Advances in Science and Technology Research Journal

Advances in Science and Technology Research Journal, 2026, 20(7), 82–94
https://doi.org/10.12913/22998624/218792
ISSN 2299-8624, License CC-BY 4.0

Received: 2025.12.22
Accepted: 2026.05.15
Published: 2026.06.01

Efficient radar signal classification using wavelet features and machine learning for embedded systems

Anna Ślesicka[1]

[1] Institute of Navigation, Polish Air Force University, 08-521 Dęblin, Poland
E-mail: a.slesicka@law.mil.pl

ABSTRACT

This paper presents a computationally efficient approach for classifying moving road traffic objects using FMCW radar. The method operates directly on raw radar IQ signals, applying Continuous Wavelet Transform (CWT) and extracting simple statistical features (mean, standard deviation, and maximum) to form compact feature vectors. Classification is performed using lightweight algorithms such as Random Forest, SVM, or kNN, achieving high accuracy on the test set (94–95% for six object classes) while maintaining minimal computational overhead. Unlike image-based CNN methods, the approach eliminates time-consuming spectrogram generation, enabling fast training and prediction suitable for real-time and embedded applications. Comparisons with selected deep learning–based methods reported in the literature indicate that the proposed framework can achieve comparable classification accuracy while requiring substantially lower computational resources. FMCW radar provides robustness against adverse weather and lighting conditions and enables detection independent of optical visibility. This work demonstrates a practical approach for efficient radar-based signal processing and pattern recognition, offering an alternative to vision-based classification for intelligent traffic monitoring and autonomous systems.

Keywords: object classification, machine learning, radar measurements, signal processing, time–frequency analysis, wavelet transform.

INTRODUCTION

Modern traffic monitoring systems and intelligent transportation systems (ITS) increasingly demand effective and reliable methods for detecting and classifying road users. Among available sensing technologies, frequency-modulated continuous wave (FMCW) radars operating in the millimeter-wave band are gaining prominence due to their robustness under adverse weather conditions, low cost, and suitability for real-time operation [1]. Traditionally, radar systems have been primarily used to estimate physical parameters such as distance or velocity, while object classification has been largely addressed using vision-based sensors.

Recent advances in artificial intelligence, particularly deep learning, have led to the development of highly effective classification models, which are predominantly applied to data obtained from cameras and LiDAR sensors [2]. Although radar sensors are increasingly integrated into perception systems, their standalone use as input for machine learning-based classification remains limited, with classical signal processing approaches still dominating radar applications.

Most existing radar-based object classification studies rely on deep learning architectures such as convolutional neural networks, recurrent neural networks, or transfer learning approaches [3–6]. These methods typically operate on time–frequency representations generated from radar signals using STFT or FFT and analyzed by deep architectures including VGG-16, ResNet-50, or DarkNet-53. Representative studies report high classification accuracy for pedestrians, vehicles, cyclists, and composite object classes

[3–6]. However, despite their effectiveness, such approaches are computationally demanding, require time-consuming spectrogram generation, and often rely on large publicly available datasets [5,6], which limits their applicability in real-time and embedded radar systems.

As an alternative, classical machine learning methods applied to radar-derived features offer more computationally efficient solutions suitable for resource-constrained platforms. Examples include kNN-based classification of micro-Doppler features [7] and probabilistic approaches such as Bayesian networks [8]. Other studies employ multi-sensor fusion by combining FMCW radar with RGB or infrared cameras to enhance detection performance [2]. While sensor fusion improves robustness in challenging conditions, it significantly increases system complexity, computational requirements, and implementation costs.

Motivated by these limitations, this study investigates whether road traffic participants can be reliably classified directly from raw FMCW radar IQ signals without generating time–frequency images. The proposed approach evaluates the effectiveness of continuous wavelet transform (CWT)-based feature extraction for capturing relevant signal dynamics and applies lightweight machine learning methods, particularly Random Forest, to distinguish six common traffic object classes. The method is validated using real-world radar measurements collected across multiple traffic scenarios and compared with kNN and SVM classifiers, demonstrating competitive classification performance with substantially reduced computational complexity and suitability for real-time and embedded applications.

The main contributions and achievements of this work include the development of a computationally efficient radar-based object classification framework operating directly on raw FMCW IQ signals without spectrogram generation, the application of CWT combined with simple statistical features (mean, standard deviation, and maximum) to obtain compact and interpretable feature vectors, a comprehensive experimental evaluation using real-world traffic data collected across multiple scenarios, and a comparative analysis of lightweight machine learning classifiers demonstrating high classification accuracy with low computational overhead, making the proposed approach suitable for real-time and embedded systems.

MATERIALS AND METHODS

This section describes the FMCW radar system used for data acquisition, the signal preprocessing procedure including CWT-based feature extraction, and the six traffic object classes considered in the study. It also presents the machine learning classifiers (Random Forest, kNN, and SVM) and the evaluation metrics used to assess classification performance.

FMCW radar

The frequency-modulated continuous wave (FMCW) radar is an advanced radar system that continuously transmits an electromagnetic wave with a frequency that changes linearly over time. Unlike traditional pulsed radars, which emit short bursts at a constant frequency, the FMCW radar uses a continuous signal with time-varying frequency – typically in the form of a sawtooth or triangular waveform [9]. An example of a sawtooth waveform generated by an FMCW radar is shown in the Figure 1.

The fundamental principle of FMCW radar operation is based on comparing the transmitted signal frequency f_T with the reflected signal frequency f_R received from a target. Since the reflected signal returns after a certain delay, the frequency difference between the transmitted and received signals – known as the beat frequency f_b – is proportional to the time delay, and thus to the distance to the object [11]:

$$R = \frac{f_B c}{2\beta} \tag{1}$$

where: c denotes the speed of light, and the modulation slope β is given by the following relationship [12]:

$$\beta = \frac{B}{T} \tag{2}$$

where: B denotes the bandwidth, defined as the difference between the maximum frequency f_{max} and the minimum frequency f_{min} [12]:

$$B = f_{max} - f_{min} \tag{3}$$

Additionally, if the object is moving relative to the radar, a Doppler frequency shift f_D occurs due to the Doppler effect, which enables the simultaneous estimation of the object's velocity [12]:

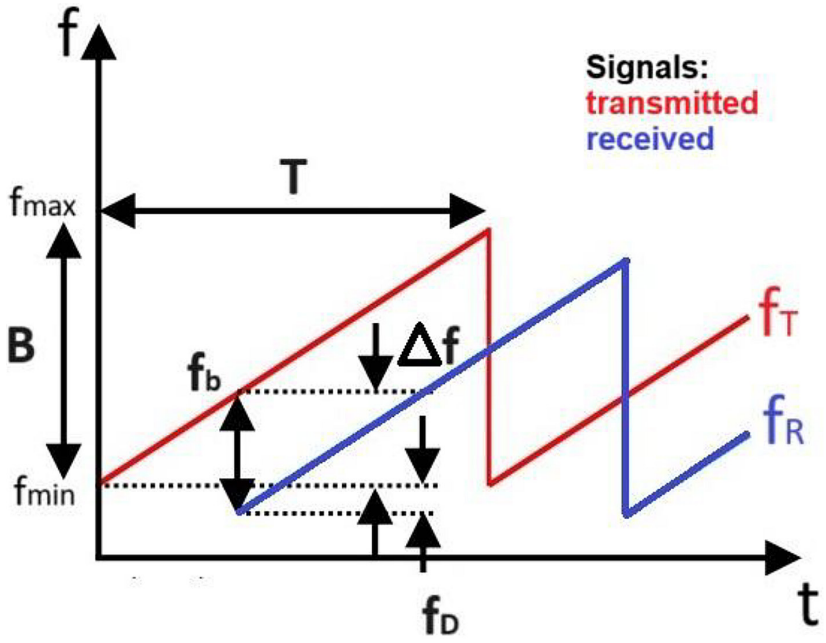

Figure 1. Sawtooth waveform of the signal in an FMCW radar [10]

$$V = \frac{f_D c}{2 f_c} \tag{4}$$

The physical experiments were conducted using the uRAD USB v1.2 radar module manufactured by Anteral, operating at a frequency of 24 GHz. A dedicated software application was used for radar control and data acquisition.

The uRAD USB v1.2 is a fully functional microwave FMCW radar operating in the 24 GHz band (24.005–24.245 GHz), designed for both research and educational purposes. A visual representation of the radar device is shown in the Figure 2.

It supports four operating modes – continuous wave, sawtooth, triangular, and dual mode – enabling the measurement of speed, distance, or both simultaneously. The radar offers a detection range from 0.45 to 100 m, with a distance resolution of ±0.04 m or ±0.3%. According to the manufacturer's datasheet, the uRAD USB v1.2 radar enables velocity measurements within a range of approximately ±0.7 to ±75 m/s in continuous-wave mode. The specified velocity resolution of 3 m/s refers to the minimum velocity separation required to distinguish two targets with similar reflectivity as separate objects. For single-target measurements, the velocity accuracy is significantly higher and is specified as ±0.05 m/s. The maximum signal sampling rate depends on the mode and ranges from 69 to 103 samples per second. The device is powered by 5 V via USB, and its compact dimensions (0.056 × 0.078 × 0.005 m) and low weight (0.008 kg) make it suitable for integration into mobile systems. The manufacturer also provides dedicated software for real-time data control and visualization [13].

Signal preprocessing

The uRAD radar enables recording of raw received signals in the form of text files containing the in-phase (I) and quadrature (Q) components of the reflected signal. In this study, the sawtooth FMCW operating mode was used, generating files I_FMCW.txt and Q_FMCW.txt. Each file consists of consecutive measurement frames containing 200 signal samples, while the final two columns store timestamp information, which is discarded during preprocessing. The I and Q components were combined to form a complex IQ signal [14]:

$$IQ(t) = I(t) + jQ(t) \tag{5}$$

Instead of processing individual frames, consecutive frames were concatenated to form a single input sample. Based on preliminary experiments with different sequence lengths, ten consecutive frames were selected as a compromise between temporal representativeness, classification performance, and computational efficiency. Shorter sequences provided insufficient temporal information, while longer sequences increased

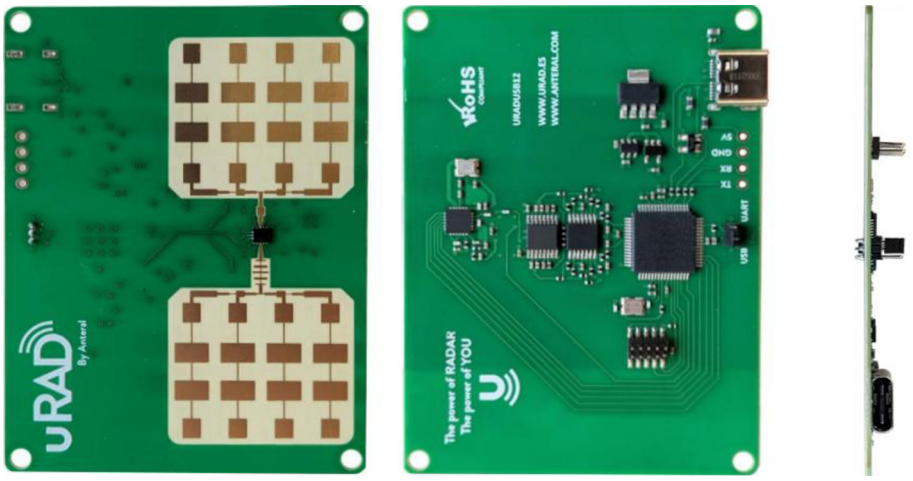

Figure 2. The uRad USB v1.2 radar used in the study [13]

redundancy and computational cost, particularly for fast-moving targets.

The resulting IQ signal was analyzed using the CWT, which provides joint time–scale representation suitable for non-stationary radar signals [18]. The CWT of a signal $x(t)$ is defined as [18]:

$$CWT_x(s,\tau) = \\ = \int_{-\infty}^{\infty} x(t)\,\psi^*_{s,\tau}(t)dt. \tag{6}$$

Unlike the Fourier Transform, which represents a signal solely in the frequency domain without time localization, the Continuous Wavelet Transform preserves both temporal and frequency information, making it particularly suitable for analyzing non-stationary signals. In CWT, the signal is convolved with scaled and shifted versions of a base wavelet, resulting in a two-dimensional time–scale representation that enables observation of how frequency components evolve over time. Such properties are especially relevant for radar signals reflected from moving objects such as pedestrians, cyclists, and vehicles.

In this study, the signal was analyzed using the Continuous Wavelet Transform with a complex Morlet wavelet (Python, cmor1.5-1.0) over a scale range from 1 to 63. From the resulting matrix of wavelet coefficients, three statistical features – the mean, standard deviation, and maximum – were extracted. These features were selected because they are computationally lightweight and have been shown to perform well in radar signal classification tasks [15]. Preliminary experiments evaluating alternative feature sets confirmed that these measures effectively capture key characteristics of the radar signal, including average energy distribution, variability of motion patterns, and peak-related micro-Doppler effects.

The selected statistical features provide a compact representation of the time–frequency structure of the signal while preserving discriminative information across different object classes. By replacing full time–frequency representations with fixed-length feature vectors, this approach enables efficient dimensionality reduction and facilitates the application of classical machine learning algorithms such as Random Forest and SVM, supporting the development of a lightweight and real-time classification framework.

Dataset

The study considered six classes of moving objects representing practical traffic scenarios: a single pedestrian (O1), two pedestrians (O2), a car (O3), a pedestrian with a car (O4), a cyclist (O5), and a cyclist accompanied by a pedestrian (O6). These classes reflect increasing motion complexity and typical operating conditions of roadside radar systems.

During data acquisition, object motion included both approximately perpendicular and radial components with respect to the radar line of sight, depending on the traffic scenario and road geometry. The radar was mounted at ~1 m above ground and remained stationary during all experiments. Object distances ranged from ~5 m to 25 m, resulting in incidence angles of ±60° relative to the radar boresight, representative of realistic roadside deployment.

Precise velocities were not measured due to the absence of external tracking systems. Pedestrians and cyclists moved at natural speeds, while vehicles followed scenario-specific traffic conditions. Multiple recordings per class captured diverse motion realizations and radar signatures despite the lack of explicit velocity annotations. Each measurement lasted ~10 s, providing comparable signal durations across classes. The ~10 s duration refers to the total data recording time for a given object pass used to build the dataset. The actual classification operates on much shorter signal segments and does not require long observation windows. Raw complex IQ data were stored as text files corresponding to consecutive measurement frames.

Experiments were conducted in four real-world scenarios:

- P1: Urban street in a small town (population <20,000) with low-to-medium traffic; daytime, cloudy summer conditions.
- P2: National road with high traffic, including cars and trucks; sunny daytime summer conditions capturing high-speed dynamics.
- P3: Supermarket parking lot with mixed traffic; evening, artificial lighting, low-light autumn conditions.
- P4: Roundabout with pedestrian crossings at a major intersection; predominantly pedestrians, daytime in spring.

All four scenarios are illustrated in Figure 3.

Table 1 summarizes the number of radar measurements collected for each of the six object classes across the four experimental scenarios (P1–P4). The table illustrates the diversity of the dataset in terms of object types, traffic density, motion characteristics, and environmental conditions.

As shown in Table 1, the dataset exhibits a significant class imbalance, with several classes (e.g., Pedestrian + Car (O4) and Cyclist + Pedestrian (O6)) being notably underrepresented. To address class imbalance, the dataset was first divided into training and test subsets. The test set was kept unchanged and was not involved in any resampling procedure. Class balancing was performed exclusively on the training data using the Synthetic Minority Over-sampling Technique (SMOTE) applied at the feature level. SMOTE generates synthetic minority samples by interpolating between training instances, improving class balance without using test data and thus preventing data leakage.

Table 2 presents the number of training samples per object class and per scenario after applying SMOTE.

While SMOTE improves class balance and stabilizes the training process, it may influence class decision boundaries by introducing synthetic samples interpolated between minority-class instances. Such samples do not always correspond to physically observed radar signatures and may affect boundary smoothness. In this study, SMOTE was applied exclusively to the training set to avoid data leakage. Alternative strategies, such as class-weighted learning or training without resampling, could be considered in future work to further assess robustness and generalization of the proposed approach.

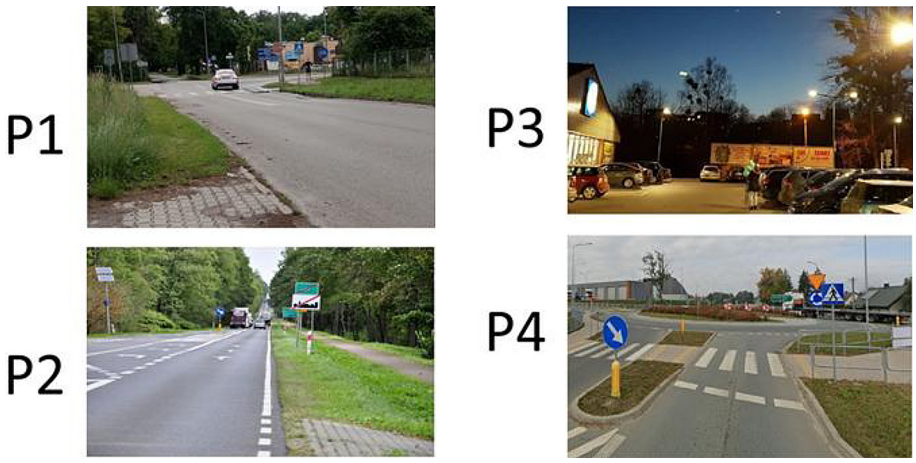

Figure 3. Illustration of the four experimental scenarios used for data collection:
P1 – urban street in a small town; P2 – national road with high traffic volume;
P3 – supermarket parking lot under artificial lighting; P4 – roundabout with pedestrian crossings

Table 1. Number of radar measurements per object class in each experimental scenario (P1–P4)

Parameter		Scenario			
		P1	P2	P3	P4
Class	O1	24	1	64	27
	O2	16	2	35	21
	O3	53	121	67	86
	O4	4	1	21	6
	O5	15	3	19	12
	O6	9	1	8	4

Table 2. Number of training samples per object class in each scenario (P1–P4) after applying SMOTE

Parameter		Scenario			
		P1	P2	P3	P4
Class	O1	68	3	181	75
	O2	71	9	155	92
	O3	53	121	67	86
	O4	41	10	215	61
	O5	100	20	127	80
	O6	135	15	119	58
Summary		468	178	864	452

Synthetic samples were generated for minority classes until all classes in the training set reached the same total number of samples as the original majority class, while preserving the original proportional distribution across scenarios. The test dataset retained its original class distribution as reported in Table 2 and was used exclusively for performance evaluation.

Classification methods and metrics

As part of the evaluation of the proposed classification framework, three widely used machine learning algorithms were applied and compared: Random Forest (RF), k-nearest neighbours (kNN), and support vector machine (SVM).

Random Forest is an ensemble learning method that constructs multiple decision trees using random subsets of samples and features, improving robustness to noise and reducing overfitting. Final classification is performed by majority voting across all trees, yielding stable and well-generalized predictions [19]. The performance of RF depends primarily on the number of trees, maximum tree depth, minimum number of samples required for node splitting, and minimum leaf size, which collectively control model complexity and generalization ability [20].

K-nearest neighbours is a distance-based classifier in which the label of a test sample is determined by the majority class among its K closest neighbors in the feature space. The method does not involve explicit training, and its performance depends mainly on the choice of K and the distance metric (e.g., Euclidean distance). While kNN is simple and effective for low-dimensional data, it is computationally demanding at inference time due to the need to search the entire training set [20].

For reproducibility, the hyperparameters of the applied classifiers are specified below. For the Random Forest model, hyperparameter selection was performed using a grid search over the number of trees (n_estimators ∈ {10, 50, 100, 200}), maximum tree depth (max_depth ∈ {None, 10, 20}), minimum number of samples required for node splitting (min_samples_split ∈ {2, 5}), and minimum number of samples per leaf (min_samples_leaf ∈ {1, 2}), with random_state set to 42. The Gini impurity criterion was used for node splitting, following the default configuration of the scikit-learn library. The best-performing hyperparameter combination on the validation split was subsequently used to report the final results.

The k-Nearest Neighbours classifier was configured with k = 5 nearest neighbors and the Euclidean distance metric (corresponding to the

Minkowski distance with p = 2). All remaining parameters were retained at their default values.

Support Vector Machine aims to find an optimal separating hyperplane by maximizing the margin between classes. For nonlinearly separable data, kernel functions enable mapping into higher-dimensional spaces where linear separation is possible. SVM is particularly effective for high-dimensional feature spaces and limited training data, offering strong generalization performance at the cost of increased training complexity [21–22].

The classification performance was evaluated using standard metrics, including accuracy, precision, recall, and the F1-score. Accuracy represents the proportion of correctly classified samples, while precision and recall characterize classification reliability and sensitivity, respectively. The F1-score provides a balanced measure by combining precision and recall.

In order to assess whether the observed differences between classifiers were statistically significant, McNemar's test was applied to paired predictions on the test set. The test was performed separately for RF vs SVM and RF vs kNN at a significance level of $\alpha = 0.05$.

EXPERIMENTAL RESEARCH

The measurements were conducted under real-world conditions across four traffic scenarios representing different road types, including intersections, roundabouts, and streets with varying traffic densities. Environmental variability was considered through different radar incidence angles, while the radar mounting height was kept constant. An example of the measurement setup deployed at the roadside in Dęblin is shown in Figure 4.

The experimental setup was based on the FMCW uRAD USB v1.2 radar, mounted in a stationary position and connected to a computer via a USB interface. A dedicated software tool was used to record raw IQ data, consisting of the in-phase (I) and quadrature (Q) components of the reflected signal. The measurements were performed passively, without interfering with pedestrian, cyclist, or vehicle movement, and only objects belonging to one of the six predefined classes were considered. Each measurement was stored in two separate files, I_FMCW_sawtooth.txt and Q_FMCW_sawtooth.txt.

Figure 4. The measurement setup

Figure 5 presents representative examples of raw IQ signals recorded for scenario P1 for all six object classes. The time-domain I/Q signals and their corresponding CWT-based time–frequency representations reveal distinct dynamic patterns associated with different motion types, enabling effective feature extraction and subsequent classification.

For each signal segment, the continuous wavelet transform was applied to obtain a time–scale representation, from which simple statistical features (mean, standard deviation, and maximum) were extracted to form fixed-length feature vectors. These vectors were labeled according to the object class and used to train Random Forest, kNN, and SVM classifiers. During testing, new IQ data underwent the same preprocessing, and the trained models were used to predict the object category.

For each scenario (P1–P4), the classifiers were trained and tested using data from the same scenario to assess feature-level discriminability under comparable measurement geometry and traffic conditions. This within-scenario protocol was adopted to isolate the effect of the proposed lightweight CWT-based feature representation and to ensure a fair comparison between classifiers. Explicit cross-scenario validation (e.g., training on P1–P3 and testing on P4) was not performed in this study and is discussed as a limitation and future research direction.

Wavelet parameter analysis

To investigate how the choice of wavelet type and the range of applied scales affect the classification performance of radar-detected objects, two parameter lists were defined: types of wavelets

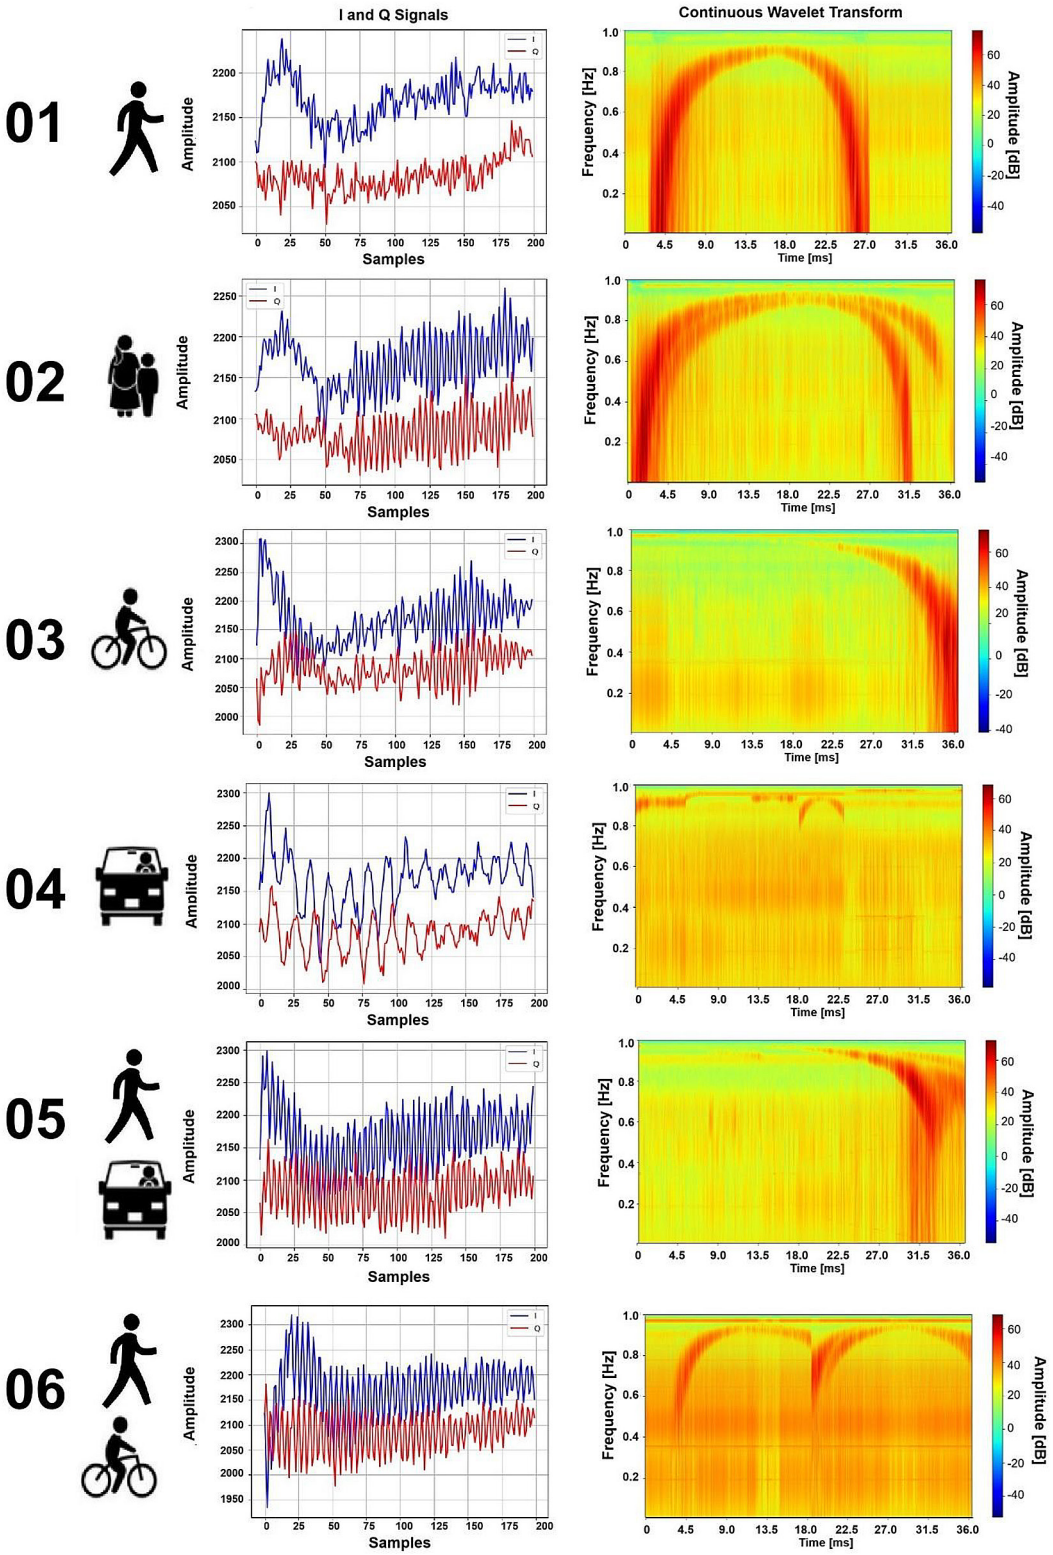

Figure 5. Sample waveforms of raw IQ signals recorded by the FMCW radar for each of the six object classes (scenario P1)

and scale ranges. For this study, radar measurements collected in scenario P3 were used, due to the largest number of samples across all object classes (O1–O6). The experiment involved iteratively testing all possible combinations of these two parameters. The results of the computations conducted are visualized in Figure 6.

The calculations were performed using four types of wavelets: complex Morlet ("cmor"), real Morlet ("morl"), Shannon ("shan"), and Gaussian

("gaus"). The scale ranges were defined in three intervals: 1–31, 1–63, and 1–127. The best overall classification performance – measured by the combined accuracy across all six object classes – was achieved using complex Morlet wavelets with a scale range of 1–63.

This setting provided a balanced trade-off between time and frequency resolution. A narrower scale range (e.g., 1–31) led to a loss of important signal features, while a wider range (e.g., 1–127) introduced redundancy and noise. The real wavelets yielded lower classification accuracy, confirming that IQ radar signals are more effectively analyzed using complex wavelets.

Comparison of classifier performance

The performance of three classifiers was compared: Random Forest, support vector machine, and k-nearest neighbours. All models used the same data representation – feature vectors extracted using the CWT – to ensure a fair comparison, of course, across the four different scenarios (P1–P4). The SVM classifier employed a radial basis function (RBF) kernel with parameters set to C = 10 and *gamma* = 'scale'. The parameter C controls the trade-off between fitting the training data closely and maintaining the model's ability to generalize – higher values of C (such as 10) reduce the tolerance for misclassifications, which can increase the risk of overfitting. On the other hand, *gamma*='scale' automatically adjusts the influence of individual training samples on the decision boundary based on the number of input features and the variance of the data. The kNN model was based on 5 nearest neighbors (*n_neighbors* = 5), with classification decisions made according to the majority class among these neighbors.

The Tables 3, 4, present the detailed performance metrics for the three classifiers: Random Forest, SVM, and k-Nearest Neighbours for each of the six classes.

Random Forest demonstrated the highest effectiveness among the methods compared. The F1-scores for all classes remained at a very high level (above 0.84), reaching as high as 0.99 for class O6. This model showed both high precision and recall, meaning it not only correctly identified many samples but also rarely misclassified other classes as the target class. Thanks to its ensemble structure based on multiple decision trees, RF handled imbalanced data distribution and noise effectively. An additional advantage was its stability and fast execution time, making it the most reliable solution for this task.

The support vector machine (SVM) exhibited lower overall performance compared to the Random Forest, with results highly dependent on the class. For the minority classes O2 and O5, precision, recall, and F1-score were near zero, indicating that the model failed to correctly classify these samples. This behavior is typical for SVMs on imbalanced datasets, where the model tends to favor majority classes. In contrast, well-represented classes such as O1 and O6 achieved reasonable F1-scores of 0.71 and 0.75, demonstrating that SVM can be stable when class features are clearly separable.

K-nearest neighbours (kNN) demonstrated moderate performance – better than SVM for underrepresented classes, but lower overall compared to Random Forest. The model performed reasonably well for the well-represented classes O1 and O6, achieving F1-scores of 0.77–0.82 and 0.70–0.74, respectively. However, its results for minority classes O2 and O4 remained low, with

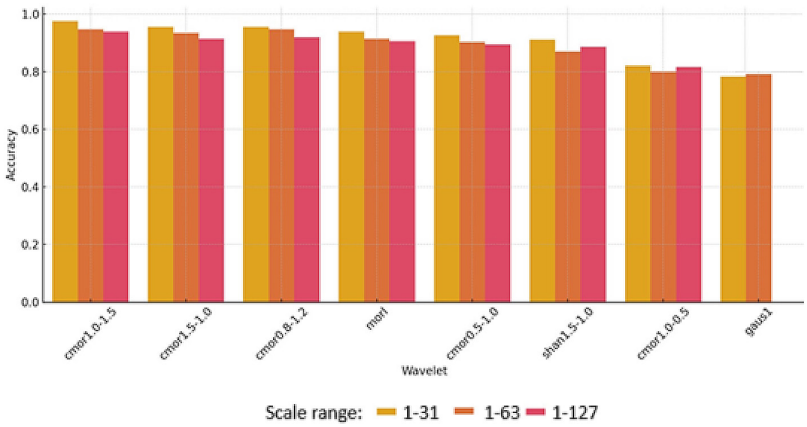

Figure 6. Impact of wavelet type and scale range

Table 3. Classification performance for random forest

Parameter		O1	O2	O3	O4	O5	O6
P1	Pre.	0.96	0.76	0.93	0.94	1.00	1.00
	Rec.	1.00	0.96	0.93	0.88	0.67	0.97
	F1	0.98	0.85	0.93	0.91	0.80	0.99
P2	Pre.	0.85	0.70	0.88	0.82	0.90	0.92
	Rec.	0.90	0.78	0.91	0.80	0.60	0.95
	F1	0.87	0.74	0.89	0.81	0.72	0.93
P3	Pre.	0.92	0.82	0.91	0.93	0.95	0.96
	Rec.	0.95	0.90	0.90	0.92	0.75	0.97
	F1	0.94	0.86	0.91	0.93	0.84	0.97
P4	Pre.	0.90	0.78	0.89	0.88	0.92	0.95
	Rec.	0.93	0.85	0.87	0.85	0.70	0.96
	F1	0.91	0.81	0.88	0.86	0.79	0.95

Table 4. Classification performance for SVM

Parameter		O1	O2	O3	O4	O5	O6
P1	Pre.	0.55	0.00	0.29	0.50	0.00	0.66
	Rec.	1.00	0.00	0.17	0.03	0.00	0.86
	F1	0.71	0.00	0.21	0.06	0.00	0.75
P2	Pre.	0.48	0.10	0.32	0.45	0.05	0.60
	Rec.	0.85	0.12	0.20	0.08	0.07	0.78
	F1	0.61	0.11	0.25	0.14	0.06	0.68
P3	Pre.	0.50	0.15	0.35	0.48	0.10	0.63
	Rec.	0.80	0.18	0.25	0.12	0.12	0.75
	F1	0.62	0.16	0.29	0.19	0.11	0.68
P4	Pre.	0.53	0.12	0.33	0.50	0.08	0.61
	Rec.	0.82	0.15	0.28	0.14	0.10	0.72
	F1	0.65	0.13	0.30	0.22	0.09	0.66

F1-scores of 0.19–0.23 and 0.46–0.52. The kNN algorithm relies on local similarity and is particularly sensitive to noise and uneven data distribution in the feature space. Its effectiveness depends strongly on the density and distribution of the training data. Additionally, kNN exhibits the longest prediction time among the compared methods, making it less suitable for real-time applications.

McNemar's test, applied to paired predictions on the test set at a significance level of $\alpha = 0.05$, indicated that the Random Forest classifier significantly outperformed both SVM and kNN ($p < 0.05$). The asymmetry in misclassification patterns confirmed that RF correctly classified a significantly higher number of samples misclassified by the alternative models, while the opposite situation occurred considerably less frequently. These results demonstrate that the observed performance differences were not attributable to random variation in the test data.

Comparison of the developed method with selected reference methods

The final stage of the study focused on evaluating computational efficiency, model complexity, and the contribution of individual features to classification performance. Scenario P3 was selected for CNN-based experiments due to its relatively large number of radar measurements (864 samples after SMOTE) and a more balanced class distribution compared to the remaining scenarios.

To avoid data leakage, the dataset was split at the measurement-session (trajectory) level, ensuring that samples derived from the same recording were assigned exclusively to either the training or

test set. An 80/20 split with a fixed random seed was applied to ensure reproducibility.

A lightweight CNN architecture was adopted to reflect deployment-oriented constraints. The network consisted of three convolutional layers with 32, 64, and 128 filters (3 × 3 kernels), each using ReLU activation. Max-Pooling2D layers followed the first two convolutional layers. The convolutional block was followed by a Flatten layer, a Dense layer with 128 neurons and ReLU activation, a Dropout layer (rate = 0.3), and a softmax output layer with six neurons. The network was trained using the Adam optimizer (learning rate 0.001) and the sparse categorical crossentropy loss function (Table 5).

The total number of trainable parameters was approximately 0.48 million, significantly lower than in standard deep architecture such as VGG-16 (~138 million) or ResNet-50 (~25.6 million). Preliminary experiments showed fast convergence, with no meaningful accuracy improvement beyond four epochs; therefore, training was limited to four epochs to reduce computational cost.

Generating CWT-based time–frequency images (1200 × 600 pixels) and storing them as PNG files introduced substantial preprocessing overhead. Training the CNN for four epochs required approximately 428 s, whereas classical classifiers such as Random Forest and SVM, trained directly on statistical CWT features, completed training in under one second. Despite this difference, the CNN achieved a test accuracy of 94.41%, comparable to the best classical classifiers (94–95%).

Most existing radar-based object classification methods rely on deep learning models operating on STFT- or FFT-based time–frequency images and employ large architectures such as VGG-16, ResNet-50, or DarkNet-53 [3–6]. Although high accuracies (88.6%–96.56%) have been reported [3–4], these methods typically require substantial computational resources and large training datasets.

In contrast, the proposed approach uses lightweight classifiers (Random Forest, kNN, and SVM) trained on compact statistical features extracted from CWT representations of raw radar signals. The method achieves comparable classification accuracy while substantially reducing preprocessing effort, training time, and model complexity, which makes it suitable for real-time and embedded radar applications. A quantitative comparison with selected reference methods is provided in Table 6.

Feature importance analysis

To enhance the interpretability of the proposed approach, feature importance analysis was performed using the Random Forest classifier. Feature importance scores were extracted based on the decrease in impurity criterion, allowing estimation of the relative contribution of individual features to the classification decision.

The results indicate that the standard deviation and maximum of CWT coefficients contribute more strongly to classification than the mean, highlighting the importance of signal variability and peak-related micro-Doppler characteristics in distinguishing object classes.

Table 5. Classification performance for kNN

Parameter		O1	O2	O3	O4	O5	O6
P1	Pre.	0.75	0.20	0.45	0.60	0.15	0.70
	Rec.	0.90	0.25	0.40	0.45	0.10	0.78
	F1	0.82	0.22	0.42	0.52	0.12	0.74
P2	Pre.	0.70	0.18	0.42	0.55	0.12	0.68
	Rec.	0.85	0.20	0.38	0.40	0.09	0.73
	F1	0.77	0.19	0.40	0.46	0.10	0.70
P3	Pre.	0.73	0.22	0.48	0.58	0.14	0.72
	Rec.	0.88	0.24	0.42	0.45	0.12	0.77
	F1	0.80	0.23	0.45	0.51	0.13	0.74
P4	Pre.	0.72	0.21	0.46	0.57	0.13	0.71
	Rec.	0.87	0.23	0.41	0.44	0.11	0.76
	F1	0.79	0.22	0.43	0.50	0.12	0.73

Table 6. Performance comparison of the developed method with selected reference methods

Ref.	Method	Input data	Object classes	Acc. (%)	Remarks
Angelov [3]	CNN + LSTM	FFT/STFT images	Pedestrian, car, cyclist, 2 pedestrians	88.6	VGG-16 network; 200 epochs; limited dataset
Hadhrami [4]	Transfer learning (VGG16 + SVM)	VGG-16 features	6 classes	96.56	First demonstration of using pre-trained CNN as feature extractor
Dadon [5]	CNN	STFT images	Pedestrians, animals	93.62	Micro-Doppler signature; augmented training dataset
Buchman [6]	CNN	STFT images	Pedestrians, animals	94.04	Specialized CNN; MAFAT Radar Challenge dataset
Proposed method	Lightweight classifier (RF / kNN /SVM) + CWT	Raw CWT features (mean, std, max)	6 classes	94–95	No image generation; fast training; suitable for real-time applications

CONCLUSIONS

The conducted research represents a significant contribution to the development of methods for classifying road traffic participants using radar data. Unlike commonly used, computationally demanding image-based methods – such as CNNs applied to spectrograms – the proposed approach is lightweight, efficient, and operates directly on raw IQ signals from an FMCW radar. The method utilizes the CWT for time-scale analysis and extracts simple statistical features (mean, standard deviation, and maximum), which serve as effective inputs for classifiers such as Random Forest, SVM, and kNN.

The proposed approach also offers significant time efficiency advantages. Unlike image-based CNN methods, numerical CWT features eliminate the need for time-consuming image generation and reduce training and prediction times from several minutes to fractions of a second, making the method particularly suitable for real-time and embedded applications.

Comparison with existing literature shows that the proposed lightweight method achieves classification accuracy comparable to complex deep learning architectures (e.g., VGG-16, ResNet-50), while substantially reducing computational overhead and dataset requirements. Experiments were primarily conducted on the P3 scenario, which provides the largest and most balanced dataset after SMOTE, ensuring robust and reproducible results.

The developed method can be applied in intelligent traffic monitoring systems, advanced driver-assistance systems (ADAS), autonomous vehicles, as well as in pedestrian crossing surveillance and urban traffic monitoring. Its independence from lighting conditions also makes it suitable for military and emergency applications, enabling reliable real-time detection and classification of objects. For real-time automotive applications, the proposed method can be applied using short sliding windows, enabling near-instantaneous classification that is independent of the longer data acquisition durations used during offline dataset construction.

Additionally, very slow-moving pedestrians with velocities close to zero may be more challenging to detect due to the lower velocity measurement limit of the radar, which constitutes a hardware-related limitation of the sensing system.

An analysis of classification errors indicates that most misclassifications occur between composite object classes and their corresponding single-object counterparts, such as pedestrian vs. pedestrian with cyclist or cyclist vs. cyclist with pedestrian. This confusion can be attributed to overlapping micro-Doppler signatures and similar motion dynamics, especially when one object dominates the radar return. Additionally, partial occlusion and variability in relative object motion may further blur class boundaries. These observations highlight inherent challenges in radar-based classification of multi-object scenarios and explain the reduced performance observed for composite classes.

Despite the promising results, the proposed approach has several limitations that should be acknowledged. A limitation of the present study is that model evaluation was performed within individual scenarios, and explicit cross-scenario validation (e.g., training on P1–P3 and testing on P4) was not conducted. Cross-scenario experiments introduce additional variability related to scene layout, traffic composition, and radar–target geometry, which constitutes a distinct generalization problem. Addressing this aspect represents an important direction for future work and will be pursued using a larger and more diverse dataset.

A statistical significance analysis using McNemar's test was performed to compare classifier performance. The results confirmed that the Random Forest classifier significantly outperformed the alternative models ($p < 0.05$). As a non-parametric test applied to paired predictions on the test set, McNemar's test provides a robust assessment under class imbalance conditions. Future research may extend this analysis by incorporating repeated cross-validation or larger multi-scenario datasets to further evaluate generalization stability.

Extending the approach to a broader taxonomy of road users will require additional labeled data and may benefit from hierarchical classification as well as open-set mechanisms (confidence-based rejection) to handle previously unseen categories.

Future research directions include expanding the number of object classes (e.g., motorcycles, scooters), integrating radar data with additional sensors such as cameras, leveraging pre-trained models via transfer learning, and exploring hybrid approaches combining traditional statistical features with deep learning techniques.

REFERENCES

1. Jankiraman M. FMCW radar design. Norwood, MA, USA: Artech House; 2018.
2. Kulhandjian H., et al. AI-based pedestrian detection and avoidance at night using multiple sensors. J. Sens. Actuator Netw. 2024; 13(3): 34. https://doi.org/10.3390/jsan13030034
3. Angelov A., Robertson A., Murray-Smith R., Fioranelli F. Practical classification of different moving targets using automotive radar and deep neural networks. IET Radar Sonar Navig. 2018; 12(10): 1082–1089. https://doi.org/10.1049/iet-rsn.2018.0103
4. Hadhrami E.A., Mufti M.A., Taha B., Werghi N. Ground moving radar targets classification based on spectrogram images using convolutional neural networks. In: Proceedings of the 19th International Radar Symposium (IRS); Bonn, Germany, June 2018; 1–9. https://doi.org/10.23919/IRS.2018.8447897
5. Dadon Y.D., et al. Moving target classification based on micro-Doppler signatures via deep learning. In: Proceedings of the IEEE Radar Conference (RadarConf21); Atlanta, GA, USA, May 2021; 1–6.https://doi.org/10.1109/RadarConf2147009.2021.9455270
6. Buchman D., Drozdov M., Krilavičius T., Maskeliūnas R., Damaševičius R. Pedestrian and animal recognition using Doppler radar signature and deep learning. Sensors. 2022; 22(9): 3456. https://doi.org/10.3390/s22093456
7. Zhang Y., Peng Y., Wang J., Lei P. Human motion recognition based on radar micro-Doppler features using Bayesian network. In: Proceedings of the IEEE International Conference on Signal, Information and Data Processing (ICSIDP); Dongguan, China, December 2019. https://doi.org/10.1109/ICSIDP47821.2019.9173502
8. Wei Y., Zhang Y., Xu Z., Li X. Classification of pedestrian motion based on micro-Doppler feature with LFMCW radar. In: Proceedings of the IEEE International Conference on Signal, Information and Data Processing (ICSIDP); Dongguan, China, December 2019. https://doi.org/10.1109/ICSIDP47821.2019.9173363
9. Skolnik M.I. Radar handbook. 3rd ed. New York, NY, USA: McGraw-Hill; 2008.
10. Mahafza B.R., Elsherbeni A.Z. MATLAB simulations for radar systems design. New York, NY, USA: Chapman & Hall/CRC; 2004.
11. Richards M.A. Fundamentals of radar signal processing. 2nd ed. New York, NY, USA: McGraw-Hill Education; 2014.
12. Budge M., German S. Basic radar analysis. 2nd ed. Norwood, MA, USA: Artech House; 2020.
13. Anteral. uRAD USB v1.2 24 GHz radar – velocity and 1D positioning: datasheet. Pamplona, Spain; 2022.
14. Zieliński T. Digital signal processing – from theory to applications (in Polish). Warsaw, Poland: WKiŁ; 2005.
15. Dobrowolski A. Signal transformations: from theory to practice (in Polish). Legionowo, Poland: BTC; 2018.
16. Ślesicki B., Ślesicka A. A new method for traffic participant recognition using Doppler radar signature and convolutional neural networks. Sensors. 2024; 24(12): 3832. https://doi.org/10.3390/s24123832
17. Ślesicki B., Ślesicka A., Kawalec A., Walenczykowska M. Improving recognition of road users via Doppler radar data and deep learning convolutional networks. Electronics. 2024; 13(20): 4070. https://doi.org/10.3390/electronics13204070
18. Walenczykowska M., Kawalec A. Application of continuous wavelet transform and artificial neural network for automatic radar signal recognition. Sensors. 2022; 22(19): 9743. https://doi.org/10.3390/s22197434
19. Goodfellow I., Bengio Y., Courville A. Deep learning. Cambridge, MA, USA: MIT Press; 2016.
20. Osowski S. Neural networks for information processing (in Polish). Warsaw, Poland: Oficyna Wydawnicza PW; 2020.
21. Cortes C., Vapnik V. Support-vector networks. Mach. Learn. 1995; 20(3): 273–297. https://doi.org/10.1007/BF00994018
22. Nalepa K., Kawulok M. Selecting training sets for support vector machines: A review. Artif. Intell. Rev. 2019; 52: 857–900. https://doi.org/10.1007/s10462-017-9611-1