



# DL-enabled Road Traffic Monitoring Services using 6G mmWave Channel State Information

Marcus Haferkamp, Simon Häger, and Christian Wietfeld

Communication Networks Institute (CNI), TU Dortmund University, 44227 Dortmund, Germany

E-mail: {Marcus.Haferkamp, Simon.Haeger, Christian.Wietfeld}@tu-dortmund.de

**Abstract**—Future wireless networks will integrate sensing functionality to improve communication performance and offer services for verticals. An example use case of 6G sensing could be in the context of future intelligent transportation systems (ITSs) aiming for versatile benefits, including environmental and socioeconomic factors. To this end, accurate road traffic information acquired in real-time by ubiquitous cellular networks will be paramount to impose instant measures and develop data-driven long-term solutions. This paper showcases the wireless radio channel as a privacy-preserving environmental sensor intrinsically affected by its surroundings, facilitating edge intelligence for road traffic monitoring. Notably, we consider millimeter-wave (mmWave) deployments owing to their propagation and communication system characteristics for sensing, i.e., low penetration and diffraction combined with directional beamforming antennas and high bandwidth. We employ real-world mmWave channel traces gathered in an urban road traffic scenario to demonstrate its suitability for two use cases well suited for future 6G sensing services. Utilizing machine learning (ML) and deep learning (DL) models, we demonstrate road user classification accuracy of up to 99.8 % and 97.8 % for binary and three-class traffic counting, respectively. In the second use case, we outline the feasibility of mmWave-enabled estimation of passing vehicles' moving direction with up to 92.3 %.

**Index Terms**—6G, integrated sensing and communication, road traffic monitoring, millimeter-wave, channel state information, machine learning.

## I. MOTIVATING 6G SENSING FOR ITSS

Cellular networks offer a versatile, high-speed, low-latency, and robust communication infrastructure. Whereas there is an almost ubiquitous 5G coverage at sub-6 GHz frequencies, ultra-high connectivity may be provided in urban hotspots using mmWave carriers. Turning towards future 6G, the network infrastructure will become perceptive of its environment using inherent radar-like sensing functionalities [1] to improve communication performance and serve verticals within the scope of integrated communication and sensing (ICAS). High-performance sensing services are expected at mmWave frequencies because of changed propagation characteristics, the broad bandwidth for high time, and the large beamforming antenna arrays for high angular resolution [2, 3].

The seamless integration of commercial-grade user positioning introduced with 5G can be seen as the forerunner of the 6G ICAS vision [4]. By integrating multiple such services, the spectrum and cost efficiency of 6G networks can be improved while making some dedicated systems obsolete [1, 5] or enhancing their performance by sensor fusion [6]. With numerous research questions regarding bistatic radar-like sensing capabilities in 6G yet to be

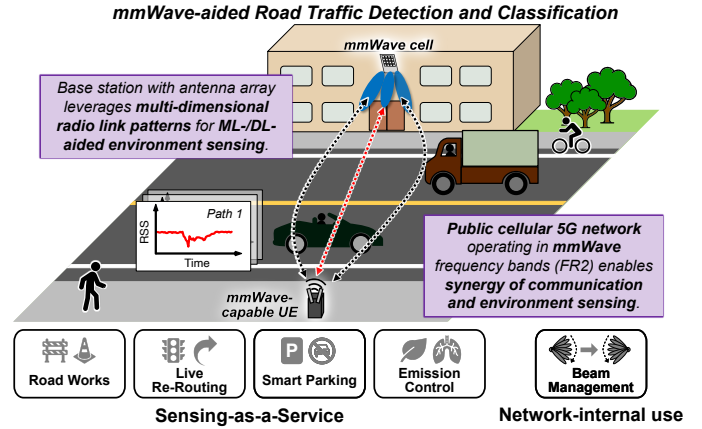


Fig. 1. 6G networks shall become environment-perceptive, making dedicated solutions obsolete. For example, mmWave frequencies facilitate privacy-preserving urban road traffic monitoring in the scope of ITS.

addressed [7–9], exploiting regular channel statistics in monostatic deployments has become a hot research topic. Recent research has established that communication-centric, bistatic system deployments are also suitable for ICAS applications. For example, the authors in [10] demonstrate that a two-node mmWave-based system can predict link blockages as a form of environment monitoring within an urban vehicular scenario to re-align the beam for a robust communication link. In this paper, we utilize extensive mmWave measurement data provided in [10] to perform DL-aided vehicle type classification and movement direction estimation for ITSs.

Similarly, so-called radio-based sensing techniques (radio tomography) at sub-6 GHz frequencies leverage wireless channel state information (CSI) between base stations (BSs) and connected user equipments (UEs) to successfully detect road users [11] with favorable system properties in terms of accuracy, privacy, cost, and robustness compared to alternative approaches [12]. Contrary to vision-based solutions, its advantage for vehicle classification is the synergy of communication and sensing services, potentially facilitating a cost-efficient, large-scale and privacy-preserving operation. Overall, the design of such systems is a trade-off between economic and performance-related needs, making dedicated solutions indispensable for some use cases. Transferring this concept to future 6G mmWave networks, Fig. 1 illustrates road traffic surveillance in an urban setting by leveraging the received signal strength (RSS) metric. Whenever the line-of-sight (LOS) link between the BS and the road-side unit (RSU)

UE is (partially) obstructed by passing road users, the BS gathers characteristic channel patterns caused by signal attenuation and reflection along the distinct spatial direction defined by the beamforming antenna array state. The acquired multidimensional data enable BSs or downstream core network (CN) components to deduce comprehensive insights into regional road traffic patterns [3]. Real-time data, such as traffic volume and speed, allow instant measures, such as dynamic traffic re-routing or vulnerable road user warnings [13]. Furthermore, this research has the potential to significantly support urban planners, providing them with the data they need to make informed decisions and implement targeted traffic regulations and mobility concepts. This, in turn, could lead to significant environmental, economic, and public health benefits.

Against this background, this work contributes the following practical aspects to the current state of research on mmWave-aided ICAS for road traffic monitoring.

- We utilize a large-scale, real-world mmWave measurement dataset [10] to showcase the suitability of spatiotemporal RSS link patterns for DL-enabled, network-external ITS services. This contrasts with [14], which focused on LOS link blockage prediction as a network-internal sensing service.
- We present a multistep ML/DL-aided classification and direction estimation approach for different road users, leveraging only in-band RSS patterns.
- We evaluate convolutional neural network (CNN), deep neural network (DNN), recurrent neural network (RNN), and Random Forest (RF) models to classify heterogeneous road users into (i) two and (ii) three classes. In addition, we provide an outlook on estimating moving directions.

The remainder of this paper is as follows. In Sec. II, we describe the methodology for acquiring and pre-processing the mmWave channel measurement data, followed by the neural architecture search (NAS) process to identify suitable ML models for the different tasks. We then discuss the prediction performance of the best-performing models in Sec. III for (i) the two classification tasks and (ii) their suitability for estimating the driving direction of vehicles. Finally, Sec. IV concludes with the key findings of our work.

## II. METHODOLOGY

This section describes the mmWave dataset, data pre-processing, and estimator training. In Sec. II-A, we briefly introduce the dataset and data pre-processing following [14] to preserve maximum performance comparability, followed by the taxonomy used for vehicle classification. Sec. II-B explains the model architecture search applied to find high-performance yet less complex DL models.

### A. Real-world mmWave Dataset

This work uses the open-data and real-world mmWave measurement datasets provided in [10]. For gathering mmWave data, the authors used a fixed multi-modal sensor setup including 60 GHz BS and UE, closely positioned on opposite sides of a two-way road, to obtain LOS link blockage patterns induced by passing road users. The BS uses a 16-element uniform

linear array (ULA) with a 64-beam codebook, performing beam steering with  $1.4^\circ$  azimuth steps. Tab. I summarizes relevant information about the measurement setup, the dataset, and the applied data pre-processing steps.

Fig. 2 shows an example of a normalized mmWave power pattern over time gathered during a vehicle passage (red-colored region), prepared as a heatmap (*top*) and line plot (*bottom*). The heatmap reveals that the center beams carry much of the total received signal power due to the ULA performing horizontal beam sweeping. The center beams point towards the UE, resulting in higher RSS levels. The bottom plot depicts the mean received signal power over time. The mean RSS is relatively constant before and after the link blockage caused by the passing vehicle. But imminent to the blockage, the power pattern exposes a peak. Consistent with [14], we call this phenomenon the *pre-blockage signature*, mainly due to constructive and destructive interference induced by passing vehicles. During the passage, road users function as scatterers, focusing the electromagnetic (EM) waves on the direct propagation path between BS and UE before turning this effect into the opposite by blocking the path. Since this signature is characteristic of different vehicle types, depending on their specific shapes, materials, etc., we leverage this to classify vehicle types (*radio-fingerprinting*).

Before the classification step (cf. Sec. III-A), we identify and label vehicle passages in the continuous mmWave data streams using synchronized image data from the dataset. Because we intend to analyze the significance of different

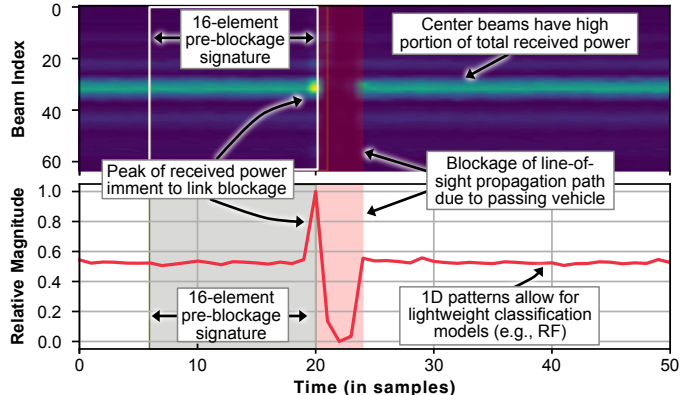


Fig. 2. Example of a 2D multi-beam mmWave power pattern affected by a passing road user (*top*). The 1D representation (*bottom*) reveals the two characteristic components of such a power pattern: the actual blockage and the imminent pre-blockage signature.

TABLE I. OUTDOOR MMWAVE MEASUREMENT SETUP [14] AND DATA PRE-PROCESSING FOR THE MACHINE LEARNING.

Parameter	Description
Deployment	Stationary receiver (RX) and transmitter (TX)
Data modalities	<b>60 GHz mmWave RSS</b> , GPS, Images
mmWave RX	16-element phased array (ULA)
mmWave TX	using <b>64-beam codebook</b>
Scenario data	Omni-directional
	Scenarios 17-21 (1,731 passage sequences)
Data pre-processing	Standardization, Beam-wise averaging, Edge-padding (only complete sequences), Data augmentation (10 dB)

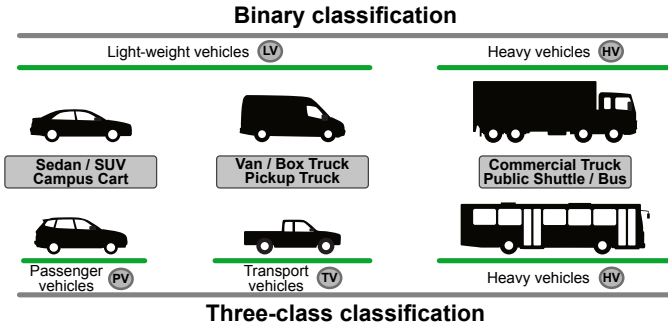


Fig. 3. Taxonomy used for binary and three-type vehicle classification in accordance with [15]. The former task focuses on LV and HV (top), whereas the latter conducts a more fine-grained distinction of PV, TV, and HV (bottom).

sequence types and beam sets in Sec. III, we prepare six dataset variations with different input sequence lengths and beam combinations for each classification task. In particular, input data differs regarding the input sequence type (complete vs. pre-blockage signature) and number of beams (64, 54, and 10), respectively. Hereafter, we standardize this data for an efficient model learning by subtracting its mean  $\mu$  and dividing by its standard deviation  $\sigma$ . Using edge-padding, we gather equal-sized sequences required by models like DNN and RF. Because the original dataset has sets of varying cardinality for each vehicle class, we mitigate this imbalance using data augmentation. For this, we duplicate random samples of the minority classes and add a 10 dB white Gaussian noise. This leads to datasets with 3,210 samples for the binary and 4,023 samples for the three-type classification tasks, respectively. We average the multi-beam magnitude data to enhance the model evaluation, resulting in flattened sequences suitable for ML models such as RFs. Similar to [14], we group the various identified vehicle-type passages into  $N_{BC} = 2$  and  $N_{TC} = 3$ , leaving out human (non-vehicle) road users, as depicted in Fig. 3. For the binary classification task ( $N_{BC} = 2$ ), we divide all labeled vehicle passages into the classes light-weight vehicles (LV) and heavy vehicles (HV). Analogously, we group all labeled passages into  $N_{TC} = 3$  classes, splitting the LV class into passenger vehicle (PV) and transport vehicles (TV) classes.

### B. Neural Architecture Search for Deep Learning Models

Finding high-performance yet low-complexity model architectures is time-consuming and computationally expensive. Thus, we utilize the NAS framework provided by [16], which builds up on *TensorFlow*, to discover suitable CNN, DNN, and RNN models. To enhance the model search, we apply the hyperparameter optimization approach *Hyperband* [17], speeding up randomized hypertuning through adaptive resource allocation and early stopping. As a performance baseline, we also provide the results for RF, utilizing *scikit-learn*'s model and randomized parameter search. Fig. 4 illustrates the NAS process: a search algorithm finds model architectures by incrementally stacking neural network building blocks  $B$  (e.g., recurrent or convolutional layers) to a set of architecture candidates  $C$  before applying random mutations.

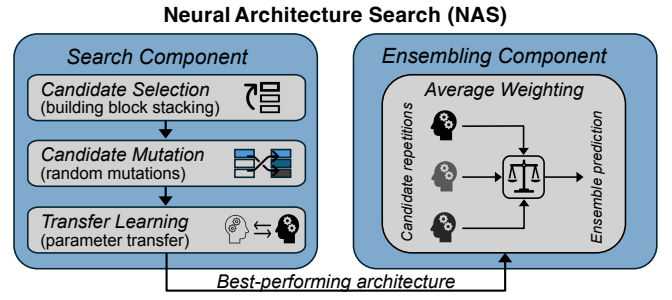


Fig. 4. We use NAS to find well-performing DL models. The process incorporates two main components: the search component for building and mutating candidate architectures and the ensembling component for average-weighting multiple candidate repetitions.

More specifically, the NAS algorithm incorporates a search and a mutation component. The former component finds new candidates  $c \in C$  by stacking up to  $k$  blocks and iteratively mutating them. The mutation component invokes random mutations to the best-performing candidate  $c$  with the minimum loss, possibly increasing its network depth to a user-defined depth limit  $k$ . Before the mutation algorithm starts training a new architecture, transfer learning ensures that those parameters of blocks found in the original architecture and its mutated derivative are preserved for further exploitation. Last, an ensembling module  $A$  combines multiple copies of a specific candidate, retraining those with different shuffled data and initialization parameters. The final model is the average weighted ensemble of the model instances.

Fig. 5 depicts the best-performing DNN model architectures when using complete sequences with signal power data of 64 antenna beams. The figure's top row (a) shows the DNN architecture found for the binary classification task, whereas the bottom row (b) depicts the corresponding one for the three-type vehicle classification task. Both models stack multiple fully connected layer (FCN) blocks, each incorporating a flatten layer, a regular densely connected layer, and a rectified linear unit (ReLU) layer. Both models mainly differ regarding the count of FCN blocks and the intermediate dropout layers.

### III. MODEL PERFORMANCE EVALUATION

This section presents the best-performing models found and their specific evaluation results for different input data. First, we present the vehicle type classification results for the binary




Fig. 5. Identified DNN model architectures for classifying road users into (a) two and (b) three classes, utilizing complete passage signatures with mean power samples of 64 antenna beams.



and the three-type classification tasks in Sec. III-A. Then, we give a short outlook on the suitability of the proposed approach for estimating the motion direction of passing vehicles in Sec. III-B.

### A. Multi-class Vehicle Type Classification

Tab. II summarizes the accuracy results of the best-performing models, using pre-blockage and complete signatures and different sets of beam magnitudes as input data.

TABLE II. ACCURACY RESULTS FOR BINARY AND THREE-TYPE VEHICLE CLASSIFICATION USING DL AND ML MODELS FOR DIFFERENT SEQUENCE TYPES AND ANTENNA BEAM CONFIGURATIONS.

# Vehicle Classes	Sequence Type	# Antenna Beams	[14] <sup>1</sup>	Model Accuracy [%]			
				CNN	DNN	RNN	RF <sup>2</sup>
LV HV	Pre-blockage Signature <sup>1</sup> (16 time steps)	10 (inner)	-	55.3	92.6	55.3	98.8
		54 (outer) <sup>1</sup>	81.5	56.5	96.8	62.5	99.0
LV HV	Complete Signature <sup>2</sup> (87 time steps)	10	-	79.1	93.8	89.3	99.2
		54	-	90.7	97.4	90.8	99.3
PV TV HV	Pre-blockage Signature <sup>1</sup>	10	-	63.6	82.9	65.8	91.3
		54	63.7	61.2	87.1	65.3	91.1
PV TV HV	Complete Signature <sup>2</sup>	10	-	83.7	96.1	87.7	93.6
		54	-	86.2	97.8	88.0	93.8
		64	-	88.3	96.3	89.8	93.6

Best model. Worst model.

<sup>1</sup>Highest accuracy according to related work [14] <sup>2</sup>Used in own work [11, 12]

We can determine a performance gap for the binary and three-type classification tasks depending on the sequence type data (complete vs. pre-blockage signature) rather than the number of labels to predict. Specifically, there is a notable gap for CNN and RNN, where the model-specific maximum accuracy scores differ by at least 20 percentage points, respectively. On the contrary, the performances vary less significantly for DNN and RF. With the specific model architectures in mind, the former models try to exploit the timely relationship between sequence steps, whereas the latter are not aware of this. This implies that the architectures of the CNN and RNN models are inattentive to the given input data. Regarding the binary classification task, the evaluated model architectures achieve maximum prediction accuracies using complete signatures with 64 beam power data. All models exhibit decreased performances when this number of antenna beam data is reduced. This performance drop is significant for CNN (13.6%) and quite notable for DNN (6%), whereas it is negligible for RNN (1.8%) and RF (0.2%). Compared to the binary classification, some models' maximum prediction performances in the three-class classification task are slightly lower. Specifically, this applies to CNN and RNN, which achieve up to 88.3% and 89.8% accuracy. The prediction performances of DNN (97.8%) and RF (93.8%) are 2 and 5.6 percentage points lower than for the binary task, yet well above 90% accuracy.

A possible reason is that the CNN and RNN models holistically grasp the relationship between the duration of RSS drops and the inherent vehicle dimensions. Otherwise, it seems difficult for them to extract such meaningful correlations from the shorter 16-step pre-blockage signature sequences, characterized by large shares of constant RSS values, followed

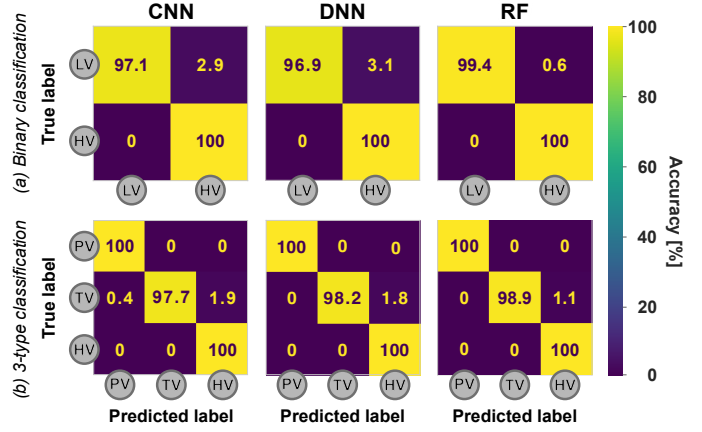


Fig. 6. Confusion matrices of ML models found for (a) binary and (b) three-type vehicle type classification. The normalized predicted true conditions for both tasks show generally high prediction performances.

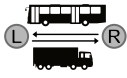
by a short-time peak of signal power (cf. Fig. 2).

The confusion matrices in Fig. 6 show the class-specific accuracy performances of the CNN, DNN, and RF models for both classification tasks, using complete signature traces with power data of 64 beams. For the binary classification task (a), the matrices confirm the high classification performance of the provided CNN, DNN, and RF models. All models tend to misclassify some LVs as HVs regarding the normalized actual labels. This effect is more pronounced for the DL models than for RF, possibly due to its ensemble nature, fostering an improved predictive accuracy and overfitting control. Otherwise, all models classify HVs correctly. This prediction accuracy regarding HVs can also be observed for all models in the three-type classification task, depicted in the bottom row of Fig. 6. Contrary to the binary task for LVs, all models achieve 100% accuracy in identifying PVs. But now, a few vehicles related to the class TV are challenging for all models, leading to a degraded normalized performance of 97.7% (CNN), 98.2% (DNN), and 98.9% (RF). In all these cases, the models misclassify the TVs as HVs, revealing uncertainties when classifying vehicles belonging to either one of these classes.

### B. Exploring Potentials: Moving Direction Estimation

In addition to the previously shown vehicle type classification, trajectory determination of passing vehicles, such as motion direction and traveled lane, is also of primary interest in modern ITSs. With the help of real-time high-resolution road traffic data, traffic models may help improve current and future road traffic by empowering optimization services like traffic re-routing and planning. This subsection gives an outlook on the predictive performances of the previously discussed model architectures regarding estimating passing vehicles' moving directions. Since we know from previous work that determining the motion direction using a single radio link is challenging, we focus on specific beam configurations and vehicle types for this task. Specifically, we choose three different subsets of beam configurations, each incorporating ten adjacent beams aligned to the left, center, and right, to exploit potential spatial effects. Moreover, we limit this

TABLE III. ACCURACY OF VEHICLE MOVEMENT DIRECTION ESTIMATION USING ML-BASED MMWAVE TOMOGRAPHY AS PROPOSED IN [3].

# Passage Directions	Sequence Type	# Antenna Beams	Accuracy [%]		
			CNN	DNN	RF
	Pre-blockage Signature (16 time steps)	10 (inner)	76.9	84.6	61.5
		10 (left)	69.2	84.6	53.8
		10 (right)	61.5	84.6	57.0
	Complete Signature (87 time steps)	10 (inner)	76.9	76.9	69.2
		10 (left)	84.6	76.9	61.5
		10 (right)	76.9	92.3	76.9

analysis to HVs because they exhibit more pronounced signal reflection and scattering characteristics due to their specific large-scale front faces. We omit the data augmentation step, as the used dataset includes 125 passage sequences with a balanced share of both driving directions (0: *left*  $\rightarrow$  *right*, 1: *right*  $\rightarrow$  *left*). Tab. III lists the prediction accuracy for CNN, DNN, and RF models, using complete and pre-blockage signature sequences for different sets of beam power data.

We generally observe that the maximum accuracies achieved are significantly lower than those achieved for the vehicle classification studied in Sec. III-A. This applies, in particular, to RF and DNN, which now yield less than 90% accuracy, with one exception when using DNN (92.3%). Moreover, the accuracy performance of the models fluctuates for the considered antenna beam configurations, regardless of whether using complete or pre-blockage signature data. This behavior differs from the prior classification task, where we observe that using complete signature traces is beneficial in most cases. Similar to the classification task, we cannot determine a preferred antenna beam configuration, leading to superior prediction performances for pre-blockage and complete sequence data. These promising insights confirm the newly considered direction estimation with a single mmWave link is a more challenging task, which is worth further investigation in our future work.

#### IV. CONCLUSIONS AND OUTLOOK

In the scope of ongoing 6G research in the field of ICAS, the exploitation of native wireless channel information is expected to enable perceptive networks whose sensing services are, for example, crucial for future ITSs. This paper uses real-world channel data to examine the potential of multi-beam mmWave CSI for privacy-preserving multi-class road user classification and movement direction estimation.

The takeaways of this paper are as follows. Contrasting previous works that leverage channel data for radio access network (RAN)-internal optimization, e. g., by mmWave link blockage predictions [14], we highlight sensing-as-a-service perspectives. Specifically, we evaluate DL and ML models for multi-class vehicle classification and provide an outlook on passage direction estimation. The results show up to 99.8% and 97.8% accuracy for binary and three-type road user classification. Regarding movement direction estimation, prediction accuracy is achieved by up to 92.3% for heavy road users (e. g., buses).

Generally, the high performance of both tasks underlines

the suitability and, thus, the importance of beamforming-based angular information available at mmWave frequencies. The results prove that a multi-class vehicle type classification allows for very high prediction performances.

In contrast, vehicle moving direction estimation is promising but challenging when leveraging only a single radio link. To improve direction estimation performance, future work should address distributed multi-node system configurations or combining different approaches (sensor fusion).

These promising results for both tasks encourage further experimental performance analyses for selected commercial use cases covering more complex settings like multi-lane roads. Besides the performance aspects, practical factors are of interest, tackling computational and, thus, economic costs.

#### ACKNOWLEDGMENT

This work has been funded by the German Federal Ministry of Education and Research (BMBF) in the course of the 6GEM Research Hub under the grant number 16KISK038.

#### REFERENCES

- [1] F. Liu, Y. Cui, C. Masouros, J. Xu *et al.*, "Integrated sensing and communications: Toward dual-functional wireless networks for 6G and beyond," *IEEE J. Sel. Areas Commun.*, vol. 40, no. 6, pp. 1728–1767, Jun. 2022.
- [2] A. Ali, N. Gonzalez-Prelcic, R. W. Heath, and A. Ghosh, "Leveraging sensing at the infrastructure for mmWave communication," *IEEE Commun. Mag.*, vol. 58, no. 7, pp. 84–89, Jul. 2020.
- [3] S. Häger, M. Haferkamp, and C. Wietfeld, "Beam-based 6G networked sensing architecture for scalable road traffic monitoring," in *Proc. IEEE SysCon*, Apr. 2023.
- [4] S. Häger, N. Gratza, and C. Wietfeld, "Characterization of 5G mmWave high-accuracy positioning services for urban road traffic," in *Proc. IEEE VTC-Spring*, Jun. 2023.
- [5] J. A. Zhang, F. Liu, C. Masouros, R. W. Heath *et al.*, "An overview of signal processing techniques for joint communication and radar sensing," *IEEE J. Sel. Top. Signal Process.*, vol. 15, no. 6, pp. 1295–1315, Nov. 2021.
- [6] S. Wu, C. Chakrabarti, and A. Alkhateeb, "Proactively predicting dynamic 6G link blockages using LiDAR and in-band signatures," *IEEE Open J. Commun. Soc.*, vol. 4, pp. 392–412, Jan. 2023.
- [7] A. Liu, Z. Huang, M. Li, Y. Wan, *et al.*, "A survey on fundamental limits of integrated sensing and communication," *IEEE Commun. Surv. Tutor.*, vol. 24, no. 2, pp. 994–1034, May 2022.
- [8] A. Kadelka, G. Zimmermann, J. Plachý, and O. Holschke, "A CSP's view on opportunities and challenges of integrated communications and sensing," in *Proc. IEEE JC&S*, Mar. 2023.
- [9] W. Yuan, Z. Du, X. Meng, F. Liu, and C. Masouros, *Integrated Sensing and Communication for Vehicular Networks*. Springer Nature Singapore, Jul. 2023, pp. 419–444.
- [10] A. Alkhateeb, C. Charan, T. Osman, A. Hredzak *et al.*, "DeepSense 6G: A large-scale real-world multi-modal sensing and communication dataset," *IEEE Commun. Mag.*, vol. 61, no. 9, pp. 122–128, Sep. 2023.
- [11] M. Haferkamp, B. Sliwa, and C. Wietfeld, "A low cost modular radio tomography system for bicycle and vehicle detection and classification," in *Proc. IEEE SysCon*, Apr. 2021.
- [12] B. Sliwa, N. Piatkowski, and C. Wietfeld, "The channel as a traffic sensor: Vehicle detection and classification based on radio fingerprinting," *IEEE IoT J.*, vol. 7, no. 8, pp. 7392–7406, Mar. 2020.
- [13] Ericsson. [Online]. Available: [www.ericsson.com/en/blog/2021/10/joint-sensing-and-communication-6g](http://www.ericsson.com/en/blog/2021/10/joint-sensing-and-communication-6g) (Accessed 2024-06-14).
- [14] S. Wu, M. Alrabeiah, C. Chakrabarti, and A. Alkhateeb, "Blockage prediction using wireless signatures: Deep learning enables real-world demonstration," *IEEE Open J. Commun. Soc.*, vol. 3, pp. 776–796, Mar. 2022.
- [15] United Nations (UN) Economic and Social Council (ECE), "World forum for harmonization of vehicle regulations; resolution on the construction of vehicles; ECE/TRANS/WP.29/78/Rev.6," Jul. 2017.
- [16] H. Mazzawi, X. Gonzalvo, A. Kracun, P. Sridhar *et al.*, "Improving keyword spotting and language identification via neural architecture search at scale," in *Proc. ISCA Interspeech*, Sep. 2019.
- [17] L. Li, K. Jamieson, G. DeSalvo, A. Rostamizadeh, and A. Talwalkar, "Hyperband: A novel bandit-based approach to hyperparameter optimization," *J. Mach. Learn. Res.*, vol. 18, no. 1, p. 6765–6816, Jan. 2017.