

Generalizing across Temporal Domains with Koopman Operators

Qiu hao Zeng¹, Wei Wang¹, Fan Zhou², Gezheng Xu¹, Ruizhi Pu¹, Changjian Shui³,
Christian Gagné⁴, Shichun Yang², Charles X. Ling¹, Boyu Wang^{1*}

¹University of Western Ontario

²Beihang University

³Vector Institute

⁴Université Laval

Abstract

In the field of domain generalization, the task of constructing a predictive model capable of generalizing to a target domain without access to target data remains challenging. This problem becomes further complicated when considering evolving dynamics between domains. While various approaches have been proposed to address this issue, a comprehensive understanding of the underlying generalization theory is still lacking. In this study, we contribute novel theoretic results that aligning conditional distribution leads to the reduction of generalization bounds. Our analysis serves as a key motivation for solving the Temporal Domain Generalization (TDG) problem through the application of Koopman Neural Operators, resulting in Temporal Koopman Networks (TKNets). By employing Koopman Operators, we effectively address the time-evolving distributions encountered in TDG using the principles of Koopman theory, where measurement functions are sought to establish linear transition relations between evolving domains. Through empirical evaluations conducted on synthetic and real-world datasets, we validate the effectiveness of our proposed approach.

1 Introduction

Modern machine learning techniques have achieved unprecedented success over the past decades in numerous areas. However, one fundamental limitation of most existing techniques is that a model trained on one dataset cannot generalize well on another dataset if it is sampled from a different distribution. Domain generalization (DG) aims to alleviate the prediction gap between the observed source domains and an *unseen* target domain by leveraging the knowledge extracted from multiple source domains (Blanchard, Lee, and Scott 2011; Muandet, Balduzzi, and Schölkopf 2013; Arjovsky et al. 2019; Li et al. 2018a).

Existing DG methods can be roughly categorized into three groups: data augmentation / generation, disentangled / domain-invariant feature learning, and meta-learning (Wang et al. 2021). In many real-world applications, the temporal dynamics across domains are common and can be leveraged to improve accuracy for the unseen target domain (Kumar, Ma, and Liang 2020; Liu et al. 2020; Wang, He, and

Katabi 2020). However, one intrinsic problem with these existing DG methods is that most of them treat all the domains equally and ignore the relationship between them, implicitly assuming that they are all sampled from a stationary environment. For example, it is common that source domains are constituted of images collected over the last few years and the target domain is the unseen future. For geological applications, the source samples can be collected along different altitudes, longitude, and latitude, while the target is to generalize to some regions where the data is absent due to inaccessibility. Medical data is also often collected with age or other indicators as intervals, and we hope the model can perform well on younger or elder age groups where the samples may be rare. As a more concrete example, Fig. 1(a) shows several instances from the rotated MNIST (RMNIST) dataset, a widely used benchmark in the DG literature, where the digit images of each subsequent domain are rotated by 15° . Fig. 1(b) reports the generalization performances of several state-of-the-art DG algorithms on the data set, from which it can be clearly observed that the performances drop when deploying the models on outer domains (i.e., domains of 0 and 75 degrees). The results indicate that the algorithms ignore the evolving pattern between the domains. Consequently, they are good at “interpolation” but not at “extrapolation”.

In this paper, we address this learning scenario as *temporal domain generalization* (TDG) (Zeng et al. 2023; Bai, Ling, and Zhao 2022; Nasery et al. 2021; Qin, Wang, and Li 2022), which aims to capture and exploit the temporal dynamics in the environment. TDG aims to generalize to a target domain along a specific direction by extracting and leveraging the relations between source domains. Specifically, we develop a novel theoretical analysis that highlights the importance of modeling the relation between two consecutive domains to extract the evolving pattern of the environment. Koopman theory (Koopman 1931) states that any complex dynamics can be modeled by a linear Koopman operator acting on the space of measurement functions. Inspired by our theoretical results, we propose to capture the temporal dynamics using the Koopman operator and align the distribution of data in the Koopman space. As a comparison, Fig. 1(c) shows the performance improvement of TKNets over the other algorithms on RMNIST dataset. It can be observed that the performance gap between TKNets

*Corresponding authors: Fan Zhou, Charles X. Ling, Boyu Wang.

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: (a) Evolving manner among RMNIST domains. (b) Accuracy of traditional DG methods on evolving domains. These methods cannot generalize well on outer domains (0° and 75°). (c) Comparison between the performance of our method and baselines on outer domains. The proposed method outperforms all the baselines.

and the other baselines has widened as the domain distance increases. More details can be found in Sec. 5.

Here, we emphasize the key difference between Temporal Domain Generalization (Bai, Ling, and Zhao 2022; Nasery et al. 2021; Qin, Wang, and Li 2022) and Temporal Domain Adaptation (Kim et al. 2020; Liu et al. 2020; Wang, He, and Katabi 2020). While both learning paradigms aim to tackle the issue of evolving domain shifts, the latter still requires unlabeled instances from the target domain. In this sense, TDG is more challenging, and existing theoretical and algorithmic results cannot be applied to this problem directly.

We summarize our **contributions** as three-fold: (1) We derive theoretical analysis to the generalization bound of the TDG problem, which highlights the importance of learning the transition function to mitigate the temporal domain shifts; (2) Motivated by the theoretic results, we propose a novel algorithm TKNeTs, which learns the complex and non-linear dynamics based on Koopman theory; (3) We conduct experiments on synthetic and real-world datasets, and the empirical results suggest an improved performance.

2 Related Work

Domain Generalization (DG). Domain generalization aims to train a model which generalizes on all domains. Existing DG methods can be classified into three categories. The first and most popular category is representation learning, which focuses on learning a common representation across domains. It can be achieved by domain-invariant representation learning (Blanchard, Lee, and Scott 2011; Ghifary et al. 2015; Ganin and Lempitsky 2015; Arjovsky et al. 2019) and feature disentanglement (Xu et al. 2014; Ilse et al. 2020). The former focuses on aligning latent features across domains, and the latter tries to distill domain-shared features. Secondly, data manipulation can also empower the model with generalization capability. Data manipulating techniques include data augmentation (Yue et al. 2019; Shankar et al. 2018), which usually extends the dataset by applying specific transformations on existing samples, and data generation (Rahman et al.

2019), which often applies neural networks to generate new samples. (Nguyen et al. 2021) converts DG to an infinite-dimensional constrained statistical learning problem under a natural model of data generation. The theoretically grounded method proposed in (Nguyen et al. 2021) leverages generating model among domains to learn domain-invariant representation. The third most commonly used method is meta-learning. The meta-learning framework (Li et al. 2018a; Balaji, Sankaranarayanan, and Chellappa 2018; Li et al. 2019) is used to improve the generalizing capability by simulating the shift among domains. Apart from the above three categories, (Mancini et al. 2018) tries to ensemble multiple models into a unified one that can generalize across domains. Besides, gradient operation (Huang et al. 2020a), self-supervision (Carlucci et al. 2019) and random forest (Ryu et al. 2019) are also exploited to improve generalizing capability. Different from the existing DG methods that focus on learning one unified model for all domains, our approach tries to train a prediction model for the target domain by leveraging the evolving pattern among domains.

Temporal Domain Adaptation (TDA) / Temporal Domain Generalization (TDG). Many previous works in the domain adaptation area notice the evolving pattern of domains and leverage it to improve performance in different settings. (Liu et al. 2020) proposes a meta-adaptation framework that enables the learner to adapt from one single source domain to continually evolving target domains without forgetting. (Kumar, Ma, and Liang 2020) focuses on adapting from source domains to the target domain with large shifts by leveraging the unlabeled intermediate samples. (Wang, He, and Katabi 2020) combines the traditional adversarial adaptation strategy with a novel regression discriminator that models the encoding-conditioned domain index distribution. (Chen and Chao 2021) investigate how to discover the sequence of intermediate domains without index information and then adapt to the final target. TDG recently has been actively studied. (Qin, Wang, and Li 2022) introduce a probabilistic model to mitigate the covariate shift and concept shift. (Nasery et al. 2021) design gradient interpola-

tion (GI) loss to penalize the curvature of the learned model along time and design TReLU activation whose parameters are tuned w.r.t. time. (Bai, Ling, and Zhao 2022) proposed a novel Bayesian framework to explicitly model the concept drift over time. (Zeng et al. 2023) proposed to learn the evolving patterns in the framework of meta-learning. However, existing TDG methods have not explored the theoretic generalization bound of TDG problem. In this paper, our theoretical results are based on proposed λ -consistency, an intuitive and realistic measurement of evolving levels in the environments.

3 Background

3.1 Koopman Theory

In discrete-time dynamical systems, the formulation of state transitions is as $v_{t+1} = F(v_t)$, where $v \in \mathcal{V} \subseteq \mathbb{R}^d$ represents the system state and F characterizes the vector field governing the system dynamics. However, it's difficult to directly capture the system dynamics due to the presence of nonlinearity or noisy data. To address this concern, Koopman theory (Koopman 1931) hypothesizes that the system state can be effectively projected onto an infinite-dimensional Hilbert space defined by measurement function $\mathcal{G}(\mathcal{V}) := \{g : \mathcal{G} \rightarrow \mathbb{R}\}$. This projected space can then be advanced forward in time through an infinite-dimensional linear operator \mathcal{K} , hence

$$\mathcal{K} \circ g(v_t) = g(F(v_t)) = g(v_{t+1}). \quad (1)$$

The Koopman operator facilitates the transition of observations of the state to subsequent time steps by mapping between function spaces (Li et al. 2020). In this work, we employ measurement functions to effectively map encoded features to the functional space. By adopting Koopman theory, we can learn the transition model, which characterizes the temporal dynamics governing the system's behavior.

3.2 Temporal Domain Generalization

We aim to propose a formal understanding of the generalization bound in predicting the unseen target environment. Let $\{\mathcal{D}_1(z, y), \mathcal{D}_2(z, y), \dots, \mathcal{D}_m(z, y)\}$ be m observed source domains sampled from an environment \mathcal{E} , where $z \in \mathcal{Z}$ and $y \in \mathcal{Y}$ are, respectively, the data Koopman embedding and its corresponding label, and $\mathcal{D}_i(z, y)$ characterizes the joint probability over the i -th domain. The embeddings z are mapped by a composite embedding function such that $\mathcal{G} \circ \phi(\mathcal{X}) \rightarrow \mathcal{Z}$, where $x \in \mathcal{X}$ is the input raw data, ϕ is a normal embedding neural network, and \mathcal{G} is the measurement function. The goal of TDG is to learn a hypothesis $h \in \mathcal{H}$ so that it can have a low risk on an unseen but time-evolving target domain \mathcal{D}_t :

$$R_{\mathcal{D}_t}(h) \triangleq \mathbb{E}_{(z,y) \sim \mathcal{D}_t}[\ell(h(z), y)]$$

where $\ell : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}_+$ is a non-negative loss function, and \mathcal{H} is a hypothesis class that maps \mathcal{Z} to the set \mathcal{Y} . The objective of TDG is to generalize the model on \mathcal{D}_t along a specific direction when there is an underlying evolving pattern between the source domains and $\mathcal{D}_t = \mathcal{D}_{m+1}$.

4 Theoretical Analysis and Methodology

4.1 Theoretical Motivations

Existing theoretical studies of temporal domain (Wang, He, and Katabi 2020; Kumar, Ma, and Liang 2020; Liu et al. 2020; Wang, Li, and Zhao 2022) make assumptions that "consecutive domains are similar", where the similarity is characterized by some assumption on a distributional distance $d(\mathcal{D}_t, \mathcal{D}_{t+1})$. For instance, (Kumar, Ma, and Liang 2020; Qin, Wang, and Li 2022) assume that $d(\mathcal{D}_t, \mathcal{D}_{t+1}) < \epsilon$, and the assumption in (Liu et al. 2020) is $d_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{D}_{t_1}, \mathcal{D}_{t_2}) \leq \alpha|t_1 - t_2|$. While such assumptions and the corresponding theoretical frameworks are intuitive, the similarity assumption does not properly characterize the temporal dynamics. Instead, we argue that we should focus on *consistency* of the environment instead of *similarity* of consecutive domains. With the Koopman theory applied, we ensure that consistency holds between the temporal domains (Azencot et al. 2020).

To capture the consistent evolving pattern in the Koopman space of \mathcal{E} , it is reasonable to assume that such a pattern can be *approximately* captured by the governing function $\mathcal{K} : \mathcal{Z} \rightarrow \mathcal{Z}$ in a way such that the *forecasted* domain $\mathcal{D}_{i+1}^{\mathcal{K}} \triangleq \mathcal{D}_i(\hat{z}, y)$ is close to \mathcal{D}_{i+1} as much as possible, where $\hat{z} = \mathcal{K} \circ \mathcal{G} \circ \phi(x)$. Then, the forecasted domain $\mathcal{D}_t^{\mathcal{K}}$ (transformed from the last source domain \mathcal{D}_m) can be adopted as an alternative to \mathcal{D}_t for domain generalization if \mathcal{K} can properly capture the evolving pattern. Intuitively, capturing the evolving pattern in \mathcal{E} is hopeless if it varies arbitrarily. On the other hand, if the underlying pattern is consistent over domains in the Koopman space, it is reasonable to assume that there exists a Koopman operator \mathcal{K}^* that would perform consistently well over all the domain pairs. To this end, we introduce the notion of λ -consistency of an environment \mathcal{E} .

Definition 4.1 (λ -Consistency). We define the ideal mapping function in the worst-case domain: $\mathcal{K}^* = \arg \min_{\mathcal{K}} \max_{\mathcal{D}_i \in \mathcal{E}} d(\mathcal{D}_i || \mathcal{D}_i^{\mathcal{K}})$, where $\mathcal{D}_{i+1}^{\mathcal{K}} \triangleq \mathcal{D}_i(\hat{z}, y)$. Then, an evolving environment \mathcal{E} is λ -consistent if the following holds:

$$|d(\mathcal{D}_i || \mathcal{D}_i^{\mathcal{K}^*}) - d(\mathcal{D}_j || \mathcal{D}_j^{\mathcal{K}^*})| \leq \lambda, \quad \forall \mathcal{D}_i, \mathcal{D}_j \in \mathcal{E}.$$

where $d(\mathcal{D}_i || \mathcal{D}_i^{\mathcal{K}})$ is the Kullback–Leibler (KL) divergence between $\mathcal{D}_i^{\mathcal{K}}$ and \mathcal{D}_i (Fuglede and Topsoe 2004).

Our definition of λ -consistency offers a fundamentally new perspective to depict the temporal pattern. The distance is between forecasted domains and real domains, without any assumptions about the distance of domain pairs in \mathcal{E} . It is worth mentioning that λ -consistency is not an assumption. Instead, it is only a value that depicts how *consistently* environment evolves in \mathcal{E} . A small λ means that the environment evolves steadily, which means the future domains are predictable and can be leveraged for performance improvement. Reversely, there are some cases in which the evolving mechanism is too complex, non-Markov, or too random, which will result in a large λ .

We show the generalization bound between $R_{\mathcal{D}_t}(h)$ and $R_{\mathcal{D}_t^{\mathcal{K}^*}}(h)$ as follows, the complete proof is deferred to the Section A in Appendix due to space limitations:

Theorem 4.2. Let $\{\mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_m\}$ be m observed source domains sampled sequentially from an evolving environment \mathcal{E} , and \mathcal{D}_t be the next unseen target domain: $\mathcal{D}_t = \mathcal{D}_{m+1}$. \mathcal{G} is the range of the interval of the loss function. Then, if \mathcal{E} is λ -consistent, we have

$$R_{\mathcal{D}_t}(h) \leq R_{\mathcal{D}_t^{\mathcal{K}^*}}(h) + \frac{\mathcal{G}}{\sqrt{2(m-1)}} \times \left(\sqrt{\sum_{i=2}^m \mathbb{E}_{y \sim \mathcal{D}_i(y)} d(\mathcal{D}_i(z|y) \| \mathcal{D}_i^{\mathcal{K}^*}(z|y))} + \sqrt{(m-1)\lambda} \right) \quad (2)$$

Discussion We note two key theoretical differences between TDG and previous studies of DA in evolving environments (Ben-David et al. 2010; Liu et al. 2020). (1) In Domain Adaptation (DA), the target risk is bounded in terms of source and target domains (e.g., $\mathcal{H}\Delta\mathcal{H}$ -divergence), while our analysis relies on the distance between forecasted and real domains. (2) DA theories are built upon the assumption that there exists an ideal joint hypothesis that achieves a low combined error on both domains, while our assumption is the λ -consistency in the Koopman space of \mathcal{E} . The theoretical result motivates us to design algorithms with Koopman operators, which directly minimizes the KL-term in Eq. (2). We further bridge the gap between the theory and the method in Section 4.3.

4.2 Proposed Methods

We propose Temporal Koopman Networks (TKNets), a deep neural network model based on Koopman theory to capture the evolving patterns between the Temporal Domains. Our analysis reveals two strategies for designing TKNets:

- (i) Learning the Koopman operator \mathcal{K} to capture the evolving pattern by minimizing the distance between the distributions of forecasted and real domains in the Koopman space.
- (ii) Note that $\mathcal{D}_{i+1}^{\mathcal{K}} = \mathcal{K} \circ \mathcal{G} \circ \phi(\mathcal{D}_i)$ is produced from \mathcal{D}_i , but its quality is evaluated on \mathcal{D}_{i+1} . Minimizing $d(\mathcal{D}_i^{\mathcal{K}} \| \mathcal{D}_i)$ naturally leads to the adoption of the Koopman Theory for learning the linear infinite-dimensional operator \mathcal{K} .

Specifically, the input samples from both \mathcal{D}_i and \mathcal{D}_{i+1} are first transformed to a representation space \mathcal{V} induced by an embedding function $\phi : \mathcal{X} \rightarrow \mathcal{V}$. Then we use measurement functions $\mathcal{G} := [g_1, \dots, g_n]$ that span the Koopman space function $\mathcal{G} : \mathcal{V} \rightarrow \mathcal{Z}$ (for example, the first dimension of measurement functions can be $g_1(x) = \sin x$) which maps the samples into the Koopman space so that we can find a Koopman operator \mathcal{K} leading to distribution of $\mathcal{K} \circ \mathcal{G} \circ \phi(\mathcal{D}_i)$ aligned with the distribution of $\mathcal{G} \circ \phi(\mathcal{D}_{i+1})$. In all, TKNets consists of three components: a embedding function $\phi : \mathcal{X} \rightarrow \mathcal{V}$, measurement functions $\mathcal{G} : \mathcal{V} \rightarrow \mathcal{Z}$, and Koopman operator $\mathcal{K} : \mathcal{G}(\mathcal{V}) \rightarrow \mathcal{G}(\mathcal{V})$. The key idea of TKNets is to use $\{\phi, \mathcal{G}, \mathcal{K}\}$ to capture the evolving pattern of \mathcal{E} by mapping the data in the nonlinear dynamic system into the Koopman space such that the transition function can be modeled by an infinite-dimensional linear Koopman operator. We actually align the distribution of $\mathcal{D}_{i+1}^{\mathcal{K}}(z|y)$ and $\mathcal{D}_{i+1}(z|y)$, as suggested by Theorem 4.2. We implicitly minimize the distance between class-conditional semantic centroids (Xie et al. 2018; Shui et al. 2021), which is

an approximation of the semantic conditional distribution of each class. We illustrate the relationship between theory and practical implementation in Sec. 4.3.

Let $\mathcal{S}_i = \{(x_n^i, y_n^i)\}_{n=1}^{N_i}$ be the data set of size N_i sampled from \mathcal{D}_i , and \mathcal{S}_i^k be the subset of \mathcal{S}_i with class $k \in \{1, \dots, K\}$, where K is the total number of classes. In TKNets, the semantic centroids are computed from the support set \mathcal{S}_i through the composite function $\mathcal{K} \circ \mathcal{G} \circ \phi$. Since the Koopman operator captures the transition relationship between two consecutive domains, the computed centroids from \mathcal{S}_i are the forecasted centroids in the domain $i+1$. Then, the forecasted centroid of domain $i+1$ is the mean vector of the *support* instances belonging to \mathcal{S}_i^k :

$$c_{i+1}^k = \frac{1}{|\mathcal{S}_i^k|} \sum_{(x_n^i, y_n^i) \in \mathcal{S}_i^k} \mathcal{K} \circ \mathcal{G} \circ \phi(x_n^i)$$

The query instances are from \mathcal{S}_{i+1} and are passed through the composite embedding function $\mathcal{G} \circ \phi$. The predictive distribution for \mathcal{D}_{i+1} is given by

$$\mathcal{D}(y^{i+1} = k | x^{i+1}) = \frac{\exp(-d_{\text{eu}}(\mathcal{G} \circ \phi(x^{i+1}), c_i^k))}{\sum_{k'=1}^K \exp(-d_{\text{eu}}(\mathcal{G} \circ \phi(x^{i+1}), c_i^{k'}))}, \quad (3)$$

where $d_{\text{eu}} : \mathcal{Z} \times \mathcal{Z} \rightarrow [0, +\infty)$ is a euclidean distance function of embedding space, and we adopt squared Euclidean distance in our implementation, as suggested in (Snell, Swersky, and Zemel 2017). During the training stage, at each step, we randomly choose the data sets $\mathcal{S}_i, \mathcal{S}_{i+1}$ from two consecutive domains as support and query sets, respectively. Then, we sample N_B samples from each class k in \mathcal{S}_i , which is used to compute centroid c_i^k for the query data in \mathcal{S}_{i+1} . Model optimization proceeds by minimizing the negative log probability:

$$J = \sum_{i=1}^{m-1} \sum_{k=1}^K \sum_{n=1}^{N_B} -\frac{1}{KN_B} \log \mathcal{D}(y_n^{i+1} = k | x_n^{i+1}) \quad (4)$$

The pseudocode to compute J for a training episode is shown in Algorithm 1. In the testing stage, we pass the instances from \mathcal{S}_m and \mathcal{S}_t through $\mathcal{K} \circ \mathcal{G} \circ \phi$ and $\mathcal{G} \circ \phi$ respectively as support and query sets and then make predictions for the instances in \mathcal{D}_t using Eq. (3).

4.3 From the Theory to TKNets

We further demonstrate the connection between the theory and our proposed method TKNets (the complete proof is deferred Section A.9 in Appendix):

Theorem 4.3. The optimization loss J defined in Eq. (4), is the approximation of the upper bound of KL term in Def. 4.1 and an inter-class distance loss, which implies

$$\underbrace{-\sum_{i=2}^m \mathbb{E}_{y' \sim \mathcal{D}_i(y')} \mathbb{E}_{y \sim \mathcal{D}_i^{\mathcal{K}^*}(y), y \neq y'} d(\mathcal{D}_i(z|y') \| \mathcal{D}_i^{\mathcal{K}^*}(z|y))}_{\text{Inter-Class Distance loss}} + \underbrace{\sum_{i=2}^m \mathbb{E}_{y \sim \mathcal{D}_i(y)} d(\mathcal{D}_i(z|y) \| \mathcal{D}_i^{\mathcal{K}^*}(z|y))}_{\text{KL-Divergence loss in Eq. (2)}} \leq \frac{J}{K-1} \quad (5)$$

Algorithm 1: TKNNets (one episode)

Input: $\{\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_m\}$: m data sets from consecutive domains. N_B : the number of support and query instances for each class. $\text{RANDOMSAMPLE}(\mathcal{S}, N)$: a set of N instances sampled uniformly from the set \mathcal{S} without replacement.

Output: The loss J for a randomly generated training episode.

$t \leftarrow \text{RANDOMSAMPLE}(\{1, \dots, m\})$

for k in $\{1, \dots, K\}$ **do**

$S^k \leftarrow \text{RANDOMSAMPLE}(\mathcal{S}_i^k, N_B)$

$c^k = \frac{1}{|S^k|} \sum_{(x_j, y_j) \in S^k} \mathcal{K} \circ \mathcal{G} \circ \phi(x_j)$

$Q^k \leftarrow \text{RANDOMSAMPLE}(\mathcal{S}_{i+1}^k, N_B)$

end for

$J \leftarrow 0$

for k in $\{1, \dots, K\}$ **do**

for (x, y) in Q **do**

$J \leftarrow J + \frac{1}{KN_B} \left[d(\mathcal{G} \circ \phi(x), c^k) \right.$

$\left. + \log \sum_{k'} \exp(-d(\mathcal{G} \circ \phi(x), c^{k'})) \right]$

end for

end for

Discussion Minimizing the above KL term, which is the distance between the conditional distributions of forecasted and real domains, corresponds to minimizing the KL term in Eq. (2), thereby reducing the generalization bound. The first term can be regarded as a regularizer term that achieves maximization of the inter-class dissimilarity.

5 Experiments

5.1 Experimental Setup

We evaluated our algorithm on 6 datasets, 4 of which were collected in real-world scenarios (RMNIST (Ghifary et al. 2015), Portrait (Kumar, Ma, and Liang 2020; Chen and Chao 2021), Cover Type (Kumar, Ma, and Liang 2020)), and FMoW (Christie et al. 2018)).

(1) **Evolving Circle (EvolCircle, Fig. 2)** consists of 30 evolving domains, where the instances are generated from 30 2D-Gaussian distributions with the same variances but different centers uniformly distributed on a half-circle. (2) **Rotated Plate (RPlate, Fig. 3)** consists of 30 domains, where the instances of each domain are generated by the same Gaussian distribution but the decision boundary rotates from 0° to 348° with an interval of 12° . (3) **Rotated MNIST (RMNIST)** We randomly select only 2400 instances in the raw MNIST dataset and split them into 12 domains equally. Then we apply the rotations with degree of $\theta = \{0^\circ, 10^\circ, \dots, 110^\circ\}$ on each domain respectively. The amount of samples in each domain is only 200, which makes this task more challenging. (4) **Portrait** data was originally proposed in (Ginosar et al. 2015) and has been used as a benchmark dataset for studying evolving domain adaptation (Chen and Chao 2021; Kumar, Ma, and Liang 2020) and other problems related to evolving domains (Chen et al.

2021; Lei, Hu, and Lee 2021; Mancini et al. 2019; Zhou et al. 2022). It aims to predict gender based on the photos of high school seniors across different decades. We divided the dataset into 12 domains by year. (5) **Cover Type** data set of geology aims to predict cover type (the predominant kind of tree cover) from 54 strictly cartographic variables. To generate evolving domains, we sort the samples by the ascending order of the distance to the water body, as proposed in (Kumar, Ma, and Liang 2020). Then we equally divided the data set into 10 domains by distance. (6) **FMoW** A large satellite image dataset with target detection and classification tasks (Christie et al. 2018). We select 5 common classes to compose a classification task. The dataset is divided into 19 domains by time.

We compared the proposed method with the following baselines: (1) **GroupDRO** (Sagawa et al. 2019); (2) **MLDG** (Li et al. 2018a); (3) **MMD** (Li et al. 2018b); (4) **SagNet** (Nam et al. 2021); (5) **VREx** (Krueger et al. 2021); (6) **SD** (Pezeshki et al. 2020); (7) **IRM** (Arjovsky et al. 2019); (8) **Mixup** (Yan et al. 2020); (9) **CORAL** (Sun and Saenko 2016); (10) **MTL** (Blanchard et al. 2021); (11) **RSC** (Huang et al. 2020b); (12) **DIRL** (Nguyen et al. 2021); (13) **ERM** (Vapnik 1998); (14) **ERM-Near Only** using the most adjacent domains for training; (15) **ERM-W** ERM with weighted loss where the samples from domains closer to the target have higher weights; (16) **GI** (Nasery et al. 2021); (17) **LSSAE** (Qin, Wang, and Li 2022); (18) **CIDA** (Wang, He, and Katabi 2020); (19) **EAML** (Liu et al. 2020). All the baselines and experiments were implemented with DomainBed package (Gulrajani and Lopez-Paz 2020) under the same setting, which guarantees extensive and sufficient comparisons.

5.2 Results and Analysis

Overall Evaluation The performances of our proposed method and baselines are reported in Table 1. It can be observed that TKNNets consistently outperforms extensive DG baselines over all the data sets and achieves 85.7% on average which is significantly higher than the other algorithms ($\approx 8\% - 20\%$). The results indicate that existing DG methods cannot deal with evolving domain shifts well while TKNNets can properly capture the evolving patterns in the environments. ERM-Near is a simple solution for TDG as the most recent domain shift least with respect to the target. It performs well on synthetic data as the samples of one domain are enough for training a classifier while the low sample utilization makes it hard to handle real datasets. Section B.2 in the appendix presents an experiment on ERM with domain index, which shows that directly incorporating it does not improve learning in the TDG problem. We also implement CIDA (Wang, He, and Katabi 2020), as a competitive TDA method. With the fair random hyper-parameter search strategy, the accuracy of TKNNets is 14.5% higher on average. The results illustrate the superiority of TKNNets.

Evolving $P(X)$ (EvolCircle). The decision boundaries on the unseen target domain \mathcal{D}_{30} learned by ERM and TKNNets are shown in Fig. 2(c) and Fig. 2(d) respectively. We can observe that TKNet fits the ground truth significantly better than ERM. This indicates that our approach can cap-




Figure 2: Visualization of the EvolCircle dataset. (a) 30 domains indexed by different colors, where the left bottom one is the target domain. (b) Positive and negative instances are denoted by red and blue dots respectively. (c) The decision boundaries learned by ERM. (d) Decision boundaries of the last model on all domains. (e) Decision boundaries of models in each domain.




Figure 3: Visualization of the RPlate data set. (a) True decision boundaries evolve over domains. (b) & (c) The decision boundaries learned by ERM and TKNNets on the target domain.

Algorithm	EvolCircle	RPlate	RMNIST	Portrait	Cover Type	FMoW	Average
ERM	72.7 ± 1.1	63.9 ± 0.9	79.4 ± 0.0	95.8 ± 0.1	71.8 ± 0.2	54.6 ± 0.1	74.7
ERM-Near	90.6 ± 1.4	94.1 ± 1.4	70.2 ± 1.3	80.9 ± 0.3	65.4 ± 0.2	44.2 ± 0.1	74.2
ERM-W	80.4 ± 1.1	73.4 ± 0.5	82.0 ± 1.1	94.9 ± 0.3	71.5 ± 0.2	50.2 ± 0.1	75.4
GroupDRO	75.5 ± 1.0	70.0 ± 4.9	76.5 ± 0.2	94.8 ± 0.1	66.4 ± 0.5	57.3 ± 0.1	73.4
MLDG	91.5 ± 2.0	66.9 ± 1.8	75.0 ± 0.3	66.2 ± 1.7	68.4 ± 0.7	43.8 ± 0.0	68.6
MMD	86.7 ± 5.7	59.9 ± 1.4	35.4 ± 0.0	95.4 ± 0.1	69.8 ± 0.4	60.0 ± 0.0	67.8
SagNet	78.7 ± 3.2	63.8 ± 2.9	79.4 ± 0.1	95.3 ± 0.1	65.3 ± 2.2	56.2 ± 0.1	73.1
VREx	82.9 ± 6.6	61.1 ± 2.6	79.4 ± 0.1	94.3 ± 0.2	66.0 ± 0.9	61.2 ± 0.0	73.3
SD	81.7 ± 4.3	65.3 ± 1.4	78.8 ± 0.1	95.1 ± 0.2	69.1 ± 0.9	55.2 ± 0.0	74.2
IRM	86.2 ± 3.0	67.2 ± 2.1	47.5 ± 0.4	94.4 ± 0.3	66.0 ± 1.0	58.8 ± 0.0	70.0
Mixup	91.5 ± 2.6	66.8 ± 1.8	81.3 ± 0.2	96.4 ± 0.2	69.7 ± 0.6	59.5 ± 0.0	77.5
CORAL	86.8 ± 5.1	61.9 ± 1.4	78.4 ± 0.1	95.1 ± 0.1	68.1 ± 1.3	56.1 ± 0.0	74.4
MTL	77.7 ± 2.4	66.0 ± 1.2	77.2 ± 0.0	95.4 ± 0.1	69.2 ± 0.9	51.7 ± 0.0	72.9
RSC	91.5 ± 2.1	67.9 ± 4.2	74.7 ± 0.1	95.5 ± 0.1	69.4 ± 0.3	55.7 ± 0.1	75.8
DIRL	53.3 ± 0.2	56.3 ± 0.4	76.3 ± 0.3	93.2 ± 0.2	61.2 ± 0.3	43.4 ± 0.3	64.0
CIDA	68.0 ± 2.8	91.6 ± 2.5	86.5 ± 1.2	94.5 ± 1.2	70.2 ± 0.3	63.4 ± 0.2	70.9
EAML	90.1 ± 1.3	91.2 ± 0.3	83.1 ± 0.8	95.2 ± 0.1	69.3 ± 0.1	61.1 ± 0.2	81.7
LSSAE	91.5 ± 2.3	91.7 ± 1.3	86.0 ± 0.2	96.0 ± 0.2	71.3 ± 0.5	60.5 ± 0.1	82.8
GI	92.1 ± 1.7	94.3 ± 0.4	86.3 ± 0.2	96.3 ± 0.1	71.7 ± 0.3	64.5 ± 0.1	83.3
TKNNets (Ours)	94.2 ± 0.9	95.0 ± 0.5	87.5 ± 0.1	97.2 ± 0.0	73.8 ± 1.0	66.8 ± 0.1	85.7

Table 1: Comparison of accuracy (%) among different methods.

ture the evolving pattern of $P(X)$ according to source domains and then learn a better classifier for the target domain. Furthermore, we can observe that the decision boundary learned by ERM achieves better performance on the observed source domains. This is because it focuses on improving generalization ability on all source domains, which leads

the poor performance on the outer target domain \mathcal{D}_{30} . Instead, TKNNets can “foresee” the centroids for the target domain, which guarantees a good generalization performance, though it may not perform well on source domains.

Evolving $P(Y|X)$ or $P(X|Y)$ (RPlate). By visualizing the datasets, we can observe that the predicted boundary

# Domains	7	9	11	13	15	17	19
Mixup	83.8 ± 0.6	83.3 ± 0.3	80.0 ± 0.3	78.3 ± 0.3	80.0 ± 0.3	77.3 ± 0.3	72.5 ± 0.3
IRM	35.6 ± 0.3	46.5 ± 0.3	40.4 ± 0.4	49.6 ± 0.3	46.0 ± 0.1	46.9 ± 0.3	41.3 ± 0.1
MLDG	82.7 ± 0.3	80.0 ± 0.6	79.0 ± 0.1	74.2 ± 0.1	77.9 ± 0.3	71.7 ± 0.1	68.8 ± 0.3
ERM	81.3 ± 0.1	79.7 ± 0.2	79.7 ± 0.3	75.6 ± 0.3	77.8 ± 0.3	69.1 ± 0.3	74.4 ± 0.1
TKNets (Ours)	81.1 ± 0.1	82.8 ± 0.3	88.1 ± 0.3	87.3 ± 0.5	86.6 ± 0.1	85.6 ± 0.3	86.3 ± 0.3

Table 2: Comparison of accuracy (%) of different methods on RMNIST data set with different number of domains.

Domain Distance	3°	5°	7°	10°	15°	20°
Mixup	92.5 ± 0.1	91.9 ± 0.3	88.4 ± 0.3	81.3 ± 0.2	73.1 ± 0.1	59.4 ± 0.3
IRM	69.4 ± 0.2	63.4 ± 0.1	49.7 ± 0.0	47.5 ± 0.4	35.6 ± 0.3	31.3 ± 0.3
MLDG	90.9 ± 0.1	87.5 ± 0.2	85.0 ± 0.2	75.0 ± 0.3	71.9 ± 0.1	56.9 ± 0.3
ERM	92.2 ± 0.1	88.8 ± 0.0	82.8 ± 0.1	79.4 ± 0.0	66.3 ± 0.1	53.1 ± 0.3
TKNets (Ours)	94.7 ± 0.3	93.1 ± 0.3	90.8 ± 0.3	88.3 ± 0.1	86.4 ± 0.3	84.1 ± 0.3

Table 3: Comparison of accuracy (%) of different methods on RMNIST data set with different distance between domains.

Domain	T+1	T+2	T+3	T+4
ERM	79.4 ± 0.0	65.0 ± 1.1	50.7 ± 1.2	41.0 ± 1.7
GI	86.3 ± 0.2	65.9 ± 3.1	48.5 ± 1.2	41.1 ± 1.1
EAML	83.1 ± 0.8	67.5 ± 1.5	55.6 ± 1.7	44.1 ± 1.6
TKNets	87.5 ± 0.1	83.1 ± 0.8	72.9 ± 2.4	72.2 ± 1.1

Table 4: Performance of the algorithms on multiple-step target domains.

of TKNets better approximates the ground truth, compared with the result of ERM. This indicates that our approach can also capture the $P(Y|X)$ evolving pattern. Existing DG methods perform poorly on this dataset because the ground truth labeling function varies. Under the evolving labeling functions, even the same instance can have different labels in different domains. Thus, there does not exist a single model that can perform well across all the domains. For this situation, learning a model specifically for one domain instead of all domains can be a possible solution. TKNets can capture the evolving pattern and produce a model specifically for the target domain.

When to apply TKNets? Existing DG methods assume that the distances between observed and unseen domains are not very large. However, the dissimilarity between domains is a crucial factor that can fundamentally influence the possibility and performance of generalization. To investigate the impact of variances of the environment, we create a series of variations on the raw RMNIST data by jointly varying the number of domains (Table 2) and the degree interval (Table 3) between two consecutive domains. On the one hand, the greater number of domains and larger distance between them lead to more significant differences across domains. This makes traditional DG methods harder to train one model from all domains, but instead more domains benefit our TKNets to learn the evolving pattern to achieve better performance. On the other hand, we observe TKNets sig-

nificantly outperform other baselines when the number of domains and the distance between domains increase.

In Table 2, the performance of traditional DG methods fluctuate when the number of domains increases. TKNets’ performance continuously improves when the domain number increases since it easily learns the evolving manners from more domains. Please refer to Section B.1 for more discussion about this. From Table 3, when the domain distance increases, the performance of DG methods decreases severely while the performance of TKNets drops slightly.

Domain Generalization across Temporal Domains. There are many scenarios in the real world where it is necessary for us to have the model applied to the next several domains instead of one. In this section, to test the performance of the algorithm on more distant domains, we extend the RMNIST (Ghifary et al. 2015) to 15 domains, with the last four domains as target domains. From the Table 4, as expected, the model performance drops heavily as the index increases due to the domain shifting. However, the accuracy of our algorithm decreases very slowly compared to other algorithms. Even on $T+4$ domains, the accuracy still remains at 72.2%. While the accuracy of other algorithms has dropped to 41.0% ~ 44.1%. The result demonstrates the generalization performance of our algorithm on multiple domains.

6 Conclusions

In this paper, we study the problem of domain generalization in an evolving environment and propose temporal domain generalization (TDG) as a general framework to address it. Our theoretical analysis highlights the role of learning a Koopman operator to capture the evolving pattern over domains. Motivated by our theory, we propose temporal Koopman networks (TKNets), a simple and efficient algorithm for TDG. Experiments on both synthetic and real-world datasets validate the effectiveness of our method.

Acknowledgements

We appreciate constructive feedback from anonymous reviewers and meta-reviewers. This work is supported by the Natural Sciences and Engineering Research Council of Canada (NSERC), Discovery Grants program.

References

- Arjovsky, M.; Bottou, L.; Gulrajani, I.; and Lopez-Paz, D. 2019. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*.
- Azencot, O.; Erichson, N. B.; Lin, V.; and Mahoney, M. 2020. Forecasting sequential data using consistent Koopman autoencoders. In *International Conference on Machine Learning*, 475–485. PMLR.
- Bai, G.; Ling, C.; and Zhao, L. 2022. Temporal Domain Generalization with Drift-Aware Dynamic Neural Networks. *arXiv preprint arXiv:2205.10664*.
- Balaji, Y.; Sankaranarayanan, S.; and Chellappa, R. 2018. Metareg: Towards domain generalization using meta-regularization. *Advances in Neural Information Processing Systems*, 31: 998–1008.
- Ben-David, S.; Blitzer, J.; Crammer, K.; Kulesza, A.; Pereira, F.; and Vaughan, J. W. 2010. A theory of learning from different domains. *Machine learning*, 79(1): 151–175.
- Blanchard, G.; Deshmukh, A. A.; Dogan, Ü.; Lee, G.; and Scott, C. 2021. Domain Generalization by Marginal Transfer Learning. *J. Mach. Learn. Res.*, 22: 2–1.
- Blanchard, G.; Lee, G.; and Scott, C. 2011. Generalizing from several related classification tasks to a new unlabeled sample. *Advances in neural information processing systems*, 24: 2178–2186.
- Carlucci, F. M.; D’Innocente, A.; Bucci, S.; Caputo, B.; and Tommasi, T. 2019. Domain generalization by solving jigsaw puzzles. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2229–2238.
- Chen, H.-Y.; and Chao, W.-L. 2021. Gradual Domain Adaptation without Indexed Intermediate Domains. *Advances in Neural Information Processing Systems*, 34.
- Chen, Y.; Luo, H.; Ma, T.; and Zhang, C. 2021. Active online learning with hidden shifting domains. In *International Conference on Artificial Intelligence and Statistics*, 2053–2061. PMLR.
- Christie, G.; Fendley, N.; Wilson, J.; and Mukherjee, R. 2018. Functional map of the world. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 6172–6180.
- Fuglede, B.; and Topsoe, F. 2004. Jensen-Shannon divergence and Hilbert space embedding. In *International Symposium on Information Theory, 2004. ISIT 2004. Proceedings.*, 31. IEEE.
- Ganin, Y.; and Lempitsky, V. 2015. Unsupervised domain adaptation by backpropagation. In *International conference on machine learning*, 1180–1189. PMLR.
- Ghifary, M.; Kleijn, W. B.; Zhang, M.; and Balduzzi, D. 2015. Domain generalization for object recognition with multi-task autoencoders. In *Proceedings of the IEEE international conference on computer vision*, 2551–2559.
- Ginosar, S.; Rakelly, K.; Sachs, S.; Yin, B.; and Efros, A. A. 2015. A Century of Portraits: A Visual Historical Record of American High School Yearbooks. *CoRR*, abs/1511.02575.
- Gulrajani, I.; and Lopez-Paz, D. 2020. In Search of Lost Domain Generalization. *CoRR*, abs/2007.01434.
- Huang, Z.; Wang, H.; Xing, E. P.; and Huang, D. 2020a. Self-challenging improves cross-domain generalization. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*, 124–140. Springer.
- Huang, Z.; Wang, H.; Xing, E. P.; and Huang, D. 2020b. Self-challenging improves cross-domain generalization. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*, 124–140. Springer.
- Ilse, M.; Tomczak, J. M.; Louizos, C.; and Welling, M. 2020. Diva: Domain invariant variational autoencoders. In *Medical Imaging with Deep Learning*, 322–348. PMLR.
- Kim, J.; Yoo, S.; Park, G.; and Kim, J. 2020. Continual Unsupervised Domain Adaptation with Adversarial Learning. *CoRR*, abs/2010.09236.
- Koopman, B. O. 1931. Hamiltonian systems and transformation in Hilbert space. *Proceedings of the National Academy of Sciences*, 17(5): 315–318.
- Krueger, D.; Caballero, E.; Jacobsen, J.-H.; Zhang, A.; Binas, J.; Zhang, D.; Le Priol, R.; and Courville, A. 2021. Out-of-distribution generalization via risk extrapolation (rex). In *International Conference on Machine Learning*, 5815–5826. PMLR.
- Kumar, A.; Ma, T.; and Liang, P. 2020. Understanding self-training for gradual domain adaptation. In *International Conference on Machine Learning*, 5468–5479. PMLR.
- Lei, Q.; Hu, W.; and Lee, J. 2021. Near-optimal linear regression under distribution shift. In *International Conference on Machine Learning*, 6164–6174. PMLR.
- Li, D.; Yang, Y.; Song, Y.-Z.; and Hospedales, T. M. 2018a. Learning to generalize: Meta-learning for domain generalization. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Li, H.; Pan, S. J.; Wang, S.; and Kot, A. C. 2018b. Domain generalization with adversarial feature learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 5400–5409.
- Li, Y.; He, H.; Wu, J.; Katabi, D.; and Torralba, A. 2020. Learning Compositional Koopman Operators for Model-Based Control. In *International Conference on Learning Representations*.
- Li, Y.; Yang, Y.; Zhou, W.; and Hospedales, T. 2019. Feature-critic networks for heterogeneous domain generalization. In *International Conference on Machine Learning*, 3915–3924. PMLR.
- Liu, H.; Long, M.; Wang, J.; and Wang, Y. 2020. Learning to Adapt to Evolving Domains. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M. F.; and Lin, H., eds., *Advances in Neural Information Processing Systems*, volume 33, 22338–22348. Curran Associates, Inc.
- Mancini, M.; Bulò, S. R.; Caputo, B.; and Ricci, E. 2018. Best sources forward: domain generalization through source-specific nets. In *2018 25th IEEE international conference on image processing (ICIP)*, 1353–1357. IEEE.
- Mancini, M.; Bulò, S. R.; Caputo, B.; and Ricci, E. 2019. Adagraph: Unifying predictive and continuous domain adaptation through graphs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6568–6577.
- Muandet, K.; Balduzzi, D.; and Schölkopf, B. 2013. Domain generalization via invariant feature representation. In *International Conference on Machine Learning*, 10–18. PMLR.
- Nam, H.; Lee, H.; Park, J.; Yoon, W.; and Yoo, D. 2021. Reducing Domain Gap by Reducing Style Bias. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8690–8699.
- Nasery, A.; Thakur, S.; Piratla, V.; De, A.; and Sarawagi, S. 2021. Training for the Future: A Simple Gradient Interpolation Loss to Generalize Along Time. *Advances in Neural Information Processing Systems*, 34: 19198–19209.

- Nguyen, A. T.; Tran, T.; Gal, Y.; and Baydin, A. G. 2021. Domain Invariant Representation Learning with Domain Density Transformations. In *Thirty-Fifth Conference on Neural Information Processing Systems*.
- Pezeshki, M.; Kaba, S.-O.; Bengio, Y.; Courville, A.; Precup, D.; and Lajoie, G. 2020. Gradient starvation: A learning proclivity in neural networks. *arXiv preprint arXiv:2011.09468*.
- Qin, T.; Wang, S.; and Li, H. 2022. Generalizing to Evolving Domains with Latent Structure-Aware Sequential Autoencoder. In *ICML*.
- Rahman, M. M.; Fookes, C.; Baktashmotlagh, M.; and Sridharan, S. 2019. Multi-component image translation for deep domain generalization. In *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*, 579–588. IEEE.
- Ryu, J.; Kwon, G.; Yang, M.-H.; and Lim, J. 2019. Generalized convolutional forest networks for domain generalization and visual recognition. In *International Conference on Learning Representations*.
- Sagawa, S.; Koh, P. W.; Hashimoto, T. B.; and Liang, P. 2019. Distributionally robust neural networks for group shifts: On the importance of regularization for worst-case generalization. *arXiv preprint arXiv:1911.08731*.
- Shankar, S.; Piratla, V.; Chakrabarti, S.; Chaudhuri, S.; Jyothi, P.; and Sarawagi, S. 2018. Generalizing Across Domains via Cross-Gradient Training. In *International Conference on Learning Representations*.
- Shui, C.; Li, Z.; Li, J.; Gagné, C.; Ling, C. X.; and Wang, B. 2021. Aggregating From Multiple Target-Shifted Sources. In *Proceedings of the International Conference on Machine Learning*, 9638–9648.
- Snell, J.; Swersky, K.; and Zemel, R. S. 2017. Prototypical Networks for Few-shot Learning. *CoRR*, abs/1703.05175.
- Sun, B.; and Saenko, K. 2016. Deep coral: Correlation alignment for deep domain adaptation. In *European conference on computer vision*, 443–450. Springer.
- Vapnik, V. N. 1998. *Statistical Learning Theory*. Bantam.
- Wang, H.; He, H.; and Katabi, D. 2020. Continuously Indexed Domain Adaptation. In *ICML*.
- Wang, H.; Li, B.; and Zhao, H. 2022. Understanding Gradual Domain Adaptation: Improved Analysis, Optimal Path and Beyond. *arXiv preprint arXiv:2204.08200*.
- Wang, J.; Lan, C.; Liu, C.; Ouyang, Y.; and Qin, T. 2021. Generalizing to Unseen Domains: A Survey on Domain Generalization. *CoRR*, abs/2103.03097.
- Xie, S.; Zheng, Z.; Chen, L.; and Chen, C. 2018. Learning semantic representations for unsupervised domain adaptation. In *International conference on machine learning*, 5423–5432. PMLR.
- Xu, Z.; Li, W.; Niu, L.; and Xu, D. 2014. Exploiting low-rank structure from latent domains for domain generalization. In *European Conference on Computer Vision*, 628–643. Springer.
- Yan, S.; Song, H.; Li, N.; Zou, L.; and Ren, L. 2020. Improve unsupervised domain adaptation with mixup training. *arXiv preprint arXiv:2001.00677*.
- Yue, X.; Zhang, Y.; Zhao, S.; Sangiovanni-Vincentelli, A.; Keutzer, K.; and Gong, B. 2019. Domain randomization and pyramid consistency: Simulation-to-real generalization without accessing target domain data. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2100–2110.
- Zeng, Q.; Wang, W.; Zhou, F.; Ling, C.; and Wang, B. 2023. Foresee What You Will Learn: Data Augmentation for Domain Generalization in Non-stationary Environment. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(9): 11147–11155.
- Zhou, S.; Zhao, H.; Zhang, S.; Wang, L.; Chang, H.; Wang, Z.; and Zhu, W. 2022. Online Continual Adaptation with Active Self-Training. In *International Conference on Artificial Intelligence and Statistics*, 8852–8883. PMLR.