doi: 10.5937/str2600006S

An Active Learning Framework for Drone Classification in Radio Frequency Domain

Boban Sazdić-Jotić [1]
Velibor Ilić [2]
Ninko Miletić [1]
Darko Mikanović [1]

Radio-frequency–based drone classification is a critical capability for modern antidrone systems. However, the development of dependable artificial intelligence models in this domain is hindered by the high cost and complexity of expert data labeling. This challenge is particularly pronounced in radio-frequency signal analysis, where labeled datasets are typically small, partially unlabeled, and continuously evolving during system deployment. This paper proposes a tailored active learning framework for drone classification in the radio-frequency domain, integrating human expertise into an iterative learning process that selectively queries the most informative unlabeled samples. By prioritizing sample informativeness, the proposed framework aims to achieve high detection performance while significantly reducing labeling effort. The approach is evaluated using the VTI_USRP_DroneSET dataset, comprising radio-frequency spectrograms acquired in realistic outdoor conditions within the 2.4 GHz frequency band. Experimental results demonstrate that the proposed active learning strategy achieves mAP50–95 performance comparable to conventional supervised learning while requiring only one quarter of the labeled data. The results confirm that active learning enables data-efficient radio-frequency based drone classification without compromising detection accuracy. Furthermore, near-optimal performance is consistently obtained with an optimal training duration of 80 epochs, reducing both annotation and computational costs. These findings confirm that active learning provides a data-efficient and cost-effective solution for RF-based drone classification and is well suited for real-time deployment in operational antidrone systems where labeled data are scarce and continuously acquired.

Key words: **active learning, antidrone, artificial intelligence, classification, drone.**

Introduction

ACTIVE learning is a subfield of machine learning (ML) within artificial intelligence (AI) that focuses on improving the labeling process by allowing an AI model to query the most informative data points selectively. In traditional supervised learning, substantial amounts of labeled data are needed, which can be both expensive and time-consuming to obtain, especially in domains such as medical imaging, legal document analysis, or sentiment annotation. (Ren et al., 2022b; Settles, 2009, 2012) At its core, active learning enables the AI model to choose the data it learns from. The key concept is that an AI model can achieve better results with fewer labeled samples if it can actively select the most informative data to label. The main ideas behind active learning can be summarized in the following key points. The first is labeling cost awareness. Labeling data is expensive (e.g., satellite images, legal documents), and active learning helps reduce this burden. The second is the faster learning loop, which involves starting with a small, labeled dataset, training an AI model on this data, using the AI model to identify uncertain or informative unlabeled samples, querying a human (or an expert) to label those samples, and adding them to the training set for retraining. The third aspect concerns the various query strategies controlled by the operator or human. This approach allows for selecting samples where the AI model is least confident, facilitating uncertainty sampling. By utilizing query-by-committee, it is possible to employ an ensemble of AI models to find disagreements. (Wang et al., 2017) The human operator can choose points that will most significantly alter the AI model and select samples that would most effectively reduce generalization error. Furthermore, it ensures a diverse choice of examples, not just those considered uncertain.

Categorization of active learning is proposed in (Settles, 2009) into three primary types: pool-based, stream-based, and membership query synthesis. In pool-based active learning, the AI model has access to a large dataset (“pool”) of unlabeled data and selectively queries the most informative samples based on a predefined strategy, such as uncertainty or diversity sampling. Stream-based active learning, on the other hand, processes data instances sequentially as they arrive, making real-time decisions on whether to query the label of each instance or discard it. This approach is instrumental in online or time-sensitive applications. Lastly, membership query synthesis involves the AI model generating synthetic data instances from the input space and requesting labels.

[1] Military Technical Institute, Ratka Resanovića 1, 11000 Belgrade, Serbia
[2] Institute for Artificial Intelligence, 21000 Novi Sad, Serbia
Correspondence to: Boban Sazdić-Jotić, e-mail: boban.sazdic.jotic@vs.rs

Although it is theoretically powerful, this method is rarely used in practice due to the risk of producing unrealistic or uninformative instances that are difficult for human annotators to label accurately. Each type offers different trade-offs in terms of applicability, efficiency, and labeling cost, depending on the task and data availability. The predominant approach to active learning is the pool-based type, primarily because the challenge of labeling data is a foundational issue faced by researchers in this domain. (Zhan et al., 2021) The typical pool-based active learning workflow involves training an initial AI model on a small, labeled dataset, then using that AI model to evaluate an unlabeled data pool. Based on a selection strategy, such as uncertainty sampling, query-by-committee, or expected error reduction, the AI model shows the most uncertain or diverse examples and queries a human expert to obtain labels for them. These newly labeled instances are then added to the training data, and the AI model is retrained, thereby continuing the iterative cycle. (Lewis & Gale, 1994; Tong & Koller, 2001)

Active learning offers several advantages in AI applications, most notably the reduction of annotation costs and improved training efficiency by enabling AI models to selectively query the most informative or uncertain samples for labeling. (Ilić, 2024) This characteristic is particularly beneficial in domains where labeled data are expensive or time-consuming to obtain, such as medical diagnostics and legal document analysis. In the context of applications based on the radio-frequency (RF) domain, active learning further enables rapid adaptation to newly observed signal patterns, which is essential in dynamic operational environments where characteristics may change due to varying transmission protocols or environmental conditions. By achieving competitive performance with fewer annotated samples, active learning also facilitates efficient learning under evolving data distributions. Primary contribution of the proposed approach lies in reducing annotation effort, expert involvement, and system scalability constraints, rather than in substantially improving absolute detection accuracy.

Despite these advantages, active learning also presents several limitations. The sample selection process may introduce bias toward regions of high uncertainty, and the annotation of ambiguous or borderline samples can be challenging even for domain experts. Furthermore, the iterative training and querying cycles inherent to active learning increase computational overhead. The effectiveness of the approach additionally depends on the availability of a sufficiently representative initial labeled dataset to initialize the learning process. These limitations underscore the importance of carefully designed active learning strategies that are specifically tailored to RF signal analysis and drone detection tasks.

The research paper is organized as follows: Section 2 presents a literature review in the corresponding area, Section 3 outlines the Methodology used in the research, including an explanation of the performed experiments, Section 4 illustrates the results with a discussion and explanation, and finally, Section 5 concludes the research paper.

Literature

2.1 General Scope

Active learning has demonstrated effectiveness across a wide range of AI applications, including natural language processing (NLP), such as semantic analysis, and computer vision (CV), particularly in image segmentation and object detection. By selectively querying the most informative samples for labeling, active learning provides a principled framework for improving learning efficiency and reducing annotation costs across diverse AI tasks. (Sener & Savarese, 2018)

In healthcare and medical diagnostics, active learning enables AI models to prioritize diagnostic uncertain cases for expert annotation, thereby reducing the annotation burden on clinicians while maintaining high diagnostic accuracy. In NLP, it improves tasks such as sentiment analysis and named entity recognition by selectively querying the most informative or ambiguous text samples, leading to enhanced AI model performance with fewer labeled instances. Similarly, in CV applications, including image classification and object detection, active learning reduces the need for extensive manual annotation, which is often both labor-intensive and expensive. These advantages are directly relevant to RF-based drone classification, where expert annotation of spectrogram data requires specialized knowledge and where new signal patterns may emerge during system operation. By selectively identifying the most informative RF samples for annotation, active learning enables efficient AI model adaptation while significantly reducing labeling effort, making it particularly well-suited for practical antidrone systems operating under constrained annotation and computational resources. (Ren et al., 2022a) Combining active learning with deep learning (DL), through approaches such as Bayesian deep active learning, has been effective for reducing data requirements in complex tasks. (Gal et al., 2017) Recent work has also explored active learning in the context of object detection under weak supervision. Authors in (Y. Wang et al., 2023) proposed ALWOD, an active learning framework for weakly supervised object detection, demonstrating that selective sample querying can substantially improve detection performance while reducing annotation effort.

Active learning is a powerful technique that minimizes labeling costs by selecting the most informative data points for annotation. In (Kumar et al., 2020), the authors explore various query strategies for classification, regression, and clustering, emphasizing how these strategies refine training datasets and improve learning efficiency. They categorize classification query strategies into informative-based, representative-based, and hybrid approaches, while also discussing advanced methods incorporating reinforcement learning and DL. The study provides a comparative analysis of these strategies, highlighting their effectiveness in reducing generalization errors and improving AI model performance. Additionally, the survey outlines practical applications, implementation guidelines, and challenges associated with active learning in real-world scenarios. Motivated by these findings, this study investigates whether an active learning framework can reduce manual labeling effort in RF-based drone classification while preserving detection accuracy under real-world conditions.

2.2 Computer Vision (CV) Applications

The study on hyperspectral image classification is presented in (Zhao et al., 2023) and represents the new framework combining a multi-attention Transformer (MAT) with adaptive super pixel segmentation-based active learning (MAT-ASSAL). Traditional DL methods, such as convolutional neural networks (CNNs), struggle to capture long-range spectral-spatial relationships due to limited receptive fields. The MAT AI model enhances contextual dependency between spectral-spatial embeddings using selfattention mechanisms, while an outlook-attention module improves local feature extraction. To improve training with a limited number of labeled samples, the study introduces an active learning strategy based on superpixel segmentation, selecting only the most informative samples. An adaptive super pixel segmentation algorithm also preserves edge details in complex regions while reducing redundancy in uninformative areas. Experimental results show that MAT-ASSAL outperforms seven state-of-the-art methods on various hyperspectral image datasets, achieving better classification accuracy even with small sample sizes.

In the study by (Cevik et al., 2016), active learning was applied to the University of Wisconsin breast cancer simulation AI model, where traditional calibration required evaluating 378,000 parameter combinations, but only sixty-nine of them closely matched observed data. By integrating active learning, researchers were able to identify all sixty-nine best combinations by evaluating just 5,620 of the original set, showing that only 1.49% of all combinations needed to be evaluated for effective calibration. This method enables simulation AI models to efficiently find promising parameter values, improving accuracy while significantly reducing computational costs and time.

The study in (Menke et al., 2024) introduces active learning as a bridge between unsupervised and supervised domain adaptation for object detection. It proposes three novel selection strategies, leveraging domain discriminators and multi-model scoring to improve annotation efficiency. Experimental results show a 2.4% performance improvement while reducing annotation costs by 13%, thereby enhancing adaptation effectiveness.

The authors in (Ilić & Tadić, 2022) exploited a small subset of manually labeled data from different datasets for the training process, integrated with an active learning methodology. Such an approach involves two distinct networks: a CNN and a Self-Correcting Neural (SCN) Network. Initially, the CNN is trained solely on manually labeled data and then employed to predict labels for previously unlabeled samples. Subsequently, the SCN is trained on the complete dataset, incorporating both manually labeled examples and samples automatically labeled by the CNN. After training, the SCN generates new predictions for all available samples, which are then compared to the original training labels. The authors prove that their approach of active learning enables achieving comparable results using only 6.11% to 59.4% of manually labeled data from different datasets, as opposed to utilizing the entire dataset for training.

The authors in (Mittal et al., 2025) examined a broad range of active learning approaches and identified key aspects relevant to the selection of effective active learning strategies. Their study provides comprehensive guidelines for achieving optimal performance in image classification and semantic segmentation tasks.

2.3 Drone Related Applications

Despite the growing interest in active learning across various CV and AI domains, its application to drone-related problems remains limited. This can be attributed to the complexity of drone-acquired data, the diversity of operational scenarios, and the high cost of expert labeling, all of which pose challenges for the direct adoption of active learning methodologies. Consequently, only a small number of studies have explicitly investigated the integration of active learning into drone-based applications.

Among these, (Hu et al., 2025) incorporated active learning into drone-related tasks by proposing a hierarchical active learning approach (HALD) for low-altitude drone-view object detection. The proposed method was evaluated on the VisDrone and CityPersons datasets, demonstrating that the effectiveness of sample informativeness estimation can be significantly improved through the use of four hierarchical levels.

In a related study, (Miao et al., 2024) applied active learning to the accurate identification of vegetation types using drone imagery. Their results showed that, under limited training data conditions, labeling costs could be reduced to approximately 20% while maintaining an average classification accuracy of 93.2%. The authors concluded that active learning effectively filters high-value samples that contribute most to AI model training, thereby substantially reducing annotation effort.

Methodology

This section outlines the system model or methodology employed in this research paper. Figure 1 presents the experimental procedures implemented in the study for drone classification in the RF domain.

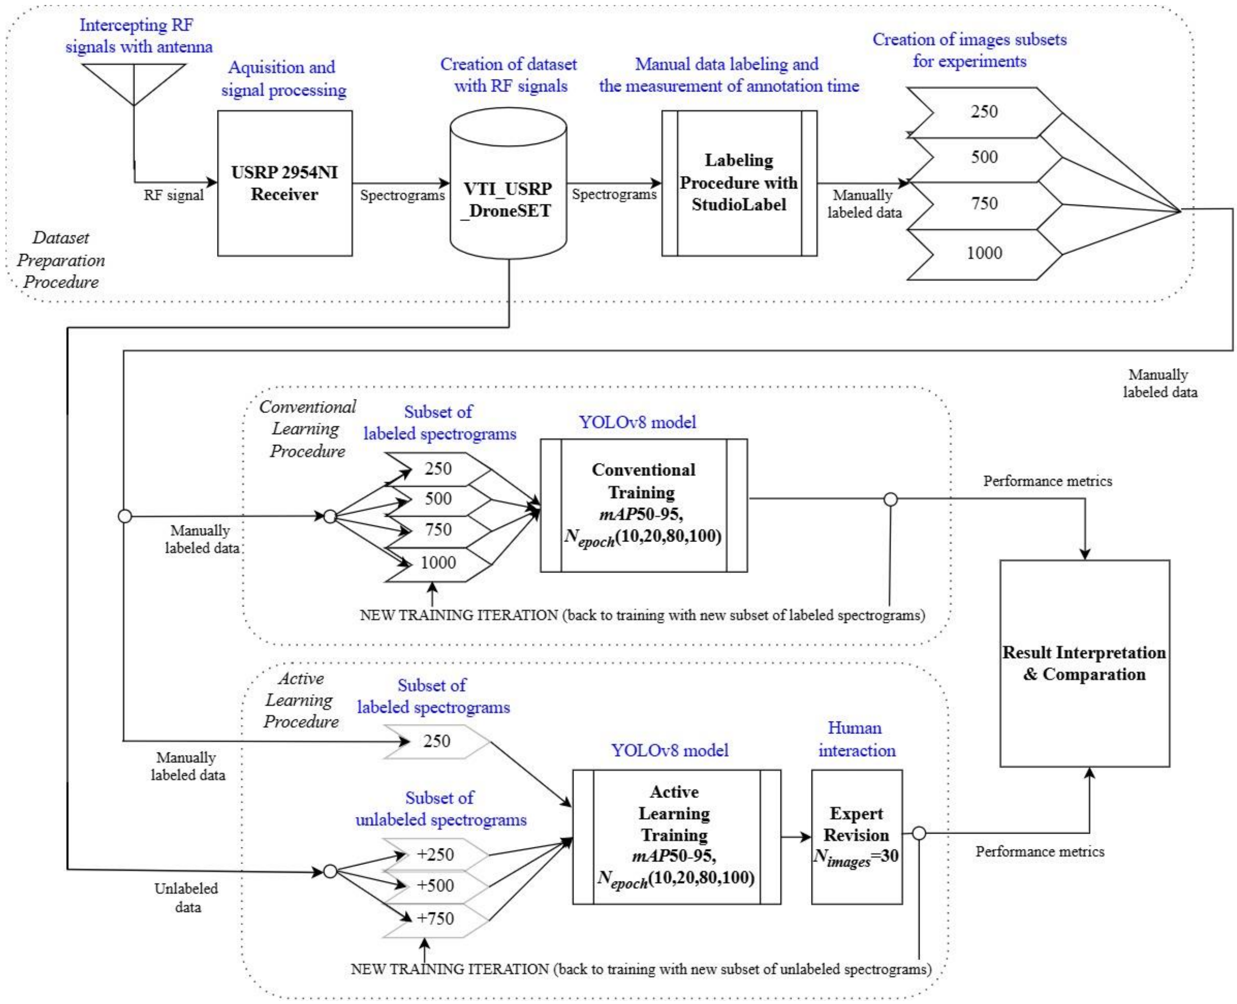

Figure 1. Methodology used in the research.

The described methodology involves three crucial steps: the procedure for dataset preparation with images – spectrograms of radio signals from drones (interception and acquisition of RF signals, creation of the dataset, labeling procedure and creation of images subsets for experiments), the training procedure (conventional supervised learning and active learning procedure), and the procedure of recording and explaining metrics for the best YOLO version 8 model which was used in this study (YOLOv8). (Redmon et al., n.d.)

3.1 Dataset preparation

Images – spectrograms were obtained from the receiver Universal Software Radio Peripheral (USRP) NI-2954 in realistic environments (outdoor) in the 2.4 GHz Industrial, Scientific, Medical (ISM) frequency band. It is important to note that the data were collected under outdoor operating conditions, capturing real-world signal environments and associated noise characteristics. The receiver performed the drones' RF signals and processed the raw in-phase and quadrature (I/Q) data to produce an output spectrogram matrix. Frequency acquisition was performed at a bandwidth of IB = 200 MHz with a sampling frequency of F_s = 150 MS/s. After acquiring the I/Q data, the field-programmable gate array (FPGA) module performed spectrogram calculations and morphological operations to enhance drone classification capabilities. For spectrogram calculation, the Short-Time Fourier Transform (STFT) with N_{DFT} = 2048 and a frame size of F_{size} = 1000 was used. For the morphological preprocessing of the spectrogram images, dilation and erosion operations were applied.The parameters associated with both operations, including the structuring element size and shape, were empirically determined based on the characteristics of the acquired spectrograms, with the objective of effectively suppressing noise while preserving salient signal structures. It is noteworthy to show that, in every situation, the drone operator was positioned more than two kilometers from the receiver. In contrast, the drones were in flight conditions towards the USRP NI-2954 receiver. These conditions were engaged on purpose to create more realistic scenarios in the outdoor environment. The receiver USRP NI-2954 produces three spectrograms per second, which are sent to the next step in the methodology used, the creation of the VTI_USRP_DroneSET dataset[3] with RF signals from drones and the labeling step.

The proposed approach utilizes six types of signals: "0" indicates ambient noise when there are no drones present; "LB" refers to signals from drones using the Lightbridge communication protocol; "OC20" and "OC40" describe OcuSync protocol signals with bandwidths of 20 MHz and 40 MHz, respectively. In addition, "2" stands for signals generated by two drones working simultaneously, while "3" denotes signals produced when three drones are flying at the same time. The ambient recordings include background noise (absent of intentional signals), as well as radio frequency emissions from Wi-Fi, Bluetooth, and 3G mobile phone systems. The classes labeled as "LB", "OC20", or "OC40"

[3] The data supporting the findings of this study are not publicly available but can be obtained from the authors upon reasonable request.

include radio signals originating from one DJI commercial drone using the respective communication protocols. Meanwhile, the classes for two and three drones comprise the radio signals emitted by two and three drones operating simultaneously. Finally, the VTI_USRP_DroneSET dataset includes 2,500 binary spectrogram images, each with a spatial resolution of 2048 × 1000 pixels, acquired using a USRP NI-2954 receiver. The spectrograms are represented as binary images by applying morphological operations to the original spectrograms. The resulting spectrograms stand for time–frequency signal distributions and serve as input data for subsequent analysis and classification tasks with the You Only Look Once (YOLO)model. Figure 2 presents two examples of manually labeled samples created with the open-source data labeling platform, Label Studio. (HumanSignal, 2025)

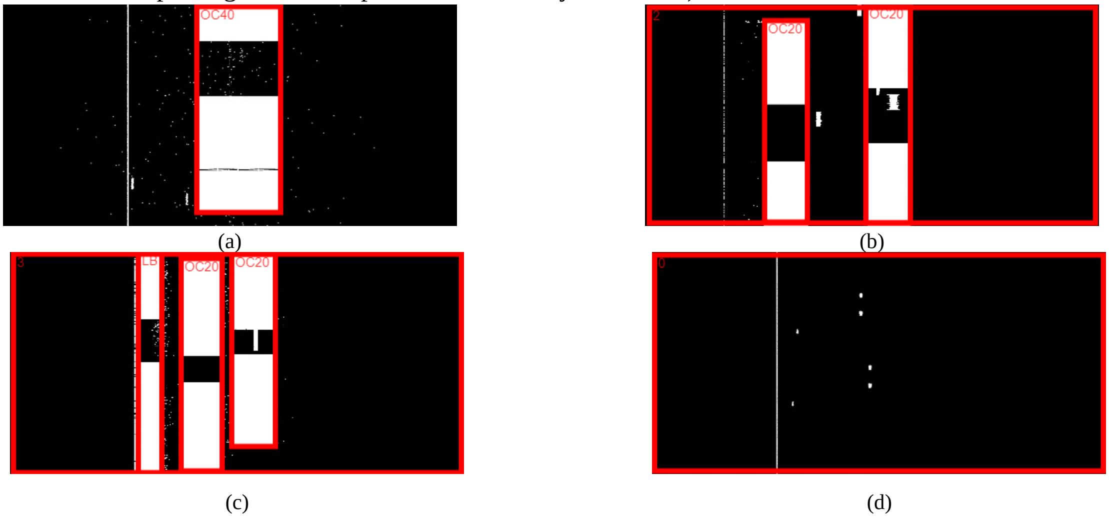

Figure 2. An example of spectrograms manually labeled:

(a) – one RF signal from drone (class “OC40”), (b) – two RF signals from drones (class “2”), (c) – three RF signals from drones (class “3”), and (d) – ambient (class “0”).

It should be noted that in occurrences where multiple RF signals appeared concurrently, distinct bounding boxes were assigned to each signal if they could be visually distinguished within the time–frequency domain, including situations involving partial overlap. This separation was further supported by morphological preprocessing, which enhanced the dominant structures of individual signals. During the sample labeling process, particular attention was devoted to the measurement of annotation time, as this metric is essential for an objective assessment of labeling effort. Accordingly, the time required to label samples belonging to different classes was systematically recorded. These measurements provide a necessary basis for the subsequent quantitative comparison of conventional supervised learning and active learning approaches. Four subsets of the VTI_USRP_DataSet, including 250, 500, 750, and 1,000 labeled images, respectively, were employed in both experimental scenarios. The total annotation time was empirically determined through direct measurement of the experts’ annotation activity. Two domain experts, referred to as E1 and E2, were assigned an equal number of images for labeling. The dataset consisted of spectrograms of radio-frequency signals emitted by unmanned aerial vehicles. For the ambient class (“0”), the entire spectrogram was annotated. For the “LB”, “OC20”, and “OC40” classes, only the video transmission signal was labeled. In the case of classes “2” and “3”, both individual radio signals and groups of signals corresponding to two or three simultaneous transmissions were annotated. As a result of this increased annotation complexity, the labeling time for classes “2” and “3” was higher than that needed for the ambient class or for classes involving a single drone transmission. The overall outcomes of the annotation procedure per image are reported in Table 1.

Table 1. Average annotation time per sample across different classes.

Classes	“0” ambient *[s]*	“LB”/ “OC20”/ “OC40” one drone *[s]*	“2” two drones *[s]*	“3” three drones *[s]*
Expert E1	4	9.5	15.2	22.1
Expert E2	4.3	9.6	15.7	23.0

The results summarized in Table 1 show that the labeling procedure is both time-consuming and labor-intensive, highlighting the practical challenges associated with manual annotation in real-world radio-frequency signal classification tasks. The average labeling time per sample varies across classes, ranging from 4.0 s to 23.1 s, depending on the number of bounding boxes required to be drawn on each spectrogram. The mean labeling time averaged over all classes is 12.93 s per sample across different classes, which corresponds to a total annotation time of approximately 12,938 s (3.6 h) for a subset of 1,000 images. Such a workload is difficult to accomplish within a single continuous annotation session due to the sustained level of attention required for detailed inspection of RF spectrograms. Consequently, in practical settings, the effective labeling duration is longer, requiring approximately three working days for expert annotators to complete the task alongside their regular professional duties.

3.2 Conventional Supervised versus Active Learning Procedure

Following the acquisition steps and the creation of properly labeled subsets of spectrograms, we conducted two distinct experiments: one focused on conventional supervised training and the other on an active learning methodology. In both cases, YOLOv8 was employed to ensure a fair comparison. The objective was to demonstrate the effectiveness of active learning using a limited labeled dataset, rather than to compare the influence of different AI model architectures on performance metrics such as precision, recall, or confidence. In the active learning setting, unlabeled samples were iteratively selected for expert annotation based on model uncertainty, prioritizing instances with the lowest prediction confidence.

Both the conventional supervised learning and active learning experiments were conducted using identical evaluation protocols to ensure a fair and consistent comparison. The conventional supervised training experiments employed 250, 500, 750, and 1,000 labeled spectrograms, while the active learning framework used 250 labeled spectrograms for initial training, followed by progressively increasing the number of unlabeled samples through three iterative querying (250, 500, and 750 unlabeled spectrograms). Both experiments were conducted with the training and test sets in a ratio of 80% to 20% of labeled images on the same GPU machine (NVIDIA RTX 4500 Ada Generation Graphics Card with 24GB graphics memory).

After each training iteration cycle, object detection performance metrics were systematically evaluated and recorded. In both experimental settings, detection performance was assessed using the mean Average Precision averaged over intersection over union (IoU) thresholds from 0.50 to 0.95 (mAP50–95) metric, which represents the mean average precision (mAP) averaged over multiple IoU thresholds ranging from 0.50 to 0.95 in increments of 0.05. This metric was selected due to its widespread adoption as a standard benchmark for object detection, particularly within the COCO evaluation protocol. The mAP50–95 value was computed by calculating the average precision (AP) for each object class at each IoU threshold and subsequently averaging the results across all thresholds and classes. The IoU metric quantifies the spatial overlap between predicted bounding boxes and corresponding ground-truth annotations, where an IoU threshold of 0.50 represents a lenient localization criterion, and 0.95 corresponds to a highly stringent requirement. As a result, the mAP50–95 metric provides a comprehensive evaluation of detection performance by jointly assessing classification accuracy and localization precision. In the context of spectrogram-based drone detection, higher mAP50–95 values indicate accurate detection and localization of RF signal regions, whereas comparatively higher AP50 values combined with lower mAP50–95 scores indicate correct signal detection with reduced bounding-box localization accuracy.

In all experiments, a consistent set of training parameters for YOLO was employed. The learning rate was set to η = 0.01, controlling the step size for AI model weight updates. A batch size of N_{batch_sizre} = 16, samples were used for each training iteration cycle, $N_{iteration}$ = {1, 2, 3}. The number of training epochs varied across experiments and was set to N_{epoch} = {10, 20, 80, 100}, corresponding to the total number of passes through the training dataset. Model optimization was performed using the stochastic gradient descent (SGDM) optimizer. In addition, thirty images were selected for expert review following each training procedure.

3.3 Results recording and comparison

The final stage of the proposed methodology consisted of recording, analyzing, and interpreting the performance metrics of the best-performing YOLO model. This stage was conducted after the completion of the training process and involved the systematic storage of training parameters (arguments), the learned YOLO model parameters (weights), and the distribution of class labels. In addition, label correlograms, both normalized and non-normalized confusion matrices, and a comprehensive set of performance curves-including the F1–confidence, precision–confidence, precision–recall, and recall-confidence curves-were generated and analyzed. Furthermore, prediction batches corresponding to each training iteration were preserved to enable detailed qualitative assessment. All mAP50–95 results obtained across training iterations and for different values of N_{epoch} were also stored in tabular form to facilitate systematic comparison and subsequent quantitative analysis.

Results

The overall experimental results are reported in Tables 2 and 3. Table 2 presents the performance obtained using classical learning as a function of the number of labeled spectrogram images employed for YOLO model training and the number of epochs per training iteration cycle, thereby enabling an assessment of performance scalability with respect to annotation effort.

Table 2. Overall mAP50–95 results for conventional supervised training across different training set sizes.

N_{epoch}	Number of labeled images			
	250	500	750	1,000
10	0.96	0.97	0.98	0.98
20	0.97	0.98	0.98	0.98
80	0.98	**0.99**	**0.99**	**0.99**
100	**0.99**	**0.99**	**0.99**	**0.99**

The results in Table 2 demonstrate that conventional supervised training achieves progressively higher *mAP50–95* values with increasing numbers of labeled images and training epochs (N_{epoch}). For a fixed number of epochs, performance improves as the training set size increases from 250 to 1,000 labeled images, while for a fixed dataset size, longer training leads to higher detection accuracy. At N_{epoch} = 80 and 100, performance approaches saturation, with *mAP50–95* ≈ 0.99 achieved across all dataset sizes, including the smallest set of 250 labeled images. The marginal gains observed beyond five hundred images and 80 epochs indicate diminishing returns in performance relative to increased labeling effort and training duration. These results highlight the inefficiency of purely supervised learning in scenarios with limited labeling resources, motivating the use of more data-efficient approaches such as active learning.

Table 3 reports the results achieved using the proposed active learning framework, in which 250 labeled spectrogram images were used for initial YOLO model training, and performance is evaluated as a function of the number of labeled images incorporated during subsequent active learning iteration cycles.

Table 3. Overall mAP 50-95 results for active training with 250 labeled images for initial training.

N_{epoch}	Initial training with 250 labeled im-ages	Number of unlabeled images		
		250	500	750
10	0.96	0.94	0.97	0.96
20	0.97	0.97	0.98	0.98
80	0.98	**0.99**	**0.99**	**0.99**
100	**0.99**	**0.99**	**0.99**	**0.99**

The results in Table 3 demonstrate the effectiveness of the proposed active learning framework when initialized with 250 labeled images and progressively augmented with additional samples selected from the unlabeled pool. Overall, the results show that active learning enables rapid performance improvements with a limited number of labeled samples, achieving *mAP50–95* values comparable to those of conventional supervised training. At lower training epochs (N_{epoch} = 80), performance exhibits minor variability depending on the number of queried unlabeled images, with *mAP50–95* values ranging between 0.94 and 0.97. As the number of training epochs increases to 20, performance stabilizes and improves consistently, reaching *mAP50–95* ≈ 0.98 for larger queried sets. For N_{epoch} = 80 and 100, near-optimal performance (*mAP50–95* = 0.99) is achieved across all configurations, regardless of the number of unlabeled images incorporated during active learning iterations.

The high detection accuracy observed for the presented subsets can be attributed, in part, to the quality of the VTI_USRP_DroneSET spectrograms. The dataset exhibits minimal interference and well-defined signal patterns, which are further enhanced through the application of morphological operations during preprocessing. These operations suppress background noise and emphasize relevant signal structures, thereby facilitating effective feature learning and contributing to the excellent classification performance achieved with a limited number of labeled samples. Consequently, enhanced data quality and the proposed training strategy play a crucial role in achieving high accuracy under constrained labeling conditions. These results indicate that the proposed active learning strategy efficiently exploits informative samples from the unlabeled data pool, allowing high detection performance to be attained with reduced labeling effort. This behavior is particularly advantageous in practical deployment scenarios, where labeled data is scarce, and new samples become available incrementally.

A comparison of Tables 2 and 3 shows that active learning achieves performance comparable to conventional supervised training while requiring fewer labeled samples. Although both approaches converge to similar mAP50–95 ≈ 0.99 values at higher numbers of training epochs, the active learning framework reaches this performance more efficiently by selectively incorporating informative samples from the unlabeled pool. These results highlight the advantage of active learning in reducing annotation effort without compromising detection accuracy. Figure 3 presents a comparative graphical analysis of mAP50–95 performance obtained using conventional supervised learning and active learning methodologies across different dataset configurations and training epochs.

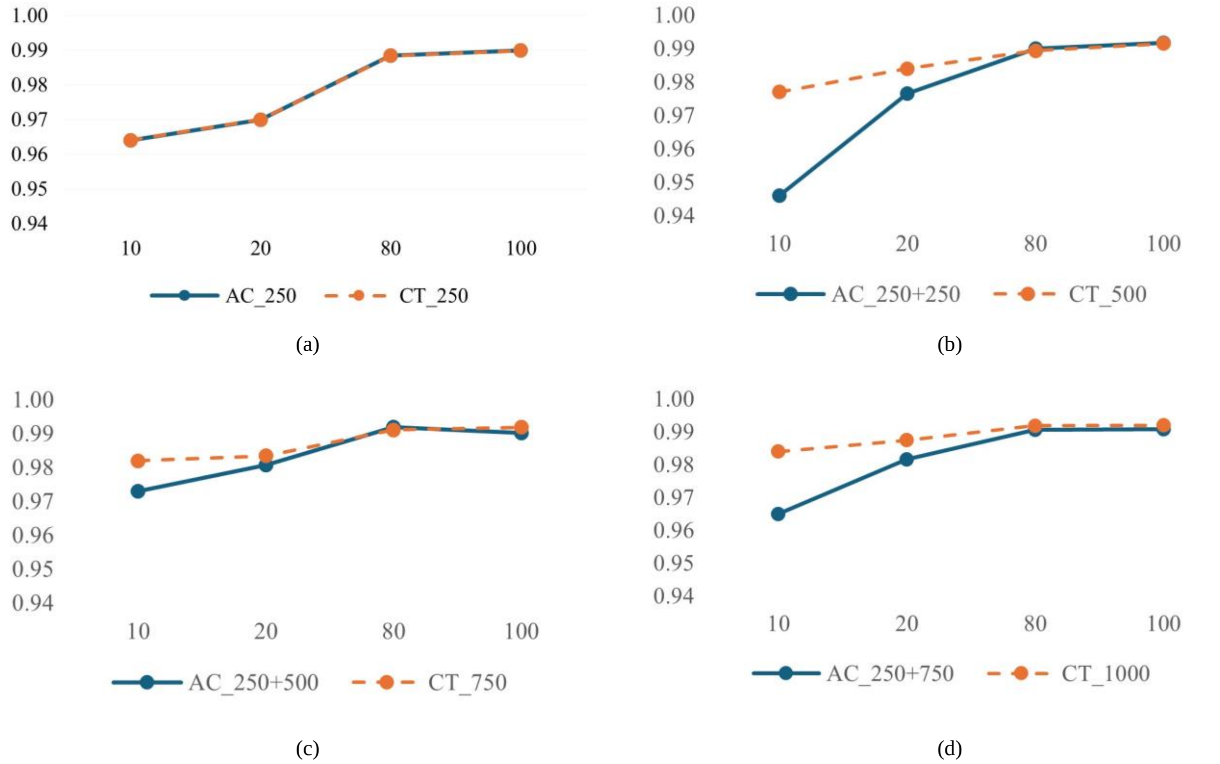

Figure 3. Comparison of mAP50–95 results for conventional supervised and active learning approaches:[4]

(a) – the first subset (AC_250 and CL_250 obtained with 250 labeled images), (b) – the second subset (AC_250+250 with additional 250 unlabeled images and CL_500 with 500 labeled images), (c) – the third subset (AC_250+5-0 with additional 500 unlabeled images and CL_750 with 750 labeled images), and (d) – the fourth subset (AC_250+750 with additional 750 unlabeled images and CL_1000 with 1000 labeled images).

Figure 3(a) shows the results for the initial subset containing 250 labeled images, where both conventional supervised learning and active learning achieve the same performance. Figures 3(b) to (d) illustrate scenarios in which the active learning framework is initialized with 250 labeled images and subsequently augmented with 250, 500, and 750 additional samples, respectively, selected from the unlabeled pool, while the conventional supervised learning approach is trained using 500, 750, and 1,000 labeled images. In all cases, performance is reported as a function of the number of training epochs. The results presented in Figure 3 show that detection performance improves with increasing training duration for both conventional supervised learning and active learning approaches, followed by a clear saturation trend. Across all dataset configurations, N_{epoch} = 80 consistently emerges as an optimal training duration, as extending the training to 100 epochs results in only negligible improvements in mAP50–95. From a practical perspective,

[4] AL denotes active learning, while CL refers to conventional supervised learning.

this observation has important implications for training cost reduction. Limiting the training process to 80 epochs allows near-optimal detection performance to be achieved while reducing computational time, energy consumption, and overall training cost. When combined with the reduced annotation effort enabled by the active learning framework, this optimal epoch selection further enhances the efficiency of the proposed methodology. Consequently, the results demonstrate that high classification performance can be attained with both fewer labeled samples and lower training costs, making the proposed approach particularly suitable for real-world RF-based drone classification systems.

Finally, we performed the real-time testing of the trained YOLO model with the USRP NI-2954 receiver in realistic environments (outdoor) conditions. Figure 4 shows the output results of the YOLO model for the corresponding spectrograms.

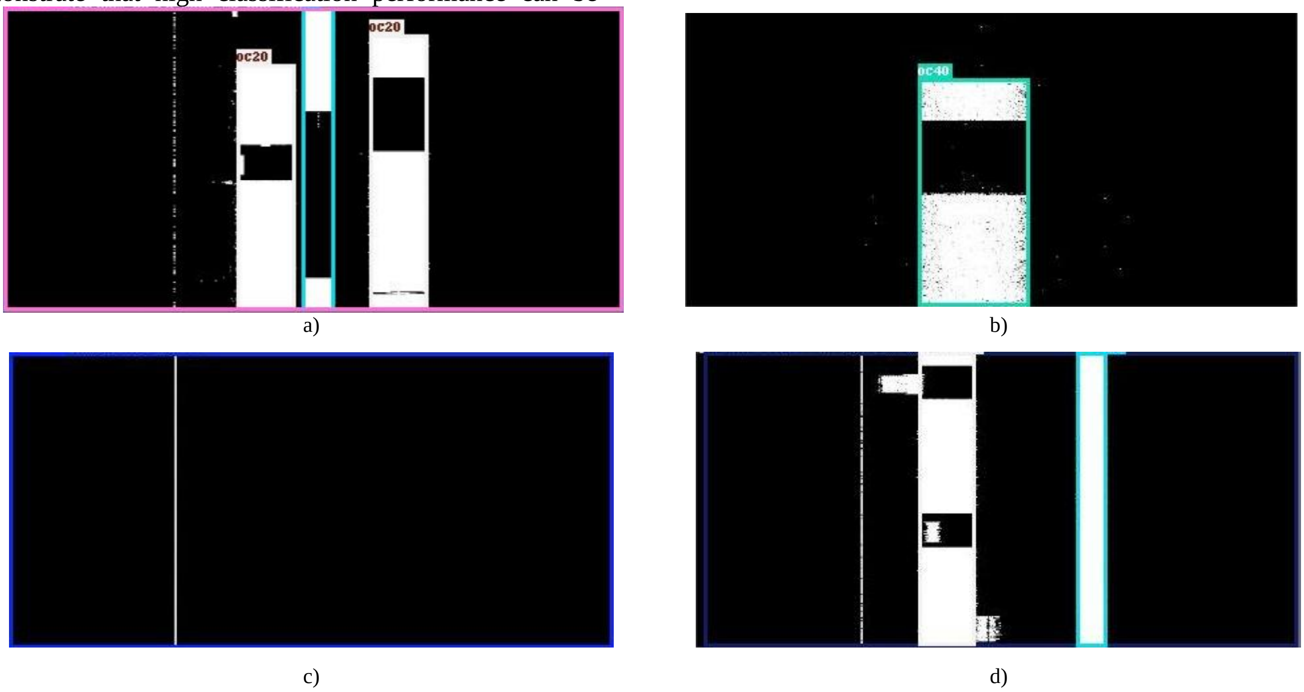

Figure 2. An example of spectrograms manually labeled:

a – three RF signals from drones (classes “3”, “OC20”, “LB”, and “OC40”), b – one RF signal from drone (class “OC40”), c – ambient (class “0”), and d – two RF signals from drones (classes “2”, “OC20”, and “LB”)

It should be noted that the practical implementation leveraged the existing methodology, with the distinction that the acquired spectrograms were obtained through separate experimental measurement sessions and were immediately forwarded to the trained model obtained via active learning. Specifically, the USRP NI-2954 receiver directly transmitted the spectrograms to the computing platform used for both model training and inference, enabling real-time evaluation of the detection and identification (classification) capabilities for drone radio-frequency signals. Figure 4(a) illustrates an example in which the best-performing trained YOLO model classifies a scenario involving three simultaneously operating drones. In this case, the model correctly detects the presence of three drones, indicated by a purple bounding region enclosing the entire spectrogram containing the three RF signals. Additionally, two instances of the “OC20” class (outlined in white) and one instance of the “LB” class (outlined in light blue) are identified. Figure 4(b) presents a scenario in which a single drone is active. The YOLO model correctly identifies one instance of the “OC40” class, highlighted by a light green bounding box. Figure 4(c) shows a case in which no drone RF signals are present. The model correctly classifies the spectrogram as belonging to the “0” class, indicated by a blue bounding region surrounding the entire spectrogram, where interference at a single frequency is observed. Finally, Figure 4(d) illustrates a scenario with two drones operating simultaneously. In this case, the model correctly identifies one instance of the “OC20” class (outlined in white) and one instance of the “LB” class (outlined in light blue).

The examples presented in Figure 4 demonstrate the practical effectiveness of the proposed active learning–based framework for RF drone detection and classification. The trained YOLO model accurately identifies scenarios involving single and multiple drones, as well as cases in which drone RF signals are absent, under real-world measurement conditions. These results confirm the capability of the proposed approach to support reliable real-time detection and classification of drone RF signals in operational environments.

From a time–cost perspective, these findings have important practical implications. The active learning framework achieves detection performance comparable to that of conventional supervised learning while requiring only one quarter of the labeled samples. This leads to a substantial reduction in annotation time and human effort, directly improving cost efficiency. Furthermore, constraining the training process decreases computational time and energy consumption, thereby improving both the efficiency and deployability of the proposed approach. The findings demonstrate a distinct trade-off between training duration and annotation costs; although extended training slightly enhances detection accuracy, performance stops at approximately 80 epochs. By integrating this optimal training period with the active learning framework's reduced labeling requirements, the proposed approach substantially reduces both computational and annotation expenses without dropping mAP50–95 performance. In practical RF-based drone classification scenarios, this balance enables fast deployment, lower energy demands, and minimized expert labeling efforts, ultimately offering a viable and cost-effective solution for real-world antidrone implementation.

Conclusion

This paper presented a tailored active learning framework for drone classification in the radio-frequency domain, addressing the fundamental challenges of limited labeled data and high annotation costs in practical antidrone systems. Moreover, the iterative expert-in-the-loop process enables post-hoc inspection of selected samples, which may facilitate the identification of potential labeling inconsistencies from earlier annotation stages. By integrating expert knowledge into an iterative sample-selection process, the proposed methodology enables efficient training of a YOLO-based detector while reducing labeling requirements. Experimental evaluation on the VTI_USRP_DroneSET dataset, collected under realistic outdoor conditions in the 2.4 GHz frequency band, demonstrated that the active learning approach achieves detection performance comparable to conventional supervised learning as measured by the mAP50–95 metric. Notably, equivalent accuracy was attained using only 25% of the labeled samples, and an optimal training duration of 80 epochs was identified, resulting in significant reductions in annotation effort, computational time, and energy consumption. Real-time testing with the USRP NI-2954 receiver further confirmed the robustness and operational viability of the proposed framework in detecting and classifying single and multiple drone scenarios, as well as correctly identifying the absence of drone RF signals. To the authors' knowledge, this work represents one of the first systematic investigations of active learning for RF-based drone classification under real measurement conditions. The proposed framework therefore offers a scalable, data-efficient, and cost-effective solution for next-generation antidrone systems. Despite these promising results, the study is limited to a single dataset and frequency band, and the active learning process still relies on expert-in-the-loop annotation. Future work will focus on extending the framework to multi-band RF environments, exploring fully automated or hybrid query strategies, and validating performance in more congested and adversarial electromagnetic conditions.

Acknowledgement

This work was supported by the Ministry of Science, Technological Development and Innovations (Serbia), Contract No.451-03-137/2025-03/200325 and Contract No.451-03-136/2025-03/200135.

References

[1] Cevik, M., Ergun, M. A., Stout, N. K., Trentham-Dietz, A., Craven, M., & Alagoz, O. (2016). Using active learning for speeding up calibration in simulation models. *Journals.Sagepub.ComM Cevik, MA Ergun, NK Stout, A Trentham-Dietz, M Craven, O AlagozMedical Decision Making, 2016•journals.Sagepub.Com*, *36*(5), pp. 581–593. https://doi.org/10.1177/0272989X15611359

[2] Gal, Y., Islam, R., & Ghahramani, Z. (2017). *Deep Bayesian Active Learning with Image Data* (pp. 1183–1192). PMLR. https://proceedings.mlr.press/v70/gal17a.html

[3] Hu, H., Han, T., Wang, Y., Zhong, W., Yue, J., & Zan, P. (2025). Hierarchical Active Learning for Low-Altitude Drone-View Object Detection. *International Journal of Computer Vision*, *133*(3), pp. 1140–1152. https://doi.org/10.1007/s11263-024-02228-y

[4] HumanSignal, Inc. (2025). *Label Studio*. https://labelstud.io/

[5] Ilić, V. (2024). Reducing Manual Labeling Effort by Identifying the Most Informative Unlabeled Data via Active Learning. *International Conference on Applied Internet and Information Technologies (AIIT 2024)*, pp. 35–43.

[6] Ilić, V., & Tadić, J. (2022). Active learning using a self-correcting neural network (ALSCN). *Applied Intelligence*, *52*(2), 1956–1968. https://doi.org/10.1007/S10489-021-02515-Y/METRICS

[7] Kumar, P., Technology, A. G.-J. of C. S. and, & 2020, undefined. (2020). Active learning query strategies for classification, regression, and clustering: A survey. *SpringerP Kumar, A GuptaJournal of Computer Science and Technology, 2020•Springer*, *35*(4), pp. 913–945. https://doi.org/10.1007/S11390-020-9487-4

[8] Lewis, D. D., & Gale, W. A. (1994). A Sequential Algorithm for Training Text Classifiers. Proceedings of the 17th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 1994, pp. 3–12. https://doi.org/10.1007/978-1-4471-2099-5_1

[9] Menke, M., Wenzel, T., & Schwung, A. (2024). Bridging the gap: Active learning for efficient domain adaptation in object detection. *Expert Systems with Applications*, p. *254*, 124403. https://doi.org/10.1016/J.ESWA.2024.124403

[10] Miao, S., Wang, C., Kong, G., Yuan, X., Shen, X., & Liu, C. (2024). Utilizing active learning and attention-CNN to classify vegetation based on UAV multispectral data. *Scientific Reports*, *14*(1), 31061. https://doi.org/10.1038/s41598-024-82248-3

[11] Mittal, S., Niemeijer, J., Çiçek, Ö., Tatarchenko, M., Ehrhardt, J., Schäfer, J. P., Handels, H., & Brox, T. (2025). Realistic Evaluation of Deep Active Learning for Image Classification and Semantic Segmentation. *International Journal of Computer Vision*, *133*(7), pp. 4294–4316. https://doi.org/10.1007/s11263-025-02372-z

[12] Redmon, J., Divvala, S., Girshick, R., & Farhadi, A. (n.d.). *You Only Look Once: Unified, Real-Time Object Detection*. http://pjreddie.com/yolo/

[13] Ren, P., Xiao, Y., Chang, X., Huang, P. Y., Li, Z., Gupta, B. B., Chen, X., & Wang, X. (2022a). A Survey of Deep Active Learning. *ACM Computing Surveys*, *54*(9), pp. 1–40. https://doi.org/10.1145/3472291/REFERENCES

[14] Ren, P., Xiao, Y., Chang, X., Huang, P.-Y., Li, Z., Gupta, B. B., Chen, X., & Wang, X. (2022b). A Survey of Deep Active Learning. *ACM Computing Surveys*, *54*(9), pp. 1–40. https://doi.org/10.1145/3472291

[15] Sener, O., & Savarese, S. (2018, February 15). Active Learning for Convolutional Neural Networks: A Core-Set Approach. *ICLR 2018 Conference Track 6th International Conference on Learning Representations*.

[16] Settles, B. (2009). Computer Sciences Department Active Learning Literature Survey.

[17] Settles, B. (2012). *Active Learning*. Springer International Publishing. https://doi.org/10.1007/978-3-031-01560-1

[18] Tong, S., & Koller, D. (2001). Support Vector Machine Active Learning with Applications to Text Classification. *Journal of Machine Learning Research*, pp. 45–66.

[19] Wang, H., Jin, Y., & Doherty, J. (2017). Committee-Based Active Learning for Surrogate-Assisted Particle Swarm Optimization of Expensive Problems. *IEEE Transactions on Cybernetics*, *47*(9), pp. 2664–2677. https://doi.org/10.1109/TCYB.2017.2710978

[20] Wang, Y., Ilic, V., Li, J., Kisačanin, B., & Pavlovic, V. (2023). ALWOD: Active Learning for Weakly-Supervised Object Detection. *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 6459–6469.

[21] Zhan, X., Liu, H., Li, Q., & Chan, A. B. (2021). A Comparative Survey: Benchmarking for Pool-based Active Learning.

[22] Zhao, C., Qin, B., Feng, S., Zhu, W., Sun, W., Li, W., & Jia, X. (2023). Hyperspectral Image Classification with Multi-Attention Transformer and Adaptive Superpixel Segmentation-Based Active Learning. *IEEE Transactions on Image Processing*, *32*, pp. 3606–3621. https://doi.org/10.1109/TIP.2023.3287738

Received: 21.01.2026.
Accepted: 10.02.2026.

Primena aktivnog učenja za klasifikaciju bespilotnih vazduhoplova u radio frekvencijskom domenu

Klasifikacija bespilotnih vazduhoplova (dronova) u radio frekvencijskom domenu je kritična sposobnost za moderne antidron sisteme. Međutim, razvoj pouzdanih modela veštačke inteligencije u ovom domenu otežan je složenošću i potrebama za određeni nivo stručnog domenskog znanja prilikom označavanja podataka. Ovaj izazov je posebno izražen u analizi radio frekvencijskih signala, gde su označeni skupovi podataka obično mali, delimično neoznačeni i kontinuirano se povećavaju tokom primene antidron sistema. Ovaj rad predlaže prilagođenu strategiju aktivnog učenja za klasifikaciju bespilotnih vazduhoplova u radio frekvencijskom domenu, integrišući eksperte – operatere u iterativni proces učenja koji selektivno ispituje najznačajnije neoznačene uzorke. Davanjem prioriteta uzorcima, predloženi okvir ima za cilj postizanje visokih performansi klasifikacije uz značajno smanjenje vremena za označavanje. Pristup je testiran korišćenjem VTI_USRP_DroneSET skupa podataka, koji se sastoji od radio frekvencijskih spektrograma snimljenih u realnim uslovima na otvorenom unutar 2,4 GHz frekvencijskog opsega. Eksperimentalni rezultati pokazuju da predložena strategija aktivnog učenja postiže performanse mAP50–95 uporedive sa konvencionalnim nadgledanim učenjem, dok zahteva samo jednu četvrtinu označenih podataka. Rezultati potvrđuju da aktivno učenje omogućava efikasnu klasifikaciju vazduhoplova u radio frekvencijskom domenu bez ugrožavanja tačnosti klasifikacije. Osim toga, optimalne performanse predloženog pristupa se dobijaju sa trajanjem obuke od 80 epoha, smanjujući zahtevano vreme i računarske potrebe za obučavanje. Ovi rezultati potvrđuju da predložena strategija aktivnog učenja pruža efikasno i isplativo rešenje za klasifikaciju bespilotnih vazduhoplova (dronova) u radio frekvencijskom domenu i da je pogodna za primenu u realnom vremenu u stvarnim antidron sistemima inicijalno obučenim na malom skupu podataka, ali sa tendencijom da se stalno povećavaju tokom operativne upotrebe.

Ključne reči: **aktivno učenje, antidron, veštačka inteligencija, klasifikacija, dron.**