Domain Generalization via Selective Consistency Regularization for Time Series Classification

Wenyu Zhang
Institute for Infocomm Research, A*STAR

Mohamed Ragab and Chuan-Sheng Foo
Institute for Infocomm Research, A*STAR
Centre for Frontier AI Research, A*STAR

Abstract—**Domain generalization methods aim to learn models robust to domain shift with data from a limited number of source domains and without access to target domain samples during training. Popular domain alignment methods for domain generalization seek to extract domain-invariant features by minimizing the discrepancy between feature distributions across all domains, disregarding inter-domain relationships. In this paper, we instead propose a novel representation learning methodology that selectively enforces prediction consistency between source domains estimated to be closely-related. Specifically, we hypothesize that domains share different class-informative representations, so instead of aligning all domains which can cause negative transfer, we only regularize the discrepancy between closely-related domains. We apply our method to time-series classification tasks and conduct comprehensive experiments on three public real-world datasets. Our method significantly improves over the baseline and achieves better or competitive performance in comparison with state-of-the-art methods in terms of both accuracy and model calibration.**

I. INTRODUCTION

Increasing accessibility to data has spurred the use of data-driven deep learning methods. In practical deployments, models need to be robust to data distribution shifts between training *(source domains)* and test *(target domain)* data, when data are collected for the same task but in different environments [1], [2]. Such domain shift may occur as the collection of training data is subject to resource constraints and may not provide sufficient coverage, or is conducted in controlled settings that do not fully assimilate real environments [3]–[5].

In this work, we consider the *domain generalization* (DG) setting where we leverage data from multiple source domains to generalize to unseen target domains. This differs from *domain adaptation* (DA) which assumes access to unlabeled target samples for training. We further focus on time series classification; while there are many works studying DG in image classification [1]–[3], [6]–[9] and natural language tasks [10]–[12], there is limited literature and evaluation for time series classification. We therefore aim to address this gap. Example use-cases include generalization across human subjects for activity recognition [13], and across operating conditions for sensor-based machine fault detection [14].

A popular DG approach is domain alignment, that is, enforcing feature distributions to be the same across all source domains [1], [13], [15]–[19] to learn domain-invariant features that potentially generalize to unseen domains. However, some works in the related domain adaptation and transfer learning literature suggest that domains can have different informative


Fig. 1. **Left**: Existing domain alignment methods align all source domains equally. **Right**: Our proposed method, by considering inter-domain relationships, selectively aligns closely-related domains and allows greater diversity in output predictions. Best viewed in colour.

characteristics and transferring between less related domains can inhibit the learning of these characteristics [11], [20], [21], causing negative transfer and harming generalization.

In this paper, we propose a novel DG approach for time series classification where we further consider cross-domain relationships during domain alignment, and regularize subsets of source domains separately based on their similarity (as illustrated in Figure 1). We propose selective cross-domain consistency regularization to encourage invariant class-conditional predictions amongst similar source domains, so that model predictions are less dependent on domain-specific information, while still being allowed to be diverse to capture cross-domain differences. Our consistency regularization is applied on the logits (pre-softmax classifier outputs); as far we know, existing DG methods regularize on either features or soft label predictions [1], [10], [16], [17], [22]–[27] or on logits in conjunction with features [28]. Features are challenging to align in high-dimensional space, and soft label predictions are normalized logits and have less variability. We leverage auxiliary domain metadata to directly infer inter-domain relationships based on application-specific knowledge when such metadata is available. Otherwise, we estimate latent relationships using a clustering-based measure of domain divergence amongst source domains. Figure 2 provides an overview of our proposed method. We summarize our contributions as follows:

- We propose a new DG methodology for time series classification that selectively enforces prediction consistency between source domains estimated to be closely-related, which helps to calibrate the model from being over and/or under-confident in its predictions. We combine this with domain-wise augmentation to generate more diverse samples in closely-related domains for regularization;
- The proposed method is easy to implement with data augmentation and logit regularization on top of empirical risk minimization;
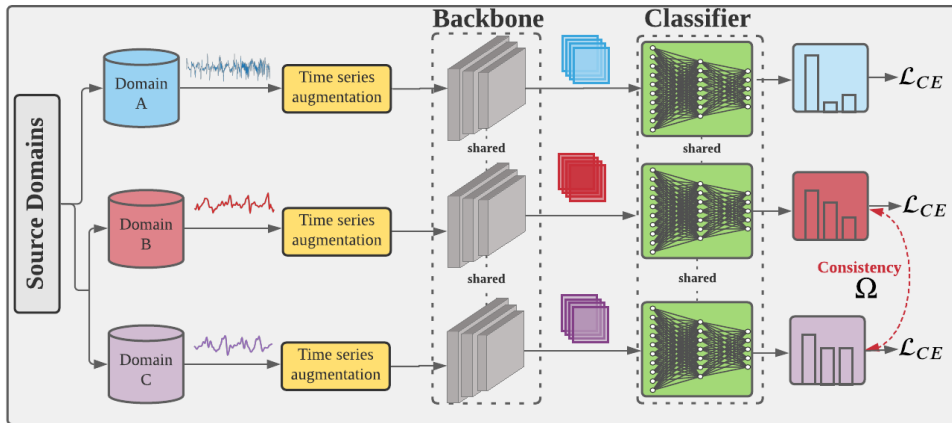
Fig. 2. Method overview. Considering multiple source domains, our proposed method enforces prediction consistency between similar domains through selective regularization of logits (pre-softmax classifier outputs). We encourage similar domains (B and C in illustration) to share predicted class relationships, while allowing diverse predicted class relationships across dissimilar domains (A vs B and C in illustration).

- We perform extensive evaluation of existing methods on 3 public datasets for classification accuracy and calibration, and our proposed method demonstrates better or competitive performance compared to state-of-the-art methods.

II. RELATED WORKS

We discuss related works to learn more robust representations with a single model. Ensembles [7], [11], [29], [30] may not scale to multiple source domains and not discussed here.

Data augmentation and generation: Enriching training sample diversity naturally helps generalization. [31]–[33] generates samples by adversarial perturbation, and [34], [35] interpolate between samples of different domains. Due to advancements in style transfer [36] and generative adversarial networks [37], many image augmentation techniques are developed [2], [38]–[40]. However, advanced techniques for image augmentation may not readily apply to time series.

Learning domain-invariant features: An approach to learn invariant features is to train the feature extractor such that the same classifier is optimal for all domains [41], [42]. Many works learn invariant features by aligning representations from all source domains. Some works directly minimize the distance between features or soft labels by a distance measure or adversarial networks [13], [15]–[18], [43]–[45], and some others use meta-learning to simulate domain shift [3], [12], [24]. However, distribution alignment between less related domains can cause negative transfer [11], [20], [21].

Robustness: Models can learn to be more robust to domain shifts through perturbations during training. [6] episodically switches the feature extractor or classifier to domain-specific counterparts, and [46] zeros out features associated with the highest gradient in the classification layer to learn more diverse features. Another approach to robustness minimizes worst-case risk over all domains, resulting in a training objective that is a weighted average of source domain losses [47], [48].

III. PROPOSED METHOD

A. Preliminaries: Notation

We denote total N observed samples from M source domains as $\{(\boldsymbol{x}_n, \boldsymbol{y}_n, d_n)\}_{n=1}^{N}$, where for the n-th sample,

$\boldsymbol{x}_n$ and $\boldsymbol{y}_n$ are predictor and response, and $d_n \in \{1, \dots, M\}$ is domain label and is unavailable at test-time. $\boldsymbol{y}_n$ is a one-hot vector of the true class label in L classes. Samples in each domain d are independently and identically distributed according to a domain-dependent data distribution as $(\boldsymbol{x}, \boldsymbol{y}) \sim P_d(\mathcal{X}, \mathcal{Y})$.

A neural network model is composed of feature extractor $f(\cdot)$ parameterized by Θ that yields learned features $\boldsymbol{z} = f(\boldsymbol{x}; \Theta)$, and classifier $h(\cdot)$ parameterized by Ψ that yields logits output $\boldsymbol{g} = h(\boldsymbol{z}; \Psi)$. Soft labels (a vector of estimated class probabilities) are obtained by $\boldsymbol{s} = softmax(\boldsymbol{g})$ where $\boldsymbol{s}[i] = \frac{\exp(\boldsymbol{z}[i])}{\sum_{\ell=1}^{L} \exp(\boldsymbol{z}[\ell])}$. The final predicted class is the one with the highest probability in $\boldsymbol{s}$ i.e. $\arg\max_i \boldsymbol{s}[i]$.

B. Selective Cross-Domain Consistency Regularization

We hypothesize that domains can have different informative characteristics and consequently class relationships which do not exactly match. That is, there exists domain $d^{(i)}$ and $d^{(j)}$ with $P_i(\mathcal{X}) \neq P_j(\mathcal{X})$ where there exists sample $\boldsymbol{x}^{(d^{(i)})}$ such that $P_i(\mathcal{Y}|\boldsymbol{x}^{(d^{(i)})}) \neq P_j(\mathcal{Y}|\boldsymbol{x}^{(d^{(j)})})$ for all $\boldsymbol{x}^{(d^{(j)})}$ under ground truth distributions; note $P_i(\mathcal{Y}|\boldsymbol{x}) = P_j(\mathcal{Y}|\boldsymbol{x})$ for the same $\boldsymbol{x}$. Aligning these two domains can result in inaccurate class relationships. Specifically, objectives such as DANN [10] align $d^{(j)}$ to $d^{(i)}$ assuming that for every $\boldsymbol{x}^{(d^{(i)})}$, there exists a matching $\boldsymbol{x}^{(d^{(j)})}$ where $P_i(\mathcal{Y}|\boldsymbol{x}^{(d^{(i)})}) = P(\mathcal{Y}|f(\boldsymbol{x}^{(d^{(i)})}; \Theta)) = P(\mathcal{Y}|f(\boldsymbol{x}^{(d^{(j)})}; \Theta)) = P_j(\mathcal{Y}|\boldsymbol{x}^{(d^{(j)})})$. When this assumption is invalid, aligning implies that a $\boldsymbol{x}^{(d^{(j)})}$ where $P_i(\mathcal{Y}|\boldsymbol{x}^{(d^{(i)})}) \neq P_j(\mathcal{Y}|\boldsymbol{x}^{(d^{(j)})})$ will be aligned to $\boldsymbol{x}^{(d^{(i)})}$ in the feature space, resulting in inaccurate class relationships since $P_i(\mathcal{Y}|\boldsymbol{x}^{(d^{(i)})}) = P(\mathcal{Y}|f(\boldsymbol{x}^{(d^{(i)})}; \Theta)) = P(\mathcal{Y}|f(\boldsymbol{x}^{(d^{(j)})}; \Theta)) \neq P_j(\mathcal{Y}|\boldsymbol{x}^{(d^{(j)})})$ [20]. Instead, we align only similar domains to balance between domain-invariance and diversity. We align logits instead of features to directly calibrate task-specific predictions through both feature extractor and classifier.

More formally, we aim to learn model parameters such that the conditional distribution $P(\mathcal{G}|\mathcal{Y}) = P(h(f(\mathcal{X}; \Theta); \Psi)|\mathcal{Y})$ is invariant for closely-related domains according to latent inter-domain relationships. By encouraging invariant class-conditional predictions, model predictions can be less reliant on domain-specific information and more robust to domain

shifts at test-time. Specifically, for closely-related domains $d^{(i)}$ and $d^{(j)}$ with $g^{(d^{(i)},\ell)} \sim P_i(\mathcal{G}|\mathcal{Y} = \ell)$ and $g^{(d^{(j)},\ell)} \sim P_j(\mathcal{G}|\mathcal{Y} = \ell)$ for class ℓ, we apply squared maximum mean discrepancy (MMD) [49] for distribution matching. Squared MMD between two distributions $\mathbb{Q}_i$ and $\mathbb{Q}_j$ is defined as

$$MMD^2(\mathbb{Q}_i, \mathbb{Q}_j) = \|E_{\mathbb{Q}_i}(\phi(\mathcal{G})) - E_{\mathbb{Q}_j}(\phi(\mathcal{G}))\|_{\mathcal{B}}^2 \quad (1)$$

for reproducing kernel Hilbert space $\mathcal{B}$ and mapping function ϕ. We take $\mathcal{B}$ as Euclidean space and ϕ as identity function, for which MMD is equivalent to the L_2 distance between distribution means. This formulation is computationally simple and empirically outperforms a more complex alternative (see Appendix B).

The objective function for model learning is

$$L(\Theta, \Psi) = L_{CE}(\Theta, \Psi) + \lambda\Omega(\Theta, \Psi) \quad (2)$$

where $L_{CE}(\Theta, \Psi) = -\sum_{n=1}^{N} \boldsymbol{y}_n \cdot \log(\boldsymbol{s}_n)$ is the supervised cross-entropy loss, and $\Omega(\Theta, \Psi)$ is the proposed regularizer:

$$\Omega(\Theta, \Psi) = \sum_{d^{(i)}=1}^{M} \sum_{d^{(j)}=1}^{M} w(d^{(i)}, d^{(j)}) \sum_{\ell=1}^{L} \|\bar{g}^{(d^{(i)},\ell)} - \bar{g}^{(d^{(j)},\ell)}\|_2^2 \quad (3)$$

where $\bar{g}^{(d^{(i)},\ell)}$ is the mean logit vector for domain $d^{(i)}$ class ℓ, referred to as the class-conditional domain centroid. Weights $w(d^{(i)}, d^{(j)}) \geq 0$ depend on pairwise domain similarity between $d^{(i)}$ and $d^{(j)}$ to impose greater regularization on more similar domains. We describe estimation procedures for $w(d^{(i)}, d^{(j)})$ in Section III-C. The learned model is encouraged to preserve prediction consistency between similar domains. This helps to prevent the model from being over-confident in its predictions due to overfitting on spurious features of individual domains, thus improving calibration, in addition to increasing robustness to domain shifts. We summarize the proposed method in Algorithm 1.

Algorithm 1 Proposed method

Input: Training data $(\boldsymbol{x}, \boldsymbol{y}, d)$; regularization hyperparameter λ, model parameterized by (Θ, Ψ); learning rate η; # iterations I; cluster update interval U
Additional input for metadata sim.: Metadata T
Additional input for learned sim.: RBF kernel parameter ξ

1: **for** iteration $i = 0 : I - 1$ **do**
2:     Augment time series input $x \leftarrow \text{Augment}(x)$
3:     **if** $(i \bmod U) = 0$ **then**
4:         Update regularization weights w by Eqn 4 (metadata sim.) or 6 (learned sim.)
5:     Get regularization $\Omega(\Theta, \Psi)$ by Eqn 3
6:     Get overall objective $L(\Theta, \Psi)$ by Eqn 2
7:     Update model $(\Theta, \Psi) \leftarrow (\Theta, \Psi) - \eta\nabla L(\Theta, \Psi)$

C. Defining Domain Similarity for Selective Regularization

Inter-domain relationships need to be determined in order to select closely-related domains for consistency regularization. We consider two scenarios: when domain metadata is available, and when it is not; domain metadata are descriptions of source domains to provide (possibly limited) context of the environments in which data is collected.

1) Metadata Based Similarity: With expert knowledge of the application, users can directly use available metadata to infer relationships and group the domains into clusters. Domains across clusters are assumed to not share class relationships and hence are not regularized. We denote the function $clust : \{1, \ldots, M\} \times \{T\} \to \{1, \ldots, K\}$ as the map from domain index and metadata T to cluster index for K clusters, and $\mathcal{S}_c = \{d | clust(d, T) = c, d \in \{1, \ldots, M\}\}$ the set of domains in cluster c. We set

$$w_{meta}(d^{(i)}, d^{(j)}) = \sum_{c=1}^{K} \frac{1}{2|\mathcal{S}_c|} \mathbb{1}_{\{clust(d^{(i)},T)=clust(d^{(j)},T)=c\}}$$

$$(4)$$

and the consistency regularization function in Equation 3 can be restated as

$$\Omega_{meta}(\Theta, \Psi) = \sum_{c=1}^{K} \sum_{d \in \mathcal{S}_c} \sum_{\ell=1}^{L} \|\bar{g}^{(d,\ell)} - \gamma^{(c,\ell)}\|_2^2 \quad (5)$$

where we refer to $\gamma^{(c,\ell)} = \frac{1}{|\mathcal{S}_c|} \sum_{d \in \mathcal{S}_c} \bar{g}^{(d,\ell)}$ of cluster c class ℓ as the class-conditional cluster centroid. The scaling factor in w_{meta} ensures that when domain centroids are equidistant to their cluster centroids, the amount of contribution each cluster makes to the regularization is proportional to its size.

2) Learned Similarity: When no domain metadata is available, we propose estimating domain relationships by inter-domain divergences during training. We measure the distance between two domains for class ℓ as the squared L_2 distance between their class-conditional domain centroids, and $d^{(j)}$ is defined as the nearest neighbor domain to $d^{(i)}$ if it is nearest to $d^{(i)}$ for the most number of classes. We estimate the nearest neighbor domain as the most similar domain to $d^{(i)}$ at fixed intervals during training (every 100 iterations in our experiments), and we enforce prediction consistency between each domain and its nearest neighbor domain. That is, we set the weights $w_{learned}(d^{(i)}, d^{(j)})$ as per Equation 6 if $d^{(j)}$ is the nearest neighbor and 0 otherwise, where

$$w_{learned}(d^{(i)}, d^{(j)}) = \frac{1}{L} \sum_{\ell=1}^{L} \exp\left(\frac{-\|\bar{g}^{(d^{(i)},\ell)} - \bar{g}^{(d^{(j)},\ell)}\|_2^2}{2\xi^2}\right)$$

$$(6)$$

by applying an RBF kernel on the inter-domain distance with hyperparameter ξ. We block gradients to prevent w and inter-domain distance from updating in opposing directions.

D. Domain-wise Time Series Augmentation

To achieve additional robustness to data perturbations, we apply time series augmentations with 0.5 probability. For each source domain, we sample an augmentation function from a pre-defined distribution at each iteration. The domain-wise augmentation simulates potential test-time domain shifts. We consider 3 time series augmentations, namely mean shift, scaling and masking, in Table I. The choice of augmentations depends on the dataset to avoid perturbing characteristics known to be important for classification. We provide augmentation details for each dataset in Section IV.

TABLE I
TIME SERIES AUGMENTATIONS.

Augmentation	General Expression
mean shift	$a_{mean}(x) = x - \mu + \mu_{new}$
scaling	$a_{scale}(x) = \left(\frac{x-\mu}{\sigma}\right) * \sigma_{new} + \mu$
masking	$a_{mask}(x[i]) = \begin{cases} x[i] & \text{w.p. } 0.9 \\ \mu & \text{w.p. } 0.1 \end{cases}$

TABLE II
DOMAIN ATTRIBUTES.

Loc.	Loading torque				User	Phone model			
	0	1	2	3		Nexus	S3	S3 mini	S+
Drive	A	B	C	D	User 1	A	B	C	D
Fan	E	F	G	H	User 2	E	F	G	H
					User 3	I	J	K	L

(a) Bearings   (b) HHAR

IV. EXPERIMENTS

We compare with baseline ERM and state-of-the-art DG methods: GroupDRO [47], VREx [48], IRM [41], Interdomain Mixup [34], RSC [46], MTL [50], MLDG [3] and Correlation [45]. We also reformulate 4 popular domain adaption methods for DG following [1]: MMD-DG [17], CORAL-DG [22], DANN-DG [10] and CDANN-DG [18], [51].

We evaluate the proposed method on three real-world time series datasets, and use leave-one-domain-out evaluation following [1]. We provide details of the evaluation procedure, backbone networks and hyperparameters in the Appendix A. We evaluate on classification performance and expected calibration error (ECE), where the latter is a metric that measures how closely model confidence match probabilities of correct predictions and zero ECE means perfect calibration. Specifically, ECE is the expected difference between confidence and accuracy empirically approximated by

$$ECE = \sum_{j=1}^{J} \frac{\|B_j\|}{N} \|acc(B_j) - conf(B_j)\|$$

where $B_j = \{n | n \in \{1, \ldots, N\}, max(\boldsymbol{s}_n) \in \left(\frac{j-1}{J}, \frac{j}{J}\right]\}$ is a bin containing indices of samples whose confidence for the predicted class falls in the corresponding interval, for $j \in \{1, \ldots, J\}$ and total $J = 15$ bins following [52]. For each bin, $acc(B_j) = \frac{1}{|B_j|} \sum_{n \in B_j} \mathbb{1}(\hat{\boldsymbol{y}}_n = \boldsymbol{y}_n)$ and $conf(B_j) = \frac{1}{|B_j|} \sum_{n \in B_j} max(\boldsymbol{s}_n)$.

A. Results on Fault Detection

The Bearings dataset [53] from Case Western Reserve University is widely used for predictive maintenance. It contains 12kHz vibration signals to detect bearings faults in rotating machines. We extract length-4096 samples by a sliding window with stride 290 [54]. There are 1 healthy class and 9 fault classes: inner-race fault (IF), outer-race fault (OF), and ball fault (BF) with each subdivided into dimensions 0.007, 0.014 and 0.021 inches. We apply mean shift, scaling and masking augmentations by setting $\mu = \bar{x}$, $\mu_{new} = 0$, $\sigma = sd(x)$ and $\sigma_{new} = 1$ in Table I. Samples are augmented with probability 0.5, and additional augmentations are applied with probability

0.5 to allow a mixture of perturbations. There are 8 domains: drive end and fan end location with each operated at 0, 1, 2, and 3 loading torques as in Table II. For metadata based similarity, consistency regularization is applied on domains with the same location. From Table III, the proposed method improves over baseline ERM on almost all target domains. On average, the proposed method attains the best accuracy across all methods with 87.9% and 89.1% given metadata based and learned similarity, respectively. The proposed method also attains lowest calibration errors of 9.2% and 6.2%.

B. Results on Human Activity Recognition

The HHAR dataset [55] consists of multi-channel sensor readings to classify six activities: Biking, Standing, Sitting, Walking, Stair down, and Stair up. Following a recent DA work [13], we focus on smartphone accelerometer readings in the x, y and z direction and extract length-128 samples. Samples are scaled by $\frac{1}{20}$ so that readings for all 3 channels approximately fall between -1 and 1. For this application, mean and standard deviation are known to be important classification features and activities such as Standing are sensitive to abrupt sensor reading changes [56], hence we apply limited augmentation i.e. scaling with $\mu = 0, \sigma = 1$ and $\sigma_{new} \sim Unif(0.8, 1.2)$. To keep domain number to a suitable level for leave-one-domain-out evaluation, we use 12 domains: the first 3 users each with 4 phone models as in Table II. For metadata based similarity, consistency regularization is applied on domains with the same user. From Table IV, the proposed method improves over the baseline ERM in almost all cases, and has the best accuracy of 88.5% on average. The second-best performing method RSC encourages learning more diverse features by feature masking, and applying the proposed method on RSC by alternating between the two methods further improves average accuracy to 88.9%. We chose the alternating procedure [57] so that strategies from the two methods do not directly interfere with each other.

We note that although RSC performs similarly to our proposed method, it utilizes masking instead of alignment strategies, and our proposed method significantly improves over all alignment methods evaluated namely Correlation, DANN-DG, CDANN-DG, CORAL-DG and MMD-DG. Furthermore, our proposed method (ECE=$4.1\%, 4.9\%$) learns better calibrated models than RSC (ECE=5.3%).

C. Results on Mortality Prediction

MIMIC-III [58], [59] is a public clinical dataset for binary in-hospital mortality prediction. We select 11 vital signs that are hourly aggregated, with 24 hours length sample from each patient [60]. We apply the same data augmentations as in Section IV-B. Different age groups represents different domains with Group A: 20-45, Group B: 46-65, Group C: 66-85, Group D: over 85. Due to class imbalance, we downsample the negative class so that both classes have equal sample size. We evaluate performance with AUC following [60]. From Table V, RSC underperforms even ERM on this task, while alignment methods improve performance over baseline ERM.

TABLE III

BEARINGS: CLASSIFICATION ACCURACY AND EXPECTED CALIBRATION ERROR (ECE) ON TARGET DOMAIN USING LEAVE-ONE-DOMAIN-OUT TESTING.

Method	Accuracy (%) ↑									ECE (%) ↓
	A	B	C	D	E	F	G	H	Avg	Avg
ERM	65.4	93.8	96.0	71.2	68.4	83.0	94.0	86.2	82.2	14.2
IRM	60.7	87.0	89.2	76.0	62.8	80.9	92.8	88.5	79.7	15.9
GroupDRO	55.7	70.0	77.8	74.8	60.7	59.2	65.6	50.5	64.3	31.9
Interdomain Mixup	62.0	86.5	96.8	76.0	82.0	95.4	97.7	87.2	85.4	13.8
MLDG	62.8	77.6	85.9	72.8	63.3	58.5	60.7	55.0	67.1	30.8
MTL	35.3	64.5	66.6	48.3	47.6	48.2	36.7	44.9	49.0	44.4
Correlation	46.5	79.0	90.0	69.1	71.5	85.5	80.5	83.9	75.7	18.2
VREx	63.8	90.5	97.2	81.6	70.3	83.9	92.0	84.3	83.0	14.3
RSC	62.4	94.4	98.0	86.5	73.4	87.4	97.3	85.7	85.6	12.0
DANN-DG	56.2	84.7	92.2	80.2	70.0	79.1	89.1	90.5	80.3	15.1
CDANN-DG	56.0	80.8	94.8	80.2	70.7	81.4	90.3	84.0	79.8	15.7
CORAL-DG	62.5	77.9	90.0	76.0	63.1	79.2	74.5	83.0	75.8	19.7
MMD-DG	53.9	67.7	84.2	67.6	63.2	74.3	74.7	56.7	67.8	28.7
Ours (Metadata sim.)	86.8	95.3	97.6	79.8	77.4	82.7	93.4	90.8	87.9	9.2
Ours (Learned sim.)	89.1	97.9	97.1	75.8	81.5	85.3	94.4	91.8	**89.1**	**6.2**

TABLE IV

HHAR: CLASSIFICATION ACCURACY AND EXPECTED CALIBRATION ERROR (ECE) ON TARGET DOMAIN USING LEAVE-ONE-DOMAIN-OUT TESTING.

Method	Accuracy (%) ↑													ECE (%) ↓
	A	B	C	D	E	F	G	H	I	J	K	L	Avg	Avg
ERM	86.4	91.6	81.0	91.7	71.3	96.9	96.4	85.9	85.0	88.1	86.6	89.5	87.5	6.0
IRM	87.2	92.0	80.0	90.5	71.7	96.5	96.6	85.3	84.3	88.4	87.3	89.6	87.4	6.2
GroupDRO	80.4	76.7	52.4	74.6	63.6	77.3	76.6	75.0	86.3	86.8	82.8	70.3	75.2	15.2
Interdomain Mixup	80.2	68.9	61.4	69.7	55.3	71.1	81.4	64.7	87.8	85.0	84.9	71.5	73.5	13.5
MLDG	81.7	75.8	58.8	79.3	58.4	70.7	70.9	68.0	86.8	87.7	84.7	70.4	74.4	18.2
MTL	79.6	75.8	60.9	77.1	62.8	75.7	80.2	72.7	85.7	81.1	79.3	71.9	75.2	17.9
Correlation	80.3	91.0	81.7	87.9	69.1	95.4	95.8	88.5	85.1	85.6	88.6	88.9	86.5	5.8
VREx	87.1	90.6	80.5	92.2	71.0	96.5	96.7	85.5	85.5	88.7	87.5	90.2	87.7	5.7
RSC	87.3	90.5	84.4	92.2	73.9	96.7	96.9	86.2	86.8	87.5	88.5	90.1	88.4	5.3
DANN-DG	84.7	89.5	72.4	92.8	71.2	95.1	94.8	84.2	81.6	84.3	84.9	86.7	85.2	7.5
CDANN-DG	85.6	86.0	79.8	89.6	72.4	93.6	95.6	83.0	81.3	87.0	83.4	85.8	85.2	8.2
CORAL-DG	80.5	76.8	58.4	74.3	62.5	77.5	85.8	74.2	86.8	79.7	86.2	69.1	76.0	15.4
MMD-DG	81.9	74.0	52.9	75.4	60.3	76.8	79.0	74.6	86.9	85.6	83.7	68.9	75.0	17.2
Ours (Metadata sim.)	87.4	91.0	80.7	94.6	75.7	96.5	97.1	86.2	85.0	89.2	89.1	89.8	88.5	**4.1**
Ours (Learned sim.)	87.3	90.6	85.3	93.5	76.0	96.1	96.7	86.0	85.1	88.1	88.6	88.9	88.5	4.9
RSC														
+ Ours (Metadata sim.)	87.2	89.9	86.0	93.6	75.6	96.6	96.4	86.1	85.7	88.6	90.4	90.6	**88.9**	4.5
+ Ours (Learned sim.)	86.3	89.0	84.3	93.3	74.4	96.8	96.9	86.8	87.0	86.7	90.4	89.2	88.5	4.8

TABLE V

MIMIC-III: AUC AND EXPECTED CALIBRATION ERROR (ECE).

Method	AUC (%) ↑					ECE (%) ↓
	A	B	C	D	Avg	Avg
ERM	84.9	78.0	77.4	73.1	78.4	7.6
RSC	83.5	78.1	77.5	73.5	78.2	8.5
DANN-DG	84.9	79.5	76.2	74.0	78.7	7.4
CDANN-DG	85.2	79.2	76.1	75.1	78.9	8.2
CORAL-DG	86.3	78.0	77.1	74.7	79.0	7.5
MMD-DG	85.5	77.8	77.0	73.6	78.5	8.2
Ours (Learned sim.)	86.5	78.1	77.5	74.4	**79.1**	**6.4**

TABLE VI

EFFECT OF REGULARIZATION AND AUGMENTATION STRATEGIES.

Strategy		Avg Acc (%)	
Reg	Aug	Bearings	HHAR
✗	✗	82.2	87.5
✗	✓	86.5	88.1
(Metadata sim.)			
✓	✗	87.1	88.5
✓	✓	87.9	88.5
(Learned sim.)			
✓	✗	86.8	88.3
✓	✓	89.1	88.5

(a) Regularization and augmentation

Aug	Avg Acc (%)
None	82.2
Mean shift	83.0
Scale	82.4
Mask	82.4
All	86.5

(b) Bearings: Time series augmentation without consistency regularization

The proposed method has highest AUC (79.1%), and although CORAL-DG achieves similar AUC (79.0%), our proposed method's calibration error is lower by 1.1%.

V. FURTHER ANALYSIS

Ablation study: In Table VI, we see that applying data augmentations and consistency regularization individually improves model performance over ERM for both datasets. Performance increases further when both strategies are applied.

Effect of design choices for regularizer: To further study the regularizer in isolation, we apply the proposed method with metadata based selection and no time series augmentations.

Distance functions: We experiment with 3 distance functions: squared L_2 distance, cosine distance and KL-divergence, either between individual samples and cluster centroids or between domain and cluster centroids. Regularization is applied on the features z, logits g or soft labels s. Comparing the generalization performance in Table VII, domain-level regularization tends to have higher accuracy, possibly because it allows greater diversity of representations in each domain.

Logits vs Soft labels vs Features: Regularizing on logits

TABLE VII
BEARINGS: ACCURACY USING DIFFERENT REGULARIZATION FUNCTIONS.

| | Avg Accuracy (%) | | | | | |
| | Sample-level | | | Domain-level | | |
Regularize on	L2	cos	KL	L2	cos	KL
Features z	83.9	84.9	N/A	85.4	82.3	N/A
Logits g	86.2	84.2	N/A	**87.1**	85.9	N/A
Soft labels s	82.9	83.3	82.8	82.5	83.8	83.8

TABLE VIII
BEARINGS: ACCURACY GIVEN DIFFERENT CLUSTER ASSIGNMENTS.

Cluster assignment	# clusters	Avg Acc (%)
{A},{B},{C},{D},{E},{F},{G},{H}	8 (ERM)	82.2
{A,B},{C,D},{E,F},{G,H}	4	84.4
{A,B,C,D},{E,F,G,H}	2	**87.1**
{A,B,C,D,E,F,G,H}	1	85.6

results in higher accuracy in most cases (Table VII). It allows more flexibility as both feature extractor and classifier are directly regularized while preserving class-relationships. Soft labels are normalized logits and have limited variability across source domains for further alignment. The feature space is generally much larger than the logit or label space, and hence possibly more difficult to effective align. We observe that all choices of distance functions and representations attain better performance than ERM (82.2%), with regularization of logits at the domain-level achieving the best accuracy of 87.1%.

Cluster assignments: Next, we study the effect of cluster assignments on DG performance for each target domain in the Bearings dataset in Table VIII. Consistency regularization tends to improve performance over ERM across all target domains. While regularization with 4 clusters and 1 cluster both improve accuracy over ERM by 2.2% and 3.4% respectively, regularization with 2 clusters achieves the highest improvement of 4.9%. Domains in each of the two clusters have the same machine location, and hence can be expected to be closely-related with similar class relationships. This shows good cluster assignments are critical for good performance.

Training process: Figure 3a shows loss components using regularization with 2 clusters, where the model initially prioritized learning the classification task and then increased domain alignment as training progresses. Figure 3b and 3c show loss and computation time per iteration (total 3000 iterations) for different number of clusters ran on Tesla V100-SXM2 GPU. Hyperparameter is fixed at $\lambda = 0.01$. Training loss converges and computation time is low for all cluster numbers tested.

(a) Loss components   (b) Training loss   (c) Training time

Fig. 3. Bearings: Loss and average computation time with target H. ERM is equivalent to no regularization or assigning each domain to its own cluster.

Analysis of target risk bound: In multi-source DA, the target domain risk is bounded by a sum consisting empirical source domain risks and divergence between source and target domains, as stated in Theorem 1.

Theorem 1: [61] Let $\mathcal{H}$ be a hypothesis class with VC dimension v. For source domain $\{\hat{S}_m\}_{m=1}^M$ with $\frac{N}{M}$ i.i.d. samples from each domain, and target domain $\hat{T}$ with N samples, with probability $1 - \delta$, for all $\gamma \in \mathcal{H}$ and $\forall a \in \mathbb{R}_+^M$ such that $\sum_{m=1}^M a_m = 1$,

$$\epsilon_T(\gamma) \leq \sum_{m=1}^M a_i \left(\hat{\epsilon}_{S_m}(\gamma) + \frac{1}{2} div_{\mathcal{H}}(\hat{T}; \hat{S}_m) \right) + r_a$$
$$+ O\left(\sqrt{\frac{1}{N} \left(\log \frac{1}{\delta} + v \log \frac{N}{v} \right)} \right) \quad (7)$$

where $\epsilon(\gamma)$ and $\hat{\epsilon}(\gamma)$ are risk and empirical risk of γ w.r.t. the true labeling function respectively, r_a is upper bound of $\inf_{\gamma \in \mathcal{H}} \left(\sum_{m=1}^M a_m \epsilon_{S_m}(\gamma) + \epsilon_T(\gamma) \right)$, and $\mathcal{H}$-divergence $div_{\mathcal{H}}(T; S) := 2\sup_{A \in \mathcal{A}_{\mathcal{H}}} \|P_T(A) - P_S(A)\|$ with $A_{\mathcal{H}} := \{\gamma^{-1}(\{1\}) | \gamma \in \mathcal{H}\}$.

In DA, adversarial learning and regularization are used to align source and target domains to directly reduce $\mathcal{H}$-divergence. Unlike DA, target domain $\hat{T}$ is unavailable during training for DG. Aligning source domains helps to prevent learning spurious features, but aligning all source domains as in existing domain alignment methods does not guarantee reducing $\mathcal{H}$-divergence.

We create a two-channel toy dataset for binary classification of sawtooth and square-shaped signals to compare the bound terms and empirical target risk for different methods; we leave out the r_a term since we use the same backbone network for all methods. Domain A and B have shaped signals on channel 1. Domain C and D additionally have square signals on channel 2 for the square signal class. Noise added to signal is Gaussian distributed with standard deviation 2 for A and B and 4 for C and D. We use D as target and the rest as source.

We implement our proposed method with regularization on domain A and B since C differs by having additional channel 2 features. In Table IX, lower $\mathcal{H}$-divergence indicates lower target risk. Our proposed method has lower $\mathcal{H}$-divergence and target risk than DANN-DG and CDANN-DG. By not aligning domain C, we also allow the model to retain channel 2 features which target domain D can use for classification. More details are in Appendix B.

TABLE IX
TOY DATASET: COMPARISON OF BOUND TERMS AND TARGET RISK.

Method	Source Risk	$\mathcal{H}$-divergence	Target Risk
DANN-DG	0.196	1.597	0.399
CDANN-DG	0.198	1.591	0.333
Ours	0.196	1.536	0.111

VI. CONCLUSION

In this work, we introduced a representation learning method for domain generalization for time series classification. We applied augmentations to improve sample diversity, and selective consistency regularization to enforce similar predictions for similar domains. From comprehensive experiments, we showed that the proposed method significantly improves over ERM and performs competitively compared to state-of-the-art methods in both classification accuracy and model calibration.

REFERENCES

[1] I. Gulrajani and D. Lopez-Paz, "In Search of Lost Domain Generalization," in *ICLR*, 2021.
[2] D. Hendrycks, S. Basart, N. Mu, S. Kadavath, F. Wang, E. Dorundo, R. Desai, T. L. Zhu, S. Parajuli, M. Guo, D. Song, J. Steinhardt, and J. Gilmer, "The many faces of robustness: A critical analysis of out-of-distribution generalization," in *ArXiv*, 2020.
[3] D. Li, Y. Yang, Y.-Z. Song, and T. M. Hospedales, "Learning to generalize: Meta-learning for domain generalization," in *AAAI*, 2018.
[4] W. Zhang, S. Seto, and D. K. Jha, "Cazsl: Zero-shot regression for pushing models by generalizing through context," in *IROS*, 2020.
[5] A. Gupta, A. Murali, D. Gandhi, and L. Pinto, "Robot learning in homes: Improving generalization and reducing dataset bias," in *NIPS*, 2018.
[6] D. Li, J. Zhang, Y. Yang, C. Liu, Y.-Z. Song, and T. M. Hospedales, "Episodic training for domain generalization," in *ICCV*, 2019.
[7] M. Mancini, S. R. Bulò, B. Caputo, and E. Ricci, "Best sources forward: Domain generalization through source-specific nets," in *ICIP*, 2018.
[8] N. Somavarapu, C.-Y. Ma, and Z. Kira, "Frustratingly simple domain generalization via image stylization," in *ArXiv*, 2020.
[9] F. Carlucci, A. D'Innocente, S. Bucci, B. Caputo, and T. Tommasi, "Domain generalization by solving jigsaw puzzles," in *CVPR*, 2019.
[10] Y. Ganin, E. Ustinova, H. Ajakan, P. Germain, H. Larochelle, F. Laviolette, M. Marchand, and V. Lempitsky, "Domain-adversarial training of neural networks," in *Journal of Machine Learning Research*, vol. 17, 2016, pp. 59:1–59:35.
[11] J. Guo, D. Shah, and R. Barzilay, "Multi-source domain adaptation with mixture of experts," in *EMNLP*, 2018.
[12] Y. Balaji, S. Sankaranarayanan, and R. Chellappa, "MetaReg: Towards domain generalization using meta-regularization," in *NeurIPS*, 2018.
[13] G. Wilson, J. R. Doppa, and D. J. Cook, "Multi-source deep domain adaptation with weak supervision for time-series sensor data," in *SIGKDD*, 2020.
[14] H. Zheng, R. Wang, Y. Yang, Y. Li, and M. Xu, "Intelligent fault identification based on multisource domain generalization towards actual diagnosis scenario," in *IEEE Transactions on Industrial Electronics*, vol. 67, no. 2, 2020, pp. 1293–1304.
[15] S. Motiian, M. Piccirilli, D. A. Adjeroh, and G. Doretto, "Unified deep supervised domain adaptation and generalization," in *ICCV*, 2017.
[16] T. M. and T. H., "Domain generalization using a mixture of multiple latent domains," in *AAAI*, 2020.
[17] H. Li, S. J. Pan, S. Wang, and A. C. Kot, "Domain generalization with adversarial feature learning," in *CVPR*, 2018.
[18] M. Long, Z. Cao, J. Wang, and M. I. Jordan, "Conditional adversarial domain adaptation," in *NeurIPS*, 2018.
[19] J. Wang, C. Lan, C. Liu, Y. Ouyang, and T. Qin, "Generalizing to unseen domains: A survey on domain generalization," in *IJCAI*, 2021.
[20] Z. Wang, Z. Dai, B. Póczos, and J. Carbonell, "Characterizing and avoiding negative transfer," in *CVPR*, 2019.
[21] K. Bousmalis, G. Trigeorgis, N. Silberman, D. Krishnan, and D. Erhan, "Domain separation networks," in *NIPS*, 2016.
[22] B. Sun and K. Saenko, "Deep coral: Correlation alignment for deep domain adaptation," in *ECCV Workshops*, 2016.
[23] S. Ben-David, J. Blitzer, K. Crammer, A. Kulesza, F. Pereira, and J. Vaughan, "A theory of learning from different domains," in *Machine Learning*, vol. 79, no. 1, 2010, pp. 151–175.
[24] Q. Dou, D. Coelho de Castro, K. Kamnitsas, and B. Glocker, "Domain generalization via model-agnostic learning of semantic features," in *NeurIPS*, 2019.
[25] M. Wang and W. Deng, "Deep visual domain adaptation: A survey," in *Neurocomputing*, vol. 312, 2018, pp. 135–153.
[26] Z. Wang, M. Loog, and J. Gemert, "Respecting domain relations: Hypothesis invariance for domain generalization," *ICPR*, 2020.
[27] F. Ahmed, Y. Bengio, H. v. Seijen, and A. C. Courville, "Systematic generalisation with group invariant predictions," in *ICLR*, 2021.
[28] D. Kim, S. Park, J. Kim, and J. Lee, "Selfreg: Self-supervised contrastive regularization for domain generalization," in *ArXiv*, 2021.
[29] O. Vinyals, C. Blundell, T. Lillicrap, D. Wierstra *et al.*, "Matching networks for one shot learning," in *NeurIPS*, 2016.
[30] A. D'Innocente and B. Caputo, "Domain generalization with domain-specific aggregation modules," in *Pattern Recognition*, 2019, pp. 187–198.
[31] R. Volpi, H. Namkoong, O. Sener, J. Duchi, V. Murino, and S. Savarese, "Generalizing to unseen domains via adversarial data augmentation," in *NeurIPS*, 2018.
[32] S. Shankar, V. Piratla, S. Chakrabarti, S. Chaudhuri, P. Jyothi, and S. Sarawagi, "Generalizing across domains via cross-gradient training," in *ICLR*, 2018.
[33] D. Stutz, M. Hein, and B. Schiele, "Disentangling adversarial robustness and generalization," in *CVPR*, 2019.
[34] S. Yan, H. Song, N. Li, L. Zou, and L. Ren, "Improve unsupervised domain adaptation with mixup training," in *ArXiv*, 2020.
[35] Y. Wang, H. Li, and A. Kot, "Heterogeneous domain generalization via domain mixup," in *ICASSP*, 2020.
[36] Y. Jing, Y. Yang, Z. Feng, J. Ye, Y. Yu, and M. Song, "Neural style transfer: A review," in *IEEE Transactions on Visualization and Computer Graphics*, vol. 26, no. 11, 2020, pp. 3365–3385.
[37] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *NIPS*, 2014.
[38] K. Zhou, Y. Yang, T. Hospedales, and T. Xiang, "Learning to generate novel domains for domain generalization," in *ECCV*, 2020.
[39] H. Nam, H. Lee, J. Park, W. Yoon, and D. Yoo, "Reducing domain gap via style-agnostic networks," in *ArXiv*, 2019.
[40] R. Gong, W. Li, Y. Chen, and L. Van Gool, "Dlow: Domain flow for adaptation and generalization," in *CVPR*, 2019.
[41] M. Arjovsky, L. Bottou, I. Gulrajani, and D. Lopez-Paz, "Invariant risk minimization," in *ArXiv*, 2019.
[42] E. Rosenfeld, P. Ravikumar, and A. Risteski, "The risks of invariant risk minimization," in *ArXiv*, 2020.
[43] I. Albuquerque, J. Monteiro, M. Darvishi, T. H. Falk, and I. Miltliagkas, "Generalizing to unseen domains via distribution matching," in *ArXiv*, 2020.
[44] D. Mahajan, S. Tople, and A. Sharma, "Domain generalization using causal matching," in *Uncertainty and Robustness in Deep Learning (ICML Workshop)*, 2020.
[45] D. Arpit, C. Xiong, and R. Socher, "Predicting with high correlation features," *ArXiv*, 2019.
[46] Z. Huang, H. Wang, E. Xing, and D. Huang, "Self-challenging improves cross-domain generalization," in *ECCV*, 2020.
[47] S. Sagawa, P. W. Koh, T. Hashimoto, and P. Liang, "Distributionally robust neural networks for group shifts: On the importance of regularization for worst-case generalization," in *ArXiv*, 2019.
[48] D. Krueger, E. Caballero, J. Jacobsen, A. Zhang, J. Binas, R. L. Priol, and A. C. Courville, "Out-of-distribution generalization via risk extrapolation (REx)," in *ArXiv*, 2020.
[49] A. Gretton, K. M. Borgwardt, M. J. Rasch, B. Schölkopf, and A. Smola, "A kernel two-sample test," *Journal of Machine Learning Research*, vol. 13, no. 25, pp. 723–773, 2012.
[50] G. Blanchard, A. A. Deshmukh, U. Dogan, G. Lee, and C. Scott, "Domain generalization by marginal transfer learning," *Journal of Machine Learning Research*, vol. 22, no. 2, pp. 1–55, 2021.
[51] Y. Li, X. Tian, M. Gong, Y. Liu, T. Liu, K. Zhang, and D. Tao, "Deep domain generalization via conditional invariant adversarial networks," in *ECCV*, 2018.
[52] C. Guo, G. Pleiss, Y. Sun, and K. Q. Weinberger, "On calibration of modern neural networks," in *ICML*, 2017.
[53] W. A. Smith and R. B. Randall, "Rolling element bearing diagnostics using the case western reserve university data: A benchmark study," in *Mechanical Systems and Signal Processing*, vol. 64. Elsevier, 2015, pp. 100–131.
[54] W. Zhang, G. Peng, C. Li, Y. Chen, and Z. Zhang, "A new deep learning model for fault diagnosis with good anti-noise and domain adaptation ability on raw vibration signals," in *Sensors*, vol. 17, no. 2. Multidisciplinary Digital Publishing Institute, 2017, p. 425.
[55] A. Stisen, H. Blunck, S. Bhattacharya, T. S. Prentow, M. B. Kjærgaard, A. Dey, T. Sonne, and M. M. Jensen, "Smart devices are different: Assessing and mitigating mobile sensing heterogeneities for activity recognition," in *Proceedings of the 13th ACM conference on embedded networked sensor systems*, 2015, pp. 127–140.

[56] S. Seto, W. Zhang, and Y. Zhou, "Multivariate time series classification using dynamic time warping template selection for human activity recognition," *IEEE Symposium Series on Computational Intelligence*, pp. 1399–1406, 2015.

[57] W. Zhang, M. Ragab, and R. Sagarna, "Robust domain-free domain generalization with class-aware alignment," *ICASSP*, 2021.

[58] A. Johnson, T. Pollard, and M. Roger, "MIMIC-III clinical database," PhysioNet, 2016.

[59] A. Johnson, T. Pollard, L. Shen, L.-W. Lehman, M. Feng, M. Ghassemi, B. Moody, P. Szolovits, L. Celi, and R. Mark, "MIMIC-III, a freely accessible critical care database," *Scientific Data*, vol. 3, 2016.

[60] R. Cai, J. Chen, Z. Li, W. Chen, K. Zhang, J. Ye, Z. Li, X. Yang, and Z. Zhang, "Time series domain adaptation via sparse associative structure alignment," in *AAAI*, 2021.

[61] H. Zhao, S. Zhang, G. Wu, J. M. F. Moura, J. P. Costeira, and G. J. Gordon, "Adversarial multiple source domain adaptation," in *NeurIPS*, 2018.

[62] Y. Liu, Y.-M. Zhang, X.-Y. Zhang, and C.-L. Liu, "Adaptive spatial pooling for image classification," in *Pattern Recognition*, vol. 55. Elsevier, 2016, pp. 58–67.

[63] I. Gulrajani and D. Lopez-Paz, "Domainbed github repository," https://github.com/facebookresearch/DomainBed, accessed: 2020-10-26.

A. Implementation Details

We provide additional details on the implementation for our experiments and the datasets used.

We evaluate the proposed method on three real-world time series datasets for fault detection, human activity recognition and mortality prediction following the evaluation procedure in [1]. For each dataset, we use leave-one-domain-out evaluation where we treat each domain as the unseen target in turn and train with rest as source domains. Each source-target combination is evaluated over 3 random seeds, and each method is tuned with 20 random hyperparameter configurations with the best configuration selected using validation accuracy. All methods use the same backbone networks specific to each dataset.

1) Hyperparameters: We fix learning rate 0.001, weight decay 5×10^{-5}, and batch size 32 per domain. Models are trained for 3000 iterations, with learning rate reduced by a factor of 10 after 2400 iterations. All other hyperparameters are tuned by random sampling from distributions in Table XVI. All experiments are run using Adam optimizer with the NVIDIA container image for PyTorch, release 20.03.

2) Datasets and Network Architectures: We provide details on the sample size of the datasets. Backbone network architectures used for each dataset is given in Table XVII.

Bearings: All domains have the same number of samples. For each domain, the sample size of each class is 'normal': 416, 'IF:0.007': 371, 'BF:0.007': 409, 'OF:0.007': 417, 'IF:0.014': 387, 'BF:0.014': 408, 'OF:0.014': 398, 'IF:0.021': 407, 'BF:0.021': 383, 'OF:0.021': 404. We use a 6-layer CNN as feature extractor and a 3-layer FCN as classifier [57].

HHAR: Sample size differs across domain according to availability of data per user and device, as in Table XVIII. We use a 3-layer CNN as feature extractor and a 1-layer fully-connected network as classifier [62].

MIMIC-III: All classes are sampled to be of the same size. The sample size of each class for the 4 domains are A: 164, B: 625, C: 1179, D: 495. We use the same network architecture as for HHAR but with smaller hidden dimension size.

B. Further Experiment Results

Visualization of learned domain relationships on Bearings and HHAR: The proposed method estimates inter-domain relationships when domain metadata are not provided. Figure 4 plots the proportion of runs each pair of domain is estimated to be closest neighbors. For Bearings, the estimated clusters approximately match location groupings. For HHAR, the variation between phone models (i.e. Nexus and S3 versus S3 mini and S+) appears larger than that between some users (i.e. User 1 versus User 3). Domain relationships estimated from data can be different from those inferred from metadata descriptions and can contain finer measures of inter-domain similarity. In fact, using learned similarity obtains higher accuracy (89.1%) than using fixed metadata-inferred similarity (87.9%) for Bearings.


(a) Bearings   (b) HHAR

Fig. 4. Fraction of runs where domain j (column) is the nearest neighbor of domain i (row) at end of training. Diagonal values are 1.

Additional experiments on Bearings and HHAR: We additionally evaluate on a more challenging setting where target domain conditions are not combinations of source domain conditions. For Bearings, we use only domains from drive-end location, so each domain has a distinct loading torque. For HHAR, we use only domains from the first user, so each domain has a distinct phone model. We fix the hyperparameter of our proposed method with learned inter-domain similarity to $\lambda = 0.01$ and $\xi = 0.1$. From Table X and XI, our proposed method outperforms ERM on all target domains in Bearings, and on average in HHAR.

TABLE X
BEARINGS DRIVE-END DOMAINS: CLASSIFICATION ACCURACY AND EXPECTED CALIBRATION ERROR (ECE) ON TARGET DOMAIN.

Method	Accuracy (%) ↑					ECE (%) ↓
	A	B	C	D	Avg	Avg
ERM	84.4	94.9	98.9	86.0	91.0	5.2
Ours (Learned sim.)	**93.9**	**97.8**	**99.0**	**88.7**	**94.8**	**3.1**

TABLE XI
HHAR USER 1: CLASSIFICATION ACCURACY AND EXPECTED CALIBRATION ERROR (ECE) ON TARGET DOMAIN.

Method	Accuracy (%) ↑					ECE (%) ↓
	A	B	C	D	Avg	Avg
ERM	**79.0**	80.5	**71.5**	80.3	77.8	14.6
Ours (Learned sim.)	78.7	**81.8**	70.8	**82.0**	**78.3**	**13.4**

TABLE XII
BEARINGS: REGULARIZATION STRATEGIES FOR REGULARIZATION WITH LEARNED INTER-DOMAIN SIMILARITY.

Selected neighbor	Weight function	Avg Accuracy (%)
Random	1	80.9
Random	RBF ($\xi = 0.1$)	81.2
Nearest	1	**84.9**
Nearest	RBF ($\xi = 0.1$)	<u>83.2</u>

Ablation study: We study the effect of neighbor domain selection and weight function in regularization with learned inter-domain similarity on Bearings. Hyperparameters are fixed at $\lambda = 0.01$ and $\xi = 0.1$. In Table XII, regularizing each domain with its nearest neighbor achieves higher accuracy than regularizing random pairs of domains. While nearest neighbor selection and a fixed weight of 1 has best accuracy in this
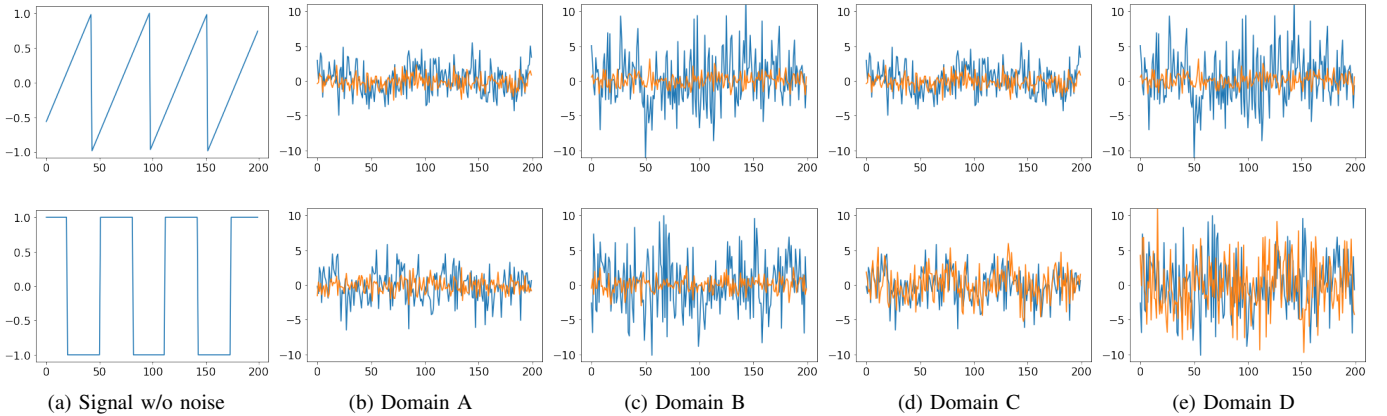
Fig. 5. Toy dataset: Example series. (a) Clean sawtooth and square signals used for data generation, (b)-(e) Samples generated for domain A to D, with channel 1 in blue and channel 2 in orange. Best viewed in color.

TABLE XIII
CHOICE OF MMD KERNEL FUNCTION.

Kernel	Avg Acc (%)	
	Bearings	HHAR
(Metadata sim.)		
Gaussian	83.1	87.6
Linear	**87.1**	**88.5**
(Learned sim.)		
Gaussian	86.5	87.6
Linear	**86.8**	**88.3**

TABLE XIV
TOY DATASET: DOMAIN ATTRIBUTES. DOMAIN A, B AND C ARE SOURCE, DOMAIN D IS TARGET.

Channel with signal	Noise std. dev.	
	2	4
Channel 1	A	B
Channel 1 and Channel 2	C	D

TABLE XV
TOY DATASET: COMPARISON OF BOUND TERMS AND TARGET RISK.

Method	Source Risk				$\mathcal{H}$-divergence				Target Risk
	A	B	C	Avg	A	B	C	Avg	D
DANN-DG	0.207	0.375	0.005	0.196	1.992	0.800	2	1.597	0.399
CDANN-DG	0.242	0.352	0	0.198	1.972	0.800	2	1.591	0.333
Ours	0.157	0.362	0.070	0.196	2	0.608	2	1.536	0.111

study, we note that the RBF hyperparameter ξ can be tuned and the RBF weight approaches 1 as $\xi \to 0$.

Lastly, we study the use of different kernel functions in MMD for distribution matching in our proposed method. The linear kernel is the default used in all our experiments, and the Gaussian kernel is implemented as in [63]. We do not apply time series augmentations in order to study the effect of regularization in isolation. From Table XIII, the linear kernel has better performance than the Gaussian kernel for both Bearings and HHAR datasets.

Analysis of target risk bound: We create a toy dataset

according to domain attributes in Table XIV for binary classification of sawtooth and square-shaped signals to compare the target risk bound terms in main manuscript Theorem 1 and empirical target risk. Each sample is a length-200 time series with two channels. Domain A and B only have sawtooth and square signals on channel 1. Domain C and D additionally have square signals on channel 2 for the square signal class. Signals are generated with number of cycles uniformly at random in $[\frac{1}{6}, 5]$, and noise added to the signals is Gaussian distributed with standard deviation 2 for domain A and B and 4 for domain C and D. We add Gaussian $N(0, 1)$ noise to channels with no shape signals. We illustrate example series from each domain in Figure 5. We use domain D as unseen target domain and the rest as source domains. Each source domain has 400 samples and target domain has 1200 samples following sample size specifications in main manuscript Equation 1.

We provide estimates of the target risk bound terms in Table XV on models trained using existing domain alignment methods (i.e. DANN-DG, CDANN-DG) and our proposed method with consistency regularization applied on domain A and B. We do not regularize domain C in our method because it is the most dissimilar from the other two source domains by having additional channel 2 features that are correlated with signal class. All regularization hyperparameters are set to 1 i.e. $\lambda = 1$, and we use the same network architecture as for HHAR with 1-layer CNN. We do not apply time series augmentations in order to study the effect of regularization in isolation. To control for possible differences due to model convergence, we stop training for all methods when average source risk falls below 0.2. We follow the procedure in [10] to estimate $\mathcal{H}$-divergence. For each source S_m and target T domain pair:

1) Construct domain classification training set with 200 samples from source domain and 200 samples randomly selected from target domain, and construct domain classification test set with the other 200 samples from source domain and 200 samples randomly selected from the remainder of target domain;

TABLE XVI
SETUP FOR HYPERPARAMETER TUNING.

Method	Hyperparameter	Distribution
IRM	Regularization λ	$10^{Unif(-1,5)}$
	Iterations of penalty annealing	$\lfloor 10^{Unif(0,4)} \rfloor$
GroupDRO	Group weight temperature η	$10^{Unif(-3,-1)}$
Interdomain Mixup	Beta shape parameter α	$10^{Unif(-1,1)}$
MTL	Embedding averaging proportion	$\{0.5, 0.9, 0.99, 1\}$
MLDG	Meta-learning loss β	$10^{Unif(-1,1)}$
Correlation	Regularization λ	10^{-5}
CORAL-DG, MMD-DG	Regularization λ	$10^{Unif(-3,-1)}$
DANN-DG, CDANN-DG	Discriminator learning rate	$10^{Unif(-5,-3.5)}$
	Discriminator weight decay	$10^{Unif(-6,-2)}$
	Discriminator Adam β_1	$\{0, 0.5\}$
	Discriminator steps	$\lfloor 2^{Unif(0,3)} \rfloor$
	Discriminator gradient penalty	$10^{Unif(-2,1)}$
	Adversarial regularization λ	$10^{Unif(-2,2)}$
VREx	Regularization λ	$10^{Unif(-1,5)}$
	Iterations of penalty annealing	$\lfloor 10^{Unif(0,4)} \rfloor$
RSC	Feature drop percentage p	$Unif(0, 0.5)$
	Batch percentage	$Unif(0, 0.5)$
Ours	Regularization λ	$10^{Unif(-3,-1)}$
	RBF kernel parameter ξ	$10^{Unif(-2,2)}$

2) Train SVMs with RBF kernel for binary domain classification with regularization hyperparameter $C \in \{0.001, 0.01, 0.1, 1, 10, 10, 100\}$;

3) Select C with lowest training error, and evaluate selected SVM on test set to obtain test error ρ;

4) Approximate the $\mathcal{H}$-divergence by $div_{\mathcal{H}}(\hat{S}_m, \hat{T}) = 2(1 - 2\rho)$.

We observe from Table XV that a lower $\mathcal{H}$-divergence tend to indicate a lower target risk. Our proposed method has the lowest $\mathcal{H}$-divergence, especially between domain B and D. Source domain B is aligned with only one domain in our method instead of two in other methods, and it is the most similar to target domain D in terms of signal noise level. By not regularizing source domain C, we also allow the model to retain channel 2 signal features for classification on target domain D.

TABLE XVII
BACKBONE NETWORK ARCHITECTURES FOR EACH DATASET. CONVOLUTION OPERATION IS ABBREVIATED AS 'CONV' AND FULLY CONNECTED OPERATION IS ABBREVIATED AS 'FC'.

Layer	Operation	Specifications
Convolution	Conv BatchNorm LeakyReLU	8 (filter: 64×1, stride: 2, pad: 1)
Convolution (3 times)	Conv BatchNorm LeakyReLU	8 (filter: 3×8, stride: 2, pad: 1)
Convolution	Conv LeakyReLU	8 (filter: 3×8, stride: 2, pad: 1)
Convolution	Conv	8 (filter: 8×8, stride: 1, pad: 1)
Fully connected	FC	32
Fully connected (2 times)	FC ReLU	32
Fully connected	FC	10

(a) Network for Bearings

Layer	Operation	Specifications
Convolution	Conv BatchNorm LeakyReLU	hidden_dim (filter: 8×3, stride: 1, pad: 1)
Convolution	Conv BatchNorm LeakyReLU	2*hidden_dim (filter: $5 \times$ hidden_dim, stride: 1, pad: 1)
Convolution	Conv BatchNorm LeakyReLU	hidden_dim (filter: $3 \times$ (2*hidden_dim), stride: 1, pad: 1)
Pooling	Average pooling	1 (filter: 121, stride:121)
Fully connected	FC	6

(b) Network for HHAR (hidden_dim = 128) and MIMIC-III (hidden_dim = 32)

TABLE XVIII
HHAR: SAMPLE SIZE DISTRIBUTION PER DOMAIN.

Domain	Class					
	Biking	Standing	Sitting	Walking	Stair down	Stair up
A	626	933	652	676	874	778
B	346	468	316	341	435	376
C	175	234	162	176	207	212
D	298	237	264	226	223	275
E	999	682	681	771	692	1013
F	487	387	312	407	370	495
G	234	196	161	199	186	245
H	385	251	267	298	253	331
I	539	817	628	768	723	857
J	293	427	312	374	358	445
K	147	213	168	211	164	244
L	275	229	264	265	248	300