NordicDat: A Cross-Border Predictive QoS Dataset

Topi Miekkala*, Pasi Pyykönen*, Georgios Drainakis†, Panagiotis Pantazopoulos†,
Tobias Muller‡, Konstantinos V. Katsaros†, Vasilis Sourlas†, Angelos Amditis†, Dimitra I. Kaklamani§
*VTT Technical Research Centre of Finland, Tampere, Finland, {topi.miekkala, pasi.pyykonen} @vtt.fi
†Institute of Communication and Computer Systems (ICCS), Athens, Greece,
{giorgos.drainakis, ppantaz, k.katsaros, v.sourlas, a.amditis} @iccs.gr
‡Robert Bosch GmbH, Leonberg, Germany, {tobias.mueller8} @de.bosch.com
§National Technical University of Athens (NTUA), Athens, Greece, {dkaklam} @mail.ntua.gr

***Abstract*—The advent of 5G and beyond systems is expected to shape the automotive vertical, as safety-critical vehicular applications rely on the network to meet their stringent Quality of Service (QoS) requirements. Predictive QoS (pQoS) has been proposed as a mechanism that allows automotive applications to proactively adapt in view of forthcoming QoS changes. Although pQoS is typically facilitated via classical (centralized) Machine Learning (ML) methods, the demand for data privacy has led to the emergence of distributed ML schemes. Efficient training of ML models however requires large volumes of (kinematic-state and connectivity) QoS data, so as to capture the involved spatio-temporal effects.**

To that end we hereby present and publicly share *NordicDat*, a QoS dataset collected during a two-week measurement campaign, driving across three European countries. NordicDat contains over 90K samples of physical layer, network and mobility-related features. Contrary to prior works, it includes multiple instances of cross-boarder roaming, diverse vehicle speed profiles and radio access technologies (generations). Further, we provide a thorough NordicDat data analysis, highlighting the dependencies between the NordicDat's features and the resulting QoS values (throughput, delay). To showcase its broad usability, we train pQoS ML models over NordicDat in classical and distributed fashion. Our results demonstrate for the first time the viability of distributed pQoS with real-word data, which achieves similar (within a margin of 10%) accuracy to that of classical ML, cropping privacy-preserving benefits.

***Index Terms*—Dataset, Predictive QoS, Automotive, Cross-Border Roaming, 5G-NSA, Distributed AI/ML**

I. Introduction

Cooperative, connected and automated mobility (CCAM) technologies are expected to revolutionize the mobility sector towards safer, efficient and sustainable transport. Typical automotive applications span from tele-operated driving, co-operative lane merge, *etc.* [1] to 5G cross-border corridors and platooning [2]. Mobile network connectivity is a key enabler towards this development, since such applications often require Quality of Service (QoS) guarantees *e.g.,* ubiquitous coverage, minimum required rates, *etc.* [3]. Despite the latest cellular generation technology aiming to provide QoS guarantees [3], connectivity remains subject to numerous environmental and technical factors affecting the achievable QoS, which in turn can threat user experience or even jeopardize safety [1].

To mitigate such risks, the notion of predictive Quality of Service (pQoS) has been introduced by the 5G Automotive Assosiation (5GAA) [3]. pQoS is a mechanism that estimates future network QoS values and informs the automotive application about a forthcoming QoS degradation event. This allows the automotive application to *proactively* adapt its functionality, by taking responsive actions to the estimated QoS changes *e.g.,* speed reduction, fail-safe maneuver, abortion/overtaking of an operation, *etc.* pQoS is typically facilitated by the mobile network operator (MNO) that has access to detailed radio and network level information *e.g.,* cell capacity, traffic prioritisation, mobility management optimization, *etc.* [3]. An alternative approach is an Over-The-Top (OTT) solution, where pQoS is provided by a third party *e.g.,* a car manufacturer, automotive supplier, *etc.* [3]. Our work focuses on OTT-based pQoS, to tackle any QoS data availability issues due to MNO confidentiality limitations.

Estimating QoS values in vehicular environments is a complex task, since radio conditions can rapidly vary in the course of time [4]. Artificial Intelligence and Machine Learning (AI/ML) is a key enabling technology towards accurate QoS prediction, so far demonstrating prominent results [5], [6], [4]. Classical AI/ML (CL) pipelines include QoS data collection in a centralized entity *e.g.,* a cloud server, where AI/ML model training occurs thereafter. Motivated by the demand for user-privacy, reduction of communication costs and scalability gains, distributed pQoS schemes have recently emerged that allow for collaborative AI/ML model training [7], [8]. A typical example is Federated Learning (FL) [7], where training is locally performed in the client-vehicles side. The server is thereafter responsible to collect and merge (aggregate) the trained local models into a collective (global) model. QoS data remains at the client-vehicles at all times, thus privacy is ensured.

As AI/ML is gaining ground to facilitate pQoS, there is an increasing demand for QoS data to effectively train AI/ML models. The process of data acquisition poses several challenges *e.g.,* high cost of measurements campaigns, quality of captured data, *etc.* [4]. In fact, QoS values exhibit high volatility and multiple inter-dependencies that are related to network and radio parameters, user mobility, spatio-temporal effects, *etc.* [4]. As such, a careful design of data collection process is required to capture a diverse set of features under various (mobility and network) scenarios [4].

Although a rich set of measurements campaigns has been performed so far [9], [10], [11], the obtained QoS datasets pose several limitations that render their application for pQoS ineffective: a) they are focused on Long-Term Evolution (LTE) radio-access network (RAN) technologies; only few of them study 5G and beyond deployments, b) they are performed in urban or indoor locations, disregarding the effect of high-speed mobility *e.g.,* in highways, a typical environment for vehicular applications, c) they are restricted to national networks, neglecting the effects of cross-border roaming, a major cause of QoS degradation [12], as shown in this work.

Seeking to overcome the above limitations, we present NordicDat, a QoS dataset obtained during a measurement campaign that spans across three European countries *i.e.,* Finland, Sweden and Norway. NordicDat comprises of 25 hours of driving runs with diverse speed patterns (speed, acceleration, direction), typically categorized under the speed profiles of highway driving. The dataset includes physical and network-layer features, as well as information on the vehicle's kinematics. The measurement area is set near each country's border, to specifically capture the effect of roaming on QoS values. Our measurements capture traces from *both* LTE and 5G RAN. As our data analysis reveals changes in terms of roaming, speed profile and RAN have major impact on the resulting QoS values. Our produced dataset is then utilized to showcase distributed pQoS via FL emulation, for throughput and delay prediction tasks. To our knowledge, this is the first attempt to present distributed pQoS using real-world public data. Our results reveal that for both tasks, FL's accuracy performance is similar to that of classical ML (by an average margin of 10%).

The captured dataset (accompanied by a short documentation) is publicly shared in an open repository [13] to support further research in the field of pQoS and fill the respective research gaps. The rest of the paper is organized as follows. State of the art (SotA) is presented in Sec. II. In Sec. III the data acquisition process is described, followed by data analysis. Preliminary results are shown in Sec. IV and we conclude in Sec. V.

II. Motivation and Related Work

Training and evaluation of AI/ML QoS prediction algorithms is often based on network simulation-based datasets [5], [6]. Such datasets however have limited credibility, since they rather fail to reveal the complex patterns that are observed in a real network environment [4]. As such, a demand for real-world QoS datasets that are able to capture the dependencies between network operations, the radio environment and the user-equipment (UE)'s behavior, emerges.

A plethora of datasets has so far been developed via measurement campaigns for various mobility scenarios and RAN technologies (see Table I). Several studies focus on driving scenarios that are in line with the automotive applications of pQoS. Existing works however are either mostly limited to LTE in terms of RAN [14], [15], [9], [11], [16], or to urban/sub-urban driving speed profiles in terms of mobility [17], [18]. Only few works provide results that include 5G measurements in highway mobility scenarios [19], [20], [21], but the involved experiments are conducted within a single country therefore fail to capture the effects of roaming. Another series of studies have investigated network properties under low-speed mobility scenarios *e.g.,* pedestrian or stationary (no-mobility) environments [22], [10], [23], [24]. Among them, a few measurement campaigns have been performed across various European countries to study the effect of roaming [12], [25], [26], however their usability is limited for pQoS, being performed in stationary or near-stationary environments. Finally, few works have focused on indoor network environments, utilizing mobility scenarios via automated guided vehicles (AGVs) [27], [28] or stationary states (office environments) [29]. Such datasets are mostly tailored to industrial networks *e.g.,* private 5G networks, device-to-device (D2D) communications*, etc.*

Our shared dataset aims to fill these research gaps and enable research in the area of pQoS, since a) it is exclusively developed via drive tests (under various mobility scenarios), b) it includes both LTE and 5G RAN technologies and most importantly, c) it is acquired in a cross-border area between three countries and therefore captures the effects of roaming on QoS values.

III. Dataset Description

The NordicDat dataset combines measurements from a 5G modem, external positioning sensors and the vehicle internal data from the Controller Area Network (CAN) protocol [30]. The aim was to collect measurement sequences which combine connectivity, positioning and kinematic data of the vehicle in geographical areas which present QoS degradation. The dataset contains 25 hours of such sequences, collected in arctic rural regions of Finland, Norway and Sweden. The sequences contain sections of low connectivity and total loss of connection, including handover events at national borders.

A. Measurement setup and data collection

The dataset was collected using a research vehicle (Martti) (see Fig. 1) offered by Valtion Teknillinen Tutkimuskeskus (VTT): a Volkswagen Touareg passenger car with modifications and external sensor installations for self-driving research purposes. All the recorded values (features) of the dataset are listed in Table II. It includes a) physical-layer parameters *e.g.,* Reference Signal Received Quality, Power, Strength Indicator (RSRQ, RSRP, RSSI, respectively), Signal to Interference plus Noise Ratio (SINR), b) network-layer parameters *e.g.,* band, RAN, serving cell, operator (coded as integer values for data anonymization), c) mobility-related values *e.g.,* position (latitude, longitude, elevation), velocity and acceleration and finally d) QoS parameters *e.g.,* downlink (DL) and uplink (UL) throughput as well as delay.

The positioning data of the vehicle was collected using two external sensors. The Global Navigation Satellite System (GNSS) data was provided by Ublox ZED-F9P Real-Time Kinematic (RTK) GNSS sensor. The latitude, longitude, altitude and GNSS service quality were recorded. The accuracy of the GNSS positioning varies in the measurement sequences, as

TABLE I: Comparison of public cellular QoS datasets

Dataset	Mobility	RAN	Location	Area Types			
				Urban	Suburban	Highway	Indoor
5G Connected Mobility [14]	driving	LTE	Nuremberg, Germany	✓	✓	✓	
5G Meas [19]	driving, walking, static	5G	Indianapolis and Chicago, US	✓	✓	✓	
5G Wild [20]	driving, walking, static	5G	2 US cities	✓		✓	
5Gophers [17]	driving, walking	5G	Minneapolis, Chicago, Atlanta, US	✓	✓		✓
BerlinV2X [15]	driving	LTE	Berlin, Germany	✓	✓	✓	
Beyond Throughput [9]	driving, static	LTE	Ireland	✓	✓	✓	✓
Lumos5G [18]	driving, walking	mmWave 5G	Minneapolis, US	✓			✓
Rome [21]	driving, walking	LTE, NB-IoT, 5G	Rome, Italy	✓			✓
SRFG [11]	driving	LTE	Salzburg, Austria			✓	
Terminal [16]	driving	LTE	Single city	✓			
5G Beams [22]	walking	mmWave 5G	Chicago, US	✓	✓		
5G Consumption [10]	walking	5G	Campus		✓		✓
5G PHY Latency [23]	walking, static	5G	Minneapolis, US	✓			
Experience [12]	mobile, static	LTE	Italy, Norway, Spain, Sweden, UK, Germany	✓			
MONROE [25]	mobile, static	LTE	Italy, Norway, Spain, Sweden, UK, Germany	✓			
Roaming [26]	static	5G	France, Italy, Spain	✓			
UE Network Traffic [24]	static, emulated driving	LTE	Volos, Greece	✓			
AGV [27]	mobile	D2D	Industrial				✓
IV2V/IV2I+ [28]	mobile	LTE	Industrial				✓
Urban Office [29]	static	LTE	Vienna, Austria				✓
NordicDat [13]	**driving**	**LTE, 5G-NSA**	**Finland, Sweden, Norway (cross-border)**		✓	✓	

Fig. 1: Martti research vehicle

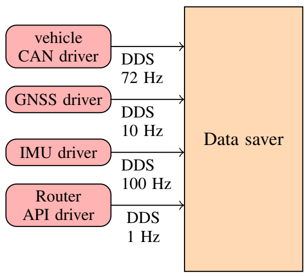

Fig. 2: Software architecture

the RTK correction signal relies on mobile connectivity [31]. The GNSS service quality ranges between Differential GNSS (DGNSS), RTK float and RTK fix. Due to this variance of the GNSS measurement accuracy, the vehicle velocity data was obtained by reading wheel speeds from the CAN bus, and converting those to vehicle speed. The vehicle orientation is obtained from Xsens MTi-680g inertia measurement unit.

The mobile connectivity data was obtained from a Teltonika RUTX50 5G modem, placed on the vehicle dashboard. AT commands ('AT' standing for 'Attention') are Application Programming Interfaces (APIs) for communicating with a cellular modem [32]. In collection of this dataset, AT commands were utilized to retrieve connectivity parameters from the device. For the inclusion of relevant network DL and UL speeds, intentional (bandwidth) strain was applied to the connection by actively downloading large files over the mobile connection during the measurements. UL and DL throughput was measured at the application-level using the ifstat API [33]. Similarly, the network delay was obtained at application level by the Linux ping utility [34]. The data saver software implemented the ping request, and reception of the response.

The software architecture of the data collection is shown in Fig. 2. The vehicle internal network communication occurs over Ethernet, with the Teltonika router using a commercially available subscription, and acting as the only gateway to the internet. Independent sensor driver software was implemented for the positioning sensors, the vehicle CAN bus, and the

TABLE II: Description of dataset values

Data source	Rate	Parameter	Unit
Teltonika RUTX50	1 Hz	timestamp	seconds (s)
		RSRQ	decibel (dB)
		RSRP	decibel (dB)
		RSSI	decibel (dB)
		SINR	decibel (dB)
		band	string
		RAN	string
		serving cell ID	integer
		delay (network ping)	milliseconds (ms)
		service status	boolean
		operator	integer
		DL throughput (ifstat in)	kilobytes per second (kb/s)
		UL throughput (ifstat out)	kilobytes per second (kb/s)
Ublox ZED-F9P	10 Hz	latitude	degrees
		longitude	degrees
		elevation	meters (m)
		GNSS mode	integer
Xsens MTi-680g	100 Hz	heading	degrees
		lateral acceleration	meters per second squared (m/s^2)
		longitudinal acceleration	meters per second squared (m/s^2)
		absolute acceleration	meters per second squared (m/s^2)
Vehicle CAN bus	72 Hz	absolute velocity	meters per second (m/s)
		longitudinal velocity	meters per second (m/s)

Teltonika router. Measurements from each sensor driver were published over the vehicle local network using Data Distribution Service (DDS) protocol [35], and received at a data saver application. As the measurement setup contained multiple sensors and devices with varying data rates, the data points were recorded in the data saver application as snapshots of the latest measurement value of each device at an 1 Hz refresh rate.

The NordicDat dataset was collected in a number of separate measurement sessions which included lengthy continuous drives with velocities ranging from low urban speeds up to 100 km/h. The data collection was conducted in the arctic rural areas of northern Finland, Norway and Sweden, highlighted in Fig. 3. While a large part of the dataset consists of dynamic driving on longer routes, the dataset also contains a number of cross-border events where handover occurs between service providers of the countries. The effect of connection loss in cross-border scenarios is reflected on weakened GNSS signal mode. The driving routes were selected with the purpose of including a variety of different levels of QoS. This includes areas of poor cellular coverage, or in some cases areas of complete connection loss. The aim of the data collection was to obtain natural dynamic measurements from challenging

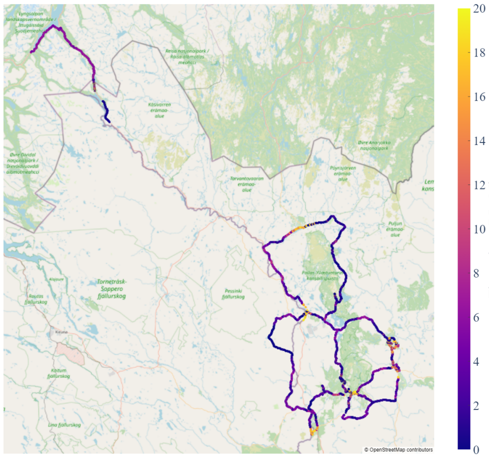

Fig. 3: Spatial effects on QoS (DL Throughput)

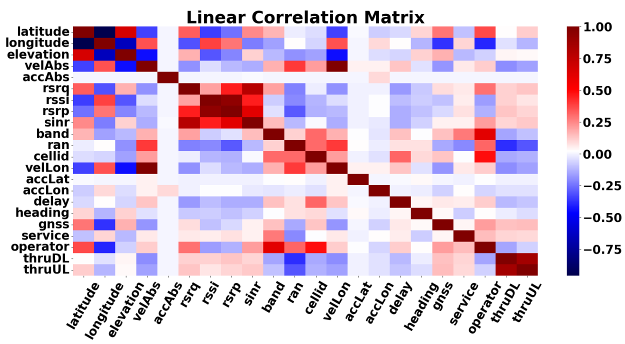

Fig. 4: NordicDat feature correlation

realistic environments. In total, the dataset contains more than 25 hours of measurements, covering close to 1200 km of driving.

B. Data analysis and statistics

Predicting QoS values is a complex task, due to the volatile nature of the involved network parameters [4]. The vehicle's location plays a key role in the resulted prediction [36], since it reflects the spatial effects due to *e.g.,* physical layer parameters, the characteristics of the surrounding environment, *etc.* These findings are also validated in our dataset. As shown in the spatial effects heatmap (see Fig. 3), DL throughput varies in range of $[0, 20]$ Mbps, across the route of the measurement campaign. Specifically throughout the campaign, the 25^{th}, 50^{th} and 75^{th} throughput percentiles values are 0.38, 4.39 and 12.03 Mbps, respectively. An overview of the (linear) relations between the pQoS values (DL and UL throughput, delay) and the respective network, spatial and mobility features is depicted in Fig. 4. A close examination shows no significant evidence of a linear correlation between the involved values. Interestingly though, DL throughput is mostly affected by mobility (velocity) and certain network-level parameters (band, RAN, service cell, operator), whereas UL throughput and delay mostly depend on physical layer parameters (RSRQ, SINR).

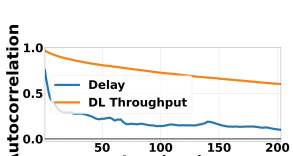

Fig. 5: Temporal autocorrelation

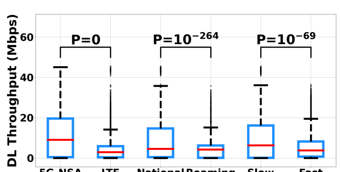

Fig. 6: Feature dependencies

In terms of temporal effects *i.e.,* to what extend previous values affect future ones, the involved QoS values exhibit high diversity (see Fig. 5). Specifically for a time interval larger than 5 seconds, consecutive delay observations exhibit an autocorrelation value of less than 0.5; linear temporal dependencies are therefore negligible. Throughput values on the other hand exhibit larger temporal effects; autocorrelation values larger than 0.8, for a horizon of up to 50 seconds. Finally, our dataset validates the results of previous studies (see Sec. II) and highlights the importance of the following factors a) RAN technology, b) Roaming effects and c) Mobility pattern on the respective QoS values. As depicted in Fig. 6, the average throughput is degraded and the density of extreme values (outliers) is higher in all the following bilateral comparisons: a) LTE *vs.* 5G, b) Roaming *vs.* national network and c) Highway *vs.* urban mobility (velocities up to 25 km/h). The respective P values that express the similarity between the bilateral distributions in all cases (calculated using the Mann–Whitney test) are almost zero for all cases, demonstrating that the above-mentioned factors have a significant effect on the resulted QoS value.

C. Potential usage and limitations

NordicDat exposes a wide range of features that include a) mobility-related metrics *e.g.,* position, speed, acceleration, b) physical layer cellular parameters *e.g.,* SNR, RSRQ, RSSI, c) network-level parameters *e.g.,* cell number, RAN, operator and finally d) QoS values, such as (appplication-level) throughput and delay. The resulted dataset is formulated as a time-series table and it can therefore be utilized to perform a series of relevant prediction tasks, besides pQoS. Relevant examples include handover (change of cell) prediction, vehicle trajectory prediction and driver intention classification.

The dataset's limitations are summarized as follows: 1) All features are obtained from the vehicle's devices, in-line with the OTT approach described in Section I. Though such an approach bypasses any MNO-related data confidentiality issues, it lacks information in regards to the overview of the network *e.g.,* cell capacity, total number of active UEs, distance to basestation, slicing policies, *etc.* [3]. 2) The 5G modem's available interfaces (see Sec. III-A) do not provide support for additional features that could enhance the accuracy of pQoS *e.g.,* resource blocks, Reference Signal Signal to Noise Ratio (RSSNR), Channel Quality Information (CQI), carriers number, coding schemes, *etc.* [4], [6]. 3) DL and UL throughput are measured via the Linux ifstat (application-level) API (see Sec. III-A). As such, no information is given in regards to the transport, network, or link layer. 4) Data collection focuses on cross-border scenarios and thus the locations under study mostly include highway road segments of low-traffic, as compared to an urban environment. The dataset thus lacks instances of QoS degradation that are

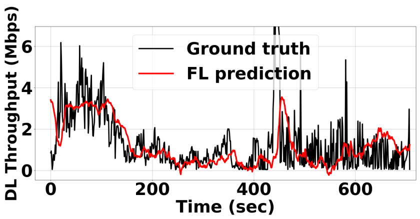

Fig. 7: DL Throughput inference

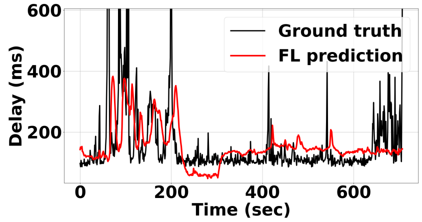

Fig. 8: Delay inference

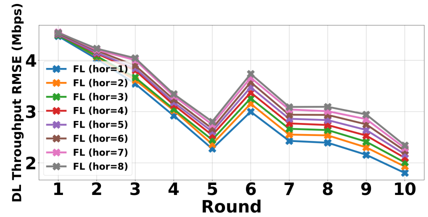

Fig. 9: FL horizons: DL Throughput

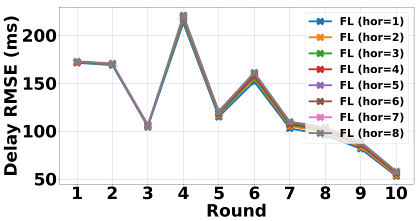

Fig. 10: FL horizons: Delay

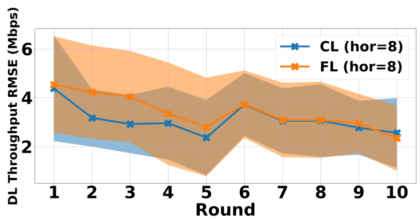

Fig. 11: CL *vs.* FL (DL Throughput)

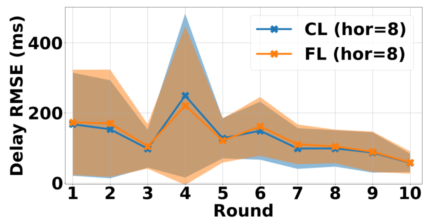

Fig. 12: CL *vs.* FL (Delay)

found in crowded urban areas, due to multiple parallel user transmissions. 5) Due to practical limitations, a single vehicle is used for the measurements. Multiple client scenarios *e.g.,* distributed AI/ML tasks can be emulated by techniques, such as dataset federation (see Sec. IV).

IV. NordicDat-Driven Experimentation

A. Problem statement and evaluation methodology

Motivated by the recent advances in distributed AI/ML (see Sec. I), we use our dataset to demonstrate pQoS via FL, also comparing against its centralized alternative. To the best of our knowledge this is the first attempt to employ distributed pQoS on a real-world public dataset. To emulate an FL setup with multiple clients, we split our dataset into 10 parts of equal size, each representing a single vehicle-client.[1] We then run a series of pQoS training tasks using both centralized (CL) and distributed (FL) ML to predict a) delay and b) DL throughput, for a total of 400 experiments (2 ML approaches $\times$ 100 repetitions $\times$ 2 QoS values). For each client, data is split at a typical 80%-20% ratio [37]. In each training round, we select 8 clients for training and 2 for testing [38]. Round duration is set to 800 secs [37], for a total of 10 rounds per experiment. We address QoS prediction as a typical time-series problem *i.e.,* we predict future QoS values based on previous QoS observations. For that cause we train a custom Long Short-Term Memory (LSTM) model with the following characteristics: 22 input features (equal to the total features of NordicDat) and 8 output features *i.e.,* a prediction horizon of 8 sec (typical for automotive applications [39]). Hyper-parameter tuning on the LSTM model via grid-search resulted in the following values: sliding window=75, hidden size=50, Min-Max normalization, decay=10^{-5}, Rectified Linear Unit activation and Mean Square Error loss function, batch size=64, learning rate=10^{-5}, epochs=50. QoS Prediction accuracy is evaluated using the *Root Mean Square Error (RMSE)* metric [37].

[1]The impact of the split strategy on prediction performance has already been examined [4] and is therefore excluded from our comparison.

B. QoS Prediction

We firstly present two representative instances (for a single vehicle-client during a single experiment) of the inference results achieved by our FL model, in terms of DL throughput (see Fig. 7) and delay prediction (see Fig. 8). These qualitative representations suggest that the FL model is able to track the complex patterns and variations of the QoS values under study (throughput and delay), across time. The quantitative results that present the mean inference values of all clients, across all experiments, for all prediction horizons (from 1 up to 8 seconds ahead) are depicted in Fig. 9 and 10. As expected, longer prediction horizons are prone to larger prediction errors (RMSE) of the QoS value, as compared to shorter ones. Interestingly, the horizon's impact on throughput is much higher as compared to delay. Specifically for throughput, increasing the horizon from 1 to 3, 5 and 8 sec, results in an increase of RMSE by 5.55%, 10.09% and 15.07% (averaged across all rounds), respectively (see Fig. 9). For delay prediction on the other hand, the respective values are 2.36%, 3.67% and 4.11% (see Fig. 10).

Having said that, we fix the horizon value to 8 sec and compare the performance of distributed pQoS (FL) to that of the classical ML approach (CL), for both throughput and delay prediction tasks. We present the mean inference values and (shady) standard deviations of all clients, across all experiments in a per-round basis (see Fig. 11, 12). For throughput prediction, FL converges similarly to CL in the first rounds of the experiments. For throughput prediction, CL outperforms FL by an average of 9.37% across all rounds (see Fig. 11). In fact, CL exhibits a maximum performance enhancement of 27.73% against FL. Interestingly though, FL achieves outperforms CL by 9.16%, during the experiment's last round. Unlike throughput, FL for delay prediction achieves a very similar performance to that of CL (see Fig. 12); on average CL outperforms FL by 2.08% across all rounds. FL tracks the performance of CL in the course of the time (rounds), even achieving better inference results (lower RMSE values) in certain instances. Overall our preliminary results suggest that FL can achieve similar accuracy levels to that of its centralized alternative, whilst preserving data privacy.

V. Conclusions

We have presented NordicDat, a dataset collected in the cross-border area of Finland, Norway and Sweden over the course of two weeks. The dataset contains a rich set of features that are related to cellular QoS values (*e.g.,* throughput, delay), network characteristics (*e.g.,* cell, operator) and vehicle kinematics (*e.g.,* speed, location). As such, it can be utilized to effectively train (classical and distributed) AI/ML models for a diverse set of use-cases *e.g.,* trajectory, handover or QoS prediction; the latter being identified as a key enabler for future automotive applications. Our data analysis has revealed non-linear correlations between the dataset's features and the involved QoS values. Nevertheless both roaming and speed, as well as radio access technology are identified as key factors that shape the QoS values' patterns across time and space. Finally, we have showcased distributed AI/ML for QoS (throughput and delay) prediction. To our knowledge that is the first demonstration of distributed pQoS on real-world data. Our preliminary results suggest that distributed pQoS achieves similar performance to that of classical ML; classical ML outperforms distributed ML by 10% and 28% (average and maximum values, respectively). Future research directions include additional measurements to capture diverse scenarios *e.g.,* urban traffic cases, as well as large-scale campaigns with multiple vehicles.

Acknowledgment

This work has been supported by the European Union's Horizon 2020 Hi-Drive Innovation Action, under grant agreement no. 101006664.

References

[1] "Predictive QoS and V2X service adaptation," Technical Report, 5GAA Automotive Association, 2022.

[2] P. Zhao *et al.*, "The case design of testing vehicle platooning in a vehicle-road cooperation system for a 5G cross-border scenario," in *IEEE 4th Int'l Conference on Civil Aviation Safety and Information Technology (ICCASIT)*. IEEE, 2022, pp. 135–139.

[3] "Making 5G proactive and predictive for the automotive industry," Industry White Paper, 5GAA Automotive Association, 2020.

[4] A. Palaios *et al.*, "Machine learning for QoS prediction in vehicular communication: Challenges and solution approaches," *IEEE Access*, 2023.

[5] S. Schwarzmann *et al.*, "Ml-based qoe estimation in 5g networks using different regression techniques," *IEEE Transactions on Network and Service Management*, vol. 19, no. 3, pp. 3516–3532, 2022.

[6] S. Barmpounakis *et al.*, "LSTM-based QoS prediction for 5g-enabled connected and automated mobility applications," in *2021 IEEE 4th 5G World Forum (5GWF)*. IEEE, 2021, pp. 436–440.

[7] S. Nakaniida and T. Fujii, "Throughput prediction by radio environment correlation recognition using crowd sensing and federated learning," in *IEEE Intl. Conf. on AI in Information and Comm.*, 2022, pp. 366–371.

[8] A. Alabbasi *et al.*, "On cascaded federated learning for multi-tier predictive models," in *IEEE ICC Workshops*, 2021, pp. 1–7.

[9] D. Raca *et al.*, "Beyond throughput: A 4g lte dataset with channel and context metrics," in *Proceedings of the 9th ACM multimedia systems conference*, 2018, pp. 460–465.

[10] D. Xu *et al.*, "Understanding operational 5G: A first measurement study on its coverage, performance and energy consumption," in *Proceedings of the Annual conference of the ACM Special Interest Group on Data Communication on the applications, technologies, architectures, and protocols for computer communication*, 2020, pp. 479–494.

[11] S. Farthofer *et al.*, "An open mobile communications drive test data set and its use for machine learning," *IEEE Open Journal of the Communications Society*, vol. 3, pp. 1688–1701, 2022.

[12] A. M. Mandalari *et al.*, "Experience: Implications of roaming in europe," in *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking*, 2018, pp. 179–189.

[13] T. Miekkala *et al.*, "Nordicdat: A cross-border predictive qos dataset," 2024. [Online]. Available: https://zenodo.org/records/10964584

[14] A. Palaios *et al.*, "Network under control: Multi-vehicle E2E measurements for ai-based qos prediction," in *IEEE 32nd Annual Int'l Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*, 2021, pp. 1432–1438.

[15] R. Hernangómez *et al.*, "Berlin v2x: A machine learning dataset from multiple vehicles and radio access technologies," in *IEEE 97th Vehicular Technology Conference (VTC2023-Spring)*, 2023, pp. 1–5.

[16] D. Schäufele *et al.*, "Terminal-side data rate prediction for high-mobility users," in *IEEE 93rd Vehicular Technology Conference (VTC2021-Spring)*, 2021, pp. 1–5.

[17] A. Narayanan *et al.*, "A first look at commercial 5G performance on smartphones," in *Procs. of The Web Conference 2020*, pp. 894–905.

[18] A. Narayanan, E. Ramadan *et al.*, "Lumos5G: Mapping and predicting commercial mmwave 5G throughput," in *Proceedings of the ACM Internet Measurement Conference*, 2020, pp. 176–193.

[19] Y. Liu and C. Peng, "A close look at 5G in the wild: Unrealized potentials and implications," in *IEEE INFOCOM 2023*, pp. 1–10.

[20] A. Narayanan *et al.*, "A variegated look at 5G in the wild: performance, power, and QoE implications," in *ACM SIGCOMM'21*, pp. 610–625.

[21] K. Kousias *et al.*, "A large-scale dataset of 4g, nb-iot, and 5G non-standalone network measurements," *IEEE Comm. Magazine*, 2023.

[22] A. Narayanan *et al.*, "A comparative measurement study of commercial 5G mmwave deployments," in *IEEE INFOCOM 2022-IEEE Conference on Computer Communications*. IEEE, 2022, pp. 800–809.

[23] R. A. Fezeu *et al.*, "An in-depth measurement analysis of 5G mmwave phy latency and its impact on end-to-end delay," in *Int'l Conf. on Passive and Active Network Measurement*. Springer, 2023, pp. 284–312.

[24] T. Tsourdinis *et al.*, "Ue network traffic time-series (applications, throughput, latency, cqi) in lte/5G networks," 2022. [Online]. Available: https://dx.doi.org/10.21227/4ars-fs38

[25] A. Safari Khatouni *et al.*, "An open dataset of operational mobile networks," in *Proceedings of the 18th ACM Symposium on Mobility Management and Wireless Access*, 2020, pp. 83–90.

[26] R. A. Fezeu *et al.*, "Roaming across the european union in the 5G era: Performance, challenges, and opportunities," in *IEEE International Conference on Computer Communications*, 2024.

[27] F. Burmeister *et al.*, "Measuring time-varying industrial radio channels for d2d communications on agvs," in *2021 IEEE Wireless Communications and Networking Conference (WCNC)*. IEEE, 2021, pp. 1–7.

[28] R. Hernangómez *et al.*, "Towards an ai-enabled connected industry: Agv communication and sensor measurement datasets," *arXiv preprint arXiv:2301.03364*, 2022.

[29] V. Raida *et al.*, "Real world performance of lte downlink in a static dense urban scenario-an open dataset," in *2020 IEEE Global Communications Conference*, pp. 1–6.

[30] W. Lawrenz, *CAN System Engineering From Theory to Practical Applications*. Springer Science Business Media, 2013.

[31] R. B. Langley, "RTK GPS," in *GPS World*, 1998, pp. 70–76.

[32] "AT commands," https://wiki.teltonika-networks.com/view/AT_Commands, accessed: 2024-03-30.

[33] "Linux ifstat," https://man7.org/linux/man-pages/man8/ifstat.8.html, accessed: 2024-03-30.

[34] "Linux ping command," https://linux.die.net/man/8/ping, accessed: 2024-04-22.

[35] "Data distribution service specification version 1.4." "https://www.omg.org/spec/DDS/1.4/About-DDS", accessed: 2024-03-30.

[36] R. Hernangomez *et al.*, "Berlin v2x," 2022. [Online]. Available: https://dx.doi.org/10.21227/8cj7-q373

[37] J. Brownlee, *Deep learning for time series forecasting: predict the future with MLPs, CNNs and LSTMs in Python*. ML Mastery, 2018.

[38] A. Reisizadeh *et al.*, "Fedpaq: A communication-efficient federated learning method with periodic averaging and quantization," in *Int'l Conf. on Artificial Intelligence and Statistics*, 2020, pp. 2021–2031.

[39] S. Mozaffari *et al.*, "Deep learning-based vehicle behavior prediction for autonomous driving applications: A review," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 1, pp. 33–47, 2020.