

Received 4 July 2023, accepted 17 July 2023, date of publication 9 August 2023, date of current version 31 August 2023.

Digital Object Identifier 10.1109/ACCESS.2023.3303528

## RESEARCH ARTICLE

# Machine Learning for QoS Prediction in Vehicular Communication: Challenges and Solution Approaches

ALEXANDROS PALAIOS<sup>1</sup>, CHRISTIAN L. VIELHAUS<sup>2</sup>, DANIEL F. KÜLZER<sup>3</sup>,  
CARA WATERMANN<sup>1</sup>, RODRIGO HERNANGÓMEZ<sup>5</sup>, (Graduate Student Member, IEEE),  
SANKET PARTANI<sup>4</sup>, (Associate Member, IEEE), PHILIPP GEUER<sup>1</sup>, ANTON KRAUSE<sup>6</sup>,  
RAJA SATTIRAJU<sup>4</sup>, MARTIN KASPARICK<sup>5</sup>, GERHARD P. FETTWEIS<sup>6</sup>, (Fellow, IEEE),  
FRANK H. P. FITZEK<sup>2</sup>, (Senior Member, IEEE), HANS D. SCHOTTEN<sup>4</sup>, (Member, IEEE),  
AND SLAWOMIR STAŃCZAK<sup>5,7</sup>, (Senior Member, IEEE)

<sup>1</sup>Ericsson Research, 52134 Herzogenrath, Germany

<sup>2</sup>Deutsche Telecom Chair of Communication Networks, Technical University of Dresden (TU Dresden), 01062 Dresden, Germany

<sup>3</sup>BMW Group, 80788 Munich, Germany

<sup>4</sup>Wireless Communications and Navigation, Rheinland-Pfälzische Technische Universität Kaiserslautern-Landau, 67663 Kaiserslautern, Germany

<sup>5</sup>Fraunhofer Heinrich Hertz Institute, 10587 Berlin, Germany

<sup>6</sup>Vodafone Chair Mobile Communications Systems, Technical University of Dresden (TU Dresden), 01062 Dresden, Germany

<sup>7</sup>Network Information Theory Group, Technische Universität Berlin, 10623 Berlin, Germany

Corresponding author: Alexandros Palaios (alex.palaios@ericsson.com)

This work was supported by the Federal Ministry of Education and Research (BMBF) of the Federal Republic of Germany as part of the A14Mobile Project under Grant 16KIS1170K.

**ABSTRACT** As cellular networks evolve towards the 6th generation, machine learning is seen as a key enabling technology to improve the capabilities of the network. Machine learning provides a methodology for predictive systems, which, in turn, can make networks become proactive. This proactive behavior of the network can be leveraged to sustain, for example, a specific quality of service requirement. With predictive quality of service, a wide variety of new use cases, both safety- and entertainment-related, are emerging, especially in the automotive sector. Therefore, in this work, we consider maximum throughput prediction enhancing, for example, streaming or highdefinition mapping applications. We discuss the entire machine learning workflow highlighting less regarded aspects such as the detailed sampling procedures, the in-depth analysis of the dataset characteristics, the effects of splits in the provided results, and the data availability. Reliable machine learning models need to face a lot of challenges during their lifecycle. We highlight how confidence can be built on machine learning technologies by better understanding the underlying characteristics of the collected data. We discuss feature engineering and the effects of different splits for the training processes, showcasing that random splits might overestimate performance by more than twofold. Moreover, we investigate diverse sets of input features, where network information proved to be most effective, cutting the error by half. Part of our contribution is the validation of multiple machine learning models within diverse scenarios. We also use explainable AI to show that machine learning can learn underlying principles of wireless networks without being explicitly programmed. Our data is collected from a deployed network that was under full control of the measurement team and covered different vehicular scenarios and radio environments.

**INDEX TERMS** Intelligent transportation systems, machine learning, quality of service, throughput prediction, vehicular communication, wireless networks.

The associate editor coordinating the review of this manuscript and approving it for publication was Chan Hwang See.

## I. INTRODUCTION

Machine learning (ML) has a strong potential to overcome challenges arising in vehicular communication and

networking, as presented in [1] and [2]. Examples of such challenges are resource allocation [3], [4] and quality of service (QoS) prediction [5]. The QoS prediction, in turn, is an enabler for a variety of use cases in the domain of vehicular communication, such as autonomous driving, platooning, cooperative maneuvering, tele-operated driving, and smart navigation. Some of these are presented in [6] and [7]. Predictions can not only enhance the existing use cases and support new ones, but also be a stepping stone for a network to become proactive. Such predictions can be an integral part of a network that proactively reacts to sustain, for example, a specific key performance indicator (KPI) target, like the maximum achievable throughput under high network load.

In this work, we focus on the prediction of QoS as a way of looking into the typical workflow for building and applying machine learning (ML) algorithms. We start with data acquisition, moving on to data analysis and statistical characterization of the data, proceeding to feature engineering and different split strategies, and concluding by testing the performance of the model. Predicting QoS is a relatively complex task since it depends on several time-varying factors. It is especially complex in the case of vehicular communication where the radio conditions might change very drastically in a short amount of time [8].

Capturing datasets that reveal the complex interdependencies between network operations, radio environment and terminal behavior is a challenging task in itself. The captured data might often lack the required quality for ML-based prediction, rendering it inappropriate. There are multiple factors that can reduce the quality of a collected dataset for training ML models, with a few examples being dataset imbalance, coverage of radio regions with bounded dynamics or limited network states gathered. Moreover, acquiring different types of data for a prediction task comes with varying acquisition costs that are typically not discussed extensively. From the theoretical perspective of statistical learning, a training dataset should often be independent and identically distributed (i.i.d.) and drawn from the same probability distribution as the test dataset [9], [10]. Radio terminals frequently experience drastic shifts in the data distribution [11] that usually degrade the performance of ML models [12].

Meaningful comparisons between prior art proposals, though desirable, can be hard. There are multiple reasons for this, for example, the pre-processing steps are not well reported, or the used performance metrics differ. Additionally, literature in the area of applying ML for wireless communications often reports the performance of the models without giving an in-depth analysis of the dataset's underlying characteristics. This raises the question of how robust such models are to changes in different environments, where the dataset's statistical properties might be different.

Our results emerge from the data analysis of a dedicated measurement campaign with the goal of capturing several characteristics of the radio environment and enabling an

in-depth discussion of ML workflows. The code for this analysis is based on customary functions from usual Python libraries for scientific computing (e.g., numpy, pandas, scikit-learn...). Our contributions include the following:

- We study the prediction of maximum throughput prediction enhancing, for example, streaming or high-definition mapping applications.
- We discuss the entire machine learning workflow highlighting less regarded, but important aspects, such as the detailed sampling procedures, the in-depth analysis of the dataset characteristics, and the effects of splits on performance. For example, we test against the data stationarity assumption and discuss methods to handle those for commercial networks. Many proposals in the literature avoid discussing such theoretical violations that many ML models require.
- We put the characteristics of wireless environments in the foreground explaining what might hinder and benefit the adoption of ML approaches for new-generation wireless networks.
- We further showcase why the lack of this type of information, missing often from the literature, can generate overconfidence in proposed ML solutions.
- We investigate the effects of diverse sets of input features specifically those that are not typically available in the prior art.
- We showcase that consumer devices and low sampling speeds can be enablers for the adoption of ML in commercial networks.
- We also use explainable AI to demonstrate that machine learning can learn underlying principles of wireless networks without being explicitly programmed.
- Our data is collected from a deployed network that avoids limitations of theoretical studies or simulation-based results.

The remaining of the paper is structured as follows: Section II presents related work in the area of predictive quality of service (pQoS), and the measurement campaign is described in Section III. In Section IV, we discuss various radio environment properties in relation to ML workflows, such as vehicle speed and stationarity. We continue the analysis on ML workflows in Section V with a careful look into train/test data splitting and feature engineering. In Section VI, we apply a set of ML algorithms to predict the maximum achievable throughput for diverse prediction models, communication directions, feature sets, splitting strategies, and prediction horizons. In Section VII we look at the topic of explainable AI (XAI) [13] showcasing what the models have learned in the absence of explicitly programmed rule sets. Finally, we draw conclusions in Section VIII.

## II. RELATED WORK

The application of ML for QoS prediction in vehicular networks has drawn attention recently. In the following, we present an overview of related works. Measurement campaigns that focus on the feasibility of teleoperated driving

(ToD) are described in [14] and [15]. The authors of [14] also include a sensitivity analysis in their work and examine the impact of factors such as speed and distance to a serving cell. When applying gathered data from measurements to QoS prediction in the current literature, a sampling interval in the order of seconds is often used (e.g., [16], [17]). This does not reflect the fast-changing wireless channel associated with high-mobility vehicular communication scenarios [18], and there is an open question as to which degree sampling schemes need to be adjusted based on the vehicle speed. Some works already have compared measurement applications running on slow-sampling consumer equipment (CE) with dedicated measurement equipment (DME) analysis [19], [20] and propagation model tuning [21]. As vehicular communication is characterized by high mobility of the terminal, it is an open question whether CE's low-cost transceiver chains might render collected data unsuitable for applying ML.

Moreover, it is known that the radio environment has dynamics that can drastically change its statistical properties over a distance of a few meters [22], [23]. Such changes in statistical properties of the data stream, where the ML model needs to reformulate its decision boundaries, are being discussed in the literature as concept drift [24], [25], [26]. Detection and analysis of concept drift using state-of-the-art algorithms have been explained in [27]. Extensive research is conducted on the management of concept drift, some of which are: using an ensemble learning approach [28], concept drift-aware federated averaging [29], or model updating mechanisms [30].

In [31] and [32] the authors present a measurement campaign and evaluate several ML models regarding their performance for downlink (DL) throughput prediction. One limitation is that the authors measure in a public network and thus cannot consider features such as the cell load. The authors of [33] consider the movement of the user and build a two-staged ML model to predict the transmission control protocol (TCP) throughput. In the first step, the movement pattern is identified. Based on this, a long short-term memory (LSTM), trained with corresponding data, is selected for the throughput prediction. Some preliminary results for uplink (UL) and DL throughput prediction using Random Forest and employing the dataset used for this paper are presented in [34]. QoS prediction in a 5G non-standalone network is examined in [35]. Besides throughput, ML is also used to predict latency [5], [36], [37] and handovers [38]. Recently, several datasets intended for ML-based studies have been made publicly available [39], [40] that typically do not provide cell-related information. Moreover, those works focus on introducing the datasets instead of applying ML methods.

An inherent limitation of ML algorithms, especially as they get more complex, is that their decision-making is difficult to understand, rendering them into black boxes. Recently, the field of explainable AI (XAI) [13] has gained traction, to ensure ML models perform as desired, not only in test environments but also in target applications.



**FIGURE 1.** The map of the areas where the measurements took place. The A9 highway is highlighted in red, while the rural and side street areas are marked in green (not all rural streets are plotted to reduce clutter). The blue area is the suburban city of Feucht.

With XAI, models can be judged in the context of human domain knowledge, e.g. by investigating the learned feature importance.

We extend prior art by providing detailed insights into the aspects mentioned above. Our dedicated measurement campaign, on a fully controlled private network, allowed us to perform a deep analysis and enable an in-depth discussion. We could control which devices could connect to the network, we could adjust the total network interference, and generate diverse traffic dynamics. We consider different radio environments, device types, sampling frequencies, measurement scenarios, and prediction horizons. We present the complete workflow towards building a pQoS model, predicting the maximum achievable throughput. We discuss several design decisions, such as the train-test split, various feature groups, and different models, and show the influence on the achieved performance. In the end, we discuss the issue of trustworthiness.

### III. MEASUREMENT CAMPAIGN

Within the project AI4Mobile, an extensive measurement campaign was performed in the 5G-ConnectedMobility test field.<sup>1</sup> In total, more than 3000 km were driven by four

<sup>1</sup><https://www.ericsson.com/en/cases/2020/5g-connectedmobility>

vehicles for one week to perform measurements from the long term evolution (LTE) network. A detailed description of the measurement campaign is available in [41]. Since we had control over the network infrastructure, we were able to measure and collect data from all parts of the network. We were also able to define custom scenarios and set parameters accordingly. Subsequently, we briefly describe the setup and highlight some scenarios.

### A. MEASUREMENT SETUP

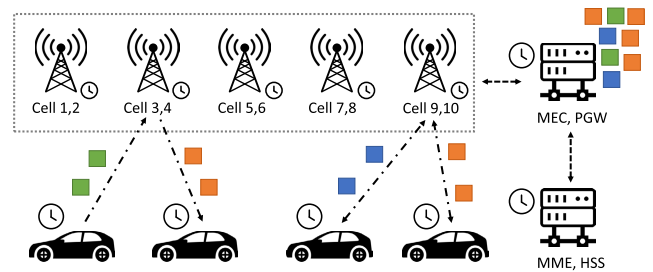
During the measurement campaign, four vehicles were equipped with several user equipments (UEs). In each vehicle, an identical DME was used for measurements with high granularity. The DME is a Linux-based PC including a cellular modem connected to a  $2 \times 2$  MIMO car antenna, which is placed on the roof of the respective car. Additionally, a global positioning system (GPS) receiver was placed on the roof and connected to the DME, providing accurate location and time synchronization via a pulse per second (PPS) signal. To capture the data from the cellular modem, we used the application MobileInsight [42].

Each vehicle was also equipped with several CE phones as pure data generators and measurement devices with a low sampling frequency for a comparison to the DME's measurements. The CE uses the application G-Net Track Pro [43] to capture radio measurements via the Android API. All devices used the application Iperf [44] to exchange data with a local server connected to the network. In addition to the UEs (CE and DME), we also collected data from the base stations and the core network, which were also time-synchronized. An overview of all captured data is given in Table 1.

The network is deployed in the south of Germany, close to Nuremberg. It consists of 10 frequency-division duplex (FDD)-LTE cells with 10 MHz of bandwidth at a carrier frequency of 700 MHz, with typical radio tower panel antennas (horizontal half-power beam width of  $65^\circ$ ). Since it is a test network, our devices were the only devices connected to the network. In Fig. 1, the area of the test field is shown, including the position of the distinct cells. While the cells were positioned to provide excellent connectivity along the highway and in the suburban city of Feucht, part of the rural area next to the highway was also covered, which allowed us to capture several radio environments.

### B. THE CONDUCTED STUDIES

The purpose of the measurements was to thoroughly examine numerous parameters through a variety of real-time scenarios in our captured dataset. This involved analyzing different driving patterns at different speeds, stationary measurements, recording different radio environments, and generating varied data traffic to create contrasting load scenarios. Furthermore, we utilized various protocols and generated traffic in both the UL and DL to the server and to other vehicles. Additional information on the parameters captured in the measurement campaign can be found in [41].



**FIGURE 2.** Simplified view of the used measurement methodology. All clocks are synchronized and all generated data is marked accordingly and captured at separate entities. Figure adapted from [41].

### C. DATA PREPROCESSING

For this study, only user datagram protocol (UDP) measurements from the dataset were considered since the protocol properties of TCP and UDP differ. Therefore the E2E throughput could not be limited by the transport layer. We used the data, where the devices were requesting maximum throughput.

Moreover, we filtered out stationary measurements, i.e., when vehicles were parked for some time. The stationary measurements contain similar radio properties for the majority of the samples, which could result in overly optimistic prediction results.

## IV. PROPERTIES OF THE RADIO ENVIRONMENT

In this section, we study properties of the radio environment that might influence ML algorithms. Radio environments have very diverse characteristics that depend on multiple factors. Some of these factors are environment-based, for example, radio propagation can be drastically different in a rural area and in a crowded city center [45], as the type and density of buildings influence the radio propagation. Another factor is, for example, the specific network deployment (e.g. antenna heights, frequency layers and the number of cells in an area). All these factors contribute to very unique properties in the collected data as UEs move. We first discuss the data collection procedures and how these might affect the quality of the collected dataset. We continue looking at the statistical properties of the radio environment. As many ML algorithms rely on multiple statistical assumptions [46], we try to highlight some of those, as they might need to be better understood by the research community that applies ML on wireless networks.

### A. SAMPLING THE RADIO ENVIRONMENT

High sampling intervals of the radio environment come at a cost, either in terms of acquiring more capable hardware, or in terms of higher power consumption or in terms of signaling, as more data need to be transferred. We specifically study the effects of different sampling intervals on the characteristics of the collected dataset. We also study the effects of the vehicular speed in the characteristics of the collected dataset to see if higher velocities call for faster sampling intervals, for sustaining specific quality characteristics in the collected



**TABLE 1. Overview of captured data from all network nodes.**

Network Node	Information Type	Sampling Interval	Features (Examples)	QoS Measurements
DME	Physical layer (PHY)	10 ms	RSRP, RSRQ, CQI	DL throughput & latency
	Medium access control (MAC)	40 ms	Buffer status report (BSR)	
	Packet data convergence protocol (PDCP)	0.5 s	PDCP UL & DL statistics	
	Radio link control (RLC)	45 ms	RLC transmissions & receptions	
	Radio resource control (RRC)	Event-based	OTA packet log	
	Built-in sensors (e.g., GPS)	1 s	Longitude, latitude, speed	
CE	Packet capturing tool (Tcpdump)	Traffic-dependent		
CE	Physical layer (PHY)	1 s	RSRP, RSRQ	DL throughput
	Measurement tool (Iperf)	1 s		
Base station	Standard operator averages	15 min	Interference, cell load	
	Reconstructed real-time information	1 ms	Interference, cell load	
Core	Mobility management entity (MME) Traces	Event-based	Handover	UL throughput & latency
	Packet capturing tool (Tcpdump)	Traffic-dependent		
Database	Weather	1 hour	Precipitation	
	Traffic APIs	5 min	Traffic flow	

dataset. We focus on LTE signal values for the analysis, i.e., reference signal received power (RSRP) and received signal strength indicator (RSSI), since these are most affected by the vehicle's speed and influence the target throughput, in turn, as we show in Section VI and in Section VII

To quantify the sensitivity of the collected measurements on these two parameters, we need to introduce the resampling error (RE) at the  $k$ -th time instance as

$$RE(k) = x_o(k) - x_{rs}(k), \quad 0 \leq k < K, \quad (1)$$

where  $x_o(k)$  represents the original signal of length  $K$  at a sampling interval of 10 ms for the DME and 1 s for the CE. The corresponding resampled value for a lower sampling interval is  $x_{rs}(k)$ . To determine  $x_{rs}(k)$ , the following procedure is applied: first, the original dataset is downsampled by averaging to a target rate to obtain  $x_{ds}(k')$ :

$$x_{ds}(k') = \frac{1}{M} \sum_{m=0}^{M-1} x_o(m + k'M), \quad 0 \leq k' < K/M, \quad (2)$$

where  $M$  accounts for the downsampling ratio for a certain pair of original and target sampling interval, e.g.,  $M = 10$  for a target rate of 10 s when considering the CE's original sampling interval of 1 s. Secondly, the downsampled dataset are then upsampled by the forward filling method, which can be expressed by  $M$ -integer division on the original index  $k$ :

$$x_{rs}(k) = x_{ds}(\lfloor k/M \rfloor) = x_{ds}\left(\frac{k - (k \bmod M)}{M}\right). \quad (3)$$

Finally, Equations (1), (2) and (3) can be combined into

$$RE(k) = \frac{1}{M} \sum_{m=0}^{M-1} x_o(k) - x_o(m + k - (k \bmod M)). \quad (4)$$

For each sample at the  $k$ -th time instance, the vehicle speed  $v(k)$  is also known, which allows to explore the relation between the mobility of the terminal and the resulting RE. Therefore, the speed range from 0 to 140 km/h is split in intervals of 5 km/h. For each interval with a lower speed boundary  $v'$ , the mean absolute resampling error  $\overline{RE}_{v'}$  is

calculated according to Eq. (5).

$$\overline{RE}_{v'} = \frac{1}{K} \sum_{v' \leq v(k) < v'+5 \text{ km/h}} |RE(k)| \quad (5)$$


The results are shown for RSRP and RSSI and measurements from the DME and CE in Fig. 3. The shaded areas indicate the central half of the distribution of the errors.

Focusing on the effect of different sampling intervals for a DME dataset (Figs. 3a and 3b), one can see that the RE increases with a larger sampling interval as expected. Interestingly, this increase seems to be below 2 dB for a sampling interval of 1 s and between 2 and 3 dB for a sampling interval of 10 s, which might still be acceptable for certain ML applications.

The mobility of the terminal seems to have some effect on the quality of the collected dataset. When the terminal is stationary, the RE is close to 0. Both RSRP and RSSI are characterized by a steep rise of the error once the terminal starts to move at an approximate speed of 10 km/h. Above that speed, the error remains stable. Of particular interest is a sampling interval of 1 s, as it shows approximately the error introduced by employing a CE instead of a DME, which collects data at a sampling interval of 10 ms. The mean RE for RSRP and RSSI are well below 2 dB, which suggests that also CE samples may enable an accurate QoS prediction. Only the sampling interval of 10 s shows larger variations across different speeds. Here we notice that the variance of the error is drastically increased for the lower sampling interval.

We also executed the same experiments with a dataset from the less granular CE. We received similar results for the DME dataset as shown in Fig. 3c and 3d, i.e., a rise of the error in the beginning, followed by a saturation of the RE. We empathize that a direct comparison of the resulting RE between the two different devices should be avoided, as the original sampling intervals are larger and thus the number of samples that are considered for each downsampling interval is smaller. This, combined with the simpler transceiver chain of the CE, is a potential explanation for why these results are less smooth.

The results indicate that the vehicle speed does not have a noticeable impact on the absolute error for different sampling



**FIGURE 3.** Relation between speed and mean absolute error (MAE) for different sampling intervals (shaded areas indicate the interquartile range).

intervals, except for vehicles moving at very low speeds. The observed effects can possibly be attributed to the Doppler effect. However, no channel measurements were carried out during the measurement campaign and thus the causes cannot be conclusively assessed on the basis of our data set. The relatively constant error for slow sampling at different speeds can further simplify the sampling procedures without the need to introduce adaptive sampling schemes. Also, lower sampling speeds can still provide useful input to ML models as the reconstruction error remains low. This type of result is necessary for optimizing data collection procedures for future ML applications. In the following sections, we use downsampled data from the DME with a sampling interval of 1s that also enables a comparison with the CE. The resulting dataset sizes for DL and UL are 18846 and 19535 samples, respectively.

## B. DATA STATIONARITY

Many theoretical results in ML literature are based on the assumption that the available dataset is i.i.d. [9], [10], [12]. This assumption can be violated in several ways, for instance, if i) samples are correlated in time (i.e., they are not independent), or ii) the underlying time series is non-stationary (i.e., not identically distributed).

There are both parametric and non-parametric hypothesis tests available in the literature to study time stationarity. Parametric tests are often more powerful, but one needs to assure that the assumptions they require (i.e., a specific well-known underlying distribution) are met by the tested data. As such, we have opted to use the parametric augmented Dickey–Fuller (ADF) test [47], which searches a unit root in the tested time series as its null hypothesis, while the alternative hypothesis is stationarity. In other words, a low  $p$ -value (for our analysis we have considered  $p < 0.05$ ) indicates strong evidence for stationarity in the data. As a parametric test, ADF requires a suitable choice of the maximum lag of the assumed autoregressive model. Since the lag order is unknown for our data, we estimate it according to the Akaike information criterion (AIC) [47].

We have run the ADF test for all physical layer (PHY) features from all measurements (c.f. Table 1, first row) in an accumulated manner. That is, for every analyzed time series of size  $T$ , we construct  $T$  sub-series each one consisting of all samples from 0 to  $t \in \{0, 1, \dots, T - 1\}$ . We then run the

ADF test against every sub-series and concatenate the output  $p$ -values into a vector of size  $T$ . Fig. 4 shows the resulting accumulated ADF test as a color code on top of some time series.

The expected behavior is that the accumulated ADF test tends to  $p \ll 0.05$  as the size of the evaluated time sub-series increases. The number of samples needed for this depends on some statistics, such as the variance (c.f. Fig. 4a & 4b), but it rarely exceeds 15 minutes for our data.

This time stationarity assumption is violated for measurements with an appreciable scenario shift, e.g., those where vehicles remain idle for some time and drive away afterward (Fig. 4c). This shows that stationarity cannot always be assumed, even though it might hold true for some datasets under a long-enough measurement time. For the measurement scenarios, we opted for a duration of 40 to 60 minutes while driving to alleviate the problem of non-stationarities.

In summary, we acknowledge the presence of non-stationary portions in our data and its possible negative impact on the performance of ML algorithms.

## C. RADIO ENVIRONMENT CORRELATIONS

Additionally, we look at an important characteristic of the radio environment, which is the correlation in time and space [48], [49]. In Fig. 5, we show an example of the autocorrelation function of the signal-to-interference-plus-noise ratio (SINR) for three distinct radio environments. We see that the radio environment tends to be highly correlated in time for all three distinct radio environments. Some higher correlations are noted for the rural environment, with the highway having the weakest of the three. This is probably explained by the faster movement of the vehicles.

In [50], the authors discuss that if correlations are not handled properly, there is a higher risk of reporting overly optimistic results. The reason is that the ML model does not learn the underlying relations between input and output clearly, and instead learns the existing dependencies of a dataset, which in our case, are the reported correlations.

## V. MACHINE LEARNING FOR WIRELESS NETWORKS

In this section, we discuss important aspects that the ML engineer needs to understand before applying ML to radio data. We discuss the notion of concept drift, different




FIGURE 4. Stationarity analysis of the data using the ADF test.

strategies for splitting the data into training and test datasets, and feature engineering.

### A. CONCEPT DRIFT

In this section, we introduce the idea of concept drift as a tool that can mitigate some of the effects of non-stationarity. ML models traditionally need a formal and precise definition of the problem to represent the decision boundaries. However, non-stationarity in datasets implies that the decision boundaries, i.e., concepts, have a higher likelihood to be relearned to accommodate variations in the underlying data distributions [51], [52]. This re-learning or concept drifts must be detected to ensure that the ML models remain accurate. Given that a sample instance  $x$  belongs to a




FIGURE 5. Autocorrelation of the SINR for excerpts of different DL measurement runs.

### Algorithm 1 Page-Hinkley Test

**Input:** A labeled dataset  $x_1, x_2, \dots, x_T$   
 Magnitude threshold  $\delta$   
 Detection threshold  $\lambda$   
 $\bar{x}_T \leftarrow 1/T \sum_{t=1}^T x_t$   
 $U_T \leftarrow \sum_{t=1}^T (x_t - \bar{x}_T - \delta)$   
 $m_T \leftarrow \min_{t=1, \dots, T} (U_t)$   
 $PH_T \leftarrow U_T - m_T$   
**if**  $PH_T > \lambda$  **then**  
 |  $\Theta \leftarrow \text{True}$   
**else**  
 |  $\Theta \leftarrow \text{False}$   
**Output:** Drift alarm  $\Theta$

class  $\omega$  ( $x \in \omega$ ), the Bayes posterior probability is:

$$P(\omega | x) = \frac{P(\omega)P(x | \omega)}{P(x)} \quad (6)$$

This indicates that the drift may occur when:

- (i) Class definition  $P(\omega | x)$  changes while  $P(x)$  remains the same, i.e., real concept drift.
- (ii) Virtual drift: Input distribution  $P(x)$  changes while  $P(\omega | x)$  remains the same, i.e., virtual concept drift.
- (iii) Prior probability of class  $P(\omega)$  changes while  $P(x | \omega)$  remains the same.

However, sometimes the class boundaries may change due to hidden contexts. This drift occurs due to ‘insufficient, unknown, or observable features in a dataset’ [53]. We aim to test how often such hidden contexts appear in real-time scenarios such as the automotive domain, where the environments change dynamically at a rapid pace. To this end, we employ the Page-Hinkley (PH) test (see Algorithm 1) to detect statistical changes in the input data stream, i.e., detect drift in data [54], [55].

The PH estimator needs a labeled dataset, a magnitude threshold ( $\delta$ ), and a detection threshold ( $\lambda$ ) as inputs. The magnitude threshold defines the degree to which noise is

**TABLE 2.** Detected drifts in test data streams for a model trained on the suburban region.

Training region	Standard deviation	Test region	
		Highway	Suburban
Suburban	$2\sigma$	56	17

permitted in the dataset. For each subsequent timestep  $t$ , a single data entry is fed to the PH estimator and a cumulative error  $U_T$  is calculated. The estimator raises an alarm if the cumulative error  $U_T$  increases beyond the minimum cumulative error  $m_T$  determined by detection threshold  $\lambda$ .


Since Algorithm 1 only returns a boolean value, we build  $T$  subseries identical to the ADF test and run the PH test against all of them. In that way, we can provide an example of detecting drifts in the input dataset as the vehicle is moving through different environments. For that example, an underlying ML model has been trained on the dataset of the suburban region. The PH estimator’s magnitude threshold  $\delta$  is set to twice the standard deviation,  $2\sigma$  of the training dataset. The dataset captured by DME 3 from Vehicle 3 in the measurement campaign is used to test for drifts. The estimator reacts to the statistical changes in the input data and raises an alarm when a drift is detected (see Table 2). The number of detected drifts on the highway is higher by a factor of 3.3 than that of the suburban region.

To verify the time and position of these drifts, an example from the dataset, where the vehicle moves through different regions is taken (see Fig. 6). Two distinct regions, i.e., suburban and highway are chosen for training the PH estimator. Furthermore, the drifts are detected for different magnitude thresholds of  $1\sigma$  and  $2\sigma$  for the two training datasets. A change in the radio environment, when trained on the suburban dataset (see Fig. 6a), can be detected when the two different thresholds are applied. We also got similar results when we trained for highway environments and the UE was moving between suburban, rural, and highway environments (see Fig. 6b).


The different types of concept drifts are expected to deteriorate ML performance. Detecting these drifts and dealing with them as early as possible will result in a resilient and accurate ML model. One way to handle the drifts is to include all the possible scenarios in the training dataset that a ML model might face during its life cycle. As this might be hard, and even unrealistic in many cases, another approach is online training, which may be based on concept drift detection [56], [57]. Hernangómez et al. [58] provide some preliminary results on this dataset using said approach.

**B. TRAIN/TEST SPLIT AND EVALUATION OF ML ALGORITHMS**

This section provides more details on splitting strategies, and it explains how different techniques can provide different insights. If not used properly, the splitting technique might



**FIGURE 6.** Detecting data drifts in distinct radio environments.



**FIGURE 7.** Visualization of different split strategies.

give skewed results, and we hope to explain some of their fundamental strengths and weaknesses in the following.

The train/test splits included in this study are random split, split by time, split by measurement run, and split by fold. Fig. 7 provides the reader with a graphical explanation of the different splitting strategies. Subsequently, we describe the four strategies.



### 1) RANDOM SPLIT

Each dataset sample has a nonzero probability of being added to the training respectively test set. This often leads to two consecutive samples from the same device being added to the training and test set. These samples, as discussed above, share similarities, explained by their correlation structure. In the coming sections, we use a split where 70 % of the data are in the training dataset and 30 % in the testing dataset. Note that this split strategy assumes samples to be i.i.d., an assumption we have shown to be violated in the radio environment.

### 2) SPLIT BY TIME

The dataset is split into two same-size parts based on the time domain. The first part of each measurement run (independent of the device) is added to the training set, and the second to the test set. As one measurement run consisted of driving the highway segment in two directions, this represents adding one direction of travel to the training, and the second to the test set. Hence, the training and test sets contain very similar radio environments but are uncorrelated in time.

### 3) SPLIT BY MEASUREMENT RUN

A measurement run (independent of the device) is considered part of either the training or the test set. The assignment and balancing of training and test set can be challenging, as the train and test datasets need to contain similar characteristics. This splitting scheme is best used when the model's generalization performance is the main focus. It is difficult to capture all dynamics (i.e., the parameter permutations described in Section III) in both training and test set. Bad performance on the test set results from the fact that the complexity and dynamics of the region are not reflected by any subset of the measurement runs. In our dataset, we ended up with a split where about 70 % of the data are in the training dataset, and 30 % in the test dataset to ensure that training contains most of the dynamics found in the test dataset.

### 4) SPLIT BY FOLDS

The last proposed splitting method aims to combine the split by time and random splitting. We divide the time domain into ten subsets of equal length for each measurement run and device. This approach promises a realistic prediction performance evaluation, as it negates many of the problems of random splitting, and combines advantages of the time splitting, where the models learn characteristics from the time-varying radio environment. Here, we assigned 70 % of the data to the training and 30 % to the test dataset.

To investigate deeper the effects of the different splitting strategies, we refer to principal component analysis (PCA) [9], an orthogonal linear transformation transforming a dataset (e.g., the training or test set) to a new basis such that the dataset's greatest variance lies in the first principal components. Therefore, it can be considered as a dimensionality reduction method and considering only a

subset of the original components, e.g., the first two for visualization purposes. We note that PCA captures only the second-order statistics and that there are alternative methods for visualization, such as t-SNE. Fig. 8 shows the differences between the training and test sets for the four splits using PCA. The random split's train and test datasets have very similar distributions (cf. Fig. 8a). The reason is that the test set does not contain unseen propagation, data traffic, and driving scheme scenarios. Because time-series data from the radio environment seems highly correlated, the training and test datasets include many similar points.

Therefore, although this seems to be the splitting scheme most used in the literature, there is a risk that reported performance is higher than for a deployed system. This should not discourage the use of random sampling. However, the data engineer needs to understand the inherent limitations of this splitting technique and be careful when reporting results.

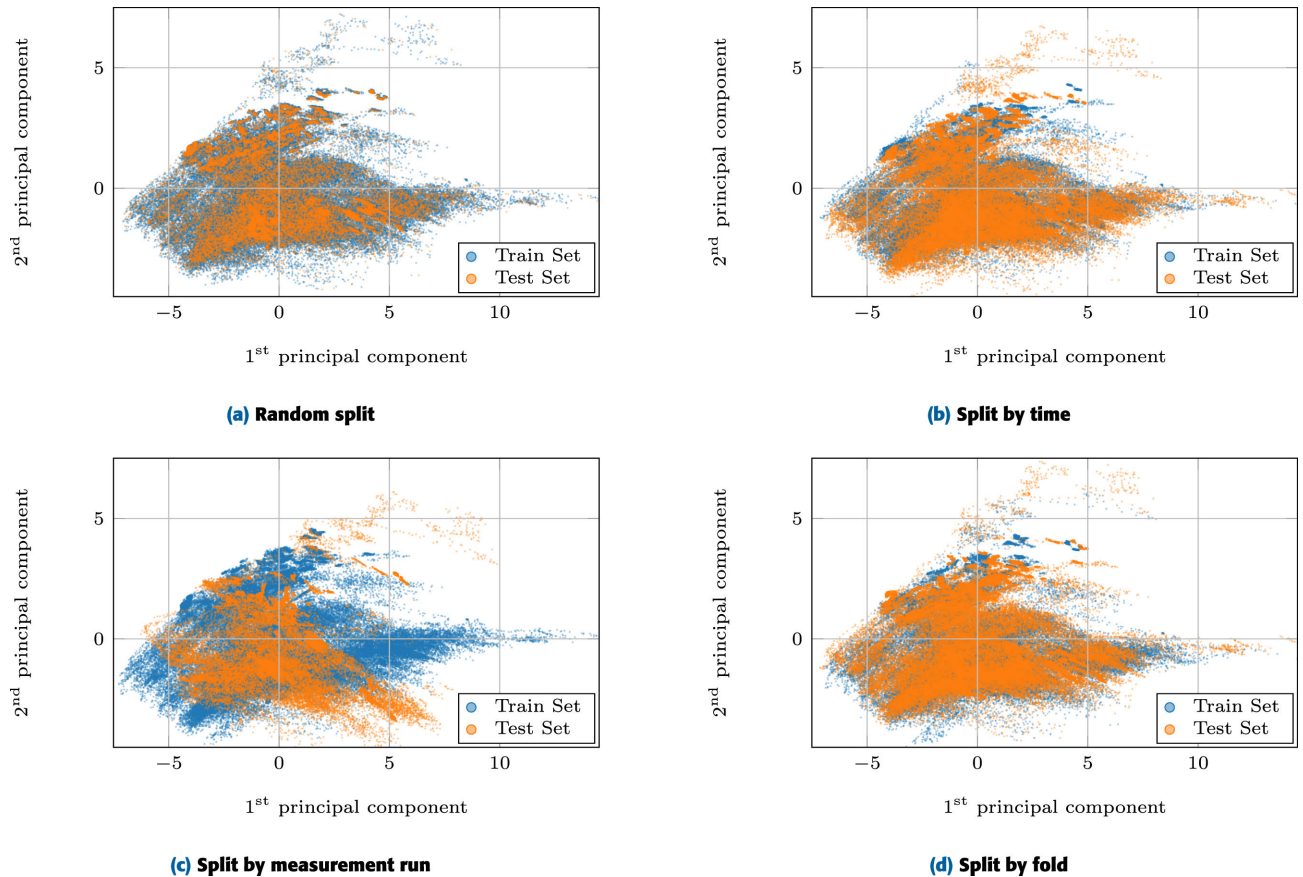
The split by time alleviates to some extent the sample correlation, as consecutive samples belong to the same set (training or test). Moreover, the varying propagation environment, due to the changing direction of travel, and differences in speed, result in distributions that are less similar than for the random split (cf. Fig. 8b). Some regions, e.g., around the value five for the second principal component, are only present in the test set. The PCA can be considered a first test on how detailed collection procedures should be and highlights how time-varying components influence differences between the training and test sets. The split by time adds some limitations on the measurement duration as for balanced train and test sets more data than, e.g., for random split is required.

The split by measurement run is more appropriate to test model performance under varying and unseen parameters in the measurement procedure, the devices' behavior, and the network. In Fig. 8c, it becomes clear that all these parameters have a substantial effect on the two datasets. Perhaps this is a good way to study the generalization performance of ML models, as a training dataset will never comprise all potential parameter permutations. The ML model makes predictions on samples that can be considered different, e.g., using our PCA evaluation.

We consider the split by folds the best splitting strategy for our dataset as it provides enough variation for ML models to be tested on unseen data while alleviating some of the disadvantages of the other splitting strategies.

## C. FEATURE ACQUISITION AND AVAILABILITY ANALYSIS

We group different types of features based on the type and the measurement capability of a device into distinct feature groups, see Table 3. The table defines a name for the feature group, a corresponding abbreviation that is used throughout the rest of the paper, and the parameters (features) that the feature group inherits. The feature group PHY includes features describing the radio environment, while channel conditions (CHAN) contains related features



**FIGURE 8.** Visualization of the first two principal components for different train/test splits. The distribution of the test and train dataset varies a lot depending on the split strategy.

describing the channel. Feature group base station data (BS) contains parameters that are aggregated per cell.

Vehicle information such as position, speed, and distance to the serving cell is included in the feature group vehicle information (VEH). Finally, the feature group radio environment map (REM) contains statistical information of PHY features and throughput in form of a radio environment map.

Since the features are captured from several different network nodes, as previously shown in Table 1, it is not realistic to assume that all features are readily available at a specific entity such as a vehicle or one of the network nodes. Data acquisition is associated with a cost function that is not constant for different sets of features. Also, confidentiality and privacy concerns could hinder the accessibility to specific features. This should be taken into account for ML models for wireless networks, instead of assuming full data availability, at least without discussing their acquisition costs.

Subsequently, we further categorized the previously defined feature groups into different access scenarios. We define these according to possible scenarios in which different sets of feature groups can be collected from exemplary entities. The defined access scenarios can be found in Table 4. The table defines a name for each access scenario,

a corresponding abbreviation that is used throughout the rest of the paper, the feature groups which can be accessed, and an exemplary entity corresponding to the access scenario.

The first access scenario is modem access (MD) and refers to features available at the modem, such as parameters from feature group PHY, which can often be easily accessed, e.g., via an API. extended modem access (EMD) refers to access to more modem data provided, e.g., by specialized software. Both MD and EMD correspond to data available at a UE, while modem and network access (MDNET), extended modem and network access (EMDNET), and modem, statistics and network access (REMNET) are available for a network operator with varying acquisition costs (in terms of computing and signaling). For example, PHY features are reported by the UE to the network during measurement reports, but these are usually sent after specific events are triggered. On the other hand, features from the CHAN group are reported regularly if the UE receives DL traffic. Since the dataset is based on vehicular measurements, we also introduce access scenarios vehicle and network access (DEVNET) and full device access (DEV) which include VEH features in addition to other feature groups. For comparison, we also include a full access (FULL) access scenario, which has access to all feature groups.

**TABLE 3. Overview of feature groups and their parameters.**

Feature Group	Abbreviation	Parameters
Physical layer	PHY	Mean RSRP, RSRQ, RSSI, SINR, RSRP margin, RSRQ margin, time since last handover
Channel conditions	CHAN	Channel matrix dominant rank index, mean CQI
Base station data	BS	Instantaneous DL/UL cell throughput (with one second resolution), number of active connected devices
Vehicle information	VEH	Position (latitude and longitude), speed, distance to serving cell (engineered from physical cell ID)
Radio environment map	REM	Radio environment map of PHY features & instantaneous device DL/UL throughput

**TABLE 4. Overview of network and device access scenarios and corresponding feature groups.**

Access Scenario	Abbreviation	Access to Feature Groups	Exemplary Entity
Modem access	MD	PHY	API of operating system (e.g., Android)
Extended modem access	EMD	PHY, CHAN	Specialist software (e.g., MobileInsight [59])
Modem and network access	MDNET	PHY, BS	Network operator
Modem access and statistics	MDREM	PHY, REM	Mapping and navigation service
Extended modem and network access	EMDNET	PHY, CHAN, BS	Network operator
Modem, statistics and network access	REMNET	PHY, BS, REM	Network operator
Vehicle and network access	DEVNET	PHY, BS, VEH	Car manufacturer (with MNO cooperation)
Full device access	DEV	PHY, CHAN, VEH	Car manufacturer
Full access	FULL	PHY, CHAN, BS, VEH, REM	-

The aim of this paper is not to define specific acquisition costs associated with the different access scenarios, but to acknowledge the potential trade-off between data acquisition costs and ML performance and compare the performance in the next section. Additionally, we would like to highlight two aspects in particular: First, there are data acquisition costs, which vary depending on the different feature groups. Second, ML models need to be tested under different assumptions on data availability including concerns of data governance and ownership.

**D. ML MODELS FOR QoS PREDICTION**

We briefly introduce different ML models, that are relevant to the prediction task. These include linear regression, ensemble methods, and neural networks. In terms of completeness, we also add a statistical model that serves as a baseline. The model hyperparameters were tuned using [60] with 5-fold cross-validation. The split by folds was used to split the dataset into a train and validation set. We used the mean absolute error (MAE) as the cost function.

**1) REM OF THE THROUGHPUT**

After building a REM of the DL/UL throughput, the predicted value is the interpolated throughput value for the current UE position. This method serves as a statistical baseline.

**2) LINEAR REGRESSION (LR)**

Linear regression is a relatively simple statistical model where a dependent variable is explained by multiple independent variables multiplied by coefficients defining their

weights. In a ML task, the weights of these coefficients are learned.

**3) RANDOM FOREST (RF)**

A random forest is an ensemble method based on decision trees. The trees are built independently and in a regression task, the mean of the individual trees is the prediction result. Our random forest (RF) is composed of 793 parallel trees with a maximum depth of 19.

**4) GRADIENT-BOOSTED DECISION TREE (GB)**

Gradient boosting is also an ensemble method based on decision trees. Here, the trees are built sequentially, minimizing the loss function with each new tree. The weighted mean according to the trees’ prediction performance is the prediction result. Our gradient-boosted decision tree (GB) ensemble is composed of 715 sequential trees with a maximum depth of 10.

**5) MULTILAYER PERCEPTRON (MLP)**

An multilayer perceptron (MLP) is a feedforward neural network that stacks together multiple layers of neurons in a sequential architecture. It is a deep learning (DL) model that has been successfully applied to regression problems with labeled tabular datasets. Our MLP consists of four hidden layers 256, 128, 64, and 32 neurons each and rectified linear unit (ReLU) activation functions. It is trained with an Adam optimizer, with a learning rate of 0.001, and a batch size of 16. The parameters of the neural network are initialized randomly. For regularization, we use early stopping with the

criterion that training ends when the error does not decrease for eight epochs on a validation dataset.

## VI. THROUGHPUT PREDICTION

In this section, we discuss the throughput prediction task based on ML. The prediction task we picked is to estimate the achieved maximum throughput under a high network load. The high network load conditions mean that the end users have to share the network resources.

We start by investigating the performance of different ML models and continue by looking at the performance when different data access scenarios apply. We then compare the influence of the split strategies on the training and test datasets. We continue looking at the effects of sampling on the prediction task. Subsequently, we give an example of how much concept drift can deteriorate performance. We conclude by showing that CE can play an important role in ML-related tasks.

For all the provided results we report multiple metrics [61], as this allows a more flexible evaluation of the prediction performance for diverse use cases and when applicable we discuss the limitations of specific metrics.

### A. MODEL PERFORMANCE

In Tables 5 and 6, we compare different models on standard regression metrics for the DL and UL direction, respectively. We assume to use all the information available for the prediction (i.e., FULL), consider all environments jointly, and refer to split by folds. Moreover, we relied on all available data samples for performance evaluation.

In each table, we highlight the best-performing model per metric in bold. For the DL direction (Table 5) the GB outperforms the other models for all metrics. RF consistently shows slightly lower performance than GB. For the mean absolute percentage error (MAPE), the MLP is on par with GB, but it shows worse performance than GB and RF for most other metrics. Considering the two simpler models, linear regression (LR) outperforms the REM interpolation for most metrics, underlining that the use of context information, which is included in the input features, is crucial for the prediction task. Moreover, as the feature space is not linear, the LR performance is poor.

The UL direction (Table 6) mainly corresponds to the performance observations in the DL prediction task. We observe that the absolute gap between GB and RF narrows for MAE, median absolute error (MedAE), and root mean square error (RMSE) (due to lower UL data rates), with RF showing similar performance as GB for the  $R^2$  score. The UL throughput range (up to 23 Mbps) is smaller than the DL throughput range (up to 70 Mbps), which is reflected in the prediction performance. While the average MAE in UL is significantly smaller than the average MAE in DL, the MAPE of the DL is comparably lower than in UL.

As there is a trade-off between maximum performance and acquisition cost, we found that performance metrics, such as MAE, converges at around 5000 training samples.

**TABLE 5. Performance comparison of different regression models in DL direction.**

Model	MAE	MAPE	MedAE	RMSE	$R^2$
RF	2.61	0.31	1.67	3.91	0.92
GB	2.46	0.29	1.48	3.77	0.93
LR	4.91	0.78	3.57	6.70	0.77
MLP	2.71	0.29	1.59	4.23	0.91
REM(y)	6.80	0.84	4.26	9.97	0.53

**TABLE 6. Performance comparison of different regression models in UL direction.**

Model	MAE	MAPE	MedAE	RMSE	$R^2$
RF	1.13	0.50	0.79	1.63	0.90
GB	1.08	0.46	0.73	1.58	0.90
LR	2.31	1.79	1.88	2.96	0.66
MLP	1.20	0.53	0.75	1.80	0.87
REM(y)	3.20	0.96	2.41	4.28	0.25

Adding further samples only leads to minor increases in prediction accuracy. However, we note that this result is specific to our dataset and may not generalize to other datasets because the minimum required dataset size depends on a multitude of factors. Such factors are the environment dynamics, the network dynamics, and the specific use case (i.e. required accuracy). Nevertheless, this result shows that the performance return degrades with an increasing number of samples and measurement campaigns should be planned carefully as there is space for optimizing the data acquisition expenditures.

In summary, GB proved to be the best-performing model considering both data traffic directions and all metrics. At the same time, the interested reader should note that there are a plethora of ML algorithms that can be applied and provide good performance. The ML engineer can always pick one based on criteria like the model's complexity and the degrees of explainability. As the ensemble methods based on decision trees offer a good degree of explainability and robustness against outliers, we pick GB for the following sections.

### B. ACCESS SCENARIOS & FEATURE GROUPS

We continue by looking at a different set of input features and how much these affect the achieved performance. In the literature, input features are very often considered quite alike in terms of acquisition costs. In a real network, there is a cost associated with collecting data from different nodes and end-users. We compare the different prediction performances for the previously defined access scenarios, from Section V, in Tables 7 and 8 for DL and UL direction, respectively.

In both DL and UL, the performance of different access scenarios is very similar, with the only difference that, in general, for DL the MAPE is significantly lower compared to UL, and the MAE, MedAE, and RMSE are higher for DL due to the larger value range. The worst-performing access scenario in both cases is MD, where only the feature group PHY is used. Prediction for EMD, modem access



**TABLE 7. Performance comparison of different access scenarios with access to different feature groups in DL using the split by folds.**

Access scenario	MAE	MAPE	MedAE	RMSE	R <sup>2</sup>
MD	6.21	0.62	3.86	9.27	0.56
EMD	5.54	0.56	3.43	8.27	0.66
MDNET	2.93	0.32	1.86	4.38	0.91
MDREM	5.81	0.53	3.46	8.85	0.62
EMDNET	2.56	0.30	1.57	3.91	0.92
REMNET	2.78	0.30	1.69	4.21	0.91
DEVNET	2.68	0.29	1.65	4.07	0.92
DEV	5.01	0.49	2.79	7.96	0.68
FULL	2.46	0.29	1.48	3.77	0.93

**TABLE 8. Performance comparison of different access scenarios with access to different feature groups in UL using the split by folds.**

Access scenario	MAE	MAPE	MedAE	RMSE	R <sup>2</sup>
MD	2.93	0.97	2.12	4.03	0.38
EMD	2.68	0.89	1.90	3.72	0.45
MDNET	1.18	0.65	0.80	1.74	0.88
MDREM	2.62	0.73	1.89	3.63	0.48
EMDNET	1.15	0.72	0.77	1.70	0.89
REMNET	1.12	0.50	0.75	1.64	0.90
DEVNET	1.14	0.56	0.80	1.66	0.89
DEV	2.48	0.76	1.75	3.49	0.53
FULL	1.08	0.46	0.73	1.58	0.90

and statistics (MDREM), and DEV performs slightly better but not by a large margin. This probably indicates that additional channel or statistical/historical information might be beneficial but does not add much useful information in addition to physical layer parameters, which can be obtained easily. The prediction with only UE features (MD/EMD) achieves the worst performance across all presented metrics and access scenarios. However, adding BS features, demonstrated in access scenarios MDNET, EMDNET, REMNET and DEVNET, significantly improves the prediction performance, resulting in a halved prediction error for DL. The same holds true for UL direction except for the MAPE, which is still significantly lower by around 20 – 30 percentage points. Also, the R<sup>2</sup> score for both DL and UL is increased to approximately 0.9 once BS features are added.

The best prediction performance is achieved when access to all features is provided, which is emphasized by the bold rows for access scenario FULL. As mentioned in the previous section, we introduced this for comparison only, as access to all features is often not realistic or associated with high acquisition costs. However, we see that the performance of FULL is improved by only a small margin compared to the access scenarios including feature group BS - even compared to MDNET, which contains only simple feature group PHY in addition to feature group BS. Thus, MDNET might be a feasible candidate if one tries to find a good trade-off between better performance and lower data acquisition costs.

**TABLE 9. Performance of GB for different splits for DL.**

Split	MAE	MAPE	MedAE	RMSE	R <sup>2</sup>
random	1.58	0.18	0.88	2.62	0.97
time	3.66	0.57	2.70	5.10	0.86
meas	4.50	0.24	3.44	6.03	0.88
folds	2.46	0.29	1.48	3.77	0.93

### C. SPLIT STRATEGIES

As discussed in Section V there are different approaches for splitting the dataset into train and test sets.

We present results for using different splitting strategies for the DL direction in Table 9. According to all analyzed metrics, the random split, as described above, yields the best prediction results.

The performance would likely decrease drastically during deployment, compared to other splits, when the model is presented with new data with a different correlation structure.

In our case, the performance of the random split is 56 % better compared to the split by folds. However, the performance of the random split is likely overly optimistic due to the temporal correlations between samples [50]. Hence, ML models trained with the random split strategy might perform worse on new uncorrelated data that appears in real deployments.

As reasoned earlier, performance results obtained with the split by folds are more likely to be more robust to other deployment scenarios. The split by folds achieves the second-best results for most metrics.

On the other hand, the performance of the splits by time and split by measurement perform worse, which supports the hypothesis that there is rather a change in the statistics over time. This change can be attributed to the radio environment, as there are different numbers and types of vehicles around during the data collection process.

The results above highlight the need for picking a single or a multitude of splitting results for assessing the ML performance for a specific use case. Failure to have a clear splitting strategy might lead to a false estimation of the performance during the deployment phase.

In Fig. 9, we show the predicted values against the measured ones. The DL throughput was predicted with the GB model, utilizing the features contained in the FULL access type and using the split by folds.

In the figure, we also include two histograms for both axes. As multiple devices were competing for higher throughput, the total capacity of the network was typically split between multiple users making the highest throughput measurements a rarity. That well explains the skewed nature of the distribution. The interested reader should note that the mean type of metrics used in this section, MAE and MAPE, are heavily influenced by the long distribution tails. Therefore we have introduced median-based metrics, such as the MedAE, that are less influenced by a few large outliers.




FIGURE 9. Visualization of the predicted throughput over the given throughput.

TABLE 10. Performance of GB for different downsampling periods for DL.

Period	MAE	MAPE	MedAE	RMSE	R <sup>2</sup>
1s	2.46	0.29	1.48	3.77	0.93
2s	2.66	0.30	1.69	4.08	0.91
3s	2.93	0.42	1.83	4.48	0.89
4s	3.11	0.37	1.97	4.68	0.88
5s	3.58	0.40	2.02	5.71	0.84
10s	4.40	0.41	2.77	6.63	0.82

D. SAMPLING INTERVALS

As discussed in Sections IV and V, the radio environment has a correlation structure that renders consecutive measurements quite alike. That means that slower sampling intervals might be used to reduce overhead, such as battery consumption and signaling. To showcase that, we compare in Table 10 results for slower sampling intervals than the GB results in Table 5. We see that a 50 percent slower sampling interval (2 s) brings a negligible drop in the performance of an ML model. Even sampling procedures that are 75 percent slower than the initial sampling (4s) might be used and still, performance could be more than adequate for some use cases, with the biggest change being that the MAE increases by 20%. We also noticed that the reduction in the sampling intervals made PHY features less relevant and increased the importance of the REM features. The correlation structure of the radio environment can be further exploited to optimize sampling procedures while keeping the ML model’s performance within the requirements of a use case. A deeper understanding of the correlation structure of the radio environment can drastically benefit the sampling procedures [45] supporting more efficient ML workflows.

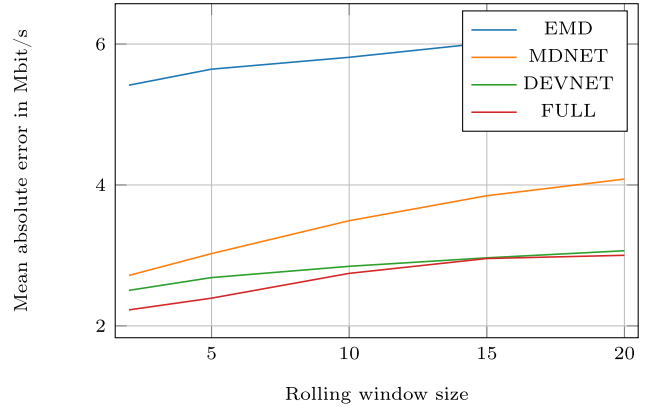


FIGURE 10. Rolling window response variable.

TABLE 11. Performance of GB model trained on suburban environment and tested on highway environment.

Direction	MAE	MAPE	MedAE	RMSE	R <sup>2</sup>
DL	9.68	0.47	5.02	15.18	0.05
UL	3.50	0.56	2.33	5.00	0.03

E. PREDICTION HORIZON

We continue by changing the prediction problem slightly. Instead of predicting the achievable instantaneous throughput, we predict the throughput in several seconds defined by the prediction horizon. This is shown in Fig. 10, for different access scenarios. We see that the access scenario EMD provides poor performance, similar to the previous subsections. The biggest improvement in the prediction error occurs when we include the MDNET access group, which includes the network information. The access scenario performance DEVNET provides marginal improvements, with the prediction performance dropping slower for longer prediction horizons. Best performance is achieved for access scenario FULL, although after a prediction horizon of 12 seconds, there is no real gain, of any extra-added features, as it performs as well as the DEVNET access scenario.

F. CONCEPT DRIFT AND ML

As was discussed in Section IV, there are concept shifts occurring frequently in the radio environment. Here, we provide some examples of the degradation of the performance for cases that ML faces concept drifts. Our example is shown in Table 11, where we trained a model on the suburban environment and afterward assumed that it was deployed on a vehicle traveling across a highway environment. We have seen already in Section IV that while driving from a suburban environment onto a highway, a large number of concept drifts is detected. The performance of the ML algorithm drops drastically, to the level of a statistical baseline as shown from the R<sup>2</sup>.

G. CONSUMER GRADE DEVICES

The presented results were calculated using data captured by DMEs. As DME is relatively expensive and comes

**TABLE 12. Performance comparison of different devices in UL and DL with MDNET features.**

Scenario	MAE	MAPE	MedAE	RMSE	R <sup>2</sup>
UL CE (Veh. 3&4)	1.13	0.38	0.74	1.65	0.82
UL DME (Veh. 3&4)	0.93	0.41	0.64	1.29	0.81
DL CE (Veh. 3&4)	3.37	0.34	2.06	5.17	0.84
DL DME (Veh. 3&4)	2.57	0.31	1.71	3.69	0.82

with increased processing and measurement set-up effort, we compare it in this section to data from CE. Since CE is significantly more accessible in price and set-up, the question arises to which degree the dataset differs from the DME dataset, and, more specifically, whether a CE dataset can be sufficient for the use case of throughput prediction.

For this comparison, we use data only from the vehicles that were equipped both with a CE and a DME (Vehicles 3 and 4) since this creates datasets of roughly the same size and diversity. Moreover, only feature groups PHY and BS were considered, (i.e., access scenario MDNET), to match the more limited feature set availability of the CE. Table 12 presents a comparison of prediction results between the CE and the corresponding DME. Both in UL and DL, the different metrics display similar performance.

The R<sup>2</sup> scores of the CE is slightly better than the DME's. The DME has higher sensitivity and we have seen larger number of outliers at the higher throughput ranges. Such outliers negatively influence the DME's performance. Overall, our findings show that CE can be part of a data collection procedure. More expensive devices might bring some small benefits but with diminishing returns.

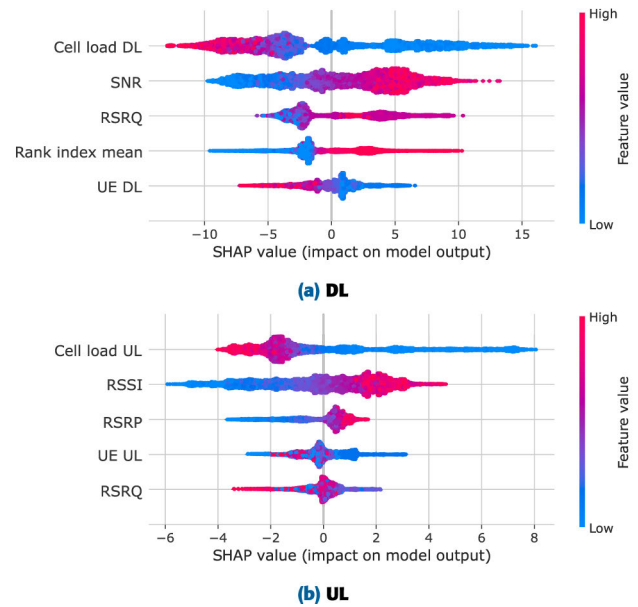
**VII. INTERPRETABILITY AND EXPLAINABILITY OF THE MODELS**

In this section, we use Shapley additive explanations (SHAP) [62] and accumulated local effects (ALE) [63] for explaining how the input features affect the prediction of the ML model.


**A. SHAPLEY ADDITIVE EXPLANATIONS**

SHAP is a framework to explain individual predictions. We use the model-agnostic variant of SHAP as described in [62]. Fig. 11 presents the SHAP values of the five most important features, in DL and UL respectively. The x-axis depicts the SHAP value that describes the impact of the input feature on the prediction. The y-axis depicts the input feature names, and the color represents the numeric input feature value.

The cell load feature has both in UL and DL a strong impact on the prediction. A higher cell load value leads to a lower value on the prediction and vice versa. On the other hand, the radio-based features (RSRP, reference signal received quality (RSRQ) and SINR) seem to have the opposite tendency. A higher radio value leads to a higher value on the prediction and vice versa.



**FIGURE 11. Feature importance in DL and UL of the five most important features in term of SHAP values.**



**FIGURE 12. ALE values for UL RSRP.**

**B. ACCUMULATED LOCAL EFFECTS**

We then continue using ALE for discovering the global effects of input features.

Fig. 12 depicts ALE values for RSRP in UL throughput prediction. The x-axis depicts the RSRP value, the y-axis the effect on the prediction, where a higher input feature value leads to a higher value of the prediction.

We note that in the ALE plot seems that there are four regions affecting the prediction differently as these are depicted in the figure. We added orange dotted lines at the perceived border of each region. The first one is with the lowest captured RSRP values. The UE needs to get at least some minimum RSRP before it can transmit data. The second region shows some linear characteristics. Higher RSRP values seem to contribute to a higher throughput significantly in this range, which demonstrates that the RSRP is contributing strongly to the prediction variable (UL throughput). The third region shows some slight saturation

in the linear trend, meaning that higher RSRP values do not contribute considerably more to a higher throughput prediction. The last region looks saturated. This probably means that either higher RSRP values do not provide any benefits for higher rates, or other factors like the cell load play a more important role on the prediction variable. We note that the ML model has discovered these regions from the input features, without being explicitly programmed, which is also close to the operating principles of the cellular networks.

The interested reader should note that the absolute numbers of the regions depend to a large extent on the characteristics of the network deployment that include for example the dynamic range of the receivers and the network states captured. Similar type of learnings have been shown in a previous paper [64], which was based on data from an operator in Asia. This increases the confidence that ML is able to learn about wireless environment characteristics on diverse networks and deployments.

## VIII. CONCLUSION

Based on a dedicated measurement campaign, we have presented insights into building reliable ML models for pQoS. Our results go beyond UE data by including network and vehicular information measurements, covering a large range of scenarios. Our measurements reveal many challenges ML models will face in real deployments.

Our first contribution discussed methods for improving sampling for the radio environment. Its correlation structure allows improved sampling procedures that reduce energy consumption and signaling for sharing collected data. The vehicle speeds do not seem to impact the statistics and characteristics of the collected data strongly, further simplifying the sampling procedures. We have also tested the data stationarity assumption, a precondition for many ML models and other theoretical approaches. Even though stationarity often holds true, especially for datasets captured over a longer duration, it should not be expected liberally. We have provided multiple examples where this assumption is violated, by focusing mostly on the aspect of concept drifts as vehicles move between radio environments. As concept drifts degrade ML performance, a reactive ML method might be needed for retraining the ML models as required, once drifts are detected. Another option is having larger datasets that cover multiple scenarios, like in our testbed, reducing the number of concept drifts that an ML model will face when deployed. A combination of the two methods might be the most feasible way forward.

We were able to predict the maximum achievable throughput under high load scenarios, where multiple users compete for maximum throughput in the network, with an MAE of 2.46 and 1.08 Mbps for down- and uplink. In particular, our findings show that the effects of data processing, different validation datasets, and sets of input features can have a very strong overall effect on the ML model performance. For that reason, simply comparing numbers from literature can be misleading. We have seen that gradient-boosted

decision tree (GB) models outperform neural networks while keeping a balance between complexity and explainability. At the same time, our results clearly show that low-cost consumer-grade devices can be part of ML processes. Their performance falls close to more expensive transceiver chains, further facilitating data collection procedures from end-user's terminals.

Moreover, we emphasized the topics of interpretability and explainability, showing that the tested ML models were able to capture the underlying principles without being explicitly programmed. It is interesting to note that the cell load was discovered as the most important feature for both communication directions, showcasing the importance of network features for such prediction tasks.

Our results indicate that more thorough testing of ML models is needed as complexities coming from the radio environment, the end users, and the effects of the network can considerably affect prediction performance, which might not always be precisely captured by any collected dataset. That can result in high performance variations of ML models, when deployed, that have never seen such effects in their training sets.

One can also draw several other conclusions and lessons from this data and its analysis, for the wireless research community and ML engineers alike. First, reporting only the ML performance of models on a specific dataset might provide overly optimistic results. Second, the data collection procedures, the handling, and processing of the data, as well as the way of reporting results, are all equally important in the long chain of ML workflows. Third, results coming from simplified analytical assumptions and simulations should be used with caution. Our results indicate that real-world-captured intricacies might hinder the further performance of ML models than hitherto understood.

Although some of these results could be regarded as intuitive, they have not been properly emphasized or taken into account in the literature, with only a few exceptions. Especially when it comes to applying ML in wireless networks, there has been a tendency to rate the usefulness of ML models based on their performance on some testing datasets. Our results clearly show that this can often be misleading, and more conservative estimates might need to be used. Moreover, the relevant acquisition costs of data are rarely discussed, and we believe that such aspects need to be part of future discussions and proposals. We hope that this publication serves as a starting point in that direction.

We believe that ML-based predictions do not only have the potential to improve specific use cases but also serve as an important enabler for a more proactive network. We made more datasets [65], [66] available to the research community that are based on similar collection principles as the ones we described. In the future, we plan to integrate the lessons learned towards methods that can innately handle some of the dynamics we noticed in the radio environment, such as non-stationarities and concept drifts. Moreover, we would like to extend our work toward other KPI metrics, such as



latency and the number of dropped packets. Finally, we would like to integrate issues of data governance since, in this study, we did not consider the acquisition cost of the different features in detail.

## ACKNOWLEDGMENT

The authors alone are responsible for the content of the paper.

## REFERENCES

- [1] F. Tang, Y. Kawamoto, N. Kato, and J. Liu, "Future intelligent and secure vehicular network toward 6G: Machine-learning approaches," *Proc. IEEE*, vol. 108, no. 2, pp. 292–307, Feb. 2020.
- [2] K. Tan, D. Bremner, J. Le Kerneec, L. Zhang, and M. Imran, "Machine learning in vehicular networking: An overview," *Digit. Commun. Netw.*, vol. 8, no. 1, pp. 18–24, Feb. 2022.
- [3] H. Ye and G. Y. Li, "Deep reinforcement learning for resource allocation in V2V communications," in *Proc. IEEE Int. Conf. Commun. (ICC)*, May 2018, pp. 1–6.
- [4] N. Zhao, Y.-C. Liang, D. Niyato, Y. Pei, M. Wu, and Y. Jiang, "Deep reinforcement learning for user association and resource allocation in heterogeneous cellular networks," *IEEE Trans. Wireless Commun.*, vol. 18, no. 11, pp. 5141–5152, Nov. 2019.
- [5] D. F. Külzer, F. Debbichi, S. Stanczak, and M. Botsov, "On latency prediction with deep learning and passive probing at high mobility," in *Proc. IEEE Int. Conf. Commun. (ICC)*, Jun. 2021, pp. 1–7.
- [6] D. F. Külzer, M. Kasparick, A. Palaios, R. Sattiraju, O. D. Ramos-Cantor, and D. Wieruch, "AI4Mobile: Use cases and challenges of AI-based QoS prediction for high-mobility scenarios," in *Proc. IEEE Veh. Technol. Conf. (VTC Spring)*, Apr. 2021, pp. 1–7.
- [7] M. Boban, A. Kousaridas, K. Manolakis, J. Eichinger, and W. Xu, "Connected roads of the future: Use cases, requirements, and design considerations for vehicle-to-everything communications," *IEEE Veh. Technol. Mag.*, vol. 13, no. 3, pp. 110–123, Sep. 2018.
- [8] L. Torres-Figueroa, H. F. Schepker, and J. Jiru, "QoS evaluation and prediction for C-V2X communication in commercially-deployed LTE and mobile edge networks," in *Proc. IEEE 91st Veh. Technol. Conf.*, May 2020, pp. 1–7.
- [9] S. Shalev-Shwartz and S. Ben-David, *Understanding Machine Learning: From Theory to Algorithms*. Cambridge, U.K.: Cambridge Univ. Press, 2014.
- [10] M. Mohri, A. Rostamizadeh, and A. Talwalkar, *Foundations of Machine Learning*. Cambridge, MA, USA: MIT Press, 2018.
- [11] J. Sarnelle, A. Sanchez, R. Capó, J. Haas, and R. Polikar, "Quantifying the limited and gradual concept drift assumption," in *Proc. Int. Joint Conf. Neural Netw. (IJCNN)*, Jul. 2015, pp. 1–8.
- [12] I. Redko, E. Morvant, A. Habrard, M. Sebban, and Y. Bennani, *Advances in Domain Adaptation Theory*. Oxford, U.K.: Elsevier, 2019.
- [13] A. Adadi and M. Berrada, "Peeking inside the black-box: A survey on explainable artificial intelligence (XAI)," *IEEE Access*, vol. 6, pp. 52138–52160, 2018.
- [14] S. Neumeier, E. A. Walelgne, V. Bajpai, J. Ott, and C. Facchi, "Measuring the feasibility of teleoperated driving in mobile networks," in *Proc. Netw. Traffic Meas. Anal. Conf. (TMA)*, Jun. 2019, pp. 113–120.
- [15] A. Gaber, W. Nassar, A. M. Mohamed, and M. K. Mansour, "Feasibility study of teleoperated vehicles using multi-operator LTE connection," in *Proc. Int. Conf. Innov. Trends Commun. Comput. Eng. (ITCE)*, Feb. 2020, pp. 191–195.
- [16] A. Narayanan, E. Ramadan, R. Mehta, X. Hu, Q. Liu, R. A. K. Fezeu, U. K. Dayalan, S. Verma, P. Ji, T. Li, F. Qian, and Z.-L. Zhang, "Lumos5G: Mapping and predicting commercial mmWave 5G throughput," in *Proc. ACM Internet Meas. Conf.*, Oct. 2020, pp. 176–193.
- [17] A. Kulkarni, A. Seetharam, A. Ramesh, and J. D. Herath, "DeepChannel: Wireless channel quality prediction using deep learning," *IEEE Trans. Veh. Technol.*, vol. 69, no. 1, pp. 443–456, Jan. 2020.
- [18] C. F. Mecklenbrauker, A. F. Molisch, J. Karedal, F. Tufvesson, A. Paier, L. Bernado, T. Zemen, O. Klemp, and N. Czink, "Vehicular channel characterization and its implications for wireless system design and performance," *Proc. IEEE*, vol. 99, no. 7, pp. 1189–1212, Jul. 2011.
- [19] M. Lauridsen, I. Rodriguez, L. M. Mikkelsen, L. C. Gimenez, and P. Mogensen, "Verification of 3G and 4G received power measurements in a crowdsourcing Android app," in *Proc. IEEE Wireless Commun. Netw. Conf.*, Apr. 2016, pp. 1–6.
- [20] H. Al-Shamisi, H. Al-Shamsi, I. Kostanic, and J. Zec, "Verifying measurements of reference signal received power (RSRP) on LTE network using an app on Android smartphones," in *Proc. IEEE 9th Annu. Inf. Technol., Electron. Mobile Commun. Conf. (IEMCON)*, Nov. 2018, pp. 372–377.
- [21] R. Enami, Y. Shi, D. Rajan, and J. Camp, "Pre-crowdsourcing: Predicting wireless propagation with phone-based channel quality measurements," *Comput. Commun.*, vol. 132, pp. 96–110, Nov. 2018.
- [22] A. Ispas, G. Ascheid, C. Schneider, and R. Thoma, "Analysis of local quasi-stationarity regions in an urban macrocell scenario," in *Proc. IEEE 71st Veh. Technol. Conf.*, May 2010, pp. 1–5.
- [23] L. Bernadó, T. Zemen, F. Tufvesson, A. F. Molisch, and C. F. Mecklenbraüker, "The (in-) validity of the WSSUS assumption in vehicular radio channels," in *Proc. IEEE 23rd Int. Symp. Pers., Indoor Mobile Radio Commun.*, Sep. 2012, pp. 1757–1762.
- [24] D. Choudhary, A. Kejariwal, and F. Orsini, "On the runtime-efficacy trade-off of anomaly detection techniques for real-time streaming data," 2017, *arXiv:1710.04735*.
- [25] A. A. Cook, G. Misirli, and Z. Fan, "Anomaly detection for IoT time-series data: A survey," *IEEE Internet Things J.*, vol. 7, no. 7, pp. 6481–6494, Jul. 2020.
- [26] J. C. Schlimmer and R. H. Granger, "Incremental learning from noisy data," *Mach. Learn.*, vol. 1, no. 3, pp. 317–354, Sep. 1986.
- [27] H. Mehmood, P. Kostakos, M. Cortes, T. Anagnostopoulos, S. Pirttikangas, and E. Gilman, "Concept drift adaptation techniques in distributed environment for real-world data streams," *Smart Cities*, vol. 4, no. 1, pp. 349–371, Mar. 2021.
- [28] A. Abbasi, A. R. Javed, C. Chakraborty, J. Nebhen, W. Zehra, and Z. Jalil, "ElStream: An ensemble learning approach for concept drift detection in dynamic social big data stream learning," *IEEE Access*, vol. 9, pp. 66408–66419, 2021.
- [29] F. E. Casado, D. Lema, M. F. Criado, R. Iglesias, C. V. Regueiro, and S. Barro, "Concept drift detection and adaptation for federated and continual learning," *Multimedia Tools Appl.*, vol. 81, no. 3, pp. 3397–3419, Jan. 2022.
- [30] B. Zhang, L. Xue, W. Wang, S. Qin, and D. Wang, "Model updating mechanism of concept drift detection in data stream based on classifier pool," *EURASIP J. Wireless Commun. Netw.*, vol. 2016, no. 1, pp. 1–7, Sep. 2016.
- [31] J. Schmid, M. Schneider, A. HöB, and B. Schuller, "A comparison of AI-based throughput prediction for cellular vehicle-to-server communication," in *Proc. 15th Int. Wireless Commun. Mobile Comput. Conf. (IWCMC)*, Jun. 2019, pp. 471–476.
- [32] J. Schmid, M. Schneider, A. HöB, and B. Schuller, "A deep learning approach for location independent throughput prediction," in *Proc. IEEE Int. Conf. Connected Vehicles Expo (ICCVE)*, Nov. 2019, pp. 1–5.
- [33] B. Wei, W. Kawakami, K. Kanai, J. Katto, and S. Wang, "TRUST: A TCP throughput prediction method in mobile networks," in *Proc. IEEE Global Commun. Conf. (GLOBECOM)*, Dec. 2018, pp. 1–6.
- [34] A. Palaios, C. Vielhaus, D. F. Külzer, P. Geuer, R. Sattiraju, J. Fink, M. Kasparick, and C. Watermann, "Effect of spatial, temporal and network features on uplink and downlink throughput prediction," in *Proc. IEEE 5G World Forum (5GWF)*, Oct. 2021, pp. 418–423.
- [35] B. Sliwa, H. Schippers, and C. Wietfeld, "Machine learning-enabled data rate prediction for 5G NSA vehicle-to-cloud communications," in *Proc. IEEE 5G World Forum (5GWF)*, Oct. 2021, pp. 299–304.
- [36] A. S. Khatouni, F. Soro, and D. Giordano, "A machine learning application for latency prediction in operational 4G networks," in *Proc. IFIP/IEEE Symp. Integr. Netw. Service Manage. (IM)*, Apr. 2019, pp. 71–74.
- [37] J. D. Mankowitz and A. J. Paverd, "Mobile device-based cellular network coverage analysis using crowd sourcing," in *Proc. Int. Conf. Comput. Tool*, Apr. 2011, pp. 1–6.
- [38] M. Feltrin and S. Tomasin, "A machine-learning-based handover prediction for anticipatory techniques in Wi-Fi networks," in *Proc. 10th Int. Conf. Ubiquitous Future Netw. (ICUFN)*, Jul. 2018, pp. 341–345.
- [39] F. Burmeister, N. Schwarzenberg, T. Höbner, and G. Fettweis, "Measuring time-varying industrial radio channels for D2D communications on AGVs," in *Proc. IEEE Wireless Commun. Netw. Conf. (WCNC)*, Mar. 2021, pp. 1–7.

- [40] S. Farthofer, M. Herlich, C. Maier, S. Pochaba, J. Lackner, and P. Dorfinger, "An open mobile communications drive test data set and its use for machine learning," *IEEE Open J. Commun. Soc.*, vol. 3, pp. 1688–1701, 2022.
- [41] A. Palaios, P. Geuer, J. Fink, D. F. Külzer, F. Götsch, M. Kasparick, D. Schäufele, R. Hernangómez, S. Partani, R. Sattiraju, A. Kumar, F. Burmeister, A. Weinand, C. Vielhaus, F. H. P. Fitzek, G. Fettweis, H. D. Schotten, and S. Stanczak, "Network under control: Multi-vehicle E2E measurements for AI-based QoS prediction," in *Proc. IEEE 32nd Annu. Int. Symp. Pers., Indoor Mobile Radio Commun. (PIMRC)*, Sep. 2021, pp. 1432–1438.
- [42] Y. Li, C. Peng, Z. Yuan, H. Deng, J. Li, and T. Wang, "MobileInsight: Analyzing cellular network information on smartphones," *GetMobile, Mobile Comput. Commun.*, vol. 21, no. 1, pp. 39–42, May 2017.
- [43] Gyokov Solutions. (2016). *G-NetTrack Pro*. Google Play Store. Accessed: May 5, 2022. [Online]. Available: <https://www.gyokovsolutions.com/G-NetTrack%20Android.html>
- [44] *IPERF 3*. Accessed: Apr. 25, 2022. [Online]. Available: <https://www.techspot.com/downloads/7303-iperf.html>
- [45] A. Palaios, J. Riihijärvi, O. Holland, A. Achtzehn, and P. Mähönen, "Measurements of spectrum use in London: Exploratory data analysis and study of temporal, spatial and frequency-domain dynamics," in *Proc. IEEE Int. Symp. Dyn. Spectr. Access Netw.*, Oct. 2012, pp. 154–165.
- [46] M. Sayed-Mouchaweh and E. Lughofer, *Learning in Non-Stationary Environments: Methods and Applications*. New York, NY, USA: Springer, 2012.
- [47] S. Ng and P. Perron, "Unit root tests in ARMA models with data-dependent methods for the selection of the truncation lag," *J. Amer. Stat. Assoc.*, vol. 90, no. 429, pp. 268–281, Mar. 1995.
- [48] N. Perpinias, J. Riihijärvi, and P. Mahonen, "Impact of model uncertainties on the accuracy of spatial interpolation based coverage estimation," in *Proc. IEEE Wireless Commun. Netw. Conf. (WCNC)*, Mar. 2017, pp. 1–6.
- [49] U. Schlicher, C. Bettstetter, and G. Brandner, "Temporal correlation of interference in wireless networks with Rayleigh block fading," *IEEE Trans. Mobile Comput.*, vol. 11, no. 12, pp. 2109–2120, Dec. 2012.
- [50] D. R. Roberts, V. Bahn, S. Ciuti, M. S. Boyce, J. Elith, G. Guillera-Arroita, S. Hauenstein, J. J. Lahoz-Monfort, B. Schröder, W. Thuiller, D. I. Warton, B. A. Wintle, F. Hartig, and C. F. Dormann, "Cross-validation strategies for data with temporal, spatial, hierarchical, or phylogenetic structure," *Ecography*, vol. 40, no. 8, pp. 913–929, Aug. 2017.
- [51] G. Widmer and M. Kubat, "Learning in the presence of concept drift and hidden contexts," *Mach. Learn.*, vol. 23, no. 1, pp. 69–101, Apr. 1996.
- [52] A. Tsymbal, "The problem of concept drift: Definitions and related work," *Comput. Sci. Dept., Trinity College Dublin*, vol. 106, no. 2, pp. 58–64, Apr. 2004.
- [53] R. Elwell and R. Polikar, "Incremental learning of concept drift in nonstationary environments," *IEEE Trans. Neural Netw.*, vol. 22, no. 10, pp. 1517–1531, Oct. 2011.
- [54] J. Lu, A. Liu, F. Dong, F. Gu, J. Gama, and G. Zhang, "Learning under concept drift: A review," *IEEE Trans. Knowl. Data Eng.*, vol. 31, no. 12, pp. 2346–2363, Dec. 2019.
- [55] L. Baier, J. Reimold, and N. Kühl, "Handling concept drift for predictions in business process mining," in *Proc. IEEE 22nd Conf. Business Informat. (CBI)*, vol. 1, Jun. 2020, pp. 76–83.
- [56] G. Ditzler and R. Polikar, "Incremental learning of concept drift from streaming imbalanced data," *IEEE Trans. Knowl. Data Eng.*, vol. 25, no. 10, pp. 2283–2301, Oct. 2013.
- [57] S. Wang, L. L. Minku, and X. Yao, "A systematic study of online class imbalance learning with concept drift," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 29, no. 10, pp. 4802–4821, Oct. 2018.
- [58] R. Hernangómez, A. Palaios, G. Guruvayoorappan, M. Kasparick, N. U. Ain, and S. Stanczak, "Online QoS estimation for vehicular radio environments," in *Proc. 30th Eur. Signal Process. Conf. (EUSIPCO)*, Aug. 2022, pp. 1701–1705.
- [59] Y. Li, C. Peng, Z. Yuan, J. Li, H. Deng, and T. Wang, "MobileInsight: Extracting and analyzing cellular network information on smartphones," in *Proc. 22nd Annu. Int. Conf. Mobile Comput. Netw.*, Oct. 2016, pp. 202–215.
- [60] T. Akiba, S. Sano, T. Yanase, T. Ohta, and M. Koyama, "Optuna: A next-generation hyperparameter optimization framework," in *Proc. 25th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, Jul. 2019, pp. 2623–2631.
- [61] M. Kuhn and K. Johnson, *Applied Predictive Modeling*. New York, NY, USA: Springer, 2013.
- [62] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in *Proc. 31st Int. Conf. Neural Inf. Process. Syst.*, Long Beach, CA, USA, Dec. 2017, pp. 4765–4774.
- [63] D. W. Apley and J. Zhu, "Visualizing the effects of predictor variables in black box supervised learning models," *J. Roy. Stat. Soc. B, Stat. Methodol.*, vol. 82, no. 4, pp. 1059–1086, Sep. 2020.
- [64] R. Zhohov, A. Palaios, and P. Geuer, "One step further: Tunable and explainable throughput prediction based on large-scale commercial networks," in *Proc. IEEE 5G World Forum (5GWF)*, Oct. 2021, pp. 430–435.
- [65] R. Hernangómez, P. Geuer, A. Palaios, D. Schäufele, C. Watermann, K. Taleb-Bouhemadi, M. Parvini, A. Krause, S. Partani, C. Vielhaus, M. Kasparick, D. F. Külzer, F. Burmeister, F. H. P. Fitzek, H. D. Schotten, G. Fettweis, and S. Stanczak, "Berlin V2X: A machine learning dataset from multiple vehicles and radio access technologies," in *Proc. IEEE 97th Veh. Technol. Conf. (VTC-Spring)*, Florence, Italy, Jun. 2023, pp. 1–5, doi: 10.1109/VTC2023-Spring57618.2023.10200750.
- [66] R. Hernangómez, A. Palaios, C. Watermann, D. Schäufele, P. Geuer, R. Ismayilov, M. Parvini, A. Krause, M. Kasparick, T. Neugebauer, O. D. Ramos-Cantor, H. Tchouankem, J. L. Calvo, B. Chen, and G. Fettweis, "Towards an AI-enabled connected industry: AGV communication and sensor measurement datasets," 2022, *arXiv:2301.03364*.

**ALEXANDROS PALAIOS** received the dual master's degrees in basic and applied cognitive science and in communication systems and networks from the University of Athens, in 2008 and 2010, respectively, and the Ph.D. degree from RWTH Aachen University, in 2016. His research interest includes the application of ML for 6G networks.

**CHRISTIAN L. VIELHAUS** received the Dipl.-Ing. degree in electrical engineering from the Technical University of Dresden (TU Dresden), Germany, in 2019, where he is currently pursuing the Ph.D. degree in electrical engineering with the Deutsche Telecom Chair of Communication Networks. His research interests include machine learning, network simulation, and congestion control.

**DANIEL F. KÜLZER** received the B.Sc. and M.Sc. degrees in electrical engineering and information technology from the Technical University of Munich, Germany, in 2016 and 2018, respectively. He is currently pursuing the Ph.D. degree, under the supervision of Prof. Sławomir Stańczak. Since 2018, he has been with BMW Group, Munich, on different topics at the intersection of artificial intelligence and wireless communication. His research interests include network optimization techniques, particularly predictive resource allocation, and machine learning for quality of service prediction and provisioning in the context of vehicular communication. In 2018, he was awarded the Engineer's degree from Télécom Paris as part of a double degree program.

**CARA WATERMANN** received the M.Sc. degree in computer engineering from the University of Duisburg-Essen, in 2021. She is currently a Researcher with Ericsson Research, Germany. She is also working in the areas of ML for wireless communications. Her research interests include ML-based pQoS and explainability.

**RODRIGO HERNANGÓMEZ** (Graduate Student Member, IEEE) received the B.Sc. degree in telecommunication technologies and services engineering and the M.Sc. degree in telecommunication engineering from Universidad Politécnica de Madrid, Spain, in 2015 and 2017, respectively. He is currently pursuing the Ph.D. degree with the Fraunhofer Heinrich Hertz Institute. He is also a Research Associate with the Fraunhofer Heinrich Hertz Institute. His research interests include the application of machine learning and signal processing to wireless communication systems and radar, and the integration of both radio technologies.

**SANKET PARTANI** (Associate Member, IEEE) received the M.Sc. degree in commercial vehicle technology from the RPTU Kaiserslautern-Landau, in 2019, where he is currently pursuing the Ph.D. degree with the Chair for Wireless Communications and Navigation. He is also a Research Assistant with the Division of Wireless Communications and Radio Navigation, University of Kaiserslautern. His research interests include V2X communications, CCAM applications, mobility modeling, and intelligence in networks.

**PHILIPP GEUER** received the master's degree in electrical engineering from RWTH Aachen University, in 2019. He is currently an experienced Researcher with Ericsson Research, Germany. His research interests include wireless networks, specifically focusing on applying ML for improving resilience, enabling proactive resource management, and predicting QoS.

**ANTON KRAUSE** received the Dipl.-Ing. degree in electrical engineering from TU Dresden, Germany, in 2022, where he is currently pursuing the Ph.D. degree with the Vodafone Chair Mobile Communications Systems.

**RAJA SATTIRAJU** received the master's degree in commercial vehicle technology from RPTU Kaiserslautern-Landau. Since 2012, he has been a Research Associate with the Chair for Wireless Communication and Navigation, working with various research projects at both the German and EU level. Since 2022, he has been a Senior Engineer with Qualcomm CDMA Technologies GmbH, assisting operators and vendors with network deployment and optimization solutions. His research interests include V2X communications, link and system level simulations, and the applications of ML for network optimization.

**MARTIN KASPARICK** received the Dipl.-Ing. and Dr.-Ing. degrees from Technische Universität Berlin, Germany, in 2009 and 2015, respectively. Since 2016, he leads the Signal and Information Processing Group, Wireless Communications Department, Fraunhofer Heinrich Hertz Institute.

**GERHARD P. FETTWEIS** (Fellow, IEEE) received the Ph.D. degree from RWTH Aachen University, in 1990. After a postdoctoral position, he joined IBM Research, San Jose, CA, USA, he joined TCSI, Berkeley, CA, USA. Since 1994, he has been a Vodafone Chair Professor with TU Dresden. Since 2018, he has been the founding Director of the Barkhausen Institute.

**FRANK H. P. FITZEK** (Senior Member, IEEE) is currently a Professor and the Head of the Deutsche Telekom Chair of Communication Networks, TU Dresden. He is the spokesman of the DFG Cluster of Excellence CeTI and the BMBF 6G Hub "6G-life."

**HANS D. SCHOTTEN** (Member, IEEE) received the Diploma and Ph.D. degrees in electrical engineering from the Aachen University of Technology, Germany, in 1990 and 1997, respectively. He was a Senior Researcher, the Project Manager, and the Head of the Research Groups, Aachen University of Technology; Ericsson Corporate Research; and Qualcomm Corporate Research and Development, respectively. During his time at Qualcomm, he has been the Director of the Technical Standards and a Coordinator of Qualcomm's activities in European research programs. Since August 2007, he has been a Full Professor and the Head of the Chair for Wireless Communications and Radio Navigation at the RPTU Kaiserslautern-Landau. Since 2012, he has been the Scientific Director of the German Research Centre for Artificial Intelligence, and heading the Intelligent Networks Department.

**SLAWOMIR STAŃCZAK** (Senior Member, IEEE) studied electrical engineering with a specialization in control theory at the Wrocław the University of Technology and the Technical University of Berlin (TU Berlin). He received the Dipl.-Ing. degree and the Dr.-Ing. degree in electrical engineering from TU Berlin, in 1998 and 2003, respectively, and the Habilitation degree, in 2006. Since 2015, he has been a Full Professor of network information theory with TU Berlin, and the Head of the Wireless Communications and Networks Department, Fraunhofer Institute for Telecommunications, Heinrich Hertz Institute (HHI).

...