

Adversarial Bi-Regressor Network for Domain Adaptive Regression

Haifeng Xia^{1*}, Pu Wang², Toshiaki Koike-Akino², Ye Wang², Philip Orlik², Zhengming Ding¹

¹Department of Computer Science, Tulane University, New Orleans, LA

²Mitsubishi Electric Research Laboratories (MERL), Cambridge, MA

{hxia, zding1}@tulane.edu, {pwang, koike, yewang, porlik}@merl.com

Abstract

Domain adaptation (DA) aims to transfer the knowledge of a well-labeled source domain to facilitate unlabeled target learning. When turning to specific tasks such as indoor (Wi-Fi) localization, it is essential to learn a cross-domain regressor to mitigate the domain shift. This paper proposes a novel method Adversarial Bi-Regressor Network (**ABRNet**) to seek more effective cross-domain regression model. Specifically, a discrepant bi-regressor architecture is developed to maximize the difference of bi-regressor to discover uncertain target instances far from the source distribution, and then an adversarial training mechanism is adopted between feature extractor and dual regressors to produce domain-invariant representations. To further bridge the large domain gap, a domain-specific augmentation module is designed to synthesize two source-similar and target-similar intermediate domains to gradually eliminate the original domain mismatch. The empirical studies on two cross-domain regressive benchmarks illustrate the power of our method on solving the domain adaptive regression (DAR) problem.

1 Introduction

Deep neural network has become popular on solving many practical scenarios whether the specific task is classification or regression [Sellami and Tabbone, 2022; Xu *et al.*, 2014]. However, considerable network parameters to be optimized means we have to collect extensive training samples with precise annotation, which tends to be time-consuming and laborious for real applications [Krizhevsky *et al.*, 2012]. The intuition is to utilize the off-the-shelf label-sufficient data to accumulate knowledge and directly adapt it to solve other specific and relevant challenges [Awais *et al.*, 2021; Xia *et al.*, 2021a]. But the knowledge transfer is negatively obstructed by the distribution discrepancy across training and test sets so that the well-learned network suffers from performance degradation [Jing *et al.*, 2021; Xia *et al.*, 2021b]. The dilemma motivates the development of unsupervised domain

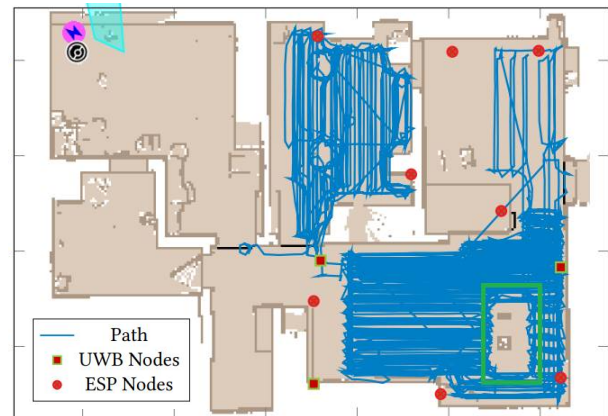


Figure 1: The floor map of the SPAWC2021 multi-modal indoor localization dataset [Arnold and Schaich, 2021], where red squares and red circles denote, respectively, UWB and Wi-Fi anchors, while blue line denotes the path of a robot. The green box denotes the area where certain furniture was moved between offline fingerprinting and online test data collection.

adaptation (UDA) to bring knowledge from source domain to predict unlabeled target instances by learning domain-invariant representations. [Pan *et al.*, 2019].

Along with this direction, abundant UDA solutions have achieved promising performance in classification task [Na *et al.*, 2021; Chang *et al.*, 2019; Xia and Ding, 2020]. To the best of our knowledge, there are a few studies exploring the domain adaptation regression while they focus on the theoretical analysis in shallow regime without the benefit of deep neural architecture [Pan *et al.*, 2010; Yamada *et al.*, 2014]. The regression task is of great importance and common in our reality, e.g., indoor wireless localization [Tong *et al.*, 2021; Zhou *et al.*, 2021; Arnold and Schaich, 2021]. Concretely, a robot moves in a multi-room environment as shown in Figure 1 and communicates with multi-modality anchors, e.g., Wi-Fi and ultra-wideband (UWB), placed at fixed locations denoted as red points. By processing the multi-modal radio frequency (RF) signals, one needs to predict the coordinates of robot in the floor map and make subsequent decisions for path planning and obstacle avoidance. Potential applications include smart home and autonomous factory where conducting the asset tracking is crucial.

The current DAR works mainly focus on the visual re-

*The work was done during the internship at MERL.

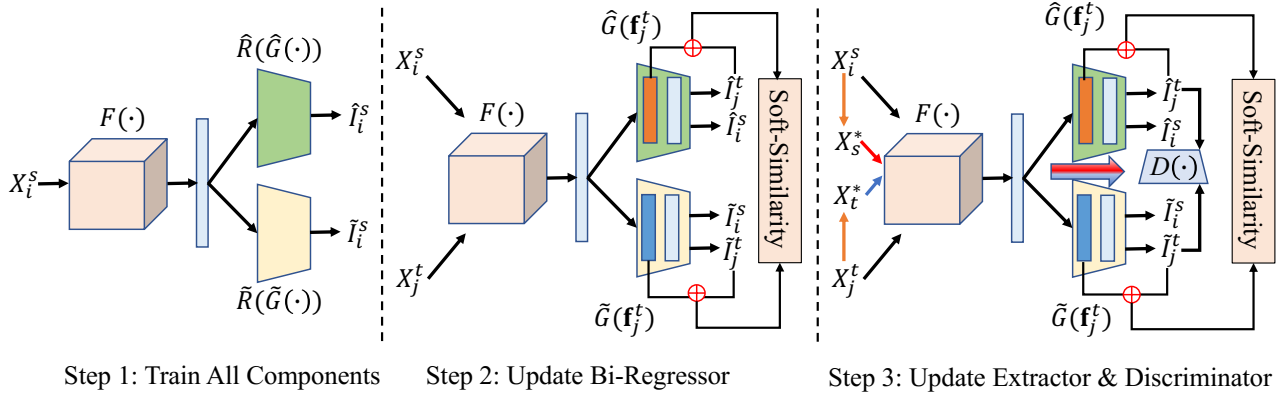


Figure 2: Overview of our proposed adversarial bi-regressors network (ABRNet) including feature generator $F(\cdot)$, two regressors $\{\hat{G}, \hat{R}\}$, $\{\tilde{G}, \tilde{R}\}$ and discriminator $D(\cdot)$. The main training process involves three stages. During the first step, ABRNet utilizes all source supervisions to train the network. For the second step, we introduce soft-similarity to maximize the bi-regressor difference. The third one aims to generate domain-invariant representations by updating feature generator and discriminator with the frozen regressors.

gression problem, where convolutional network typically detects the boundary information of object beneficial for the domain adaptation [LeCun *et al.*, 1998; Krizhevsky *et al.*, 2012]. Thus, the relatively simple alignment strategies as DANN [Ganin and Lempitsky, 2015] easily learn the domain-invariant information for regressive task. And RSD [Chen *et al.*, 2021] empirically finds the further control of feature scale learned from image data achieves better distribution alignment. However, the multi-modal RF signals are sensitive to environmental changes due to the fact that a subtle change in the environment may impact the RF signal propagation and result in signal variations at these RF anchors. This leads to a significant distribution shift across different scenarios and affecting the model generalization. And the relationship between multi-sensor RF signals and localization output is too implicit to be detected, which increases the difficulty of learning intrinsic representations across two domains. We empirically evaluate the existing UDA methods [Courty *et al.*, 2017; Li *et al.*, 2020], and notice they hardly improve the model generalization when solving indoor localization based on multi-sensor RF signals.

To overcome the practical challenges, this paper proposes a robust and effective *Adversarial Bi-Regressor Network (ABR-Net)*. Specifically, the ABRNet network consists of one feature generator and two same-architecture regressors. To identify the certain and uncertain target samples, we adopt the adversarial training strategy by maximizing the bi-regressor discrepancy to seek domain-invariant representations. However, it is difficult to directly align source and target domains in the presence of a large discrepancy. An intuitive solution is to synthesize samples in intermediate domains to bridge the original two domains and gradually eliminate their domain mismatch. The main contributions of our work are summarized three folds:

- First, we propose an adversarial bi-regressor architecture to achieve a direct alignment across the source and target domains. Specially, the disagreement of dual regressors is maximized to discover target samples outside the source support and then learn domain-invariant features.
- Second, to fight off a situation where the cross-domain

distribution discrepancy becomes relatively significant, as in the case of multi-modal indoor localization, we attempt to construct two intermediate domains between the source and target domains and gradually eliminate their mismatch to achieve distribution alignment.

- Finally, the experimental performance on multi-model RF signal and image benchmarks illustrates the effectiveness of our method on solving challenging DAR in real applications. Moreover, the empirical studies analyze the contributions of our designed modules.

2 The Proposed Method

2.1 Problem Statement and Motivation

Domain adaptive regression (DAR) aims to learn transferable knowledge from the label-sufficient source domain $\mathcal{D}_s = \{\mathbf{X}_i^s, \mathbf{I}_i^s\}_{i=1}^{n_s}$ to predict the continuous output for the target samples $\mathcal{D}_t = \{\mathbf{X}_i^t\}_{i=1}^{n_t}$, where $\mathbf{X}_i^{s/t}$, \mathbf{I}_i^s corresponds to the instances and **annotations sampled from the continuous space**, and $n_{s/t}$ is the number of instances for each domain. Moreover, \mathbf{X}_i^s and \mathbf{X}_i^t are collected from the same sensors but with varying environment settings such as the change of furniture location, but share the identical continuous label space, e.g., 2D coordinate within the same floor map.

Take the fingerprinting-based indoor localization as an example [He and Chan, 2016; Wang *et al.*, 2017; Koike-Akino *et al.*, 2020]. During the offline fingerprinting phase, a robot moves in a pre-defined area, e.g., a multi-room apartment in Figure 1, and registers locations by collecting multi-modal RF signals from a set of fixed (Wi-Fi and UWB) anchors. Denoting the fingerprinted RF signals as $\mathbf{X}_i^s \in \mathbb{R}^{m \times d}$ where m is the number of continuously fingerprinted samples during the offline fingerprinting phase, and d is the dimension of the concatenation of multi-modal RF signals. Corresponding 2D coordinates $\mathbf{I}_i = (\mathbf{I}_i^x, \mathbf{I}_i^y)$ of m fingerprinted samples can be obtained from the robot using more costly on-board sensors such as Lidar and high-precision wheel encoder, where $\mathbf{I}_i^x \in [0, L]$, $\mathbf{I}_i^y \in [0, W]$ and L, W are the length and width of the floor respectively. Then, during the online test phase,

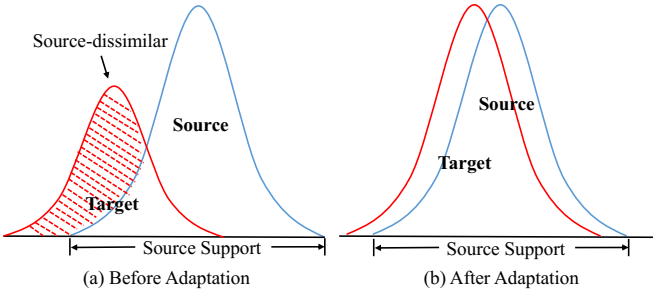


Figure 3: Before adaptation (a): target samples far from the source distribution are named as source-dissimilar ones; After adaptation (b): target features are more aligned over the source support.

new data \mathbf{X}_i^t are collected under, possibly, changed environments such as the movement of furniture and the close/open of a door. As a result, well-trained localization models using the offline fingerprinted (source domain) data \mathbf{X}_i^s may suffer from performance degradation when evaluating it on the test (target domain) samples \mathbf{X}_i^t from changed environments. The problem of interest is to predict the continuous 2D coordinates of these new test data \mathbf{X}_i^t by learning intrinsic and domain-invariant features from the offline signals.

Motivated by this consideration, this paper presents a novel method named Adversarial Bi-Regressors Network (ABR-Net) to achieve distribution alignment for DAR setting. As shown in Figure 2, ABRNet includes one feature extractor, two same-architecture regressors and one discriminator. The dual-regressor design allows the model to detect target samples outside source distribution and serves as one implicit discriminator to generate domain-invariant features. Moreover, to overcome the considerable distribution shift across source and target domains, ABRNet constructs source-similar and target-similar domains via the linear combination of original inputs and align them with the explicit discriminator to gradually eliminate the original cross-domain discrepancy.

2.2 Bi-Regressor Discrepancy Discovery

The feature generator takes source and target samples as input to yield high-level feature representations, i.e., $\mathbf{f}_i^{s/t} = F(\mathbf{X}_i^{s/t})$. For specific tasks, we can consider the CNN or LSTM network architecture as the feature extractor. However, the significant cross-domain difference results in the feature distribution mismatch. For instance, in the indoor RF localization task, source and target features corresponding to the same location may become significantly different due to environment variations nearby. Thus, the regressor learned from the source supervision fails to accurately predict the coordinate of several target instances which are shown as shaded area in Figure 3 and far from the source support. To achieve domain alignment, we first detect these target samples and adjust feature extractor to generate them within the source support. For the convenience of illustration, we name these target samples as source-dissimilar ones.

According to the voting strategy, two regressors easily produce inconsistent prediction on source-dissimilar target instances. Thus, we utilize the prediction disagreement to detect those target samples. Specifically, the introduced two regressors are equipped with the same network architec-

ture, i.e., two fully-connected layers transforming the high-level feature into the final prediction via $\hat{\mathbf{g}}_i^{s/t} = \hat{G}(\mathbf{f}_i^{s/t})$ or $\tilde{\mathbf{g}}_i^{s/t} = \tilde{G}(\mathbf{f}_i^{s/t})$ and $\hat{\mathbf{l}}_i^{s/t} = \hat{R}(\hat{\mathbf{g}}_i^{s/t})$ or $\tilde{\mathbf{l}}_i^{s/t} = \tilde{R}(\tilde{\mathbf{g}}_i^{s/t})$, where $\{\hat{G}, \hat{R}\}$ and $\{\tilde{G}, \tilde{R}\}$ are the components of two regressors, respectively. If the source and target regressors provide different predictions for the same target samples, we consider them as source-dissimilar ones with high confidence. The prediction discrepancy means that these two regressors involve different predictive abilities for target instances. For this objective, we need to maximize the discrepancy of two regressor. Different from the classification task, the final output of the regressors fail to reflect the probability distribution of samples in the continuous label space. Although the maximization of the final outputs of two regressors makes them different, such an operation easily misleads the regressor to deliberately search for solutions that are far from the ground-truth, leading to degraded performance of the regressor.

On the other hand, it is observed that the first output of regressor \mathbf{g}_i^t ($\mathbf{g} \in \{\hat{\mathbf{g}}, \tilde{\mathbf{g}}\}$) indicates the distribution of hidden feature space. Thus, we consider \mathbf{g}_i^t as the prior condition for the final output and concatenate them as $\hat{\mathbf{h}}_i^t = [\hat{\mathbf{g}}_i^t, \hat{\mathbf{l}}_i^t]$ or $\tilde{\mathbf{h}}_i^t = [\tilde{\mathbf{g}}_i^t, \tilde{\mathbf{l}}_i^t]$. To quantify the similarity between $\hat{\mathbf{h}}_i^t$ and $\tilde{\mathbf{h}}_i^t$ that reflects the bi-regressor discrepancy, one option is to use the Jaccard Similarity Coefficient (IoU score) from the field of object detection using visual sensors [Gupta *et al.*, 2014]. However, the non-differentiable nature of the IoU score leads to the optimization of the network weights problematic. Alternatively, a differentiable soft-similarity is developed here for this purpose:

$$\mathcal{L}_s = \sum_{i=1}^{n_t} \frac{\langle \hat{\mathbf{h}}_i^t, \tilde{\mathbf{h}}_i^t \rangle}{\sum_d (\hat{\mathbf{h}}_i^t + \tilde{\mathbf{h}}_i^t - \hat{\mathbf{h}}_i^t \otimes \tilde{\mathbf{h}}_i^t)}, \quad (1)$$

where $\langle \cdot, \cdot \rangle$ is the inner-product suggesting the intersection of $\hat{\mathbf{h}}_i^t$ and $\tilde{\mathbf{h}}_i^t$, \otimes is the element-wise product, and d is the dimension size of \mathbf{h}_i^t ($\mathbf{h} \in \{\hat{\mathbf{h}}, \tilde{\mathbf{h}}\}$). Note that \mathbf{g}_i^t is activated with the Sigmoid function before computing this metric. The minimization of Eq. (1) decreases the intersection between $\hat{\mathbf{g}}_i^t$ and $\tilde{\mathbf{g}}_i^t$ to further increase the difference of dual regressors, and vice versa. With Eq. (1), we can achieve the domain alignment with the bi-regressor in the following three steps.

First, due to the accessibility of source labels, it is simple to train the entire network to learn and adapt the source support set via a supervised manner as:

$$\min_{F, \hat{G}, \tilde{G}, \hat{R}, \tilde{R}} \mathcal{L}_r = \sum_{i=1}^{n_s} \left(\mathcal{L}_r^*(\hat{\mathbf{l}}_i^s, \mathbf{l}_i^s) + \mathcal{L}_r^*(\tilde{\mathbf{l}}_i^s, \mathbf{l}_i^s) \right), \quad (2)$$

where \mathcal{L}_r^* is the mean-square error (MSE) loss. This training manner likely yields two identical regressors. To avoid this issue, the second step is to maximize the difference of the bi-regressor with the aforementioned soft-similarity score to detect the source-dissimilar instances. Meanwhile, we still need to guarantee the predictive ability of model on the source samples. These considerations lead to the following objective function formulated as $\min_{\hat{G}, \tilde{G}, \hat{R}, \tilde{R}} \mathcal{L}_r + \mathcal{L}_s$. It is worth noting that the second step only updates the bi-regressor with the feature generator fixed. To this end, we can explore the

bi-regressor discrepancy to discover the target samples outside the source support. The third step optimizes the feature extractor to move these target samples into the source support set by improving the overlap between $\hat{\mathbf{h}}_i^t$ and $\tilde{\mathbf{h}}_i^t$ with the formulation as $\min_F -\mathcal{L}_s$.

2.3 Intermediate Bi-Domain Alignment

According to the above illustration, the conditional bi-regressor discrepancy maximizes the difference of two regressors to encourage the feature generator to produce domain-invariant representations. Thus, the essence of this module considers the combination of the two regressors as one discriminator to directly distinguish the source features from the target ones. However, the aforementioned training manner means these two regressors should be very similar with smaller difference which makes it invalid for considerable cross-domain shift.

An intuitive solution is to find two intermediate domains between the source and target ones and gradually mitigate the shift of two new introduced domains to align original distributions. To this end, we further explore data augmentation to achieve this motivation. First, we adopt a fixed ratio λ to linearly combine the source and target samples to synthesize source-similar and target-similar instances as:

$$\mathbf{X}_s^* = \lambda \mathbf{X}_i^s + (1 - \lambda) \mathbf{X}_j^t, \quad \mathbf{X}_t^* = (1 - \lambda) \mathbf{X}_i^s + \lambda \mathbf{X}_j^t, \quad (3)$$

where λ is set to 0.7 in all experiments. The combined instances gradually form two intermediate domains where one is closer to the source domain and the other is similar to the target domain. Second, we introduce an additional discriminator following the feature generator and utilize their adversarial relationship to mitigate the domain shift between source-similar and target-similar ones. Concretely, the augmented instances flow into the feature generator and two regressors with the corresponding output, i.e., $\mathbf{f}_*^{s/t} = F(\mathbf{X}_{s/t}^*)$, $\hat{\mathbf{I}}_*^s = \hat{R}(\hat{G}(\mathbf{f}_*^s))$ and $\tilde{\mathbf{I}}_*^t = \tilde{R}(\tilde{G}(\mathbf{f}_*^t))$. And then we adopt the adversarial loss as the following:

$$\mathcal{L}_{adv} = \mathbb{E}_{\mathbf{X}_*^s \sim \mathcal{D}_s^*} \log[D(\mathbf{f}_*^s, \hat{\mathbf{I}}_*^s)] + \mathbb{E}_{\mathbf{X}_*^t \sim \mathcal{D}_t^*} \log[1 - D(\mathbf{f}_*^t, \tilde{\mathbf{I}}_*^t)],$$

where $\mathcal{D}_{s/t}^*$ denotes the source-similar or target-similar domain and $D(\cdot)$ is the domain discriminator. Since the augmented samples are derived from the original source and target samples, their adversarial training would trigger the domain alignment of the original two domains. Thus, our augmented domain alignment can compensate the previous bi-regressor alignment.

2.4 Overall Objectives

So far, we have illustrated the details of our proposed ABR-Net to seek domain-invariant feature representation for cross-domain regression. The overall optimization iterates the following three steps until the convergence is achieved or the maximum number of iterations is reached:

$$\left\{ \begin{array}{l} \text{Step 1:} \quad \min_{F, \hat{G}, \tilde{R}, \hat{R}} \mathcal{L}_r = \sum_{i=1}^{n_s} \mathcal{L}_r^*(\hat{\mathbf{I}}_i^s, \mathbf{I}_i^s) + \mathcal{L}_r^*(\tilde{\mathbf{I}}_i^t, \mathbf{I}_i^t), \\ \text{Step 2:} \quad \min_{\hat{G}, \tilde{R}, \hat{R}} \mathcal{L}_s = \sum_{i=1}^{n_t} \frac{\langle \hat{\mathbf{h}}_i^t, \tilde{\mathbf{h}}_i^t \rangle}{\sum_a (\hat{\mathbf{h}}_i^t + \mathbf{h}_i^t - \hat{\mathbf{h}}_i^t \otimes \tilde{\mathbf{h}}_i^t)}, \\ \text{Step 3:} \quad \min_F \mathcal{L}_{adv} - \mathcal{L}_s, \quad \max_D \mathcal{L}_{adv}. \end{array} \right.$$

Method	RC	RU	RI	CU	CI	UI	Avg.
LSTM	1.39	1.08	1.12	1.39	0.98	2.43	1.40
TCA	1.36	0.97	0.93	1.32	0.98	2.22	1.30
DAN	1.39	1.02	1.03	1.34	0.98	2.29	1.34
DANN	<u>1.26</u>	0.98	<u>0.92</u>	1.26	<u>0.96</u>	<u>2.09</u>	<u>1.25</u>
JDOT	1.33	0.97	0.94	1.28	0.97	2.13	1.27
MCD	1.35	<u>0.96</u>	0.98	<u>1.21</u>	0.97	2.10	1.26
ETD	1.35	0.98	0.94	1.23	0.98	2.11	1.27
Ours	1.13	0.84	0.71	0.98	0.95	1.58	1.04

Table 1: Mean square error (MSE) computed on SPAWC2021 dataset with the combination of two sensors under Unsupervised Regressive Domain Adaptation. The best performance is highlighted with **bold** type, while the second one is emphasized with underline.

Method	RCU	RCI	RUI	CUI	RCUI	Avg.
LSTM	1.04	0.96	0.89	0.92	0.88	0.94
TCA	1.02	0.95	0.87	0.92	0.79	0.91
DAN	<u>0.96</u>	0.95	0.82	0.86	0.75	0.87
DANN	0.97	0.90	0.81	0.87	0.71	<u>0.85</u>
JDOT	0.99	<u>0.89</u>	0.88	<u>0.82</u>	<u>0.70</u>	<u>0.86</u>
MCD	1.01	0.96	<u>0.78</u>	0.85	0.74	0.87
ETD	1.04	0.94	0.85	0.86	0.80	0.90
Ours	0.85	0.70	0.57	0.68	0.54	0.67

Table 2: MSE computed on SPAWC2021 dataset with the combination of three or four sensors under Unsupervised Regressive Domain Adaptation. The best performance and the second one is highlighted with **bold** type and underline respectively.

3 Experiments

3.1 Experimental Steup

Datasets. 1) SPAWC2021 [Arnold and Schaich, 2021] is a large-scale indoor localization dataset, where a robot moves in a multi-room apartment and uses the received multi-modal RF signals from the deployed devices to predict its coordinates in this floor plan. For each location, the recorded signals for this robot consists of RSSI (**R**, scalar), CSI (**C**, 64 values over 64 sub-carrier frequencies) from 11 Wi-Fi anchors, UWB (**U**) from 3 anchors recording range and power readings, and a 9-dimensional IMU (**I**) signal. SPAWC2021 consists a well-annotated Dataset1 (source domain) with 750k samples and one unlabeled Dataset2 (target domain) with 650k instances. Their differences are the movement of furniture around the green box in Figure 1 and the change of data collection time, which results in significant distribution difference. 2) dSprites [Higgins *et al.*, 2017] is popular 2D synthetic image dataset suitable to unsupervised regressive domain adaptation and includes three domains as Color (**C**), Noisy (**N**) and Scream (**S**). And each domain involves 737,280 images. The regressive tasks in dSprites are to predict the scale and (x,y) coordinates.

Implementation details. Due to the difference of data type, we adopt distinct backbones to learn feature with corresponding optimizer on two benchmarks. For **SPAWC2021**, when moving along with the trajectory, the current location of robot is related to the status of previous t time steps. Considering

Method	C→N	C→S	N→C	N→S	S→C	S→N	Avg.
Resnet-18	0.94±0.06	0.90±0.08	0.16±0.02	0.65±0.02	0.08±0.01	0.26±0.03	0.498
TCA	0.94±0.03	0.87±0.02	0.19±0.02	0.66±0.05	0.10±0.02	0.23±0.04	0.498
DAN	0.70±0.05	0.77±0.09	0.12±0.03	0.50±0.05	<u>0.06±0.02</u>	0.11±0.04	0.377
DANN	0.47±0.07	0.46±0.07	0.16±0.02	0.65±0.05	0.05±0.00	0.10±0.01	0.315
JDOT	0.86±0.03	0.79±0.02	0.19±0.02	0.64±0.05	0.10±0.02	0.23±0.04	0.468
MCD	0.81±0.09	0.81±0.12	0.17±0.12	0.65±0.03	0.07±0.02	0.19±0.04	0.450
AFN	1.00±0.04	0.96±0.05	0.16±0.03	0.62±0.04	0.08±0.01	0.32±0.06	0.523
RSD	<u>0.32±0.02</u>	<u>0.35±0.02</u>	0.16±0.02	0.57±0.01	0.08±0.01	<u>0.09±0.02</u>	<u>0.258</u>
Ours	0.20±0.02	0.27±0.01	<u>0.14±0.01</u>	<u>0.56±0.03</u>	0.05±0.01	0.08±0.02	0.217

Table 3: Mean absolute error (MAE) computed on dSprites dataset with six tasks under Unsupervised Regressive Domain Adaptation. The best performance is highlighted with **bold** type, while the second one is emphasized with underline.

this point, we adopt Bi-LSTM as feature generator to capture the temporal dependency in the high-level representations. The inputs of Bi-LSTM correspond to the signals of 10 continuous time steps and its output is a 512-dimensional vector for each instance. In addition, the varying environmental factors have different impacts on different sensors. To evaluate the effectiveness of our method with various signals, we consider the concatenation of the multi-modal signals as the input. It is worth noting that the dimensions of the processed RSSI, CSI, UWB and IMU are 8, 40, 6 and 9, respectively. The domain-specific regressors and the discriminator both comprise of two full-connection layers. In the implementation, the Adam optimizer with a learning rate of 0.001 is used to update all network components. For **dSprites**, we following the training strategy of RSD [Chen *et al.*, 2021] utilize the pre-trained ResNet-18 as backbone to extract features and SGD with a momentum 0.95 to optimize the network architecture. The learning rate of the new added network modules is 10 times larger than that of the pre-trained module.

Baselines. To assess the power of our method on solving DAR, we compare it with the state-of-the-art domain adaptation methods which are appropriate for the regressive challenges. Specifically, the competitors are divided into two branches. The first one is metric-based solutions such as Transfer Component Analysis (**TCA**) [Pan *et al.*, 2010], Joint Distribution Optimal Transport (**JDOT**) [Courty *et al.*, 2017], Deep Adaptation Network (**DAN**) [Long *et al.*, 2015], Adaptive Feature Norm (**AFN**) [Xu *et al.*, 2019] and Enhanced Transport Distance (**ETD**) [Li *et al.*, 2020] and Representation Subspace Distance (**RSD**)¹ [Chen *et al.*, 2021]. Another is domain confusion approaches such as Domain Adversarial Neural Network (**DANN**) [Ganin and Lempitsky, 2015] and Maximum Classifier Discrepancy (**MCD**) [Saito *et al.*, 2018]. The domain adaptation baselines are deployed in the identical network (LSTM or ResNet-18) and tested on both datasets.

3.2 Comparison Results

Table 1 and Table 2 summarize the results of the mentioned methods on SPAWC2021 dataset with the concatenation of

¹Note that we attempt to evaluate **RSD** with the public code on SPAWC2021 dataset but the low-rank operation easily suffers from the vanishing gradient problem within a few iterations. Thus, we fail to report the results of RSD on SPAWC2021 dataset.

different signals. From both tables, several observations can be made as follows. **First**, our proposed method outperforms the baselines by a large margin on the average mean square error (MSE). Specially, when only two types of sensor are used, our strategy further improves the precision of localization prediction by comparing it with the second best result obtained by DANN (1.04 v.s. 1.25). It suggests that our method effectively eliminates the distribution discrepancy across Dataset1 and Dataset2 and transfers more source knowledge to target domain. Compared with DANN, our learning strategy explores the data augmentation to construct source-similar and target-similar samples and aligns them to achieve distribution matching. **Second**, it is simple to find that the model generalization is driven by the combination of different type of sensors. For example, when directly using the well-trained LSTM model into target domain, training model with CSI and IMU (CI) obtains better localization than that with UWB and IMU (UI) (0.98 v.s. 2.43). The main reason lies in that the signals UI are sensitive to varying external factors triggering considerable shift across source and target domains. Our proposed approach shows stronger mitigation of cross-domain discrepancy than the selected baselines. Especially, the comparison between our method and MCD with (1.58 v.s. 2.10) illustrates that our designed bi-regressor module better utilizes the maximization of discrepancy to discover the target samples far away from the source support and then minimize the dual-regressor difference to generate domain-invariant representations. **Third**, our method achieves the best localization precision (MSE:0.54) with the combination of all four type of sensors where the MSE of LSTM without any domain adaptations is about 0.88. The performance improvement is attributed to the efficient alignment of source and target distributions.

Table 3 reports the regressive performance on dSprites dataset. The distribution discrepancy in some tasks is smaller so that the well-learned source model (ResNet-18) shows promising prediction on target domain, e.g., N→C and S→C, where our method achieves comparable performance. With respect to more challenging tasks such as C→S, the prediction ability of our model on target domain becomes much better than that of ResNet-18 with 0.27 v.s. 0.90. Moreover, compared with RSD, our model further reduces the MAE from 0.32 to 0.20 on task C→N, which demonstrates that our model effectively eliminates the domain shift.

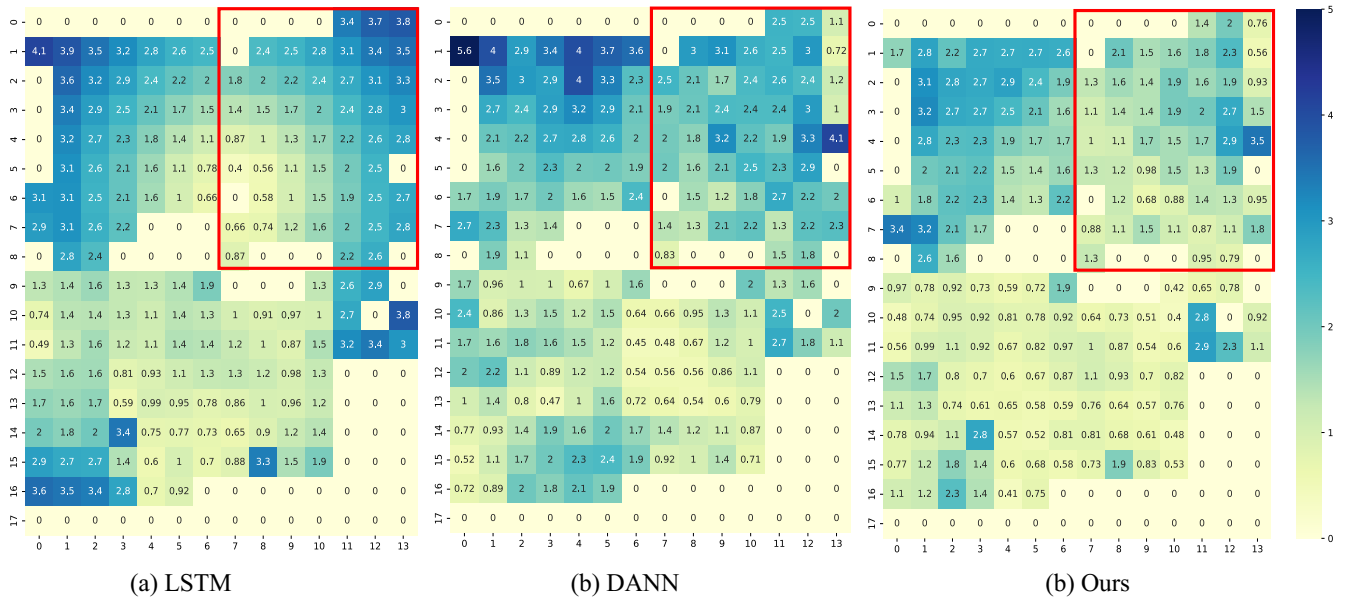


Figure 4: Visualization of average MSEs over $50 \times 50 \text{ cm}^2$ on the floor map. The x -axis and y -axis are divided several intervals with 0.5 meter, which forms the chessboard. The number of each grid is computed by taking the average MSE of all test samples falling into the grid.

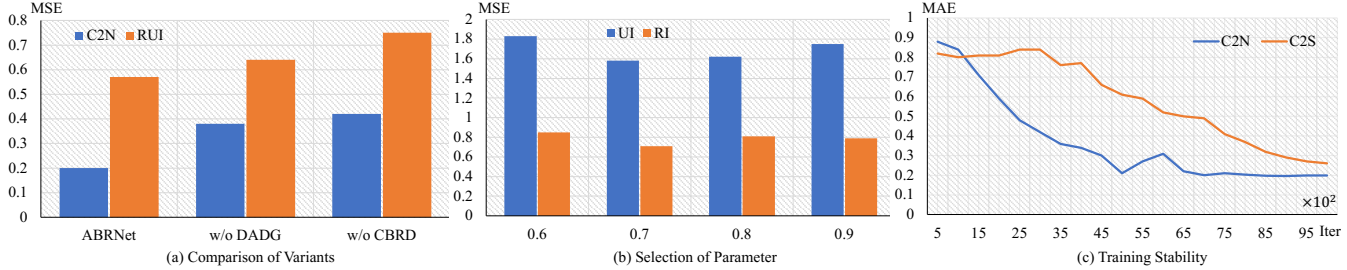


Figure 5: Empirical Analysis: (a) Comparison of ABRNet variants, (b) Selection of parameter λ , and (c) Training stability with iterations.

3.3 Empirical Analysis

Visualization. To better comprehend the performance gain of our method on SPAWC2021, Figure 4 specially visualizes the average MSE over $50 \times 50 \text{ cm}^2$ grids in the floor map and compare it with other baselines e.g., LSTM and DANN. The data description of SPAWC2021 shows for target domain, several cabinets were newly deployed in the area indicated by red box during the online data collection, leading to a distribution shift across source and target data. Thus, ResNet-18 and DANN give much worse localization performance than the proposed method for the target regressive task in the red box. Our method shows better capacity to distinguish source-dissimilar samples in the red box and significantly reduces the localization errors.

Ablation Study & Training Stability. Our method includes two important operations: conditional bi-regressor discrepancy (CBRD) and data augmentation driven game (DADG). To analyse the contribution of each module, we first remove one part while keeping the other intact to form two variants: w/o DADG and w/o CBRD. From the results in Figure 5, it is seen that CBRD contributes more to the performance gain than DADG as the amount of localization error increase more w/o CBRD than w/o DADG. In addition, we analyse the performance sensitivity to different parameter (λ)

selections on tasks UI and RI. From Figure 5, the model appears to be insensitive to λ and achieves the best performance with $\lambda = 0.7$. Finally, to assess the training stability, we record the MAE of target domain via target regressor over the number of iterations in Figure 5 (c) which suggests the training process achieves convergence.

4 Conclusion

This paper explores domain adaptation regression (DAR) aiming to transfer source knowledge to provide continuous-label prediction for unlabeled target samples in real-world applications such as indoor localization. To this end, we propose a novel adversarial bi-regressor network (ABRNet) that consists of a feature generator, two regressors and a discriminator. First, we utilize the prediction disagreement of two regressors to identify target samples outside source distribution and then seek domain-invariant features depending on the adversarial relationship between feature generator and dual regressors. Moreover, considering the difficulty of direct alignment across two domains with large domain shift, ABRNet attempts to construct two intermediate domains with smaller discrepancy to gradually eliminate the original domain mismatch. Extensive experiments verify our method outperforms the state-of-the-art on solving DAR issue.

References

- [Arnold and Schaich, 2021] Maximilian Arnold and Frank Schaich. Indoor positioning systems: Smart fusion of a variety of sensor readings. *ArXiv preprint*, 2021.
- [Awais *et al.*, 2021] Muhammad Awais, Fengwei Zhou, Hang Xu, Lanqing Hong, Ping Luo, and Sung-Ho Bae. Adversarial robustness for unsupervised domain adaptation. In *ICCV*, 2021.
- [Chang *et al.*, 2019] Woong-Gi Chang, Tackgeun You, Seonguk Seo, Suha Kwak, and Bohyung Han. Domain-specific batch normalization for unsupervised domain adaptation. In *CVPR*, 2019.
- [Chen *et al.*, 2021] Xinyang Chen, Sinan Wang, Jianmin Wang, and Mingsheng Long. Representation subspace distance for domain adaptation regression. In *ICML*, 2021.
- [Courty *et al.*, 2017] Nicolas Courty, Rémi Flamary, Amaury Habrard, and Alain Rakotomamonjy. Joint distribution optimal transportation for domain adaptation. *ArXiv preprint*, 2017.
- [Ganin and Lempitsky, 2015] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *ICML*, 2015.
- [Gupta *et al.*, 2014] Saurabh Gupta, Ross Girshick, and Pablo Arbeláez. Learning rich features from rgb-d images for object detection and segmentation. In *ECCV*, 2014.
- [He and Chan, 2016] S. He and S.G. Chan. Wi-Fi fingerprint-based indoor positioning: Recent advances and comparisons. *IEEE Communications Surveys & Tutorials*, 2016.
- [Higgins *et al.*, 2017] Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, and Shakir Mohamed. beta-vae: Learning basic visual concepts with a constrained variational framework. In *ICLR*, 2017.
- [Jing *et al.*, 2021] Taotao Jing, Bingrong Xu, and Zhengming Ding. Towards fair knowledge transfer for imbalanced domain adaptation. *IEEE Transactions on Image Processing*, 30:8200–8211, 2021.
- [Koike-Akino *et al.*, 2020] T. Koike-Akino, P. Wang, M. Pajovic, H. Sun, and P. V. Orlik. Fingerprinting-based indoor localization with commercial mmwave WiFi: A deep learning approach. *IEEE Access*, 2020.
- [Krizhevsky *et al.*, 2012] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *NIPS*, 2012.
- [LeCun *et al.*, 1998] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *IEEE*, 1998.
- [Li *et al.*, 2020] Mengxue Li, Yi-Ming Zhai, You-Wei Luo, and Peng-Fei Ge. Enhanced transport distance for unsupervised domain adaptation. In *CVPR*, 2020.
- [Long *et al.*, 2015] Mingsheng Long, Yue Cao, Jianmin Wang, and Michael Jordan. Learning transferable features with deep adaptation networks. In *ICML*, 2015.
- [Na *et al.*, 2021] Jaemin Na, Heechul Jung, and Hyung Jin Chang. Fixbi: Bridging domain spaces for unsupervised domain adaptation. In *CVPR*, 2021.
- [Pan *et al.*, 2010] Sinno Jialin Pan, Ivor W Tsang, James T Kwok, and Qiang Yang. Domain adaptation via transfer component analysis. *IEEE transactions on neural networks*, 2010.
- [Pan *et al.*, 2019] Yingwei Pan, Ting Yao, Yehao Li, Yu Wang, Chong-Wah Ngo, and Tao Mei. Transferrable prototypical networks for unsupervised domain adaptation. In *CVPR*, 2019.
- [Saito *et al.*, 2018] Kuniaki Saito, Kohei Watanabe, and Yoshitaka Ushiku. Maximum classifier discrepancy for unsupervised domain adaptation. In *CVPR*, 2018.
- [Sellami and Tabbone, 2022] Akrem Sellami and Salvatore Tabbone. Deep neural networks-based relevant latent representation learning for hyperspectral image classification. *Pattern Recognition*, 2022.
- [Tong *et al.*, 2021] Xinyu Tong, Hao Li, and Xiaohua Tian. Wi-fi localization enabling self-calibration. *IEEE/ACM Transactions on Networking*, 2021.
- [Wang *et al.*, 2017] X. Wang, L. Gao, S. Mao, and S. Pandey. CSI-based fingerprinting for indoor localization: A deep learning approach. *IEEE Trans. on Vehicular Technology*, 2017.
- [Xia and Ding, 2020] Haifeng Xia and Zhengming Ding. Structure preserving generative cross-domain learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4364–4373, 2020.
- [Xia *et al.*, 2021a] Haifeng Xia, Taotao Jing, Chen Chen, and Zhengming Ding. Semi-supervised domain adaptive retrieval via discriminative hashing learning. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 3853–3861, 2021.
- [Xia *et al.*, 2021b] Haifeng Xia, Handong Zhao, and Zhengming Ding. Adaptive adversarial network for source-free domain adaptation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9010–9019, 2021.
- [Xu *et al.*, 2014] Yong Xu, Jun Du, Li-Rong Dai, and Chin-Hui Lee. A regression approach to speech enhancement based on deep neural networks. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 2014.
- [Xu *et al.*, 2019] Ruijia Xu, Guanbin Li, Jihan Yang, and Liang Lin. Larger norm more transferable: An adaptive feature norm approach for unsupervised domain adaptation. In *ICCV*, 2019.
- [Yamada *et al.*, 2014] Makoto Yamada, Leonid Sigal, and Yi Chang. Domain adaptation for structured regression. *International Journal of Computer Vision*, 2014.
- [Zhou *et al.*, 2021] Mu Zhou, Yaohua Li, Muhammad Junaid Tahir, Xiaolong Geng, Yong Wang, and Wei He. Integrated statistical test of signal distributions and access point contributions for wi-fi indoor localization. *IEEE Transactions on Vehicular Technology*, 2021.