Domain Adaptation for Deep Learning based Signal Detection in the High Frequency Range

Yasin Ciritci*†, Ulrike Buhl†, Georg Fischer*
*Institute for Electronics Engineering, University of Erlangen-Nuremberg, Erlangen 91058, Germany
†Rohde & Schwarz GmbH & Co. KG, Munich 81671, Germany
Email of corresponding author: yasin.ciritci@fau.de

***Abstract*—Signal detection, especially in crowded frequency ranges, poses high challenges to any signal analysis solution. Therefore, in recent years, approaches apply deep learning object detection methods to the specific signal detection task. Although in principle huge similarities between *object* detection and *signal* detection exist, applying specific object detection solutions is not always straightforward though. In this work, we focus on the specific challenge of domain adaptation for deep learning signal detection. We adapt ideas from domain adaptation for object detection and introduce a novel method called Domain Adaptation with Multi-Class domain classifiers (DA-MC). This method surpasses the performance of domain adaptation with class-agnostic domain classifiers. The latter prove to be inapplicable for signal detection domain adaptation. We train and evaluate our proposed method on a synthetic test scenario. Our results prove that the domain shift for signals can effectively be solved. We furthermore evaluate an actual domain shift signal scenario using an over-the-air dataset and present promising improvements over a baseline training.**

***Index Terms*—Detection, Domain Adaptation, Spectrum Monitoring, Deep Learning, Signal Detection.**

I. Introduction

Research in object detection increased in the recent years drastically. One of the challenges is retrieving sufficient training data for the neural network. Compared to object classification, object detection, alongside the images and a class information, requires so-called *bounding boxes* corresponding to the location information of each object in the image. Hence, creating the training dataset sufficiently is expensive. Furthermore, due to limited amount of training data, the variability of trained features is limited as well. Thus, performance of the neural network might drop for unseen data. One example for such a drop are adverse weather conditions like fog that occur in unseen data but not in the trained data. In this case, the detection performance degrades compared to normal weather conditions. In principle, the performance difference between the training dataset and the test dataset, referred to as *source domain* and *target domain*, in the following, is known as *domain shift* [1], [2].

This domain shift is also relevant if deep learning detection is applied to the area of signal detection. The terminology signal detection thereby summarizes methods that identify and localize an a priori unknown amount of signals in a certain frequency band. The terminology *localization* in this sense refers to identifying the position of a signal in a time-frequency representation of the frequency band. The localization outcome contains information w.r.t. center frequency, bandwidth, and duration for each signal. One option to visualize frequency ranges over time is with a spectrogram. It represents the power of signals in a time and frequency range. Hence, the signal detection task is similar to object detection since the goal is to detect and localize signals (objects) in a spectrogram (image). While there is research regarding signal detection with deep learning [3], [4], to the best of our knowledge no solutions exist that focus on the specific field of domain shift in deep learning based signal detection. In this work, we aim at solving the domain adaptation problem for signal detection in the high frequency range. We introduce our novel method called Domain Adaptation with Multi-Class domain classifiers (DA-MC). This domain adaptation method applies a multi-class domain classifier for signal detection, which to our best knowledge is the first of its kind. We introduce a synthetic test scenario to show that the selected method is principally applicable to signal detection. We show the results for this test scenario as well as the results for an over-the-air recording scenario.

II. State of the Art

In recent years, multiple publications addressed the topic of domain adaptation for deep learning object detection. We sort these publications into four major categories and summarize their main ideas shortly in the following.

1) Adversarial Learning Methods: The first method forces the neural network to learn domain-invariant features during training [1]. This is achieved with a *Domain Classifier Network* (DCN) that is attached to the original detection network during training. On the one hand, the DCN is supposed to distinguish between input from the source domain versus input from the target domain and therefore learns domain-variant features. On the other hand, the detection network should learn domain-invariant features to detect the incoming data from the source and target domain equally well. The latter is achieved with an adversarial training strategy, which is explained in section IV with more details.

2) Self-Learning Methods: With this method, a neural network is trained with labeled data from the source domain. Then, the trained network is utilized to estimate corresponding labels for images, which are then added to the training dataset [6]. The crucial part for self-learning is the training data selection, which is selected based on a threshold [7], [8].

3) Image-To-Image Translation Methods: This approach generates target-like images from the source domain images with a Generative Adversarial Network (GAN) [9]. These target-like images are then added to the training dataset.

4) Mean Teacher Methods: This method works with two identical networks, referred to as *teacher* and *student* network. They are trained with similar, but slightly modified, images [10]. During the training, the student model is optimized with a loss function that punishes deviations from predictions of the teacher model. In turn, the teacher model is updated incrementally with the weights of the student model.

The approaches of all four categories have in common that they are developed for *object* detection networks with real world *images*. Hence, not every method is applicable for signal detection, since spectrograms cannot be treated as images and common image data augmentation methods, such as changing the saturation or rotating the image, are ineffective. The image-to-image translation and mean teacher methods modify images by means of data augmentation. Therefore, these two methods are not applicable for signal detection. Self-learning methods are suitable for signals, since they do not rely on generated data. However, the base performance of the neural network must be sufficient to obtain reliable bounding boxes for the new target domain images. In contrast, the adversarial learning method can be applied without restrictions, since it aims to generalize the feature space between source and target domain. Therefore, we focus on the adversarial learning approach, which is the most suitable method for signal detection.

III. Dataset Generation

For the dataset generation we collect multiple individual signals in time domain for each signal category. These signals are then positioned randomly w.r.t. the frequency and transformed into the frequency domain. The monochromatic power levels of the spectrogram are subsequently converted into an RGB image. In a last step, the corresponding bounding box and class labels are generated accordingly. We define these class labels from the signal categories in the high frequency range. We train with up to ten signal categories to ensure the principal functionality of the training. For the domain adaptation we focus on the following signal categories: Amplitude Modulation (AM), Single-Sideband modulation (SSB), Frequency-Shift-Keying with two tones (FSK2), signals with continuous rectangular power distribution (RECT) and signals with multi-channel characteristics (MLTC).

IV. Domain Adaptation for Signal Detection

Adversarial learning methods train a DCN such that it is able to distinguish between images from a labeled source dataset, referred to as $D_{\mathcal{S}}$ in the following, and images from a unlabeled target dataset, referred to as $D_{\mathcal{T}}$. We thereby focus on semi-supervised domain adaptation, i.e., labels only exist for images in the source dataset $D_{\mathcal{S}}$.

The domain adaptation approaches based on adversarial learning require at least one DCN [11], [12]. However, employing a class-agnostic domain classifier that does not distinguish between various classes can lead to a misalignment of features as shown in Fig. 1. Instead of matching the features of the AM class from $D_{\mathcal{S}}$ to the features of the AM class in $D_{\mathcal{T}}$, they are erroneously adapted to the features of the SSB class in the target dataset since a class-agnostic domain classifier can not learn distinct features for each class. This limitation is addressed in [13] for a classification network. There, the authors introduce a *multi-class domain classifier* where each class has its dedicated domain classifier. Therefore, the domain classifiers learn specific features for each class.

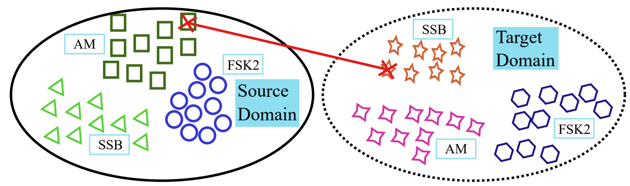

Fig. 1. Feature alignment can fail with a class-agnostic domain classifier [13].

In this work, we propose a domain adaptation network with multi-class domain classifiers for signal detection. The overview structure of the newly proposed architecture is shown in Fig. 2. We use a modified YOLOv4 architecture as backbone (marked in green in Fig. 2) and as detection head of the detection network [14]. The output of the modified YOLOv4 architecture contains beside others a varying amount of bounding box estimates and, per box, a corresponding class estimate. The boxes are subsequently fed into a Region of Interest (RoI) pooling layer. This layer, marked in light red in Fig. 2, extracts and maps the features at the positions of the bounding box estimates to predefined dimensions [12]. Its output is subsequently sorted per class by the class estimates that correspond to the respective bounding boxes. We further introduce a domain classifier for each of the K signal classes in the network. Each of the multi-class domain classifiers thereby consists of the following blocks, highlighted in light blue: the Gradient Reversal Layer (GRL), two fully connected layers (FC) and the Domain Classifier Layer (DCL). To enhance the stability of the network against domain shifts at different scales, we use three feature maps from the network at different scales for each class as in [11]. This results in $3 \times \mathrm{K}$ domain classifiers.

The loss of the described architecture in Fig. 2 consists of two parts: The domain adaptation loss and the detection loss. First, we introduce the domain adaptation loss L_{da} for all domain classifiers in (1)

$$L_{\text{da}} = -\sum_{i,j,k,l} \left[D_i \log p_{i,j}^{k,l} + (1 - D_i) \log(1 - p_{i,j}^{k,l}) \right] \quad (1)$$

Here, l, $(l = 1, 2, 3)$ represents a multi-class domain classifier at the l-th feature map scale. The k-th domain classifier, $k = 1, 2, ..., K$ represents the domain classifier of the k-th class within one multi-class domain classifier l. Then, $p_{i,j}^{k,l} \in [0, 1]$ is the predicted probability of the k-th domain classifier. This probability is per bounding box estimate j

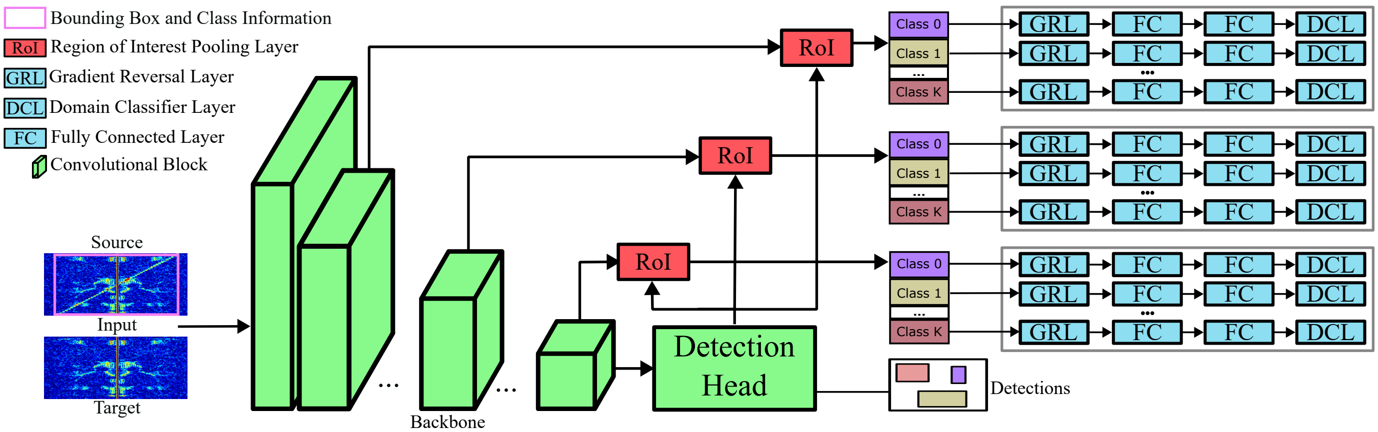

Fig. 2. Proposed Domain Adaptation with Multi-Class domain classifiers architecture. The domain adaptive part is utilized during the training process.

within the image i. Furthermore, D_i is the domain label of the i-th image that is known during train time. In particular, $D_i = 0$ marks images from the source domain and $D_i = 1$ marks the corresponding target domain images.

The domain classifier network learns domain-invariant features by minimizing the loss in (1). However, the same loss must be maximized for the detection network to learn the common features of the source and target domain. This contrary goal is achieved with the GRL [5], which is positioned between the domain classifiers and the backbone. During forward propagation, the GRL works as an identity operation by passing the incoming features through. However, during backpropagation, the domain adaptation loss L_{da} is minimized, beginning at the DCL until the GRL, to learn domain-invariant features for the domain classifier. Then, the sign of the gradient is flipped at the GRL, which effectively maximizes L_{da} for the backbone network during the backpropagation. Therefore, the backbone is adjusted to learn indistinguishable features.

The second part of the total loss consists of the detection loss L_{det}, which is the default loss from YOLOv4 [14]. The detection loss L_{det} is only trained with the dataset $D_{\mathcal{S}}$ from the source domain, while the domain adaptation loss L_{da} is trained with $D_{\mathcal{S}}$ and $D_{\mathcal{T}}$. Hence, the total loss for the domain adaptive network is $L = L_{\text{det}} + \lambda\, L_{\text{da}}$, where $\lambda \in [0, 1]$ is a balancing factor to control the influence of the domain classifier network and is set to $\lambda = 1$ during our experiments.

V. Experiments

A. Experiment Setup

We present results for our newly proposed signal detection domain adaptation network in the following.

1) Training Dataset: Besides the aforementioned datasets, there is also the so-called *oracle target dataset* $D_{\mathcal{T},\mathcal{O}}$, which combines $D_{\mathcal{T}}$ with corresponding labels. This dataset serves only to compare our network with a so called *oracle*. This oracle has perfect knowledge about the content of the target data during training time, thus representing the best result possible for the domain adaptation process. There are three combinations for the dataset utilized during the training:

- *Baseline*: Training with labeled source dataset $D_{\mathcal{S}}$ only
- Domain Adaptation: Training with labeled source dataset $D_{\mathcal{S}}$ and unlabeled target dataset $D_{\mathcal{T}}$
- Oracle: Training with labeled source dataset $D_{\mathcal{S}}$ and labeled target dataset $D_{\mathcal{T},\mathcal{O}}$

2) Dataset Content: In order to evaluate the performance of our network, we use two ways to gain the signal content. The resulting datasets strongly differ from their principle area of application. The first pair of source and target datasets, $D_{\mathcal{S}_{\text{SYN}}}$ and $D_{\mathcal{T}_{\text{SYN}}}$, respectively, is a set with signals that are synthetically generated. Its main goal is to introduce controllable signal domain shifts. Thus, any progress w.r.t. a domain shift problem can easily be verified since the expected improvements are a priori known. We use $D_{\mathcal{S}_{\text{SYN}}}$ and $D_{\mathcal{T}_{\text{SYN}}}$ to prove the principal functionality of our presented approach. A detailed description of the synthetic dataset content will follow in subsubsection V-B1. In order to show results that are comparable to results that might be achieved in a real-world use case, there is a second pair of source and target datasets, $D_{\mathcal{S}_{\text{OTA}}}$ and $D_{\mathcal{T}_{\text{OTA}}}$, respectively, with spectrograms from over-the-air recordings. Here, we give a more detailed description in subsubsection V-B2.

3) Test Dataset: For testing the performance, we need an additional pair of test datasets $D_{\mathcal{S},\text{Test}_{\text{SYN}}}$ and $D_{\mathcal{T},\text{Test}_{\text{SYN}}}$, respectively, for the synthetic dataset. While the latter is required for the actual performance evaluation, the former solely serves as reference for the graphical representation of the performance. We use an additional test dataset $D_{\text{Test}_{\text{OTA}}}$ for the over-the-air recordings. Additionally, we manually create the bounding boxes for the signals in these recordings.

B. Results of Experiments

In this section, we present the performance of our newly developed deep learning domain adaptation signal detection network. The performance of the different training runs is evaluated with two metrics, namely recall (R) and precision (P). Recall is defined as the ratio of all correct detections divided by all existing objects, while precision is defined as all correct detections divided by all detections. Furthermore, we utilize the dimension reduction method t-SNE [15] for the feature visualization of the classes to compare the different methods.

We first evaluate the results for the synthetic dataset pair $D_{\mathcal{S}_{\mathrm{SYN}}}$ and $D_{\mathcal{T}_{\mathrm{SYN}}}$, respectively. By doing so, we can ensure the principal capabilities of the domain adaptation method. Then, we test our domain adaptation approach on the over-the-air dataset $D_{\mathcal{T}_{\mathrm{OTA}}}$. The results for the latter represent achievable gains for the actual signal detection task. Before we present concrete figures for the datasets $D_{\mathcal{T}_{\mathrm{SYN}}}$ and $D_{\mathcal{T}_{\mathrm{OTA}}}$, we first give a precise description of the content of $D_{\mathcal{T}_{\mathrm{SYN}}}$.

1) Synthetic Dataset: The synthetic dataset consists of signals within the spectrograms that are analytically generated. This means, both, the analog as well as the digital signals, are generated with the communications toolbox from MATLAB. The synthetic datasets are generated with these signals as explained in section III. In order to obtain controllable signal domain shifts, there are a variety of effects that can be introduced. One option is to add artificial features to a signal class in the source dataset $D_{\mathcal{S}_{\mathrm{SYN}}}$, which do not occur in the corresponding signal class from the target dataset $D_{\mathcal{T}_{\mathrm{SYN}}}$. For example, a diagonal or vertical line can be added to a signal, as shown in Fig. 3 for an FSK2 and an AM signal.

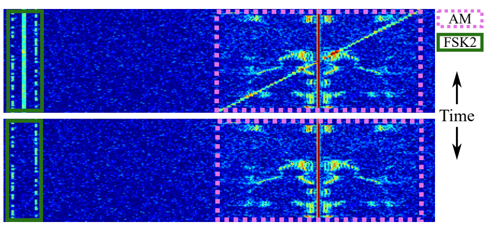

Fig. 3. Example spectrogram for $D_{\mathcal{S}_{\mathrm{SYN}}}$ (top) and $D_{\mathcal{T}_{\mathrm{SYN}}}$ (bottom).

While artificial lines exist in both signals in the spectrogram data from $D_{\mathcal{S}_{\mathrm{SYN}}}$, they are missing in the corresponding spectrogram from $D_{\mathcal{T}_{\mathrm{SYN}}}$. We choose this artificial domain shift scenario for two reasons: First, without domain adaptation, the neural network focuses on these lines during the training as these are simple features to differentiate these signals from other classes. This makes it an ideal example for proving the general functionality of our proposed approach. Second, the artificial domain shift is different for each class, which introduces an additional difficulty for a domain adaptation process with a class-agnostic domain classifier. Thus, it allows us to investigate the functionality of the *multi-class* domain adaptation mechanism.

In order to investigate the influence of the multi-class domain adaptation part, we compare the results of our approach with an approach that uses only one class-agnostic domain classifier. In particular, the results of Domain Adaptation with Class-Agnostic domain classifier (DA-CA) and DA-MC are evaluated alongside the baseline and oracle training. The results for the synthetic test dataset $D_{\mathcal{T},\mathrm{Test}_{\mathrm{SYN}}}$ are shown in Table I. We focus on the FSK2 and AM signal in this test scenario since only these two classes experience a domain shift. For the evaluation, we use recall and precision as metrics (cf. subsection V-B). The baseline training already achieves high scores for the respective signal categories. The DA-CA improves for both classes overall compared to the baseline training. However, the DA-MC surpasses even the performance of the DA-CA and reaches a score close to the oracle training.

TABLE I
BASELINE, DA-CA, DA-MC, AND ORACLE RESULTS FOR $D_{\mathcal{T},\mathrm{Test}_{\mathrm{SYN}}}$.

Method	FSK2		AM		**All**	
	R	P	R	P	R	P
Baseline	95.8	77.9	73.3	92.8	85.2	83.3
DA-CA	100	85.5	75.2	94.0	88.3	88.7
DA-MC	100	98.3	88.6	100	94.6	99.1
Oracle	100	100	94.3	100	97.3	100.0

In Fig. 4 the t-SNE plots for the various training variations (cf. subsubsection V-A1) are shown for the synthetic test datasets $D_{\mathcal{S},\mathrm{Test}_{\mathrm{SYN}}}$ and $D_{\mathcal{T},\mathrm{Test}_{\mathrm{SYN}}}$, respectively. The optimal result for a domain adaptation approach contains source and target data that cannot be distinguished by the network. For the t-SNE plot, this means that the feature positions for both, source and target dataset mix up inseparable. In Fig. 4a the plot for the baseline training is illustrated. Within each class, source and target samples can strictly be separated by their features. Then, the t-SNE plot in Fig. 4b shows the training with DA-CA. The features are not mixed for the FSK2 class and only slightly mixed for the AM class despite the results in Table I. The class-agnostic domain classifier can not learn the domain shifts of both classes. In contrast, the DA-MC method in Fig. 4c mixes up the source and target samples. This indicates that the features are not distinguishable for the neural network. Thus, the training with multi-class domain classifiers is superior compared to the training with class-agnostic domain classifier.

2) Over-the-Air Recording Dataset: The source dataset $D_{\mathcal{S}_{\mathrm{OTA}}}$ is similarly generated to the synthetic dataset. However, instead of generating the signals analytically, we filter the individual signals from wideband over-the-air recordings. In contrast, the target dataset $D_{\mathcal{T}_{\mathrm{OTA}}}$ contains wideband spectrograms with a priori unknown amount of signals. Both, the position and the category of the signals are not controllable. This leads to several challenges with the target dataset $D_{\mathcal{T}_{\mathrm{OTA}}}$ compared to the source dataset $D_{\mathcal{S}_{\mathrm{OTA}}}$. First, the uncontrollable amount of signals can lead to collisions between signals. Fig. 5 illustrates examples for the training data (top image) and the over-the-air recording (bottom image). While there are collisions in the over-the-air recording scenario, the signals are separated in the training data scenario. Therefore, the detection task is more difficult for $D_{\mathcal{T}_{\mathrm{OTA}}}$ since features of multiple signals are mixed in the spectrogram due to these collisions. Second, there are unknown signal categories, i.e., signals that are not contained in the training data, in $D_{\mathcal{T}_{\mathrm{OTA}}}$. These unknown signals confuse the neural network, which erroneously detects them as one of the existing categories. Third, in the over-the-air recordings distortions can occur that result in additional power levels in the spectrogram, i.e., due to faulty hardware. Fourth, poor filter design at the transmitter leads to visible filter edges of signals. Overall, these factors contribute to the existing domain shift between the source dataset $D_{\mathcal{S}_{\mathrm{OTA}}}$ and the target dataset $D_{\mathcal{T}_{\mathrm{OTA}}}$.

Results for the over-the-air recording test dataset $D_{\mathrm{Test}_{\mathrm{OTA}}}$

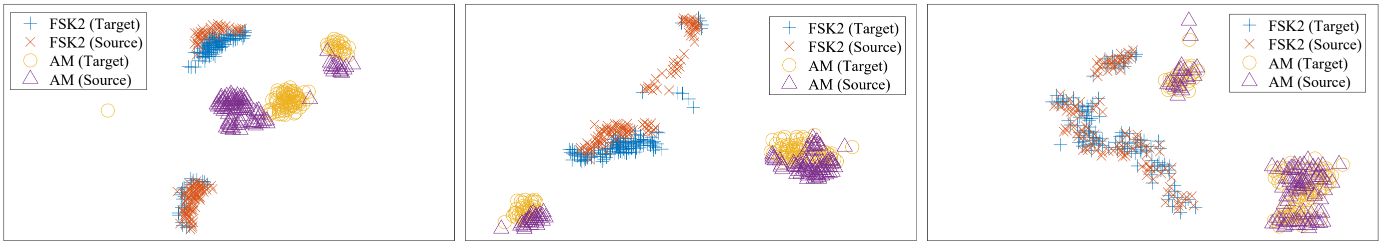

(a) Baseline (b) Class-agnostic domain adaptation (c) Multi-class domain adaptation

Fig. 4. t-SNE comparison of baseline, DA-CA and DA-MC for two different domain shifts.

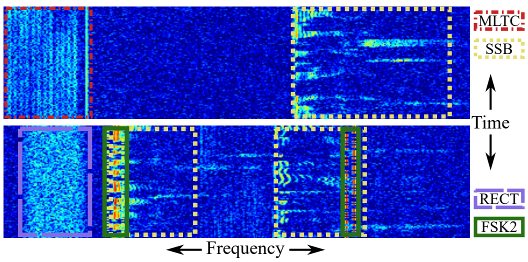

Fig. 5. Example spectrogram for $D_{\mathcal{S}_{\text{OTA}}}$ (top) and $D_{\mathcal{T}_{\text{OTA}}}$ (bottom).

are shown in Table II. In this test scenario we focus on MLTC, RECT, FSK2, and SSB as signal categories. Furthermore, we select the DA-MC method, which yields the overall best results for the synthetic dataset. The DA-MC method improves the detection performance for all classes. In particular, the score of MLTC and RECT gets close to the oracle training. Hence, our approach resolves domain shift effects for these classes without extensive labeling efforts. While FSK2 and SSB also improve the performance compared to the baseline training, there is a larger gap between DA-MC and the oracle training if compared to MLTC or RECT. One reason for the minor improvement could be their baseline score. In particular, we can observe that the recall for SSB and FSK2 in the baseline training is lower compared to the other two classes. Our domain adaptation method relies on detections, which are less reliable for SSB and FSK2. Overall, the performance increases with our domain adaptation approach for signal detection.

TABLE II
BASELINE, DA-MC, AND ORACLE RESULTS FOR $D_{\text{Test}_{\text{OTA}}}$.

Method	MLTC		RECT		FSK2		SSB	
	R	P	R	P	R	P	R	P
Baseline	75.2	83.0	89.7	85.5	58.7	89.6	55.5	95.3
DA-MC	89.9	80.2	92.7	87.9	63.5	88.7	60.6	94.4
Oracle	91.9	85.6	92.7	88.5	78.0	91.3	82.0	84.6

VI. Conclusion

In this paper, we addressed the domain adaptation challenge for signal detection in the high frequency range. As starting point, we compared four domain adaptation categories and selected the adversarial training method as the most suitable approach for the signal detection task. We introduced a novel adversarial domain adaptation network with multi-class domain classifiers for signal detection that uses a modified YOLOv4 network as its backbone. We created synthetic and over-the-air datasets for the evaluation process. Our evaluations with the synthetic test scenario show that the domain shift problem can be solved for the signal detection task while reducing the labeling efforts significantly. We surpass both, the performance of the baseline training as well as the performance of the domain adaptation method with a class-agnostic classifier and achieve a performance close to the oracle training. Furthermore, we achieve remarkable performance gains for tests with an over-the-air dataset $D_{\mathcal{T}_{\text{OTA}}}$.

References

[1] Y. Chen, W. Li, C. Sakaridis, D. Daiand L. Van Gool, "Domain adaptive faster r-cnn for object detection in the wild", pp. 3339–3348, Aug. 2018.

[2] R. Gopalan, R. Li, and R. Chellappa, "Domain adaptation for object recognition: An unsupervised approach," 2011 international conference on computer vision, pp. 999–1006, 2011.

[3] T. O'Shea, T. Roy, and T. C. Clancy, "Learning robust general radio signal detection using computer vision methods," 2017 51st asilomar conference on signals, systems, and computers, pp. 829–832, 2017.

[4] L. Wood, K. Anderson, P. Gerstoft, R. Bell, R. Subbaraman, and D. Bharadia, "Deep Learning Object Detection Approaches to Signal Identification," arXiv preprint arXiv:2210.16173, 2022.

[5] Y. Ganin and V. Lempitsky, "Unsupervised domain adaptation by backpropagation," International conference on machine learning, pp. 1180–1189, 2015.

[6] D.-H. Lee, "Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks", vol. 3, no. 2, p. 896, Aug. 2013.

[7] X. Li, "A free lunch for unsupervised domain adaptive object detection without source data", vol. 35, no. 10, pp. 8474–8481, Aug. 2021.

[8] A. RoyChowdhury et al., "Automatic adaptation of object detectors to new domains using self-training", pp. 780–790, Aug. 2019.

[9] V. F. Arruda et al., "Cross-domain car detection using unsupervised image-to-image translation: From day to night," 2019 International Joint Conference on Neural Networks (IJCNN), pp. 1–8, 2019.

[10] A. Tarvainen and H. Valpola, "Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results," Advances in neural information processing systems, vol. 30, 2017.

[11] M. Hnewa and H. Radha, "Multiscale Domain Adaptive Yolo For Cross-Domain Object Detection," 2021 IEEE International Conference on Image Processing (ICIP), Sep. 2021, doi: 10.1109/icip42928.2021.9506039.

[12] S. Zhang, H. Tuo, J. Hu, and Z. Jing, "Domain adaptive yolo for one-stage cross-domain detection," Asian conference on machine learning, pp. 785–797, 2021.

[13] Z. Pei, Z. Cao, M. Long, and J. Wang, "Multi-Adversarial Domain Adaptation," 2018, [Online]. Available: https://arxiv.org/abs/1809.02176

[14] A. Bochkovskiy, C.-Y. Wang, and H.-Y. M. Liao, "Yolov4: Optimal speed and accuracy of object detection," arXiv preprint arXiv:2004.10934, 2020.

[15] L. Maaten and G. Hinton, "Visualizing data using t-SNE.," Journal of machine learning research, vol. 9, no. 11, 2008.