca.2017.11.066 (cit. on p. 26). **SFB876-B2**

[607] V. Shpacovitch and R. Hergenröder. “Optical and surface plasmonic approaches to characterize extracellular vesicles. A review”. In: *Analytica Chimica Acta* 1005 (2018), pp. 1–15. URL: https://www.sciencedirect.com/science/article/pii/S0003267017313594 (cit. on p. 21).

[608] V. Shpacovitch et al. “Application of Surface Plasmon Resonance Imaging Technique for the Detection of Single Spherical Biological Submicron-particles”. In: *Analytical Biochemistry: Methods in the Biological Sciences* (2015) (cit. on p. 25). **SFB876-B2**

[609] V. Shpacovitch et al. “Application of the PAMONO-sensor for Quantification of Microvesicles and Determination of Nano-particle Size Distribution”. In: *Sensors* 17.2 (2017), pp. 1–14. URL: http://dx.doi.org/10.3390/s17020244 (cit. on pp. 26, 37). **SFB876-B2, SFB876-C1**

[610] S. N. Shukla and B. M. Marlin. “Interpolation-Prediction Networks for Irregularly Sampled Time Series”. In: *arXiv: Computing Research Repository* (2019). DOI: arXiv:1909.07782 (cit. on p. 13).

[611] G. L. Sicherman, W. de Jonge, and R. P. van de Riet. “Answering queries without revealing secrets”. In: *ACM Transactions on Database Systems* 8.1 (1983), pp. 41–59 (cit. on p. 411).

[612] A. J. Siddiqui, A. Mammeri, and A. Boukerche. “Towards Efficient Vehicle Classification in Intelligent Transportation Systems”. In: *Procs. of the ACM Symposium on Development and Analysis of Intelligent Vehicular Networks and Applications 2015*. DIVANet ’15. New York, NY, USA: Association for Computing Machinery, 2015, pp. 19–25. DOI: https://doi.org/10.1145/2815347.2815354 (cit. on p. 245).

[613] D. Siedhoff. “A Parameter-Optimizing Model-Based Approach to the Analysis of Low-SNR Image Sequences for Biological Virus Detection”. PhD thesis. Dortmund, Germany: TU Dortmund University, 2016. DOI: http://dx.doi.org/10.17877/DE290R-17272 (cit. on pp. 22, 29, 30, 34, 35). **SFB876-B2**

[614] D. Siedhoff, P. Libuschewski, F. Weichert, A. Zybin, P. Marwedel, and H. Müller. “Modellierung und Optimierung eines Biosensors zur Detektion viraler Strukturen”. In: *Bildverarbeitung für die Medizin 2014*. Springer Berlin Heidelberg, 2014, pp. 108–113. URL: http://dx.doi.org/10.1007/978-3-642-54111-7_24 (cit. on p. 27). **SFB876-B2**

[615] D. Siedhoff, F. Weichert, P. Libuschewski, and C. Timm. “Detection and Classification of Nano-Objects in Biosensor Data”. In: *Procs. of the Microscopic Image Analysis with Applications in Biology Workshop 2011*. 2011 (cit. on p. 30). **SFB876-B2**

[616] R. L. Siegel, K. D. Miller, H. E. Fuchs, and A. Jemal. “Cancer statistics, 2021.” In: *CA: a cancer journal for clinicians* 71.1 (2021), pp. 7–33 (cit. on p. 7).

[617] D. Siepmann. “Industrie 4.0 – Grundlagen und Gesamtzusammenhang”. In: *Einführung und Umsetzung von Industrie 4.0: Grundlagen, Vorgehensmodell und Use Cases aus der Praxis*. Ed. by A. Roth. Springer, 2016, pp. 17–82. URL: https://doi.org/10.1007/978-3-662-48505-7_2 (cit. on pp. 104, 105).

[618] J. H. Silber, P. S. Romano, A. K. Rosen, Y. Wang, O. Even-Shoshan, and K. G. Volpp. “Failure-to-rescue: comparing definitions to measure quality of care”. In: *Medical care* (2007), pp. 918–925 (cit. on p. 11).

[619] K. R. Simpson and A. Lyndon. “False alarms and overmonitoring: major factors in alarm fatigue among labor nurses”. In: *Journal of nursing care quality* 34.1 (2019), pp. 66–72 (cit. on p. 5).

[620] D. Sinwar and M. Kumar. “Anomaly detection using decision tree based classifiers”. In: *International Journal of modern Trends in Engineering and Research(IJMTER)* 3 (2016), pp. 1–4 (cit. on pp. 12, 14).

[621] B. Sliwa, R. Adam, and C. Wietfeld. “Acting selfish for the good of all: Contextual bandits for resource-efficient transmission of vehicular sensor data”. In: *Procs. of the ACM MobiHoc Workshop on Cooperative Data Dissemination in Future Vehicular Networks 2020*. Online, Oct. 2020 (cit. on p. 243). **SFB876-B4**

[622] B. Sliwa, R. Adam, and C. Wietfeld. “Client-Based Intelligence for Resource Efficient Vehicular Big Data Transfer in Future 6G Networks”. In: *IEEE Transactions on Vehicular Technology* (Feb. 2021). URL: https://arxiv.org/abs/2102.08624 (cit. on pp. 334, 338, 352). **SFB876-A4, SFB876-B4**

[623] B. Sliwa, D. Behnke, C. Ide, and C. Wietfeld. “B.A.T.Mobile: Leveraging mobility control knowledge for efficient routing in mobile robotic networks”. In: *Procs. of the IEEE GLOBECOM Workshop on Wireless Networking, Control and Positioning of Unmanned Autonomous Vehicles 2016*. Washington D.C., USA: IEEE, Dec. 2016. URL: http://dx.doi.org/10.1109/GLOCOMW.2016.7848845 (cit. on p. 273). **SFB876-B4**

[624] B. Sliwa, R. Falkenberg, T. Liebig, N. Piatkowski, and C. Wietfeld. “Boosting Vehicle-to-cloud Communication by Machine Learning-enabled Context Prediction”. In: *IEEE Transactions on Intelligent Transportation Systems* (2019) (cit. on p. 279). **SFB876-B4, SFB876-A4, SFB876-A1**

[625] B. Sliwa, R. Falkenberg, and C. Wietfeld. “Machine learning for resource-efficient data transfer in mobile crowdsensing”. In: *Machine Learning for Future Wireless Communications*. Ed. by F.-L. Luo. John Wiley & Sons Inc, Mar. 2020 (cit. on p. 334). **SFB876-A4, SFB876-B4**

[626] B. Sliwa, R. Falkenberg, and C. Wietfeld. “Towards cooperative data rate prediction for future mobile and vehicular 6G networks”. In: *Procs. of the 6G Wireless Summit 2020*. Levi, Finland: IEEE, Mar. 2020 (cit. on pp. 340, 352, 353). **SFB876-A4, SFB876-B4**

[627] B. Sliwa, S. Falten, and C. Wietfeld. “Performance evaluation and optimization of B.A.T.M.A.N. V routing for aerial and ground-based mobile ad-hoc networks”. In: *Procs. of the IEEE Vehicular Technology Conference 2019*. Kuala Lumpur, Malaysia, Apr. 2019. URL: http://arxiv.org/abs/1901.02298 (cit. on p. 273). **SFB876-B4**

[628] B. Sliwa, M. Haferkamp, M. Al-Askary, and C. Wietfeld. “A radio-fingerprinting-based vehicle classification system for intelligent traffic control in smart cities”. In: *Procs. of the IEEE Int. Systems Conference 2018*. Vancouver, Canada: IEEE, Apr. 2018 (cit. on p. 334). **SFB876-B4**

[629] B. Sliwa, T. Liebig, R. Falkenberg, J. Pillmann, and C. Wietfeld. “Efficient machine-type communication using multi-metric context-awareness for cars used as mobile sensors in upcoming 5G networks”. In: *Procs. of the IEEE Vehicular Technology Conference 2018*. Best Student Paper Award. IEEE. 2018 (cit. on p. 334). **SFB876-B4, SFB876-A4**

[630] B. Sliwa, T. Liebig, T. Vranken, M. Schreckenberg, and C. Wietfeld. "System-of-systems modeling, analysis and optimization of hybrid vehicular traffic". In: *Procs. of the IEEE Int. Systems Conference 2019*. Orlando, Florida, USA: IEEE, Apr. 2019 (cit. on p. 332). **SFB876-B4**

[631] B. Sliwa, M. Patchou, K. Heimann, and C. Wietfeld. "Simulating hybrid aerial- and ground-based vehicular networks with ns-3 and LIMoSim". In: *Procs. of the Workshop on Ns-3 2020*. Gaithersburg, Maryland, USA, June 2020 (cit. on p. 385). **SFB876-B4**

[632] B. Sliwa, M. Patchou, and C. Wietfeld. "Lightweight simulation of hybrid aerial- and ground-based vehicular communication networks". In: *Procs. of the IEEE Vehicular Technology Conference 2019*. Honolulu, Hawaii, USA, Sept. 2019 (cit. on p. 334). **SFB876-B4**

[633] B. Sliwa, N. Piatkowski, and C. Wietfeld. "LIMITS: Lightweight machine learning for IoT systems with resource limitations". In: *Procs. of the IEEE Int. Conference on Communications 2020*. Best paper award. IEEE. Dublin, Ireland, June 2020, pp. 1–7. URL: https://arxiv.org/abs/2001.10189 (cit. on pp. 277, 335). **SFB876-B4**

[634] B. Sliwa, N. Piatkowski, and C. Wietfeld. "The channel as a traffic sensor: Vehicle detection and classification based on radio fingerprinting". In: *IEEE Internet of Things Journal* (Mar. 2020) (cit. on pp. 243, 244, 247–249, 253, 255–258). **SFB876-B4**

[635] B. Sliwa, C. Schüler, M. Patchou, and C. Wietfeld. "PARRoT: Predictive ad-hoc routing fueled by reinforcement learning and trajectory knowledge". In: *Procs. of the IEEE Vehicular Technology Conference 2021*. Helsinki, Finland, Apr. 2021 (cit. on pp. 273, 279). **SFB876-A4, SFB876-B4**

[636] B. Sliwa and C. Wietfeld. "A Reinforcement Learning Approach for Efficient Opportunistic Vehicle-to-Cloud Data Transfer". In: *Procs. of the IEEE Wireless Communications and Networking Conference 2020*. Seoul, South Korea, Apr. 2020. DOI: https://doi.org/10.48550/arXiv.2001.05321 (cit. on p. 334). **SFB876-B4**

[637] B. Sliwa and C. Wietfeld. "LIMoSim: A framework for lightweight simulation of vehicular mobility in intelligent transportation systems". In: *Recent Advances in Network Simulation - The OMNeT++ Environment and its Ecosystem*. Ed. by A. Virdis and M. Kirsche. Springer Int. Publishing, May 2019. Chap. LIMoSim: A framework for lightweight simulation of vehicular mobility in intelligent transportation systems. URL: https://www.springer.com/de/book/9783030128418%5C#aboutBook (cit. on p. 338). **SFB876-B4**

[638] B. Sliwa and C. Wietfeld. "Towards data-driven simulation of end-to-end network performance indicators". In: *Procs. of the IEEE Vehicular Technology Conference 2019*. Honolulu, Hawaii, USA, Sept. 2019 (cit. on pp. 333, 335). **SFB876-B4**

[639] I. Smal, M. Loog, W. Niessen, and E. Meijering. "Quantitative Comparison of Spot Detection Methods in Fluorescence Microscopy". In: *IEEE Transactions on Medical Imaging* 29.2 (2010), pp. 282–301 (cit. on p. 27).

[640] R. Sonabend, F. J. Király, A. Bender, B. Bischl, and M. Lang. "mlr3proba: An R Package for Machine Learning in Survival Analysis". In: *Bioinformatics* (Feb. 2021) (cit. on p. 83). **SFB876-A3**

[641] S. Song, C. Li, and X. Zhang. "Turn Waste into Wealth: On Simultaneous Clustering and Cleaning over Dirty Data". In: *Procs. of the ACM SIGKDD Conference on Knowledge Discovery and Data Mining 2015*. Ed. by L. Cao, C. Zhang, T. Joachims, G. I. Webb, D. D. Margineantu, and G. Williams. ACM, 2015, pp. 1115–1124. DOI: http://doi.acm.org/10.1145/2783258.2783317 (cit. on p. 12).

[642] F. B. Sorbelli, F. Corò, S. K. Das, and C. M. Pinotti. "Energy-constrained delivery of goods with drones under varying wind conditions". In: *IEEE Transactions on Intelligent Transportation Systems* 22.9 (2020), pp. 6048–6060 (cit. on p. 213).

[643] S. Srihari, C. H. Yong, and L. Wong. *Computational Prediction of Protein Complexes from Protein Interaction Networks*. New York, NY, USA: Association for Computing Machinery and Morgan & Claypool, 2017 (cit. on p. 86).

[644] B. K. St"ocker, T. Sch"afer, P. Mutzel, J. K"oster, N. M. Kriege, and S. Rahmann. "Protein Complex Similarity Based on Weisfeiler-Lehman Labeling". In: *Procs. of the Int. Conference on Similarity Search and Applications 2021*. Ed. by G. Amato, C. Gennaro, V. Oria, and M. Radovanovic. Cham: Springer International Publishing, 2019, pp. 308–322 (cit. on p. 86).

[645] A. Stanford-Clark et al. "MQTT For Sensor Networks (MQTT-SN) Protocol Specification Version 1.2". In: *OASIS standard* (2013) (cit. on p. 266).

[646] M. R. Stavseth, T. Clausen, and J. Røislien. "How handling missing data may impact conclusions: A comparison of six different imputation methods for categorical questionnaire data". In: *SAGE open medicine* 7 (2019), p. 2050312118822912 (cit. on p. 13).

[647] D. J. Stekhoven and P. Bühlmann. "MissForest – non-parametric missing value imputation for mixed-type data". In: *Bioinformatics* 28.1 (2012), pp. 112–118 (cit. on p. 13).

[648] STMicroelectronics. *DM00037051 - STM32F405xx STM32F407xx Datasheet*. STMicroelectronics. 2020 (cit. on p. 216).

[649] B. K. Stöcker, J. Köster, E. Zamir, and S. Rahmann. "Modeling and simulating networks of interdependent protein interactions". In: *Integrative Biology* 10 (2018), pp. 290–305 (cit. on pp. 95, 102). **SFB876-C1**

[650] J. K. Stolaroff, C. Samaras, E. R. O'Neill, A. Lubers, A. S. Mitchell, and D. Ceperley. "Energy use and life cycle greenhouse gas emissions of drones for commercial package delivery". In: *Nature Communications* 9.1 (Feb. 2018) (cit. on p. 212).

[651] M. Stolpe and K. Morik. "Learning from Label Proportions by Optimizing Cluster Model Selection". In: *Procs. of the Machine Learning and Knowledge Discovery in Databases European Conference 2011*. Ed. by D. Gunopulos, T. Hofmann, D. Malerba, and M. Vazirgiannis. Springer, 2011, pp. 349–364 (cit. on pp. 139, 152, 154–156, 231). **SFB876-B3**

[652] M. Stolpe. "Distributed Analysis of Vertically Partitioned Sensor Measurements under Communication Constraints". PhD thesis. Dortmund: TU Dortmund University, 2017. URL: http://dx.doi.org/10.17877/DE290R-17839 (cit. on pp. 139, 156). **SFB876-B3**

[653] M. Stolpe. "The Internet of Things: Opportunities and Challenges for Distributed Data Analysis". In: *SIGKDD Explorations* 18.1 (June 2016), pp. 15–34 (cit. on p. 137). **SFB876-B3**

[654] M. Stolpe, H. Blom, and K. Morik. "Sustainable Industrial Processes by Embedded Real-Time Quality Prediction". In: *Computational Sustainability*. Ed. by K. Kersting, J. Lässig, and K. Morik. Springer, 2016, pp. 201–243. URL: http://link.springer.com/chapter/10.1007/978-3-319-31858-5_10 (cit. on pp. 114, 116–122, 124, 132–134, 156). **SFB876-A1, SFB876-B3**

[655] M. Stolpe, T. Liebig, and K. Morik. "Communication-efficient learning of traffic flow in a network of wireless presence sensors". In: *Procs. of the Workshop on Parallel and Distributed Computing for Knowledge Discovery in Data Bases 2015*. CEUR Workshop Proceedings. CEUR-WS, 2015, (to appear) (cit. on pp. 150, 156, 231). **SFB876-B3**

[656] J. Stoyanovich, B. Howe, and H. Jagadish. "Responsible data management". In: *Procs. of the VLDB Endowment* 13.12 (2020), pp. 3474–3488 (cit. on p. 13).

[657] C. Sturm and W. Wiesbeck. "Waveform Design and Signal Processing Aspects for Fusion of Wireless Communications and Radar Sensing". In: *Procs. of the IEEE* 99.7 (July 2011), pp. 1236–1259 (cit. on p. 387).

[658] C. Sudlow et al. "UK biobank: an open access resource for identifying the causes of a wide range of complex diseases of middle and old age". In: *PLoS medicine* 12.3 (2015), e1001779 (cit. on p. 4).

[659] S. J. Sujit, I. Coronado, A. Kamali, P. A. Narayana, and R. E. Gabr. “Automated image quality evaluation of structural brain MRI using an ensemble of deep learning networks”. In: *Journal of Magnetic Resonance Imaging* 50.4 (2019), pp. 1260–1267 (cit. on p. 6).

[660] T. Sun, D. Sheldon, and B. O’Connor. “A Probabilistic Approach for Learning with Label Proportions Applied to the US Presidential Election”. In: *Procs. of the IEEE Int. Conference on Data Mining 2017*. 2017, pp. 445–454 (cit. on p. 139).

[661] M. Sundararajan and A. Najmi. “The many Shapley values for model explanation”. In: *Procs. of the Int. Conference on Machine Learning 2020*. PMLR. 2020, pp. 9269–9278 (cit. on p. 17).

[662] T. Surmann, D. Biermann, and G. Kehl. “Oscillator model of machine tools for the simulation of self excited vibrations in machining processes”. In: *Procs. of the Int. Conference on Process Machine Interactions 2008*. 2008, pp. 23–29 (cit. on pp. 160, 173).

[663] B. Suthar, H. Patel, and A. Goswami. “A survey: classification of imputation methods in data mining”. In: *International Journal of Emerging Technology and Advanced Engineering* 2.1 (2012), pp. 309–12 (cit. on p. 13).

[664] R. S. Sutton and A. G. Barto. *Reinforcement learning: An introduction*. MIT press, 2018 (cit. on p. 222).

[665] S. Suzuki and K. Abe. “Topological structural analysis of digitized binary images by border following”. In: *Computer vision, graphics, and image processing* 30.1 (1985), pp. 32–46 (cit. on p. 176).

[666] A. Swaminathan and T. Joachims. “Batch learning from logged bandit feedback through counterfactual risk minimization”. In: *Journal of Machine Learning Research* 16 (2015), pp. 1731–1755 (cit. on pp. 236–238).

[667] A. L. Swan, A. Mobasheri, D. Allaway, S. Liddell, and J. Bacardit. “Application of machine learning to proteomics data: classification and biomarker identification in postgenomics biology”. In: *Omics: a journal of integrative biology* 17.12 (2013), pp. 595–610 (cit. on p. 9).

[668] D. Talia and P. Trunfio. *Service-oriented Distributed Knowledge Discovery*. CRC Press, 2013 (cit. on p. 116).

[669] R. M. Talukder et al. “Modifications of the PAMONO-Sensor Help to Size and Quantify Low Number of Individual Biological and Non-Biological Nano-Particles”. In: *Engineering Proceedings* 6.1 (2021). URL: https://www.mdpi.com/2673-4591/6/1/26 (cit. on p. 25).

[670] F. Tang and H. Ishwaran. “Random forest missing data algorithms”. In: *Statistical Analysis and Data Mining: The ASA Data Science Journal* 10.6 (2017), pp. 363–377 (cit. on p. 13).

[671] X. Tang, Z. Wang, Z. Xu, Y. Ma, H. Zhu, and J. Ye. “A Deep Value-network Based Approach for Multi-Driver Order Dispatching”. In: *Procs. of the Int. Conference on Knowledge Discovery and Data Mining 2019*. 2019 (cit. on p. 274).

[672] X. Tang et al. “On the origin and continuing evolution of SARS-CoV-2”. In: *National Science Review* 7.6 (2020), pp. 1012–1023 (cit. on p. 11).

[673] M. A. Tanner and W. H. Wong. “The calculation of posterior distributions by data augmentation”. In: *Journal of the American Statistical Association* 82.398 (1987), pp. 528–550 (cit. on p. 64).

[674] M. Al-Tarawneh, Y. Huang, P. Lu, and D. Tolliver. “Vehicle Classification System Using In-Pavement Fiber Bragg Grating Sensors”. In: *IEEE Sensors Journal* 18.7 (2018), pp. 2807–2815 (cit. on p. 245).

[675] D. M. J. Tax and R. P. W. Duin. “Support Vector Data Description”. In: *Machine Learning* 54.1 (Jan. 2004), pp. 45–66 (cit. on p. 125).

[676] H. Tersing et al. “Simulation of manufacturing chain of a titanium aerospace component with experimental validation”. In: *Finite Elements in Analysis and Design* 51.Supplement C (2012), pp. 10–21 (cit. on p. 158).

[677] R. Tibshirani. “The Lasso Method for Variable Selection in the Cox Model”. In: *Statistics in Medicine* 16.4 (1997), pp. 385–395 (cit. on p. 80).

[678] J. Tiemann and Y. Elmasry. *ATLAS localization framework ROS package repository*. Sept. 2018 (cit. on pp. 183, 185).

[679] J. Tiemann, Y. Elmasry, and L. Koring. *Dataset and Interactive Demo: ATLAS FaST multi-user tracking in heterogeneous logistic applications*. Sept. 2018 (cit. on pp. 183, 185).

[680] J. Tiemann, L. Koring, P. Gorczak, and C. Wietfeld. “Improving the Robustness of Control-Grade Ultra-Wideband Localization”. In: *Procs. of the IFAC Conference on Embedded Systems, Computer Intelligence and Telematics 2018*. Faro, Portugal, June 2018 (cit. on pp. 183, 187, 191).

[681] J. Tiemann, A. Ramsey, and C. Wietfeld. “Enhanced UAV Indoor Navigation through SLAM-Augmented UWB Localization”. In: *Procs. of the IEEE Int. Conference on Communications Workshops: The IEEE Int. Workshop on Communication, Computing, and Networking in Cyber Physical Systems 2018*. Kansas City, USA, May 2018 (cit. on p. 183).

[682] J. Tiemann, F. Schweikowski, and C. Wietfeld. “Design of an UWB Indoor-Positioning System for UAV Navigation in GNSS-Denied Environments”. In: *Procs. of the Int. Conference on Indoor Positioning and Indoor Navigation 2015*. Banff, Canada, Oct. 2015 (cit. on p. 182). **SFB876-A4**

[683] J. Tiemann. “Scalability, Reliability and Energy Efficiency in Ultra-Wideband Indoor Localizations”. PhD thesis. TU Dortmund University, Oct. 2016 (cit. on pp. 183, 188, 191).

[684] J. Tiemann, F. Eckermann, and C. Wietfeld. “ATLAS - An Open-Source TDOA-based Ultra-Wideband Localization System”. In: *Procs. of the Int. Conference on Indoor Positioning and Indoor Navigation 2016*. Alcala de Henares, Madrid, Spain, Oct. 2016 (cit. on pp. 182, 185, 190). **SFB876-A4**

[685] J. Tiemann, F. Eckermann, and C. Wietfeld. “Multi-User Interference and Wireless Clock Syncronization in TDOA-based UWB Localization”. In: *Procs. of the Int. Conference on Indoor Positioning and Indoor Navigation 2016*. Accepted for presentation. Alcala de Henares, Madrid, Spain, Oct. 2016 (cit. on pp. 182, 184). **SFB876-A4**

[686] J. Tiemann, J. Pillmann, and C. Wietfeld. “Ultra-Wideband Antenna-Induced Error Prediction using Deep Learning on Channel Response Data”. In: *Procs. of the IEEE Vehicular Technology Conference 2017*. Sydney, Australia, June 2017 (cit. on p. 188). **SFB876-A4**

[687] J. Tiemann and C. Wietfeld. “Scalability, Real-Time Capabilities and Energy Efficiency in Ultra-Wideband Localization”. In: *IEEE Transactions on Industrial Informatics* 15.8 (2019). forthcoming. URL: https://ieeexplore.ieee.org/document/8618347 (cit. on pp. 182, 183, 185, 188, 250). **SFB876-A4**

[688] J. Tiemann and C. Wietfeld. “Scalable and Precise Multi-UAV Indoor Navigation using TDOA-based UWB Localization”. In: *Procs. of the Int. Conference on Indoor Positioning and Indoor Navigation 2017*. Sept. 2017, pp. 1–7 (cit. on p. 182). **SFB876-A4**

[689] W. R. Tobler. “A computer movie simulating urban growth in the Detroit region”. In: *Economic geography* 46.sup1 (1970), pp. 234–240 (cit. on p. 227).

[690] M.-Q. Tran, M.-K. Liu, and Q.-V. Tran. “Milling chatter detection using scalogram and deep convolutional neural network”. In: *The International Journal of Advanced Manufacturing Technology* 107.3 (2020), pp. 1505–1516 (cit. on p. 158).

[691] T. Treppmann, K. Ickstadt, and M. Zucknick. “Integration of multiple genomic data sources in a Bayesian Cox model for variable selection and prediction”. In: *Computational and Mathematical Methods in Medicine* Vol. 2017 (2017), pp. 1–19. URL: https://www.hindawi.com/journals/cmmm/2017/7340565/ (cit. on pp. 62, 65). **SFB876-C4**

[692] O. Troyanskaya et al. “Missing value estimation methods for DNA microarrays”. In: *Bioinformatics* 17.6 (2001), pp. 520–525 (cit. on p. 13).

[693] A. J. Tudos and R. B. M. Schasfoort. “Introduction to surface plasmon resonance.” In: *Handbook of Surface Plasmon Resonance.* Ed. by A. J. Tudos and R. B. M. Schasfoort. RSC Publishing, 2008, pp. 1–13 (cit. on p. 24).

[694] N. Ueda and R. Nakano. “Generalization error of ensemble estimators”. In: *Procs. of the Int. Conference on Neural Networks 1996*. Vol. 1. IEEE. 1996, pp. 90–95 (cit. on p. 169).

[695] J. Um, I. Kim, and S. Park. “Implementation of Platform for Measurement and Analysis on LTE Traffic and Radio Resource Utilization”. In: *Procs. of the IEEE Int. Conference on Consumer Electronics 2019*. Jan. 2019, pp. 1–2 (cit. on p. 345).

[696] Z. Usfoor, K. Kaufmann, A. S. H. Rakib, R. Hergenröder, and V. Shpacovitch. “Features of Sizing and Enumeration of Silica and Polystyrene Nanoparticles by Nanoparticle Tracking Analysis (NTA)”. In: *Sensors* 20.22 (Nov. 2020), p. 6611. DOI: doi:10.3390/s20226611 (cit. on p. 26). **SFB876-B2**

[697] A. Ustundag and E. Cevikcan. *Industry 4.0: Managing the Digital Transformation*. Springer, 2017 (cit. on p. 103).

[698] T. Vaa, P. Melén, D. Andersson, and B. B. Nielsen. “NorSIKT – Nordic System for Intelligent Classification of Traffic”. In: *Procedia - Social and Behavioral Sciences* 48 (2012). Transport Research Arena 2012, pp. 1702–1712 (cit. on p. 244).

[699] S. Vaishnav, A. Agarwal, and K. Desai. “Machine learning-based instantaneous cutting force model for end milling operation”. In: *Journal of Intelligent Manufacturing* (2019), pp. 1–14 (cit. on p. 158).

[700] V. N. Vapnik. *The Nature of Statistical Learning Theory (Information Science and Statistics)*. 2nd. New York: Springer, 1999 (cit. on pp. 139, 151, 152).

[701] A. Varga and R. Hornig. “An overview of the OMNeT++ simulation environment”. In: *Procs. of the Int. Conference on Simulation Tools and Techniques for Communications, Networks and Systems & Workshops 2008*. Simutools ’08. ICST, Brussels, Belgium, Belgium: ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), 2008, 60:1–60:10 (cit. on p. 273).

[702] A. Vaswani et al. “Attention is All you Need”. In: *Advances in Neural Information Processing Systems 30: Procs. of the 2017 Conference*. Ed. by I. Guyon et al. Vol. 30. Curran Associates, Inc., 2017. URL: https : // proceedings . neurips . cc / paper / 2017 / file / 3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf (cit. on pp. 11, 13).

[703] R. G. Vaughan. *Radiance Calculator program*. 2009 (cit. on p. 201).

[704] W. Viriyasitavat, M. Boban, H. Tsai, and A. Vasilakos. “Vehicular Communications: Survey and Challenges of Channel and Propagation Models”. In: *IEEE Vehicular Technology Magazine* 10.2 (2015), pp. 55–66 (cit. on p. 274).

[705] S. Vishwanathan, N. Schraudolph, R. Kondor, and K. Borgwardt. “Graph Kernels”. In: *Journal of Machine Learning Research* 11 (2010), pp. 1201–1242 (cit. on p. 102).

[706] P. Voigt and A. Von dem Bussche. “The EU General Data Protection Regulation (GDPR)”. In: *A Practical Guide, 1st Ed., Cham: Springer Int. Publishing* 10.3152676 (2017), pp. 10–5555 (cit. on p. 14).

[707] M. Volkovs, F. Chiang, J. Szlichta, and R. J. Miller. “Continuous data cleaning”. In: *Procs. of the IEEE Int. Conference on Data Engineering 2014*. IEEE. 2014, pp. 244–255 (cit. on p. 12).

[708] T. Vranken, B. Sliwa, C. Wietfeld, and M. Schreckenberg. “Adapting a cellular automata model to describe heterogeneous traffic with human-driven, automated, and communicating automated vehicles”. In: *Physica A: Statistical Mechanics and its Applications* (Feb. 2021) (cit. on pp. 293, 295). **SFB876-B4**

[709] T. Vranken, B. Sliwa, C. Wietfeld, and M. Schreckenberg. “Performance comparison of dynamic vehicle routing methods for minimizing the global dwell time in upcoming smart

cities". In: *Procs. of the IEEE Vehicular Technology Conference 2018*. Chicago, USA, Aug. 2018 (cit. on pp. 290, 291). **SFB876-B4**

[710] T. P. E. Vranken and M. Schreckenberg. "Cellular Automata Intersection Model". In: *Collective Dynamics* 5 (2020) (cit. on pp. 290, 291). **SFB876-B4**

[711] S. Wachter, B. Mittelstadt, and C. Russell. "Counterfactual explanations without opening the black box: Automated decisions and the GDPR". In: *Harvard Journal of Law & Technology* 31 (2017), p. 841 (cit. on p. 17).

[712] G. Wang, B. Wang, X. Yang, and G. Yu. "Efficiently Indexing Large Sparse Graphs for Similarity Search". In: *IEEE Transactions on Knowledge and Data Engineering* 24.3 (2012), pp. 440–451 (cit. on p. 93).

[713] M. Wang, Q. Zhang, S. Lam, J. Cai, and R. Yang. "A review on application of deep learning algorithms in external beam radiotherapy automated treatment planning". In: *Frontiers in oncology* 10 (2020), p. 2177 (cit. on pp. 5, 8).

[714] R. Wang, F. M. Harper, and H. Zhu. "Factors Influencing Perceived Fairness in Algorithmic Decision-Making: Algorithm Outcomes, Development Procedures, and Individual Differences". In: *Procs. of the Conference on Human Factors in Computing Systems 2020*. New York, NY, USA: Association for Computing Machinery, 2020, pp. 1–14 (cit. on p. 332).

[715] S. Wang and X. Yao. "Diversity analysis on imbalanced data sets by using ensemble models". In: *2009 IEEE symposium on computational intelligence and data mining* (2009) (cit. on p. 125).

[716] S.-C. Wang. "Artificial neural network". In: *Interdisciplinary computing in java programming*. Springer, 2003, pp. 81–100 (cit. on p. 8).

[717] W. Wang, G.-S. Hong, Y. Wong, and K. Zhu. "Sensor fusion for online tool condition monitoring in milling". In: *International Journal of Production Research* 45.21 (2007), pp. 5095–5116 (cit. on p. 158).

[718] X. Wang, H. Zhao, and J. Zhu. "GRPC: A communication cooperation mechanism in distributed systems". In: *ACM SIGOPS Operating Systems Review* 27.3 (1993), pp. 75–86 (cit. on p. 223).

[719] S. Warnat-Herresthal et al. "Swarm Learning for decentralized and confidential clinical machine learning". In: *Nature* 594.7862 (2021), pp. 265–270 (cit. on pp. 14, 15).

[720] S. Watkins et al. "Ten questions concerning the use of drones in urban environments". In: *Building and Environment* 167 (2020), p. 106458 (cit. on p. 213).

[721] L. Wawrla, O. Maghazei, and T. Netland. "Applications of drones in warehouse operations". In: *Whitepaper. ETH Zurich, D-MTEC* (2019) (cit. on p. 212).

[722] C. Weis et al. "Direct Antimicrobial Resistance Prediction from clinical MALDI-TOF mass spectra using Machine Learning". In: *bioRxiv* (2021). URL: https://www.biorxiv.org/content/early/2021/07/12/2020.07.30.228411 (cit. on pp. 5–7).

[723] C. V. Weis, C. R. Jutzeler, and K. Borgwardt. "Machine learning for microbial identification and antimicrobial susceptibility testing on MALDI-TOF mass spectra: a systematic review". In: *Clinical Microbiology and Infection* (2020) (cit. on p. 7).

[724] B. Weisfeiler and A. A. Lehman. "A Reduction of a Graph to a Canonical Form and an Algebra Arising During this Reduction". In: *Nauchno-Technicheskaya Informatsia* 2.9 (1968) (cit. on pp. 88, 89).

[725] D. Westreich. "Berkson's bias, selection bias, and missing data". In: *Epidemiology (Cambridge, Mass.)* 23.1 (2012), p. 159 (cit. on p. 13).

[726] V. Weyer and H. Binder. "A Weighting Approach for Judging the Effect of Patient Strata on High-Dimensional Risk Prediction Signatures". In: *BMC bioinformatics* 16 (Sept. 2015), p. 294 (cit. on pp. 78, 79).

[727] P. Wiederkehr, F. Finkeldey, and T. Merhofe. “Augmented semantic segmentation for the digitization of grinding tools based on deep learning”. In: *CIRP Annals* 70.1 (2021), pp. 297–300 (cit. on pp. 172, 176–178). **SFB876-B3**

[728] P. Wiederkehr and T. Siebrecht. “Virtual Machining: Capabilities and Challenges of Process Simulations in the Aerospace Industry”. In: *Procedia Manufacturing* 6 (2016), pp. 80–87 (cit. on pp. 157–160).

[729] C. Williams, E. V. Bonilla, and K. M. Chai. “Multi-task Gaussian process prediction”. In: *Advances in Neural Information Processing Systems 20: Procs. of the 2007 Conference*. 2007, pp. 153–160 (cit. on p. 13).

[730] M. J. Williams, B. Werner, C. P. Barnes, T. A. Graham, and A. Sottoriva. “Identification of neutral tumor evolution across cancer types”. In: *Nature Genetics* 48.3 (Mar. 2016), pp. 238–244. URL: https://www.nature.com/articles/ng.3489 (cit. on p. 51).

[731] I. H. Witten, F. Eibe, and M. A. Hall. *Data Mining: Practical Machine Learning Tools and Techniques*. 3rd. Data Management Systems. Burlington, MA: Elsevier, Inc., 2011 (cit. on pp. 137, 148).

[732] M. Won, S. Sahu, and K. Park. “DeepWiTraffic: Low Cost WiFi-Based Traffic Monitoring System Using Deep Learning”. In: *Procs. of the IEEE Int. Conference on Mobile Ad Hoc and Sensor Systems 2019*. 2019, pp. 476–484 (cit. on p. 246).

[733] K. Wüstefeld and F. Weichert. “An Automated Rapid Test for Viral Nanoparticles Based on Spatiotemporal Deep Learning”. In: *Procs. of IEEE SENSORS 2020*. Oct. 2020, pp. 1–4. DOI: 10.1109/SENSORS47125.2020.9278935 (cit. on pp. 27, 29–33). **SFB876-B2**

[734] Y. Xie et al. “Early lung cancer diagnostic biomarker discovery by machine learning methods”. In: *Translational oncology* 14.1 (2021), p. 100907 (cit. on p. 10).

[735] H. Xu and K. Larson. “Improving the efficiency of crowdsourcing contests”. In: *Procs. of the Int. Conference on Autononomous Agents and Multi-Agent Systems 2014*. 2014 (cit. on p. 298).

[736] J. Xu, B. S. Glicksberg, C. Su, P. Walker, J. Bian, and F. Wang. “Federated learning for healthcare informatics”. In: *Journal of Healthcare Informatics Research* 5.1 (2021), pp. 1–19 (cit. on p. 14).

[737] K. Xu, W. Hu, J. Leskovec, and S. Jegelka. “How Powerful are Graph Neural Networks?” In: *Procs. of the Int. Conference on Learning Representations 2019*. 2019 (cit. on p. 89).

[738] G. Xue, J. Xu, H. Wu, and W. Lu. “Incentive Mechanism for Rational Miners in Bitcoin Mining Pool”. In: *Information Systems Frontiers* 23.10 (2021) (cit. on pp. 299, 302, 309).

[739] M. Yakout, L. Berti-Équille, and A. K. Elmagarmid. “Don’t be scared: use scalable automatic repairing with maximal likelihood and bounded changes”. In: *Procs. of the ACM SIGMOD Int. Conference on Management of Data 2013*. 2013, pp. 553–564 (cit. on p. 12).

[740] K. Yamanishi and J. Takeuchi. “A Unifying Framework for Detecting Outliers and Change Points from Non-stationary Time Series Data”. In: *Procs. of the ACM SIGKDD Conference on Knowledge Discovery and Data Mining 2002*. KDD ’02. New York, NY, USA: ACM, 2002 (cit. on p. 168).

[741] J. Yan, X. Yin, W. Lin, C. Deng, H. Zha, and X. Yang. “A Short Survey of Recent Advances in Graph Matching”. In: *Procs. of the ACM on Int. Conference on Multimedia Retrieval 2016*. Ed. by J. R. Kender, J. R. Smith, J. Luo, S. Boll, and W. H. Hsu. ACM, 2016, pp. 167–174. DOI: https://doi.org/10.1145/2911996.2912035 (cit. on p. 87).

[742] L. Yan et al. “An interpretable mortality prediction model for COVID-19 patients”. In: *Nature machine intelligence* 2.5 (2020), pp. 283–288 (cit. on p. 11).

[743] D. Yang, G. Xue, X. Fang, and J. Tang. “Crowdsourcing to smartphones: Incentive mechanism design for mobile phone sensing”. In: *Procs. of ACM MobiCom 2012*. 2012 (cit. on pp. 298, 299).

[744] K. Yang, B. Huang, J. Stoyanovich, and S. Schelter. “Fairness-Aware Instrumentation of Preprocessing Pipelines for Machine Learning”. In: *Procs. of the Workshop on Human-In-the-Loop Data Analytics 2020*. 2020 (cit. on p. 13).

[745] Q. Yang, Y. Liu, T. Chen, and Y. Tong. “Federated machine learning: Concept and applications”. In: *ACM Transactions on Intelligent Systems and Technology (TIST)* 10.2 (2019), pp. 1–19 (cit. on p. 14).

[746] M. Yayla et al. “Nanoparticle Classification Using Frequency Domain Analysis on Resource-Limited Platforms”. In: *Sensors* 19 (2019). DOI = 10.3390/s19194138. URL: https://www.mdpi.com/1424-8220/19/19/4138 (cit. on pp. 30, 31). **SFB876-B2**

[747] Z. Ye, H. Xiong, and L. Wang. “Collecting Comprehensive Traffic Information Using Pavement Vibration Monitoring Data”. In: *Computer-Aided Civil and Infrastructure Engineering* 35.2 (Jan. 2020), pp. 134–149. URL: https://doi.org/10.1111/mice.12448 (cit. on p. 245).

[748] R. Yearby. “Racial disparities in health status and access to healthcare: the continuation of inequality in the United States due to structural racism”. In: *American Journal of Economics and Sociology* 77.3-4 (2018), pp. 1113–1152 (cit. on p. 18).

[749] J. Yoon, J. Jordon, and M. Schaar. “Gain: Missing data imputation using generative adversarial nets”. In: *Procs. of the Int. Conference on Machine Learning 2018*. PMLR. 2018, pp. 5689–5698 (cit. on p. 13).

[750] F. X. Yu, S. Kumar, T. Jebara, and S. Chang. “On Learning with Label Proportions”. In: *arXiv: Computing Research Repository* (2014). DOI: arXiv:1402.5902 (cit. on pp. 142–146).

[751] F. X. Y. Yu, D. Liu, S. Kumar, T. Jebara, and S. Chang. “$\propto$SVM for Learning with Label Proportions”. In: *Procs. of the Int. Conference on Machine Learning 2013*. 2013, pp. 504–512 (cit. on p. 140).

[752] P. Yuan, A. Rezvan, X. Li, N. Varadarajan, and H. Van Nguyen. “Phasetime: Deep Learning Approach to Detect Nuclei in Time Lapse Phase Images”. In: *Journal of Clinical Medicine* 8.8 (2019) (cit. on p. 32).

[753] L. Z. Yunpu Ma Volker Tresp and Y. Wang. “Variational Quantum Circuit Model for Knowledge Graph Embedding”. In: *Advanced Quantum Technologies* (2019). DOI: DOI:10.1002/qute.201800078 (cit. on p. 388).

[754] A. Zeltins. “Construction and Characterization of Virus-Like Particles: A Review”. In: *Molecular Biotechnology* 53.1 (2013), pp. 92–107. URL: %3CGo%20to%20ISI%3E://WOS:000312640400011 (cit. on p. 23).

[755] Z. Zeng, A. K. H. Tung, J. Wang, J. Feng, and L. Zhou. “Comparing Stars: On Approximating Graph Edit Distance”. In: *Procs. VLDB Endowment* 2.1 (Aug. 2009), pp. 25–36. URL: http://dx.doi.org/10.14778/1687627.1687631 (cit. on pp. 93, 94).

[756] J. Zentgraf and S. Rahmann. “Fast lightweight accurate xenograft sorting”. In: *Algorithms for Molecular Biology* 16.1 (2021), p. 2. URL: https://doi.org/10.1186/s13015-021-00181-w (cit. on p. 46). **SFB876-C1**

[757] J. Zentgraf, H. Timm, and S. Rahmann. “Cost-optimal assignment of elements in genome-scale multi-way bucketed Cuckoo hash tables”. In: *Procs. of the Symposium on Algorithm Engineering and Experiments 2020*. SIAM, 2020, pp. 186–198. URL: https://epubs.siam.org/doi/abs/10.1137/1.9781611976007.15 (cit. on p. 46). **SFB876-C1**

[758] A. Zhang, S. Song, J. Wang, and P. S. Yu. “Time series data cleaning: From anomaly detection to anomaly repairing”. In: *Procs. of the VLDB Endowment* 10.10 (2017), pp. 1046–1057 (cit. on p. 12).

[759] K. Zhang, M. Hutter, and H. Jin. “A new local distance-based outlier detection approach for scattered real-world data”. In: *Procs. of the Pacific-Asia Conference on Knowledge Discovery and Data Mining 2009*. Springer. 2009, pp. 813–822 (cit. on p. 14).

[760] Q. Zhang, Y. Liu, G. Liu, G. Zhao, Z. Qu, and W. Yang. "An automatic diagnostic system based on deep learning, to diagnose hyperlipidemia". In: *Diabetes, metabolic syndrome and obesity: targets and therapy* 12 (2019), p. 637 (cit. on p. 6).

[761] T. Zhang and R. Ma. "Identify the cells' nuclei based on the deep learning neural network". In: *arXiv: Computing Research Repository* (2019). DOI: arXiv:1911.09830 (cit. on p. 32).

[762] X. Zhang, Z. Yang, Z. Zhou, H. Cai, L. Chen, and X. Li. "Free market of crowdsourcing: Incentive mechanism design for mobile sensing". In: *IEEE Transactions on Parallel and Distributed Systems* 25.12 (2014) (cit. on pp. 298, 299, 302).

[763] X. Zhang et al. "Incentives for Mobile Crowd Sensing: A Survey". In: *IEEE Communications Surveys and Tutorials* 18.1 (2015) (cit. on p. 298).

[764] X. Zhang, C. Gu, and J. Lin. "Support vector machines for anomaly detection". In: *Procs. of the World Congress on Intelligent Control and Automation 2006*. Vol. 1. IEEE. 2006, pp. 2594–2598 (cit. on pp. 9, 14).

[765] X. Zhang, X. Lu, S. Wang, W. Wang, and W. Li. "A multi-sensor based online tool condition monitoring system for milling process". In: *Procedia CIRP* 72 (2018), pp. 1136–1141 (cit. on p. 158).

[766] Y. Zhang, Y. Fang, and S. Zhong. "Incentive mechanism design for smartphone crowdsensing". In: *Procs. of the IEEE Int. Conference on Big Data and Cloud Computing 2015*. 2015 (cit. on pp. 298, 299, 302).

[767] B. Zhao, H. Lu, S. Chen, J. Liu, and D. Wu. "Convolutional neural networks for time series classification". In: *Journal of Systems Engineering and Electronics* (2017) (cit. on p. 130).

[768] X. Zhao, C. Xiao, X. Lin, W. Wang, and Y. Ishikawa. "Efficient processing of graph similarity queries with edit distance constraints". In: *VLDB J.* 22.6 (2013), pp. 727–752 (cit. on pp. 93, 94).

[769] X. Zhao, C. Xiao, X. Lin, W. Zhang, and Y. Wang. "Efficient structure similarity searches: a partition-based approach". In: *VLDB – The International Journal on Very Large Data Bases* 27.1 (2018), pp. 53–78 (cit. on pp. 93, 94).

[770] W. Zheng, L. Zou, X. Lian, D. Wang, and D. Zhao. "Efficient Graph Similarity Search Over Large Graph Databases". In: *IEEE Transactions on Knowledge and Data Engineering* 27.4 (2015), pp. 964–978 (cit. on p. 93).

[771] R. Zou and W. Wang. "U-CIMAN: Uncover Spectrum and User Information in LTE Mobile Access Networks". In: *Procs. of the IEEE Conference on Computer Communications 2020*. 2020, pp. 1459–1468 (cit. on p. 347).

[772] A. Zybin, Y. Kuritsyn, E. Gurevich, V. Temchura, K. Überla, and K. Niemax. "Real-time Detection of Single Immobilized Nanoparticles by Surface Plasmon Resonance Imaging". In: *Plasmonics* 5 (2009), pp. 31–35 (cit. on p. 24).

[773] A. Zybin, V. Shpacovitch, J. Skolnik, and R. Hergenröder. "Optimal conditions for SPR-imaging of nano-objects". In: *Sensors and Actuators B: Chemical* 239 (2017), pp. 338–342 (cit. on p. 25). **SFB876-B2**

[774] N. Zygouras, N. Zacheilas, V. Kalogeraki, D. Kinane, and D. Gunopulos. "Insights on a Scalable and Dynamic Traffic Management System". In: *Procs. of the Int. Conference on Extending Database Technology 2015*. 2015, pp. 653–664 (cit. on p. 235).

Index

3D Euclidean positioning error, 382
5G New Radio (NR), 385
5G networks, 224, 318, 363, 377
5G positioning, 388
6G networks, 329, 387

Active Learning (AL), 160
Age of Information (AoI), 333
Aggregate Output Classification Kernel k-Means (AOC-KK), 141
Akaike's Information Criterion (AIC), 325
Algebraic Equation Set (AES), 198
Angle-based positioning, 388
Angle-of-Arrival (AoA), 383
Angle-of-Departure (AoD), 383
Antenna
– directivity, 376
– Phased Array (PAA), 376
Anticipatory mobile networking, 275, 332
Attenuation patterns, 241
Augemented Reality (AR), 355
Autoencoder, 16
Automated vehicles (AVs), 293
AutoRegressive Integrated Moving Average (ARIMA), 325, 366

Background, 27, 34
Backpropagation, 31
Bagging, 125
Bayes
– naive Bayes classifier, 124
– optimal Bayes classifier, 143
– theorem, 58
Bayesian network, 228
Beam
– alignment, 375
– management, 385
– sweeping, 376
– tracking, 375
Bearing-based mmWave positioning, 383
Better Approach to Mobile Ad-hoc Networking (B.A.T.M.A.N.), 273
Bias, 18, 49, 78, 166, 169–172
Binary tree, 67
Biomarker discovery, 9
Black spot, 338
Blackspot-Aware Contextual Bandit (BS-CB), 331
Block Error Rates (BLER), 329
Boosting, 9, 71, 125

Cellular automaton model, 291, 293
Central Processing Unit (CPU), 35, 45, 46
Central Unit (CU), 385
Channel Impulse Response (CIR), 241
Channel phase tracking, 388
Channel Quality Indicator (CQI), 334
Channel State Information (CSI), 241
Channel-Aware Transmission (CAT), 331
Clock synchronization, 184
Clustering, 82, 86, 127, 161
Commercial-Off-The-Shelf (COTS), 356
Communicating Automated Vehicles (CAVs), 293
Communication
– Enhanced Machine Type (eMTC), 261
– massive Machine Type (mMTC), 261
– massive Machine Type Communication (mMTC), 313, 331, 355
Competition-free channel access, 189
Confidentiality, 15
– enforcement, 394
– requirement, 393
Constrained Application Protocol (CoAP), 266
Coresets, 59
Critical infrastructure communications, 358
Cross validation, 81, 130, 152, 155
Cross-validation, 80
Crowdsensing, 297, 331
Cyber-Physical System (CPS), 104, 180, 241
Cyber-physical twin, 109

Data
– acquisition, 195
– stream algorithm, 66
Data-Driven Model-Predictive Communication, 323
Data-Driven Network Simulation (DDNS), 338
Database search, 86, 88, 93

https://doi.org/10.1515/9783110785968-008

Decision tree, 124
Deep Boltzmann Trees (DBTs), 254
Deoxyribonucleic acid (DNA), 23
– fragment, 45
– sequencing, 43, 49
– transcription, 85
Detect-And-Avoid (DAA), 315
Detector, 29, 30, 32
Differential privacy, 231
Digital twin, 108
Dilution Of Precision (DOP), 384
Distributed machine learning, 137
Distributed Unit (DU), 385
Downlink Control Information (DCI), 346
Duty cycle, 315
Dynamic transit routing, 233

Embedded device, 37
Embedded system, 35
End-to-end system architecture, 384
Energy, 35
– consumption, 21, 22, 35, 36
– efficiency, 35, 185
– harvesting, 195
– saving, 35
enhanced Mobile Broadband (eMBB), 355
Ensembles, 9, 127, 162, 165, 169, 170, 172
– resampling-based, 125
Error
– mean squared, 164, 169
– measures, 161
Error Vector Magnitude (EVM), 380
Euclidean distance, 69, 168
Evolutionary algorithms, 147
Exponential families, 229
Extended Kalman Filter (EKF), 187

Fast Analysis of LTE CONtrol channels (FALCON), 344
Feature selection, 9, 58, 59, 71, 133
Field-Programmable Gate Array (FPGA), 41
Floating Car Data (FCD), 356

Gaussian processes, 228
Genetic
– analysis, 43
– optimization, 34
Genetic programming algorithm, 123
Global Navigation Satellite Systems (GNSS), 181, 242
Graphical models, 142, 228
Graphics Processing Unit (GPU), 35, 45

Hashing
– Cuckoo, 46
Hierarchical Token Bucket (HTB), 356
High accuracy 3D motion tracking, 388
Human-driven Vehicles (HVs), 293

Incentive mechanisms, 297
– data-reverse-auction, 299
– dynamic relative, 304
– static uniform, 304
– task-reverse-auction, 299
Industry 4.0, 103, 158, 179, 182, 196, 212
Inference, 65
– proofness, 391
Information
– communicating, 392
– exploiting, 392
– keeping, 392
– system, 393
Integer linear programming, 88, 95
Intelligent Transportation Systems (ITS), 241, 272, 356
Internet of Things (IoT), 137, 213, 261
– cellular, 261
– Narrowband (NB-IoT), 261

Joint Communication And Sensing (JCAS), 387

k-Means
– Fast, 151
k-Nearest Neighbour (kNN), 124
Kernel functions, 228
– graph, 87, 89
– Radial Basis Function (RBF), 335
Key Performance Indicator (KPI), 320, 354
Knowledge
– a priori, 36, 40, 52, 58, 395
– background, 32, 86, 95, 395
– epistemic state, 394
– inferred, 396

L1 norm, 59
L2 norm, 59
Label
– counts, 232
– proportions, 231
Laplace distribution, 232
Latent variable model, 48, 49, 53
Learning
– semi-supervised, 137

– supervised, 123, 137, 242, 309
– unsupervised, 124, 137, 309
Learning from Label Proportions (LLP), 136, 137, 146, 231, 233
Learning from Label Proportions by Clustering (LLPC), 147
Learning tasks, 173, 175
Lee model, 290, 293
LIghtweight Machine Learning for IoT Systems with resource limitations (LIMITS), 277, 335
Line-Of-Sight (LOS), 279, 380
– obstructed, 385
Link Expiry Time (LET), 276
Link quality, 269
Listen-Before-Talk (LBT), 315
Logistics, 110
Long Short Term Memory (LSTM), 326
Long-Range Wide-Area Network (LoRaWAN), 262, 315
Low-Power Wide Area Network (LPWAN), 313
Low-Power Wide-Area Network (LPWAN), 261, 262

M5 regression tree, 335
Machine Learning-CAT (ML-CAT), 331
Machine to Machine (M2M), 265
Machine-Type Communications (MTC), 329
Management And Network Orchestration (MANO), 356, 385
Markov chains, 50, 65
Markov Random Fields (MRFs), 229
Maturity index, 106
Maximum likelihood, 80
Maximum Power Point (MPP), 196
Median filtering, 29
Medical data analysis, 22
Medium Access Control (MAC), 315
Memory, 66, 74, 77
– utilization, 254
Merge & Reduce, 58, 59, 66
Message Queuing Telemetry Transport (MQTT), 220, 265
Message Queuing Telemetry Transport for Sensor Networks (MQTT-SN), 266
Millimeter wave (mmWave)
– communication, 375, 389
– network coverage, 385
– pencil beam, 375
– radio links, 378
Min-hashing, 47, 88, 89, 93
Mixture model, 46
Mobile Ad-hoc Networks (MANETs), 272
Mobile radio networks, 385
Mobile system, 21
Mobility-predictive routing, 274
Motion Capture (MoCap) system, 217
Multi-modal trip planning, 233
Multilayer perceptrons, 162
Nagel-Schreckenberg model, 227, 287
Nash equilibrium, 299
Network
– coverage, 263, 389
– slicing, 363, 368
Network Data Analytic Function (NWDAF), 340
Network Function Virtualization (NFV), 355
Network Interface Card (NIC), 362
Neural Network
– Convolutional Neural Networks (CNN), 6, 8, 32, 142
– Fully Convolutional Networks (FCN), 176
– Generative Adversarial Networks (GAN), 38, 40
– Recurrent Neural Networks (RNN), 326
– Residual Networks model (ResNet), 32

Offloading, 35
Operating system, 362
Optimized Link State Routing (OLSR), 273

PAMONO, 21, 31, 33, 36
Parallelism, 45
Particle, 23, 27, 30–32
– detection, 30
Passive reflectors, 385
Pearson correlation
– coefficient, 97, 127
Personalized therapy, 43
Pervasive intelligence, 332
Photovoltaics, 195
Physical Downlink Control CHannel (PDCCH), 346
Physical layer Service Data Unit (PSDU), 190

Poisson Dependency Network (PDN), 230, 289
Policy Optimizer for Exponential Models (POEM), 236
Power
– consumption, 190
Precision Time Protocol (PTP), 358
Predictive Ad-hoc Routing Fueled by Reinforcement Learning and Trajectory knowledge (PARRoT), 273
Private networks, 388
Proportional hazards regression, 63
Protein
– complex, 85, 86
– interactions, 88, 95
Proximal Policy Optimization (PPO), 223
Proximity forest, 254
Pruning, 127

Q-learning, 275
Q-routing, 274
Quality of Service (QoS), 262, 332, 365, 385

Radio
– based sensing, 388
– fingerprints, 246
– Radio Network Temporary Identifier (RNTI), 343
– smart radio environments, 385
– Software-Defined Radio (SDR), 320, 345
  – based mmWave platform, 375
Random forest, 31, 171, 254, 277, 335
Random walk Metropolis-Hastings method, 64
Read mapper, 45, 49
Real-time capability, 21
Real-Time Clock (RTC), 186
Real-Time Localization System (RTLS), 184
Reconfigurable Intelligent Surface (RIS), 385
– placement, 386
Regression, 71, 123, 162
– Bayesian, 58
– Gaussian process regression, 80
– generalized linear regression models, 66
– LASSO, 9, 80, 81
– linear, 60, 65, 67
– logic, 61
– ordinary least squares, 59
– Penalized regression models, 79
– Poisson tree, 230
– proportional hazards, 62, 76, 79
Regularization, 59, 63, 130, 161
Reinforcement learning, 127, 221, 272, 273, 283
Reinforcement Learning-CAT (RL-CAT), 331
Resource
– optimization, 22, 35
Resource Block (RB), 346
Resource Element (RE), 345
Resource-constrained, 182
Ribonucleic acid (RNA), 23
Risk classification, 43
Road-side Unit (RSU), 273
Robot Operating System (ROS), 184, 217
Root Mean Square Error (RMSE), 327
Routing
– algorithm, 272
– protocols, 272–274

Sampling, 50, 81
– informed, 162
– Markov-Chain-Monte-Carlo (MCMC), 64, 68, 139
– random, 165
Scalability, 58, 65, 188
Scheduling, 185, 364
Seasonal and Trend decomposition using LOESS (STL), 324
Security
– availability, 397
– confidentiality, 397
– rational reasoner, 396
Security mechanism, 395
– overhead, 409
Security policy, 396
– permission, 394
– prohibition, 395
Self-Organizing Maps (SOM), 133
Sequencing coverage, 47
Sequential Minimal Optimization (SMO), 335
Short-Range Device (SRD), 315
Signal
– background separation, 29
– quality measurements, 378
– Received Signal Strength Indicator (RSSI), 242, 318
– Reference Signal Received Power (RSRP), 334
– Reference Signal Received Quality (RSRQ), 334
– Signal-to-Interference-plus-Noise Ratio (SINR), 332
– Signal-to-Noise Ratio (SNR), 32, 350

Similarity
– cosine, 97, 98
– graph edit, 85, 92, 93
– Jaccard, 87, 88
– Weisfeiler-Leman, 88, 91, 93
Similarity measures, 86, 89
Simulation of Urban MObility (SUMO), 239
Sketches, 59, 89
Slice-Aware Machine Learning-Based Ultra-Reliable Scheduling (SAMUS), 365
Slice-Aware Machine Learning-based Ultra-Reliable Scheduling (SAMUS), 354
Smart Grid communications, 354
Software-Defined Networking (SDN), 355, 385
Spatio-temporal deep learning, 31
Spectral efficiency, 261
Speedup, 36
Stackelberg game, 299
Sub-sampling, 81
Support Vector Data Description (SVDD), 125
Support Vector Machine (SVM), 124, 133, 254, 277, 335
– Regression (SVR), 140, 158
Surface Plasmon Resonance (SPR), 24
Surrogate model, 71, 77
Synchronization, 187

Template matching, 31, 32
Time Of Arrival (TOA), 181
Time series data, 27, 119, 126, 133, 161, 168, 324
Time-Difference Of Arrival (TDOA), 184
Time-on-Air (ToA), 318
Timing Advance (TA), 334
Transfer learning, 32
Transit networks, 233
Transmission Time Interval (TTI), 365
Tumor, 43, 53, 64
– primary, 43
– relapse, 43
Two-Way Ranging (TWR), 190

Ultra-Low Power (ULP), 195
ultra-Reliable Low Latency Communication (uRLLC), 354, 356
Ultra-Wideband (UWB), 181
Unmanned Aerial Vehicles (UAV), 112, 184, 212, 272, 378
User Datagram Protocol (UDP), 358
User Equipment (UE), 344, 365
User positioning, 388

Variable Importance Measure, 59
Variance, 78, 166, 169, 170, 172
Variant calling, 48
Virtual Reality (VR), 355
Virus, 21

Waikato Environment for Knowledge Analysis (WEKA), 335
Weighted Euclidean distance, 151
Wireless Sensor Network (WSN), 242
Workflow management system, 44

Xenograft sorting, 46

List of Contributors

Editors

Morik, Katharina, Prof. Dr. TU Dortmund University, Department of Computer Science, Head of the chair for Artificial Intelligence & Speaker of the Collaborative Research Center 876 on Providing Information by Resource-Constrained Data Analysis, katharina.morik@tu-dortmund.de

Wietfeld, Christian, Prof. Dr. TU Dortmund University, Faculty of Electrical Engineering and Information Technology, Head of the chair for Communication Networks & Co-Speaker of the Collaborative Research Center 876, christian.wietfeld@tu-dortmund.de

Rahnenführer, Jörg, Prof. Dr. TU Dortmund University, Department of Statistics, Head of the group for Statistical Methods in Genetics and Chemometrics, CRC 876, rahnenfuehrer@statistik.tu-dortmund.de

Contributors

Arendt, Christian, M. Sc. TU Dortmund University, Faculty of Electrical Engineering and Information Technology, Communication Networks Institute, CRC 876,

Awasthi, Shrutarv, M. Sc. TU Dortmund University, Faculty of Mechanical Engineering, Chair of Materials Handling and Warehousing, FLW, CRC 876, shrutarv.awasthi@tu-dortmund.de

Bektas, Caner, M. Sc. TU Dortmund University, Faculty of Electrical Engineering and Information Technology, Communication Networks Institute, caner.bektas@tu-dortmund.de

Biskup, Joachim, Prof. Dr. TU Dortmund University, Department of Computer Sciences, Databases and Information Systems Group, CRC 876, joachim.biskup@cs.tu-dortmund.de

Borgwardt, Karsten, Prof. Dr. ETH Zurich, Department of Biosystems Science and Engineering, Machine Learning & Computational Biology Lab, karsten.borgwardt@bsse.ethz.ch

Böcker, Stefan, Dipl.- Ing. TU Dortmund University, Faculty of Electrical Engineering and Information Technology, Communication Networks Institute, CRC 876, stefan.boecker@tu-dortmund.de

Büscher, Jan, M. Sc. TU Dortmund University, Faculty of Mechanical Engineering, Institute of Production Systems (IPS), CRC 876, jan.buescher@ips.tu-dortmund.de

Cheng, Liang, Dr. University of Toledo, Electrical Engineering & Computer Science Department, EECS Department Chair & Professor, liang.cheng@utoledo.edu

Deuse, Jochen, Univ.-Prof. Dr.-Ing. TU Dortmund University, Faculty of Mechanical Engineering, Institute of Production Systems (IPS), CRC 876, jochen.deuse@ips.tu-dortmund.de

Ding, Zeyu, M. Sc. TU Dortmund University, Department of Statistics, CRC 876, zeyu.ding@tu-dortmund.de

https://doi.org/10.1515/9783110785982-009

Finkeldey, Felix, M. Sc. TU Dortmund University, Department of Computer Sciences, Chair for Software Engineering, CRC 876, felix.finkeldey@tu-dortmund.de

Gramse, Nils, M. Sc. TU Dortmund University, Faculty of Mechanical Engineering, Chair of Materials Handling and Warehousing, FLW, nils.gramse@tu-dortmund.de

Haferkamp, Marcus, M. Sc. TU Dortmund University, Faculty of Electrical Engineering and Information Technology, Communication Networks Institute, CRC 876, marcus.haferkamp@tu-dortmund.de

Heimann, Karsten, M. Sc. TU Dortmund University, Faculty of Electrical Engineering and Information Technology, Communication Networks Institute, CRC 876, karsten.heimann@tu-dortmund.de

Hergenröder, Roland, Dr. Leibniz-Institut für Analytische Wissenschaften - ISAS - e.V., CRC 876, roland.hergenroeder@isas.de

ten Hompel, Michael, Prof. Dr. Dr. h. c. Fraunhofer-Institut für Materialfluss und Logistik IML, CRC 876, michael.tenhompel@tu-dortmund.de

Ickstadt, Katja, Prof. Dr. Department of Statistics, Mathematical Statistics with Applications in Biometrics, CRC 876, ickstadt@statistik.tu-dortmund.de

Jörke, Pascal, M. Sc. TU Dortmund University, Faculty of Electrical Engineering and Information Technology, Communication Networks Institute, CRC 876, pascal.joerke@tu-dortmund.de

Jutzeler, Catherine, Prof. Dr. ETH Zurich, Department of Health Sciences and Technology, catherine.jutzeler@hest.ethz.ch

Köster, Johannes, Dr. Duisburg-Essen University, University Hospital Essen, Institute of Human Genetics, CRC 876, johannes.koester@uk-essen.de

Kriege, Nils, Ass.-Prof. Dipl.-Inf. Dr. University of Vienna, Research Group Data Mining and Machine Learning, CRC 876, nils.kriege@univie.ac.at

Krieger, Cedrik, M. Sc. TU Dortmund University, Faculty of Electrical Engineering and Information Technology, Communication Networks Institute, CRC 876, cedrik.schueler@tu-dortmund.de

Kurtz, Fabian, Dr.-Ing. TU Dortmund University, Faculty of Electrical Engineering and Information Technology, Communication Networks Institute, fabian.kurtz@tu-dortmund.de

Lang, Michel, Dr. TU Dortmund University, Department of Statistics, Statistical methods in Genetics and Chemometrics, CRC 876, lang@statistik.tu-dortmund.de

Liebig, Thomas, Juniorprof. Dr. TU Dortmund University, Department of Computer Sciences, Chair for Artificial Intelligence, CRC 876, thomas.liebig@cs.tu-dortmund.de

Machado, Maximilian, M. Sc. Lehigh University, Department of Electrical and Computer Engineering, maximillian.k.machado@gmail.com

Masoudinejad, Mojtaba, Dr.-Ing. TU Dortmund University, Department of Computer Science, Design Automation for Embedded Systems Group, CRC 876, mojtaba.masoudinejad@tu-dortmund.de

Munteanu, Alexander, Dr. TU Dortmund University, Dortmund Data Science Center, Faculties of Statistics and Computer Science, CRC 876, alexander.munteanu@tu-dortmund.de

Mutzel, Petra, Prof. Dr. Bonn University, Institute of Computer Science, Computer Science I, CRC 876, petra.mutzel@cs.uni-bonn.de

Overbeck, Dennis, M. Sc. TU Dortmund University, Faculty of Electrical Engineering and Information Technology, Communication Networks Institute, dennis.overbeck@tu-dortmund.de

Panusch, Thorben, M. Sc. TU Dortmund University, Department of Computer Science, CRC 876, thorben.panusch@ips.tu-dortmund.de

Rahmann, Sven, Prof. Dr. Saarland University, Center for Bioinformatics, CRC 876, sven.rahmann@tu-dortmund.de

Ran, Ran, M. Sc. Lehigh University, Department of Electrical and Computer Engineering, rar418@lehigh.edu

Richter, Jakob, Dr. TU Dortmund University, Department of Statistics, Statistical methods in Genetics and Chemometrics, CRC 876, richter@statistik.tu-dortmund.de

Roidl, Moritz, Dipl.-Inform. TU Dortmund University, Faculty of Mechanical Engineering, Chair of Materials Handling and Warehousing, FLW, CRC 876, moritz.roidl@tu-dortmund.de

Saadallah, Amal, M. Sc. TU Dortmund University, Department of Computer Sciences, Chair for Artificial Intelligence, CRC 876, amal.saadallah@cs.tu-dortmund.de

Schramm, Alexander, Prof. Dr. University Hospital Essen, Department of Medical Oncology, CRC 876, alexander.schramm@uk-essen.de

Schreckenberg, Michael, Prof. Dr. University of Duisburg Essen, Faculty for Physics, Physics for Transport and Traffic, CRC 876, michael.schreckenberg@uni-due.de

Shpacovitch, Victoria, Dr. Leibniz-Institut für Analytische Wissenschaften - ISAS - e.V., CRC 876, victoria.shpacovitch@isas.de

Sliwa, Benjamin, Dr.-Ing. TU Dortmund University, Faculty of Electrical Engineering and Information Technology, Communication Networks Institute, CRC 876, benjamin.sliwa@tu-dortmund.de

Stolpe, Marco, Dr. TU Dortmund University, Department of Computer Sciences, Chair for Artificial Intelligence, CRC 876, marco.stolpe@tu-dortmund.de

Stöcker, Bianca, M. Sc. Duisburg-Essen University, University Hospital Essen, Institute of Human Genetics, CRC 876, bianca.stoecker@uni-due.de

Tiemann, Janis, Dr.-Ing. TU Dortmund University, Faculty of Electrical Engineering and Information Technology, Communication Networks Institute, CRC 876, janis.tiemann@tu-dortmund.de

Vranken, Tim, M. Sc. University of Duisburg Essen, Faculty for Physics, Theoretical Physics, CRC 876, tim.vranken@uni-due.de

Weichert, Frank, Priv.-Doz. Dr. TU Dortmund University, Department of Computer Science, Intelligent Sensing in Computer Graphics, CRC 876,frank.weichert@tu-dortmund.de

Wiederkehr, Petra, Prof. Dr.-Ing. TU Dortmund University, Department of Computer Sciences, Chair for Software Engineering, CRC 876, petra.wiederkehr@tu-dortmund.de

Wüstefeld, Konstantin, M. Sc. TU Dortmund University, Department of Computer Sciences, CRC 876, konstantin.wuestefeld@tu-dortmund.de

Technical Editors

Becker, Andreas, Dr. TU Dortmund University, Department of Computer Science, CRC 876, andreas3.becker@tu-dortmund.de

Buß, Jens, Dr. TU Dortmund University, Department of Computer Science, Managing Director of the Collaborative Research Center 876, jens.buss@tu-dortmund.de

Acknowledgment

Part of the work on this book is the result of research of the Collaborative Research Center 876 "Providing Information by Resource-Constrained Analysis", which was funded from 2011–2022 by the Deutsche Forschungsgemeinschaft (DFG) under DFG project number 124020371, see: https://gepris.dfg.de/gepris/projekt/124020371?language=en.