

Unsupervised Domain Adaptation across FMCW Radar Configurations Using Margin Disparity Discrepancy

Rodrigo Hernangómez*, Igor Bjelaković*[‡], Lorenzo Servadei^{†§}, and Sławomir Stańczak*[‡]

* Fraunhofer Heinrich Hertz Institute, Berlin, Germany, {firstname.lastname}@hhi.fraunhofer.de

[†]Infinion Technologies AG, Munich, Germany, lorenzo.servadei@infineon.com

[‡]Technical University of Berlin, Berlin, Germany

[§]Technical University of Munich, Munich, Germany

Abstract—Commercial radar sensing is gaining relevance and machine learning algorithms constitute one of the key components that are enabling the spread of this radio technology into areas like surveillance or healthcare. However, radar datasets are still scarce and generalization cannot be yet achieved for all radar systems, environment conditions or design parameters. A certain degree of fine tuning is, therefore, usually required to deploy machine-learning-enabled radar applications. In this work, we consider the problem of unsupervised domain adaptation across radar configurations in the context of deep-learning human activity classification using frequency-modulated continuous-wave. For that, we focus on the theory-inspired technique of Margin Disparity Discrepancy, which has already been proved successful in the area of computer vision. Our experiments extend this technique to radar data, achieving a comparable accuracy to few-shot supervised approaches for the same classification problem.

Index Terms—radar, machine learning, deep learning, transfer learning, domain adaptation, human activity classification

I. INTRODUCTION

Radar, a well-established technology for several industrial areas, has recently gained attention for other commercial applications like human monitoring, presence detection or gesture sensing [1–3] due to the production of small and compact radar sensors [4]. Here radar offers some advantages in comparison with computer vision approaches, such as good performance under poor-lighting conditions or privacy protection (due to the difficulty to identify individuals from radar images).

Similarly as in the realm of computer vision, the use of radar often comes hand in hand with machine learning (ML) techniques, including deep learning, to overcome the burden of handcrafted feature engineering [5]. Due to the variety of system design parameters at hand, such as modulation techniques or bandwidth, these ML algorithms are required to generalize well under different radar setups. This need for inter-domain generalization is common to several ML problems and it has been studied in recent years under the paradigm of domain adaptation (DA) [6, 7].

Considered a special case of transfer learning, DA involves modifying an ML estimator that can be trained with enough data from a *source domain*, so that its performance increases when evaluated with data originated from a different *target*

domain. Unlike other transfer learning approaches, here the mismatch between source and target domain lies merely in a distinct probability measure over data rather than in different input or output spaces [7].

The reasons for DA are usually related with insufficient or incomplete data in the target domain, which can be overcome with the help of data from the source domain. In ML classification, the missing information is often the labels; this case is referred to as *unsupervised* DA. If target data are labeled, we can apply *supervised* DA techniques instead [6].

Both supervised and unsupervised DA methods have already been investigated in the Radar-ML community to overcome several problems, including individual patient differences [8], aspect angle variations [9], synthetic-to-real adaptation [10] or environmental differences [11]. In the case of cross-configuration adaptation, Khodabakhshandeh et al. [12] use supervised techniques such as Few-shot Adversarial Domain Adaptation (FADA) [13] or domain adaptation using Stochastic Neighborhood Embedding (*d*-SNE) [14] to adapt their trained human activity classifier to new frequency-modulated continuous-wave (FMCW) radar setups using few data.

In this paper, we build on the work in [12] by applying Margin Disparity Discrepancy (MDD) [15]. In that way, we confirm that this unsupervised technique, which delivers state-of-the-art results for computer vision datasets, also works for radar data and thus enables cross-configuration radar-based human activity classification based on unlabeled data.

II. PROBLEM STATEMENT

Radar-ML classification deals with the evaluation of a group of radar features $\mathbf{x} \in \mathcal{X}$ obtained from a target to find the underlying class $y \in \mathcal{Y}$ that best describes some property of the said target. The input space $\mathcal{X} \subset \mathbb{R}^m$ is characterized by a dimension m that depends on the radar technology and preprocessing steps, while the label space is defined as $\mathcal{Y} = \{1, \dots, k\}$, with k being the number of classes.

In order to achieve this classification, one has first to find a classifier h that maps \mathbf{x} into y . The ML approach assumes a sufficiently large amount of data available conveying information both about the inputs and the class so that it can be

used to train h among a restricted hypothesis class \mathcal{H} . This dataset consists of a sequence of pairs of features and labels, i.e. $\{(\mathbf{x}_i, y_i)\}_{i=1}^n$, that have been previously sampled from a certain domain \mathcal{D} , defined to be

$$\mathcal{D} = (\mathcal{X}, \mathcal{Y}, p_{\mathcal{D}}), \quad (1)$$

with an associated probability measure $p_{\mathcal{D}}$ over $\mathcal{X} \times \mathcal{Y}$. Here and hereafter, we write \mathbf{x} and y in their upright form whenever we refer to the random variables related to $p_{\mathcal{D}}$, and not to its realizations.

By choosing an objective loss function

$$\ell : \mathcal{H} \times \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}_{0+} \quad (2)$$

and minimizing it over the hypothesis class \mathcal{H} with a suitable optimization method, we can train an h that performs well for the available data. The performance of h can thus be measured by the *risk* associated with the loss for a domain \mathcal{D} . This risk $\mathcal{L}_{\mathcal{D}}$ represents the expected value of the loss of h for $p_{\mathcal{D}}$:

$$\mathcal{L}_{\mathcal{D}}(h) = \mathbb{E}_{\mathcal{D}} \ell(h, (\mathbf{x}, y)). \quad (3)$$

By assuming the indicator function $\mathbb{1}_{h(\mathbf{x}) \neq y}$ to be the loss, we obtain the *0-1 error* $\text{err}_{\mathcal{D}}(h) \triangleq \mathbb{E}_{\mathcal{D}} \mathbb{1}_{h(\mathbf{x}) \neq y}$. In practice, we do not have access to $p_{\mathcal{D}}$, so we resort to its empirical approximation $\mathcal{L}_{\hat{\mathcal{D}}}(h) \triangleq \sum_{i=1}^n \ell(h, (\mathbf{x}_i, y_i))/n$ for a dataset $\hat{\mathcal{D}}$ with n samples drawn from \mathcal{D} .

If generalization is achieved, h will also behave well for unseen data as long as it is drawn from the same domain. Unfortunately, this assumption cannot always be guaranteed. It is often the case that training data have been drawn from a *source domain* \mathcal{S} but we would like to leverage the trained classifier for a different *target domain* \mathcal{T} . Depending on how dissimilar \mathcal{S} and \mathcal{T} are, the performance of the trained classifier can degrade significantly. In our specific problem, this *domain shift* is given by the choice of different FMCW settings and presents an additional challenge in the lack of the labels for the training data from \mathcal{T} . The absence of labeled target data makes it necessary to apply *unsupervised* DA. In this paper, we explore this possibility by using Margin Disparity Discrepancy (MDD) [15].

A. Margin Disparity Discrepancy

In order to use MDD, we assume a hypothesis class induced by a space \mathcal{F} of *scoring functions* $f : \mathcal{X} \mapsto \mathbb{R}^k$. We also introduce the shorthand $f_y(\mathbf{x})$ to refer to the y -th component of $f(\mathbf{x})$. The hypothesis class is given by

$$\mathcal{H} \triangleq \left\{ h_f : \mathbf{x} \mapsto \arg \max_{y \in \mathcal{Y}} f_y(\mathbf{x}) \mid f \in \mathcal{F} \right\}. \quad (4)$$

MDD has been developed by Zhang et al. [15] as a practical algorithm based on the concept of discrepancy distance by Mansour et al. [16]. For that, they define the *margin loss* $\text{err}_{\mathcal{D}}^{(\rho)}(f)$ as

$$\text{err}_{\mathcal{D}}^{(\rho)}(f) \triangleq \mathbb{E}_{\mathcal{D}} \Phi^{(\rho)} \circ \phi_f(\mathbf{x}, y), \quad (5)$$

$$\phi_f(\mathbf{x}, y) \triangleq \frac{1}{2} \left(f_y(\mathbf{x}) - \max_{y' \neq y} f_{y'}(\mathbf{x}) \right), \quad (6)$$

$$\Phi^{(\rho)}(x) \triangleq \begin{cases} 0 & \rho \leq x \\ 1 - x/\rho & 0 \leq x \leq \rho \\ 1 & x \leq 0 \end{cases}, \quad (7)$$

and the true and empirical *margin disparity* between two scoring functions f' and f as

$$\text{disp}_{\mathcal{D}}^{(\rho)}(f', f) \triangleq \mathbb{E}_{\mathcal{D}} \Phi^{(\rho)} \circ \phi_{f'}(\mathbf{x}, h_f(\mathbf{x})), \quad (8)$$

$$\text{disp}_{\hat{\mathcal{D}}}^{(\rho)}(f', f) \triangleq \frac{1}{n} \sum_{i=1}^n \Phi^{(\rho)} \circ \phi_{f'}(\mathbf{x}_i, h_f(\mathbf{x}_i)), \quad (9)$$

to finally formulate the following minimax optimization problem:

$$\min_{f \in \mathcal{F}} \text{err}_{\hat{\mathcal{S}}}^{(\rho)}(f) + d_{f, \mathcal{F}}^{(\rho)}(\hat{\mathcal{S}}, \hat{\mathcal{T}}), \quad (10)$$

$$d_{f, \mathcal{F}}^{(\rho)}(\mathcal{S}, \mathcal{T}) \triangleq \sup_{f' \in \mathcal{F}} \left(\text{disp}_{\mathcal{T}}^{(\rho)}(f', f) - \text{disp}_{\mathcal{S}}^{(\rho)}(f', f) \right).$$

Following the principles of unsupervised DA, the MDD term $d_{f, \mathcal{F}}^{(\rho)}$ does not make use of any labels y_i . Furthermore, the solution to (10) minimizes the 0-1 error of h_f in the target domain, as the authors of [15] prove with the following theoretical bound:

$$\text{err}_{\mathcal{T}}(h_f) \leq \text{err}_{\mathcal{S}}^{(\rho)}(f) + d_{f, \mathcal{F}}^{(\rho)}(\mathcal{S}, \mathcal{T}) + \lambda, \quad (11)$$

where λ is the ideal combined margin loss:

$$\lambda = \min_{f^* \in \mathcal{F}} \left\{ \text{err}_{\mathcal{S}}^{(\rho)}(f^*) + \text{err}_{\mathcal{T}}^{(\rho)}(f^*) \right\}. \quad (12)$$

The bound in (11) can also be expressed in terms of empirical measures rather than true probability measures by the addition of Rademacher complexity terms [17, Chapter 3].

Despite the interesting properties of MDD, $\Phi^{(\rho)} \circ \phi_f$ is non-smooth, non-convex and its training causes vanishing and exploding gradients, which leads Zhang et al. [15] to work with the cross-entropy loss instead. For this, they map $f(\mathbf{x})$ to the k -simplex via the softmax function σ , as it is customary in deep learning, where the elements of $\sigma(\mathbf{z})$ are given by

$$\sigma_j(\mathbf{z}) \triangleq \frac{\exp z_j}{\sum_{i=1}^k \exp z_i}, \quad \text{for } j = 1, \dots, k. \quad (13)$$

The composition of the cross-entropy loss with the softmax yields the log-sum-exp (lse) loss \mathcal{L}_f :

$$\begin{aligned} \mathcal{L}_f(\mathbf{x}, y) &\triangleq H(\mathbb{1}_{j=y}, \sigma(f(\mathbf{x}))) = -\log \sigma_y(f(\mathbf{x})) \\ &= \log \sum_{y' \in \mathcal{Y}} \exp(f_{y'}(\mathbf{x}) - f_y(\mathbf{x})). \end{aligned} \quad (14)$$

Zhang et al. [15] propose to use \mathcal{L}_f instead of $\Phi^{(\rho)} \circ \phi_f$ for $\text{err}_{\hat{\mathcal{S}}}^{(\rho)}$ and $\text{disp}_{\hat{\mathcal{S}}}^{(\rho)}$ in (10). As for $\text{disp}_{\hat{\mathcal{T}}}^{(\rho)}$, they use the adversarial loss $\tilde{\mathcal{L}}_f$ proposed by Goodfellow et al. [18], i.e.

$$\tilde{\mathcal{L}}_f(\mathbf{x}, y) \triangleq \log(1 - \sigma_y(f(\mathbf{x}))), \quad (15)$$

so that their MDD ultimately becomes

$$\begin{aligned} \tilde{d}_{f,\psi,\mathcal{F}}^{(\gamma)}(\widehat{\mathcal{S}},\widehat{\mathcal{T}}) &\triangleq \max_{f' \in \mathcal{F}} \mathbb{E}_{\mathbf{x}^t \sim \widehat{\mathcal{T}}} \tilde{\mathcal{L}}_{f'}(\psi(\mathbf{x}^t), h_f(\psi(\mathbf{x}^t))) \\ &\quad - \gamma \mathbb{E}_{\mathbf{x}^s \sim \widehat{\mathcal{S}}} \mathcal{L}_{f'}(\psi(\mathbf{x}^s), h_f(\psi(\mathbf{x}^s))) \end{aligned} \quad (16)$$

for a *margin factor* $\gamma > 0$ and a feature extractor ψ that levels the min-player to the max-player [15] (A concrete example of ψ is given in (27), Section III). The authors explain that this is equivalent to the use of the margin loss with a margin $\rho = \log \gamma$ and that the problem is still solved for $\mathcal{S} = \mathcal{T}$ [15].

In addition to the results in [15], we observe that the use of the recently proposed soft-margin softmax [19] instead of σ in (14) provides an upper bound for $\text{err}_{\mathcal{S}}^{(\rho)}$. The entries of the soft-margin softmax $\sigma^{(\rho)}(\mathbf{z})$ are defined as

$$\begin{aligned} \sigma_j^{(\rho)}(\mathbf{z}) &\triangleq \frac{\exp(z_j - \rho)}{\exp(z_j - \rho) + \sum_{i \neq j} \exp z_i}, \\ &\quad \text{for } j = 1, \dots, k; \quad \rho \in \mathbb{R}_+ \end{aligned} \quad (17)$$

and this induces the soft-margin cross-entropy loss $\mathcal{L}_f^{(\rho)}$:

$$\begin{aligned} \mathcal{L}_f^{(\rho)}(\mathbf{x}, y) &\triangleq -\log \sigma_y^{(\rho)}(f(\mathbf{x})) \\ &= \log \sum_{y' \in \mathcal{Y}} \exp(f_{y'}(\mathbf{x}) - f_y(\mathbf{x}) + \rho \cdot \mathbf{1}_{y' \neq y}). \end{aligned} \quad (18)$$

Likewise, a soft-max adversarial loss can also be defined as

$$\tilde{\mathcal{L}}_f^{(\rho)}(\mathbf{x}, y) \triangleq \log(1 - \sigma_y^{(\rho)}(f(\mathbf{x}))). \quad (19)$$

We prove this soft-margin-based bound with the help of the following Lemma, which motivates us to investigate $\mathcal{L}_f^{(\rho)}$ further in Section III.

Lemma 1. *The soft-max cross entropy bounds the margin loss as*

$$\Phi^{(\rho)} \circ \phi_f(\mathbf{x}, y) \leq \frac{1}{2\rho} \mathcal{L}_f^{(2\rho)}(\mathbf{x}, y). \quad (20)$$

Proof. First, let us recall the generalized hinge loss [20]:

$$\tilde{h}_f^{(\theta)}(\mathbf{x}, y) \triangleq \max_{y' \in \mathcal{Y}} (f_{y'}(\mathbf{x}) - f_y(\mathbf{x}) + \theta \cdot \mathbf{1}_{y' \neq y}). \quad (21)$$

Noting that the term within $\max_{y' \in \mathcal{Y}}$ is null $\forall y' = y$, we have

$$\begin{aligned} \tilde{h}_f^{(2\rho)}(\mathbf{x}, y) &= \max \left\{ 0, \max_{y' \neq y} f_{y'}(\mathbf{x}) - f_y(\mathbf{x}) + 2\rho \right\} = \\ &= \max \{ 0, 2\rho - 2\phi_f(\mathbf{x}, y) \} = 2\rho \max \{ 0, 1 - \phi_f(\mathbf{x}, y) / \rho \}. \end{aligned} \quad (22)$$

The last expression in (22) can be derived from (7) if we set $1 - x/\rho$ instead of 1 for $x \leq 0$, hence

$$\Phi^{(\rho)} \circ \phi_f(\mathbf{x}, y) \leq \frac{1}{2\rho} \tilde{h}_f^{(2\rho)}(\mathbf{x}, y) \quad (23)$$

and the proof is concluded using the fact that

$$\log \sum_{a \in \mathcal{A}} \exp a \geq \max_{a \in \mathcal{A}} a \quad (24)$$

for any finite set \mathcal{A} . □

TABLE I
RADAR CONFIGURATION PARAMETERS

Configuration name		I	II	III	IV
Chirps per frame	n_c	64	64	64	64
Samples per chirp	n_s	256	256	128	256
Bandwidth	[GHz]	2	1	2	2
Frame period	[ms]	50	32	32	32
Chirp to chirp time	[μ s]	250	250	250	250
Range resolution	[cm]	7.5	15	7.5	7.5
Max. range	[m]	6.2	12.5	4.8	6.2
Max. speed	[m/s]	5.0	5.0	5.0	5.0
Speed resolution	[m/s]	0.15	0.15	0.15	0.15

Taking the expectation w.r.t. $p_{\mathcal{S}}$ in (20), we finally obtain

$$\text{err}_{\mathcal{S}}^{(\rho)}(f) \leq \frac{1}{2\rho} \mathbb{E}_{\mathcal{S}} \mathcal{L}_f^{(2\rho)}(\mathbf{x}, y). \quad (25)$$

Despite the gap introduced by Eqs. (23) and (24), we note that Eq. (25) delivers convincing bounds for a small $\text{err}_{\mathcal{S}}^{(\rho)}(f)$, which can be achieved by training f under enough samples from \mathcal{S} .

III. EXPERIMENTS

A. Setup

Similarly as in [12], we focus on human activity recognition using data that have been measured simultaneously with 4 different 60-GHz FMCW radar sensors. For these measurements, 2 male subjects were recorded separately while performing 5 different activities: *standing*, *waving*, *walking*, *boxing* or *boxing while walking*. Each one of the radar sensors was configured with a different set of radar parameters, presented in Table I as **I** to **IV**. Here, the divergent parameters of the different configurations (marked in bold) affect the temporal and range resolution of the range-Doppler map (RDM) sequences, as well as the maximum observable scope of the latter. From these configurations, **I** has been taken over from [12].

The input features \mathbf{x} in Fig. 1 comprise both range (\mathbf{x}_r) and Doppler (\mathbf{x}_d) information, i.e.:

$$\mathbf{x} = (\mathbf{x}_r, \mathbf{x}_d), \quad \mathbf{x}_r, \mathbf{x}_d \in \mathbb{R}^{64 \times 128}. \quad (26)$$

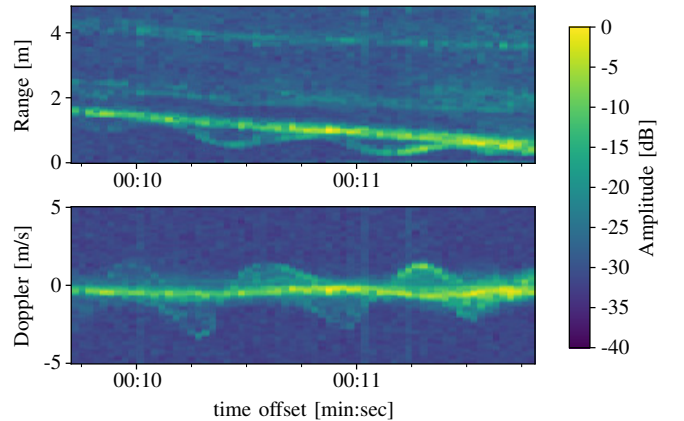


Fig. 1. Range and Doppler spectrogram for *boxing while walking*.

The radar preprocessing to produce \mathbf{x} is also based on [12], with the notable addition of cropping and resampling of the spectrograms to ensure the dimensions in (26) and the scopes of 0 to 2 s, 0.0 to 4.8 m and -5 to 5 m/s for the time, range and Doppler dimensions, respectively.

Despite all the preprocessing, the differences on resolution still yield a domain shift across configurations that we try to tackle with MDD. For that, we take both spectrograms as an input to our feature extractor ψ . Here we choose the same topology as in [12]; that is, a pair of twin branches ψ_r and ψ_d , each one consisting of 3 convolutional layers for which we concatenate the outputs:

$$\psi(\mathbf{x}) \equiv (\psi_r(\mathbf{x}_r), \psi_d(\mathbf{x}_d)). \quad (27)$$

Furthermore, we employ a bottleneck layer of 512 nodes and choose our hypothesis space \mathcal{F} to be consistent with the structure of the fully connected layers from [12].

Motivated by Lemma 1, we replace the vanilla loss terms in (16) by the soft-margin losses $\mathcal{L}^{(\rho)}$ and $\tilde{\mathcal{L}}^{(\rho)}$ with

$$\rho = 2 \log 2 \simeq 1.386 \quad (28)$$

and we set the margin factor $\gamma = 1$ since the margin ρ is already included in the loss, effectively rendering MDD as

$$\begin{aligned} \tilde{d}_{f, \psi, \mathcal{F}}^{(\rho)}(\hat{\mathcal{S}}, \hat{\mathcal{T}}) \triangleq & \max_{f' \in \mathcal{F}} \mathbb{E}_{\mathbf{x}^t \sim \hat{\mathcal{T}}} \tilde{\mathcal{L}}_{f'}^{(\rho)}(\psi(\mathbf{x}^t), h_{f'}(\psi(\mathbf{x}^t))) \\ & - \mathbb{E}_{\mathbf{x}^s \sim \hat{\mathcal{S}}} \mathcal{L}_{f'}^{(\rho)}(\psi(\mathbf{x}^s), h_{f'}(\psi(\mathbf{x}^s))). \end{aligned} \quad (29)$$

Other than that, we leave all hyperparameters to the same values as in [15] and adapt their implementation as in Fig. 2. This has been written in Pytorch as an instance of adversarial training, where a gradient reversal layer (GRL) is used to minimize the MDD loss term on ψ while maximizing on f' as the minimax formulation in (10) mandates.

The number of samples per dataset lies over 1150 samples for the train sets and over 350 samples for the test sets.

B. Results

We have run unsupervised training experiments for all possible domain pairs within configurations I-IV and summarized the resulting test accuracies on the test sets in Table II.

The figures follow the same trend as the results of MDD in computer vision datasets as reported by Zhang et al. [15]. and




Fig. 2. MDD adversarial network, adapted from [15].

TABLE II
TEST ACCURACY [%] OF MDD FOR FMCW DATA

Source configuration	Target configuration			
	I	II	III	IV
I	-	91.4	90.6	88.3
II	90.9	-	89.8	89.4
III	89.4	90.4	-	89.4
IV	92.5	85.8	90.9	-

TABLE III
MIN. AND MAX. ACCURACY [%] OF MDD FOR DIFFERENT DATASETS

Office-31	Office-Home	VisDa	FMCW
72.2-100.0	53.6-82.3	74.6 (single value)	85.8-92.5

TABLE IV
AVERAGE ACCURACY COMPARISON [%] OF THE ORIGINAL MDD IMPLEMENTATION AND THE SOFT-MARGIN VERSION

	Office-31	Office-Home	FMCW
Original MDD	88.9	68.1	89.525
Soft-margin MDD	88.3	67.6	89.9

presented in Table III. Here one can compare, for instance, the results using MDD for our FMCW data with the minimum and maximum accuracies obtained for Office-31, a dataset containing 4,652 images from three domains [21]. It is also noteworthy that the highest accuracy for FMCW exceeds both that of the Office-Home dataset (15,500 images from four domains) [22] and the VisDA dataset (280K real and synthetic images) [23].

Our results are also comparable with the FADA method for FMCW-based human activity recognition in [12], which increases the baseline accuracy of 50–60 % without DA to 88–92 %. Here it is important to note that MDD is, in contrast to FADA, an unsupervised technique and thus it presents the advantage of working with unlabeled target data.

We have also compared the average accuracy across domain combinations for the original implementation of MDD in (16) with the average accuracy for our soft-margin version in (29), taking the Office-31 and Office-Home datasets as well as our FMCW radar data. The results, which can be seen in Table IV, show little difference between both implementations.

IV. CONCLUSION

In this work, we confirm that the MDD algorithm, which has already shown promising results for unsupervised DA in the area of computer vision, is also suitable for radar data across different FMCW parameters. The obtained accuracy can become as high as for some supervised techniques [12] while using a much more limited dataset, paving thus the way for a prompt deployment of radar-based deep learning applications with custom configurations.

In our experiments, we observe that the use of the soft-margin cross entropy loss provides similar results as the original implementation by Zhang et al. [15]. Since the motivation of MDD is to bring the algorithms closer to the analytical

performance bounds of DA, we see potential in this alternative loss function to bridge the gap between theory and practice.

ACKNOWLEDGMENT

We gratefully acknowledge the support of NVIDIA Corporation with the donation of the DGX-1 used for this research. We would also like to thank Avik Santra from Infineon Technologies AG for his support throughout this work.

REFERENCES

- [1] Z. Peng, C. Li, J.-M. Muñoz-Ferreras, and R. Gómez-García, “An FMCW radar sensor for human gesture recognition in the presence of multiple targets,” in *2017 First IEEE MTT-S International Microwave Bio Conference (IMBIOC)*, May 2017, pp. 1–3.
- [2] A. Santra, R. V. Ulaganathan, and T. Finke, “Short-Range Millimetric-Wave Radar System for Occupancy Sensing Application,” *IEEE Sensors Letters*, vol. 2, no. 3, pp. 1–4, Sep. 2018.
- [3] J. Lien, N. Gillian, M. E. Karagozler, P. Amihoud, C. Schwesig, E. Olson, H. Raja, and I. Poupyrev, “Soli: Ubiquitous gesture sensing with millimeter wave radar,” *ACM Transactions on Graphics*, vol. 35, no. 4, pp. 142:1–142:19, Jul. 2016.
- [4] S. Trotta, D. Weber, R. W. Jungmaier, A. Baheti, J. Lien, D. Noppeney, M. Tabesh, C. Rumpler, M. Aichner, S. Albel, J. S. Bal, and I. Poupyrev, “2.3 SOLI: A Tiny Device for a New Human Machine Interface,” in *2021 IEEE International Solid-State Circuits Conference (ISSCC)*, vol. 64, Feb. 2021, pp. 42–44.
- [5] A. Santra and S. Hazra, *Deep Learning Applications of Short-Range Radars*. Artech House, 2020.
- [6] I. Redko, E. Morvant, A. Habrard, M. Sebban, and Y. Bennani, *Advances in Domain Adaptation Theory*. Elsevier, 2019.
- [7] W. M. Kouw and M. Loog, “An introduction to domain adaptation and transfer learning,” *arXiv:1812.11806 [cs, stat]*, Dec. 2018.
- [8] W. Yin, X. Yang, L. Li, L. Zhang, N. Kitsunan, R. Shinkuma, and E. Oki, “Self-adjustable domain adaptation in personalized ECG monitoring integrated with IR-UWB radar,” *Biomedical Signal Processing and Control*, vol. 47, pp. 75–87, Jan. 2019.
- [9] Q. Chen, Y. Liu, F. Fioranelli, M. Ritchie, and K. Chetty, “Eliminate Aspect Angle Variations for Human Activity Recognition using Unsupervised Deep Adaptation Network,” in *2019 IEEE Radar Conference (RadarConf)*, Apr. 2019, pp. 1–6.
- [10] X. Li, X. Jing, and Y. He, “Unsupervised Domain Adaptation for Human Activity Recognition in Radar,” in *2020 IEEE Radar Conference (RadarConf20)*, Sep. 2020, pp. 1–5.
- [11] M. Stephan, T. Stadelmayer, A. Santra, G. Fischer, R. Weigel, and F. Lurz, “Radar Image Reconstruction from Raw ADC Data using Parametric Variational Autoencoder with Domain Adaptation,” in *2020 25th International Conference on Pattern Recognition (ICPR)*, Jan. 2021, pp. 9529–9536.
- [12] H. Khodabakhshandeh, T. Visentin, R. Hernangómez, and M. Pütz, “Domain Adaptation Across Configurations of FMCW Radar for Deep Learning Based Human Activity Classification,” in *2021 21st International Radar Symposium (IRS)*, Berlin, Germany, Jun. 2021, pp. 1–10.
- [13] S. Motiian, Q. Jones, S. Iranmanesh, and G. Doretto, “Few-Shot Adversarial Domain Adaptation,” *Advances in Neural Information Processing Systems*, vol. 30, Nov. 2017.
- [14] X. Xu, X. Zhou, R. Venkatesan, G. Swaminathan, and O. Majumder, “D-SNE: Domain Adaptation Using Stochastic Neighborhood Embedding,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 2497–2506.
- [15] Y. Zhang, T. Liu, M. Long, and M. Jordan, “Bridging Theory and Algorithm for Domain Adaptation,” in *International Conference on Machine Learning*. PMLR, May 2019, pp. 7404–7413.
- [16] Y. Mansour, M. Mohri, and A. Rostamizadeh, “Domain Adaptation: Learning Bounds and Algorithms,” in *Proceedings of The 22nd Annual Conference on Learning Theory (COLT 2009)*, Montréal, Canada, 2009.
- [17] M. Mohri, A. Rostamizadeh, and A. Talwalkar, *Foundations of Machine Learning*, 2nd ed. Cambridge, Massachusetts; London, England: MIT Press, Dec. 2018.
- [18] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, “Generative Adversarial Nets,” in *Advances in Neural Information Processing Systems*, Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Q. Weinberger, Eds., vol. 27. Curran Associates, Inc., 2014, p. 9.
- [19] X. Liang, X. Wang, Z. Lei, S. Liao, and S. Z. Li, “Soft-Margin Softmax for Deep Classification,” in *Neural Information Processing*, ser. Lecture Notes in Computer Science, D. Liu, S. Xie, Y. Li, D. Zhao, and E.-S. M. El-Alfy, Eds. Cham: Springer International Publishing, 2017, pp. 413–421.
- [20] S. Shalev-Shwartz and S. Ben-David, *Understanding Machine Learning: From Theory to Algorithms*. Cambridge university press, 2014.
- [21] K. Saenko, B. Kulis, M. Fritz, and T. Darrell, “Adapting Visual Category Models to New Domains,” in *Computer Vision – ECCV 2010*, ser. Lecture Notes in Computer Science, K. Daniilidis, P. Maragos, and N. Paragios, Eds. Berlin, Heidelberg: Springer, 2010, pp. 213–226.
- [22] H. Venkateswara, J. Eusebio, S. Chakraborty, and S. Panchanathan, “Deep Hashing Network for Unsupervised Domain Adaptation,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 5018–5027.
- [23] X. Peng, B. Usman, N. Kaushik, J. Hoffman, D. Wang, and K. Saenko, “VisDA: The Visual Domain Adaptation Challenge,” *arXiv:1710.06924 [cs]*, Nov. 2017.