

Enhancing Evolving Domain Generalization through Dynamic Latent Representations

Binghui Xie¹, Yongqiang Chen¹, Jiaqi Wang¹, Kaiwen Zhou¹, Bo Han², Wei Meng¹, James Cheng¹

¹The Chinese University of Hong Kong

²Hong Kong Baptist University

{bhxie21, yqchen, kwzhou, jqwang23, jcheng}@cse.cuhk.edu.hk, bhanml@comp.hkbu.edu.hk

Abstract

Domain generalization is a critical challenge for machine learning systems. Prior domain generalization methods focus on extracting domain-invariant features across several *stationary domains* to enable generalization to new domains. However, in *non-stationary* tasks where new domains evolve in an underlying continuous structure, such as time, merely extracting the invariant features is *insufficient* for generalization to the evolving new domains. Nevertheless, it is non-trivial to learn both evolving and invariant features within a single model due to their conflicts. To bridge this gap, we build causal models to characterize the distribution shifts concerning the two patterns, and propose to learn both *dynamic* and *invariant* features via a new framework called Mutual Information-Based Sequential Autoencoders (MISTS). MISTS adopts information theoretic constraints onto sequential autoencoders to disentangle the dynamic and invariant features, and leverage a domain adaptive classifier to make predictions based on both evolving and invariant information. Our experimental results on both synthetic and real-world datasets demonstrate that MISTS succeeds in capturing both evolving and invariant information, and present promising results in evolving domain generalization tasks.

Introduction

Domain generalization (DG) is a critical challenge for machine learning systems that requires model to generalize beyond the assumption that training and testing data come from identical and independent distributions (Bengio et al. 2019). To address the issue, most previous DG methods focus on extracting domain-invariant features across several *stationary* source domains (Sun and Saenko 2016; Li et al. 2018; Sagawa et al. 2020; Arjovsky et al. 2019). Nevertheless, domains could also be non-stationary and evolve along with certain structures (Wang et al. 2022; Qin, Wang, and Li 2022; Yao et al. 2022a; Gagnon-Audet et al. 2022). For instance, banks assess whether a person is likely to default on a loan by examining factors such as income, career, and marital status. However, as society changes over time, it is desirable to take account of predictable trends along with the time (Bai, Ling, and Zhao 2022) (i.e., concept shift) when making predictions for new customers. Moreover, these factors,




Figure 1: An example of Evolving Domain Generalization on Portraits (Ginosar et al. 2015). The dataset consists of historical images of US high school students, and as time progresses, the visual attributes captured in the photos, such as hair type and clothing style, gradually change.

e.g., income and career types, will also change gradually due to social developments (i.e., covariate shift). Figure 1 shows another example of historical images of US high school students. Besides temporal factors, the data collected can also evolve along with geometry, transformation, and other factors (Qin, Wang, and Li 2022). The task of generalizing under such shifts, known as *Evolving Domain Generalization* (EDG), involves training models on examples from sequential source domains to generalize well to evolving unseen target domains (Wang et al. 2022; Qin, Wang, and Li 2022).

To tackle EDG, a common belief is to inherit the spirit of domain invariant learning and to learn evolving domain *invariant features* for generalizing to new domains (Gagnon-Audet et al. 2022). In contrast, plentiful empirical studies challenged the position and showed that it is also needed to additionally learn the dynamic information which is useful for predicting evolving domain patterns (Wang et al. 2022; Bai, Ling, and Zhao 2022). These phenomena raise a challenging research question:

What features do we need to learn for successful EDG?

We consider the problem from the feature learning perspective. Specifically, we build Structural Causal Models to characterize the distribution shifts caused by the underlying invariant and dynamic factors in EDG. Our causal analysis shows that it is essential to learn both dynamic and invariant features for better EDG.

Nevertheless, it is non-trivial to learn both invariant and

dynamic features within a single model due to their conflicts (Qin, Wang, and Li 2022). To this end, we propose a principled EDG method called Mutual Information-Based Sequential Autoencoders (MISTS). Specifically, MISTS adopts a variational inference strategy to identify the underlying invariant and dynamic features. To encourage more complete separation between the invariant and dynamic latent representations, MISTS employs a novel information-theoretic objective that minimizes the mutual information between them. We theoretically show the new objective of MISTS is a valid evidence lower bound (ELBO) of the data log-likelihood with respect to our causal models. Our major contributions are as follows:

- We provide theoretical evidence showing that only learning either learning invariant or dynamic features is insufficient for EDG.
- We then propose a novel framework MISTS to extract invariant and dynamic representations simultaneously and separately.
- We conduct extensive experiments on various EDG benchmarks. The results confirm that learning both invariant and dynamic features in MISTS provides a better generalization ability on unseen evolving domains.

Related Work

Domain Generalization

A rich literature is dedicated to addressing the OOD generalization challenge, which often involves introducing additional regularizations to Empirical Risk Minimization (ERM) (Vapnik 1991). Researchers such as Ganin et al. (2016); Sun and Saenko (2016); Li et al. (2018); Dou et al. (2019); Chen et al. (2023b) have explored regularization of learned features to be **domain-invariant**, while others such as Namkoong and Duchi (2016); Hu et al. (2018); Sagawa et al. (2020) have focused on regularizing models to be **robust to mild distributional perturbations** in the training distributions. Similarly, researchers such as Zhang et al. (2022b); Liu et al. (2021); Zhang et al. (2022a); Yao et al. (2022b); Chen et al. (2022) have proposed improving robustness with additional assumptions. Recently, there has been a growing interest in adopting causality theory (Pearl 2009; Schölkopf et al. 2021) and introducing causal invariance to representation learning (Peters, Bühlmann, and Meinshausen 2016; Arjovsky et al. 2019; Creager, Jacobsen, and Zemel 2021; Parascandolo et al. 2021; Wald et al. 2021; Ahuja et al. 2021; Chen et al. 2023a). These approaches require the learned representation to be **causally invariant**, such that a predictor acting on minimizes the risks of all environments simultaneously. In addition, approaches such as Koyama and Yamaguchi (2020); Krueger et al. (2021); Shi et al. (2022); Rame, Dancette, and Cord (2021) have implemented invariance by **encouraging agreements** at various levels across environments. The aforementioned studies focus on extracting the invariant features across multiple domains while disregarding other features. However, Yao et al. (2022a) provide extensive empirical evidence showing that the existing invariant learning methods may not be

suitable for non-stationarity environments. This work focuses on OOD generalization under non-stationarity environments.

Besides, some methods are proposed to utilize domain-dependent features for better OOD generalization ability, e.g., (Chattopadhyay, Balaji, and Hoffman 2020; Zhang et al. 2021; Bui et al. 2021; Zhang et al. 2023). Yet, they still treat domain index as a **discrete variable**, and can not learn the evolving drift across the domains.

Evolving Domain Generalization

Recently, many works have been proposed to tackle the challenging Evolving Domain Generalization or Temporal Domain Generalization, which can be further divided into two categories. The first line of work is primarily inspired by domain-dependent methods. For example, Wang et al. (2022) propose to learn a transformation between domains via meta-learning. However, they assumed that the sequential domains **evolve consistently**, i.e., there exists an explicit function transformations between the domains, which does not always hold for real-world datasets. Nasery et al. (2021) introduce a temporal DG algorithm with gradient interpolation (GI) that trains models to predict near-future data by learning how the activation function evolves over time. Since GI only focuses on the change of activation functions, GI has **limited power** in characterizing model dynamics (Bai, Ling, and Zhao 2022). Motivated by this, Bai, Ling, and Zhao (2022) further propose to use dynamic graphs over a recurrent structure to capture the evolving dynamics of model parameter distributions. However, the method requires adjusting the model parameters by updating the weights of edges between neurons as a graph, which can be **computationally expensive** for large datasets and models. Meanwhile, Qin, Wang, and Li (2022) propose LSSAE to model the underlying latent variables in data sample space. LSSAE disentangles the latent variables into invariant variables and dynamic variables. However, LSSAE still focuses on using the **invariant features** for prediction. In contrast, we incorporate a more rigorous analysis to demonstrate the importance of incorporating dynamic features for classification.

Continuous Domain Adaptation

The issue of continuous domain adaptation, or evolving domain adaptation, has garnered increasing attention in recent years. Various CDA methods have been developed, such as (Hoffman, Darrell, and Saenko 2014; Wang, He, and Katabi 2020; Liu et al. 2020; Lao et al. 2020). Meanwhile, some intermediate-domain-based approaches (Kumar, Ma, and Liang 2020; Chen and Chao 2021; Wang, Li, and Zhao 2022) are also known as gradual domain adaptation. However, these methods require data samples from target or intermediate domains for adaptation. Our focus is on the domain generalization task in evolving settings, where no information from target domains is accessible during training.




Figure 2: The directed acyclic graph depicting our generative model. Dashed lines indicate the causal direction is possible for either side.

Background and Motivation

Problem Definition

We consider the domain generalization tasks in which we have a sequence of evolving source domains $\mathcal{S} = \{\mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_T\}$, where each domain $\mathcal{D}_t = \{(x_i^{(t)}, y_i^{(t)})\}_{i=1}^{n_t}$, $x_i^{(t)} \in \mathcal{X}$, $y_i^{(t)} \in \mathcal{Y}$ and n_t denote the input data, label and data size of domain $t \in \{1, 2, \dots, T\}$, respectively. Our goal is to train a model $f_\theta : \mathcal{X} \rightarrow \mathcal{Y}$ on source domains \mathcal{S} to predict well on the evolving unseen target domains $\mathcal{T} = \{\mathcal{D}_{T+1}, \mathcal{D}_{T+2}, \dots\}$. For clarity, we will omit the index i and t for notations involving only a single data point or domain. EDG assumes that the domain distribution is changing following some sequential patterns (Wang et al. 2022; Qin, Wang, and Li 2022; Bai, Ling, and Zhao 2022).

Feature Learning for Successful EDG

To study what features the model should learn for successful EDG, we first build the structural causal models to characterize the distribution shifts in EDG, shown in Figure. 2. The label $y \in \{1, -1\}$ is randomly sampled from the uniform distribution at the label space. Then, the label y further controls the generation of the latent variable $z = [z_c, z_t]^T$ with respect to domain t , where z is composed of a domain invariant part z_c and a dynamic part z_t :

$$z_c \sim \mathcal{N}(y \cdot \mu_c, \sigma_c^2 I), \quad z_t \sim \mathcal{N}(y \cdot \mu_t, \sigma_t^2 I), \quad (1)$$

where $\mu_c \in \mathbb{R}^{d_c}$ and $\mu_t \in \mathbb{R}^{d_t}$ are the mean of latent invariant and dynamic features, respectively. μ_t is generated conditioned on previous $\mu_{<t}$, and the condition relationship can be parameterized by neural networks. Furthermore, the latent features z_c and z_t control the generation of invariant and dynamic patterns of the input feature x via an injective function of the latent features z_c and z_t , i.e., $x = g(z_c, z_t)$. Although we have presented the model as the distribution of z_c and z_t conditioned on y , the causal directions can be viewed either way. In general, DG frameworks suppose to learn a featurizer $\varphi : \mathcal{X} \rightarrow \mathcal{Z}$, such that there exists a classifier acting on $w : \mathcal{Z} \rightarrow \mathcal{Y}$ such that:

$$\sum_{t \in \mathcal{T}} \mathcal{L}_t(y, w \circ \varphi(x)), \quad (2)$$

where \mathcal{L}_t represents logistic or 0-1 loss. Prior invariant methods are designed to achieve robust performance on target domains by ignoring non-invariant features, i.e., by using an invariant classifier $[2\mu_c/\sigma_c^2, 0]^T$ in combination with $[z_c, 0]$ (Arjovsky et al. 2019; Rosenfeld, Ravikumar, and Risteski 2020). However, in the context of EDG, non-invariant features have a meaningful correlation with the label. Restricting the model to z_c will further limit its generalization ability in EDG. The theoretical results are informally presented in Theorem 1. Details are deferred to the Appendix.

Theorem 1 (Informal) *In the linear setting of Eq. 1, for any domain t , there exists a classifier w_t acting on z_c and z_t that achieves a lower risk than the optimal classifier w_* acting on $[z_c, 0]$.*

Theorem 1 implies that additionally using dynamic features can achieve better OOD generalization ability, which explains the success of EDG using dynamic features (Bui et al. 2021). Instead of learning sole invariant or dynamic features, Theorem 1 demonstrates that it is essential to learn both features for successful EDG.

Harnessing Dynamic Features

To effectively utilize dynamic features, we aim to find a feature space that distinguishes between z_c and z_t , enabling the model to learn the evolving pattern of z_t . One common approach is to employ a multi-classification head, as demonstrated in (Bui et al. 2021), to capture invariant and domain-specific features, respectively. However, the multi-classification head treats domains as discrete indices, hindering the ability to learn dynamic patterns. Furthermore, this method struggles to cleanly separate z_c and z_t .

In summary, there are several challenges in harnessing z_t . The first challenge is **how to capture z_t and learn the evolving dynamics from $z_{1:T}$** . Drawing inspiration from the literature on Sequential Autoencoders, we consider employing a probabilistic framework that utilizes variational inference to identify the latent structures of z_c and z_t . However, a **purely sequential autoencoder fails to cleanly separate z_c and z_t** , leading to inaccurate learning of the evolving dynamics. Therefore, we propose an information-theoretical regularizer to minimize the mutual information between z_c and z_t . Additionally, **the conditional probability $P(Y|z_c, z_t)$ varies across different domains**, making it difficult for a stationary classifier to account for this drift. Therefore, to address these variations, we employ an adaptive classifier w_t that operates on top of $C(z_c, z_t)$, where C represents a combination function. In practical implementations, the combination function is often simplified by using concatenation.

Method

In this section, we will introduce our proposed method Mutual Information-Based Sequential Autoencoders, which jointly extract *both* the invariant and dynamic features and train domain-adaptive classifiers onto the extracted features to achieve better generalization in EDG.

Probabilistic Modeling

Specifically, we can define the following probabilistic generative model for the input data of all source domains as

$$\begin{aligned} p(x_{1:T}, y_{1:T}, z_c, z_{1:T}, w_{1:T}) \\ = p(x_{1:T}, z_{1:T}, z_c) p(y_{1:T}, w_{1:T} | z_{1:T}, z_c). \end{aligned} \quad (3)$$

where x in Eq. 3 separates the generation of $(x_{1:T}, y_{1:T})$ into two parts via the decomposition of $z_{1:T}$ into the invariant feature z_c and the dynamic feature z_t . Modeling the first term $p(x_{1:T}, z_{1:T}, z_c)$ elicits proper disentanglement of z_c and z_t at the latent space. Modeling the second term $p(y_{1:T}, w_{1:T} | z_{1:T}, z_c)$ is essentially to leverage the disentangled features to predict the labels. In what follows, we will detail how to model the two terms for better EDG.

Identify Latent Representations

To model $p(x_{1:T}, z_{1:T}, z_c)$ for disentanglement, we can further decompose the first term by the chain rule of probability, as the following:

$$\begin{aligned} p(x_{1:T}, z_{1:T}, z_c) &= p(z) p(x_{1:T} | z) \\ &= \left[p(z_c) \prod_{t=1}^T p(z_t | z_{<t}) \right] \cdot \prod_{t=1}^T p(x_t | z_t, z_c). \end{aligned} \quad (4)$$

Eq. 4 shows that the generation process of domains data x_t at timestamp t depends on the corresponding dynamic latent representation z_t and invariant representation z_c . Moreover, the distribution of z_t is conditional on $z_{<t}$ from the history domains. For instance, z_c can represent the facial contour while z_t denotes the visual attributes which follow fashion trends change over time.

Our objective here is to extract the latent representations given only the observed data $x_{1:T}$. Therefore, we hope to learn a posterior distribution as follows:

$$q(z_{1:T}, z_c | x_{1:T}) = q(z_c | x_{1:T}) \prod_{t=1}^T q(z_t | z_{<t}, x_t), \quad (5)$$

which has a corresponding variational lower bound as follows:

$$\max_{p,q} \mathbb{E}_q [\log p(x_{1:T} | z) - KL[q(z | x_{1:T}) || p(z)]], \quad (6)$$

where p and q can be parameterized by recurrent neural network, and KL denotes KL divergence. The first term denotes the reconstruction term for input data x_t , the second and third terms denote KL divergence which is to align the posterior distributions z_c and z_t with the corresponding prior distributions.

Although forcing the dynamic variable z_t to predict dynamic factors can guarantee that z_t contains adequate dynamic information, it fails to guarantee that z_t excludes the static information z_c . For example, if the learned dynamic factors \hat{z}_t encompasses z_c , the invariant factor \hat{z}_c will lose some invariant information while $(\hat{z}_c, \hat{z}_{1:T})$ still reconstructs the data samples equally well as $(z_c, z_{1:T})$ up to Eq. 6. Consequently, such disentanglement results can lead to poor performance in the downstream classification tasks (Ilse et al.

2020; Qin, Wang, and Li 2022). To solve the problem, Qin, Wang, and Li (2022) used a temporal domain constraint to limit the information carried by z_t . However, the constraints will harm the encoding of z_t , and fail to capture complex dynamics.

From the information-theoretic perspective, the problem can be considered as clearly separating z_c and z_t , under the constraint that z_t and z_c should contain useful information from x_t in a specified domain (Han et al. 2021; Akuzawa, Iwasawa, and Matsuo 2021; Bai, Wang, and Gomes 2021). Specifically, we propose to additionally minimize the mutual information between z_t and z_c while maximizing the mutual information between the latent representation and the data observed, such that the z_t and z_c carry mutually exclusive meaningful information on input data:

$$\begin{aligned} \max_{p,q} \mathbb{E}_q [\log p(x_{1:T} | z) - KL[q(z | x_{1:T}) || p(z)]] \\ \text{s.t. } p, q \in \arg \min I_q(z_t; z_c), \\ p, q \in \arg \max I_q(z_t; x_t), \end{aligned} \quad (7)$$

where the mutual information terms are defined as:

$$\begin{aligned} I_q(z_c; x_t) &= \mathbb{E}_{q(z_c, x_t)} [\log q(z_c | x_t) - \log q(z_c)], \\ I_q(z_t; x_t) &= \mathbb{E}_{q(z_t, x_t)} [\log q(z_t | x_t) - \log q(z_t)]. \end{aligned} \quad (8)$$

$I_q(z_t; z_c)$ is defined similarly, and more details can be found in the Appendix. To address the constrained optimization problem stated in Equation 7, we demonstrate that it can be transformed into the optimization of a novel valid Evidence Lower Bound (ELBO) for the data log-likelihood. We provide a theoretical guarantee for this transformation, ensuring the reliability of the approach.

Theorem 2 *Let the mutual information (MI) between z_c, z_t and x_t be Eq. 8 in terms of q , and $p(z_t) = p(z_t | z_{<t}, x_t)$, then solving Problem. 7 can be converted into an evidence lower bound of $\mathbb{E}_{x_{1:T} \sim p_D} \log(x_{1:T})$ as follows:*

$$\max_{p,q} \mathbb{E}_q [\log p(x_{1:T} | z) - KL[q(z) || p(z)]],$$

In practice, we follow the spirit of (Higgins et al. 2016; Bai, Wang, and Gomes 2021; Qin, Wang, and Li 2022) to optimizing the following objective:

$$\begin{aligned} \mathcal{L}_e &= \sum_{t=1}^T \mathbb{E}_{z_c, z_t} [\log p(x_t | z_c, z_t) \\ &\quad - \alpha(KL[q(z_c | x_{1:T}) || p(z_c)] \\ &\quad + KL[q(z_t | z_{<t}, x_t) || p(z_t | p(z_{<t}))]) \\ &\quad + \beta(I_q(z_c; x_t) + I_q(z_t; x_t) - I_q(z_t; z_c))], \end{aligned} \quad (9)$$

Domain Adaptive Classifier

To model $p(y_{1:T}, w_{1:T} | z_{1:T}, z_c)$ for classification, as discussed in Sec. , the classifier is expected to evolve accordingly, e.g. the magnitude of model parameter weights for some features will evolve gradually. We use w_t to represent the classifier for domain t and consider w_t as a latent variable in the category space, inspired by Qin, Wang, and Li

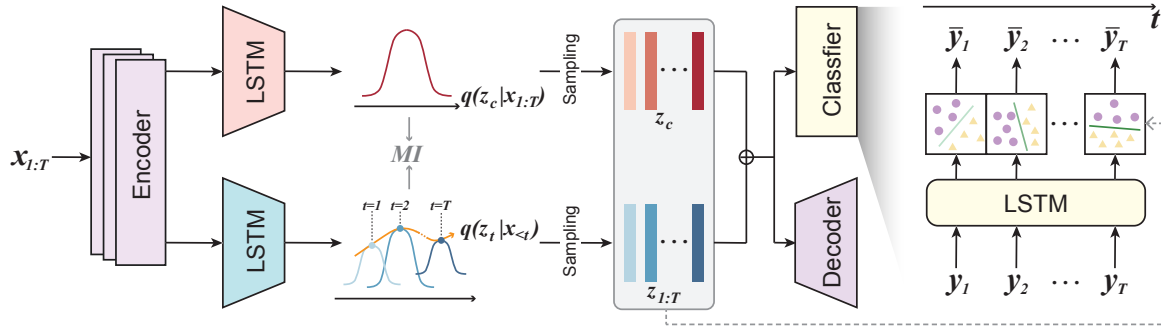


Figure 3: The MISTS framework starts by encoding the input data into the latent space, where two distinct LSTMs parameterize the corresponding posteriors to obtain the invariant and dynamic latent representations, denoted as z_c and $z_{1:T}$, respectively. These representations are then passed through a decoder and an adaptive classifier to compute the reconstruction loss and classification loss, respectively. To encourage better disentanglement, mutual information (MI) terms are applied to the invariant, dynamic latent variables, and input data. The KL-divergence terms are omitted for simplicity.

(2022). According to Eq. 3, the adaptive classifier w_t can be decomposed as:

$$p(y_{1:T}, w_{1:T} | z_{1:T}, z_c) = \prod_{t=1}^T p(w_t | w_{<t}) p(y_t | z_t, z_c, w_t),$$

where w_t can be inferred by classifiers from previous domains, and predict labels \hat{y}_t given z_c and z_t . To model the dynamic variables w_t , we parameterize $p(w_t | w_{<t})$ as a learnable categorical distribution. To optimize the prior distribution $p(w_t | w_{<t})$, the objective function can also be derived based on ELBO:

$$\mathcal{L}_c = \lambda \sum_{t=1}^T \mathbb{E}_{q(w_t | w_{<t}, y_t)} [\log p(y_t | z_c, z_t, w_t) - \alpha KL[q(w_t | w_{<t}, y_t), p(w_t | w_{<t})]], \quad (10)$$

where $q(w_t | w_{<t}, y_t)$ denotes the posterior distribution and $p(y_t | z_c, z_t, w_t)$ denotes the classification loss. Given training data \mathcal{S} , our proposed framework can be optimized through the objective function $\mathcal{L}_{\text{MISTS}} = \mathcal{L}_e + \mathcal{L}_c$. To guarantee the efficacy of our method proposed, we provide the following proposition:

Proposition 3 *Given the probabilistic generative model defined for the joint distribution of all source domains as in Eq. 3, $\mathcal{L}_{\text{MISTS}}$ is equivalent to the ELBO of the data log-likelihood $\log(x_{1:T}, y_{1:T})$ on source domains.*

Proposition 3 demonstrates that MISTS can effectively identify the latent representations of interest and model their relationship with the labels within a given dataset.

Model Instantiations and Implementation

The implementation of network architecture for MISTS is depicted in Fig. 3. It is composed of two parts: (1) encoders to extract z_c and $z_{1:T}$ (2) adaptive classifier on top of the two latent representations.

Encoder. The encoder module consists of a base-encoder θ to learn latent representation from the data samples x_t , and two LSTMs κ_c and π_q to extract z_c and z_t based on $q(z_c | x_{1:T})$ and $q(z_t | z_{<t}, x_t)$, respectively. We set the corresponding prior distribution $p(z_c)$ to be the standard Gaussian, and $p(z_t | z_{<t})$ is instantiated by π_p , which share the same architectures with π_q .

Decoder. The decoder D takes the concatenation of z_c and z_t as input and output the reconstructed data samples \hat{x}_t . For the classification module, the prior network τ_p for $p(w_t | w_{<t})$ is a LSTM network with a categorical distribution as the output, which generates linear classifiers at domain t based on the previous sequence w_t . The corresponding posterior network τ_q for $q(z_t^v | z_{<t}^v, y_t)$ share a similar structure while additionally taking the one-hot code of label y_t as the input. Finally, we can generate the classification results \bar{y}_t by using w_t and the combination of z_c and z_t . The detailed optimization procedure can be found in Algorithm 1.

Inference. To predict the label of x_{T+1} sampled from the following target domains in \mathcal{D}_{T+1} , we adopt π_q and τ_p to infer z_{T+1} and w_{T+1} and use κ_c to extract the latent representation z_c . Then, we can have the prediction results \bar{y}_{T+1} . Repeating this process, we can generalize our model into future target domains, e.g., $\mathcal{D}_{T+2}, \mathcal{D}_{T+3}$.

Experiments

Experimental Setup

To evaluate the effectiveness of MISTS, we conducted experiments on both synthetic and real-world datasets, following the setting of LSSAE (Qin, Wang, and Li 2022). Specifically, we compared our approach with invariant learning methods and the state-of-the-art EDG methods on three synthetic datasets (Circle, Sine and Rotated MNIST) and three real-world datasets (Portraits, Caltran, and Elec). We also evaluated the results on one additional variant, Sine-C, which was created for EDG settings by Qin, Wang, and Li

Algorithm 1: Optimization procedure of MISTS

```

1: Input: Source labeled datasets  $\mathcal{S}$  with  $T$  domains;
   Training epochs  $E$ ; Batch Size  $B$ .
2: Randomly initialize  $\theta, \kappa_c, \pi_p, \pi_q, \tau_p, \tau_q, D$ 
3: Assign  $\mathbf{z}_0, \mathbf{w}_0 \leftarrow \mathbf{0}$ 
4: for  $t = 1, 2, \dots, E$  do
5:   for  $i = 1, 2, \dots, T$  do
6:     Sample a batch  $B$  of data  $(x_t, y_t)$  from  $\mathcal{D}_t$ 
7:     Compute batch loss  $\mathcal{L}_e$  according to Eq. 9 for
        $\kappa_c, \pi_p, \pi_q$  and  $D$ 
8:     Compute batch loss  $\mathcal{L}_c$  according to Eq. 10 for
        $\tau_p, \tau_q$ 
9:     Update all modules by the total loss  $\mathcal{L}_{\text{MISTS}}$ 
10:   end for
11: end for

```

(2022). The domains were split into the source, intermediate, and target domains with a ratio of $1/2 : 1/6 : 1/3$, with the intermediate domains used as the validation set.

Datasets. We briefly introduce these datasets here, and leave more details on the dataset and baselines in the Appendix. (1) The **Circle** dataset (Pesaranghader and Viktor 2016) includes 30 evolving domains, where data points are sampled from 30 2D Gaussian distributions. For **Circle-C**, concept shift is introduced by gradually changing the center and radius of the decision boundary over time. (2) The **Sine** dataset (Pesaranghader and Viktor 2016) is extended to 24 evolving domains by rearranging it. To test models whether can adapt to sudden change, the labels of **Sine-C** are reversed (i.e., from 0 to 1 or from 1 to 0) from the 6th domain to the last one. (3) The **Rotated MNIST** (RMNIST) dataset (Ghifary et al. 2015) consists of MNIST digits with varying degrees of rotation. (4) The **Portraits** dataset (Ginosar et al. 2015) contains photos of high-school seniors from the 1900s to the 2000s for gender classification. We split the dataset into 34 domains based on a fixed interval over time. (5) The **Caltran** dataset (Hoffman, Darrell, and Saenko 2014) is a real-world surveillance dataset comprising images captured from a fixed traffic camera deployed in an intersection. It involves predicting the type of scene based on continuously evolving data. We divide it into 34 domains based on different times. (6) The **Elec** dataset (Dau et al. 2019) is designed for the time-section prediction of current power supply based on the hourly records of an Italian electricity company. The concept shift may arise from changes in season, weather, or price. We split it into 30 domains based on days.

Baselines. We compare MISTS with the state-of-the-art EDG methods, i.e., LSSAE and DRAIN. Additionally, we choose several representative methods from 3 main categories: classical supervised learning, continual learning, and invariant learning. **Classical Supervised Learning:** ERM (Vapnik 1991). **Continual Learning:** EWC (Kirkpatrick et al. 2017) and SI (Zenke, Poole, and Ganguli 2017). **Invariant learning:** (5) IRM (Arjovsky et al. 2019), CORAL (Sun and Saenko 2016), Mixup (Yan et al. 2020)




Figure 4: The visualization presents the decision boundaries for the Sine and Sine-C datasets. In the Sine dataset’s ground truth, positive and negative labels are denoted by green and yellow dots, respectively. Figures (b-d) illustrate the prediction results on the Sine dataset, obtained through the ERM, LSSAE, and MISTS methods, respectively.

and LISA (Yao et al. 2022b). We implemented these baselines according to (Yao et al. 2022a).

Quantitative results. The results of our proposed MISTS and the baseline methods are presented in Table 1. The experimental results reveal the significant improvement of EDG (Evolving Domain Generalization) methods over traditional DG (Domain Generalization) methods. This finding aligns with our theoretical results as well as the empirical observations in the existing EDG literature. Additionally, the substantial performance gap observed between continual learning methods (such as SI and EWC) and EDG methods highlights the importance of effectively leveraging historical knowledge to learn evolutionary patterns, which is not taken into account by methods like SI and EWC.

Notably, MISTS outperforms other EDG methods in terms of average accuracy across the seven datasets. These results highlight the importance of a clear separation of dynamic and invariant features and incorporating both features for successful EDG. By considering the evolving nature of the data, MISTS demonstrates improved performance and shows promise for enabling better adaptation to changing environments from a feature learning perspective.

Qualitative results. To assess the effectiveness of our method, we visualize the decision boundaries of our proposed approach along with two baselines, ERM and LSSAE, on the Sine and Sine-C datasets. The visualization results are presented in Figure 4. For better visualization, we apply certain smoothing and augmentation techniques to the challenging Sine-C curve, which does not affect our results.

As depicted in the figure, all methods demonstrate a good fit to the source domains (the left half of the figures). However, unlike ERM, which only adapts to the source domains, both LSSAE and MISTS exhibit desirable generalization capabilities to unseen target domains. This observation validates the effectiveness of our EDG approach in capturing underlying patterns that evolve across domains, resulting in improved performance. It is worth noting that LSSAE struggles to recover part of the desired decision boundary at the unobserved timestamps compared to MISTS. This limitation

Algorithm	Circle	Sine	Sine-C	Elec	RMNIST	Portraits	Caltran	Avg
ERM	49.3 \pm 2.1	62.7 \pm 1.1	62.3 \pm 1.2	70.8 \pm 0.6	41.6 \pm 0.7	87.9 \pm 1.4	61.2 \pm 2.3	62.3
IRM	53.6 \pm 2.7	62.3 \pm 0.9	59.4 \pm 0.7	70.5 \pm 0.1	40.2 \pm 0.4	87.2 \pm 1.7	63.8 \pm 0.8	62.4
Mixup	48.7 \pm 1.5	62.2 \pm 0.7	62.0 \pm 1.5	70.2 \pm 0.9	42.1 \pm 0.4	87.8 \pm 1.0	66.1 \pm 1.0	62.7
LISA	48.4 \pm 1.1	61.7 \pm 0.5	61.3 \pm 1.2	70.1 \pm 1.0	41.9 \pm 0.8	88.0 \pm 0.9	66.3 \pm 1.1	62.5
CORAL	51.2 \pm 4.2	58.7 \pm 1.9	60.2 \pm 2.1	70.1 \pm 0.7	42.3 \pm 0.8	86.1 \pm 1.9	65.1 \pm 1.5	62.0
GroupDRO	54.3 \pm 3.4	59.3 \pm 0.2	59.2 \pm 4.1	68.6 \pm 0.6	42.8 \pm 1.1	84.3 \pm 1.8	63.6 \pm 1.3	61.7
EWC	59.1 \pm 3.9	65.4 \pm 4.8	64.2 \pm 4.3	69.5 \pm 1.1	32.2 \pm 5.7	88.7 \pm 1.2	60.1 \pm 3.2	62.7
SI	58.3 \pm 2.7	69.7 \pm 3.9	63.6 \pm 4.7	69.8 \pm 1.3	31.4 \pm 5.3	88.7 \pm 0.8	59.4 \pm 3.7	63.0
DRAIN	54.5 \pm 4.0	69.5 \pm 2.7	71.0 \pm 1.3	71.1 \pm 0.7	44.2 \pm 1.1	-	-	-
LSSAE	63.5 \pm 4.7	68.4 \pm 3.5	63.6 \pm 3.7	71.0 \pm 0.5	45.3 \pm 1.4	88.9 \pm 1.5	68.8 \pm 3.4	67.0
MISTS	62.7 \pm 2.7	78.2 \pm 3.3	70.1 \pm 3.2	71.4 \pm 0.8	47.5 \pm 1.3	89.2 \pm 1.3	70.2 \pm 2.1	69.9

Table 1: Experimental Results (Accuracy %) on Synthetic and Real-World Datasets across Different Methods.

may stem from LSSAE’s sole reliance on an adaptive classifier built on top of invariant features, potentially leading to a loss of dynamic information. Additionally, all methods show potential for improvement when faced with abrupt changes.

Overall, the visualization results highlight the superiority of MISTS in terms of generalization to unseen target domains, confirming its ability to capture and leverage evolving patterns for enhanced performance.

Ablation Study

In this subsection, we conduct ablation studies to assess the individual contributions of different components in MISTS, using the RMNIST dataset. We systematically remove one or several components during training to create new variants of MISTS and evaluate their performance. The results are as summarized in Table 2.

To begin, we investigate the significance of leveraging additional information for improved generalization. The results from variants A and B indicate a significant degradation in performance compared to the original MISTS. This observation highlights the indispensability of both z_c and z_t components for achieving promising results.

Furthermore, we explore the impact of clean separation on the efficiency of utilizing z_c and z_t . Variant C exhibits a significant performance drop compared to Variant A due to the absence of Mutual-Information terms, suggesting that a clean separation is crucial for optimal utilization of these components. Conversely, Variant D performs slightly better than the original MISTS, indicating that it can leverage some invariant information during the extraction of dynamic features, without a clean separation. Finally, Variant E demonstrates that the absence of an adaptive classifier leads to performance similar to that of the ERM (Empirical Risk Minimization) baseline. This finding suggests that the adaptive classifier plays a vital role in enhancing the model’s performance, distinguishing it from the conventional ERM approach.

Overall, these ablation studies provide valuable insights into the importance and interplay of different components within MISTS, shedding light on the factors that contribute

	z_c	z_t	w_t	MI	Accuracy
Variant A	✓	✗	✓	✓	45.7 \pm 1.1
Variant B	✗	✓	✓	✓	10.2 \pm 0.2
Variant C	✓	✗	✓	✗	44.3 \pm 1.4
Variant D	✗	✓	✓	✗	11.4 \pm 0.4
Variant E	✓	✓	✗	✓	42.1 \pm 1.3
MISTS	✓	✓	✓	✓	47.5 \pm 1.3

Table 2: Ablation study of MISTS on dataset RMNIST.

to its superior performance compared to the variants and baselines examined.

Conclusion

This paper introduces a novel approach, Mutual Information Based Sequential Autoencoders (MISTS), for addressing the challenges of Evolving Domain Generalization (EDG). Our empirical findings demonstrate the significant impact of MISTS on enhancing the performance of downstream classification tasks within the EDG framework. The theoretical analysis further supports the effectiveness of our method in modeling the joint distribution of data and labels under evolving distribution shifts.

Future Work. While our work makes significant contributions to the understanding and advancement of Evolving Domain Generalization (EDG), it is crucial to acknowledge its limitations. One important aspect to consider is the varying importance of invariant and dynamic features across different datasets as their distributions change. Therefore, it becomes essential to find a method that can automatically adapt and adjust the relative importance of these features.

Overall, As the first to focus on simultaneously learning dynamic and invariant features, we aspire to inspire further research in exploring the essential factors that impact the performance of EDG models and effectively leveraging these factors.

Acknowledgements

We thank reviewers and meta-reviewers for their valuable comments. This work was supported by CUHK direct grant 4055146. BH was supported by the NSFC Young Scientists Fund No. 62006202, NSFC General Program No. 62376235, Guangdong Basic and Applied Basic Research Foundation No. 2022A1515011652, and HKBU Faculty Niche Research Areas No. RC-FNRA-IG/22-23/SCI/04.

References

- Ahuja, K.; Caballero, E.; Zhang, D.; Gagnon-Audet, J.-C.; Bengio, Y.; Mitliagkas, I.; and Rish, I. 2021. Invariance Principle Meets Information Bottleneck for Out-of-Distribution Generalization. In *Conference on Neural Information Processing Systems*.
- Akuzawa, K.; Iwasawa, Y.; and Matsuo, Y. 2021. Information-theoretic regularization for learning global features by sequential VAE. *Machine Learning*.
- Arjovsky, M.; Bottou, L.; Gulrajani, I.; and Lopez-Paz, D. 2019. Invariant Risk Minimization. *arXiv preprint arXiv:1907.02893*.
- Bai, G.; Ling, C.; and Zhao, L. 2022. Temporal Domain Generalization with Drift-Aware Dynamic Neural Networks. *arXiv preprint arXiv:2205.10664*.
- Bai, J.; Wang, W.; and Gomes, C. P. 2021. Contrastively disentangled sequential variational autoencoder. In *Conference on Neural Information Processing Systems*.
- Bengio, Y.; Deleu, T.; Rahaman, N.; Ke, R.; Lachapelle, S.; Bilaniuk, O.; Goyal, A.; and Pal, C. 2019. A meta-transfer objective for learning to disentangle causal mechanisms. *arXiv preprint arXiv:1901.10912*.
- Bui, M.-H.; Tran, T.; Tran, A.; and Phung, D. 2021. Exploiting domain-specific features to enhance domain generalization. In *Conference on Neural Information Processing Systems*.
- Chattopadhyay, P.; Balaji, Y.; and Hoffman, J. 2020. Learning to balance specificity and invariance for in and out of domain generalization. In *European Conference on Computer Vision*.
- Chen, H.-Y.; and Chao, W.-L. 2021. Gradual Domain Adaptation without Indexed Intermediate Domains. In *Conference on Neural Information Processing Systems*.
- Chen, Y.; Bian, Y.; Zhou, K.; Xie, B.; Han, B.; and Cheng, J. 2023a. Does Invariant Graph Learning via Environment Augmentation Learn Invariance? In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Chen, Y.; Zhang, Y.; Bian, Y.; Yang, H.; Kaili, M.; Xie, B.; Liu, T.; Han, B.; and Cheng, J. 2022. Learning Causally Invariant Representations for Out-of-Distribution Generalization on Graphs. In Koyejo, S.; Mohamed, S.; Agarwal, A.; Belgrave, D.; Cho, K.; and Oh, A., eds., *Advances in Neural Information Processing Systems*.
- Chen, Y.; Zhou, K.; Bian, Y.; Xie, B.; Wu, B.; Zhang, Y.; KAILI, M.; Yang, H.; Zhao, P.; Han, B.; and Cheng, J. 2023b. Pareto Invariant Risk Minimization: Towards Mitigating the Optimization Dilemma in Out-of-Distribution Generalization. In *International Conference on Learning Representations*.
- Creager, E.; Jacobsen, J.; and Zemel, R. S. 2021. Environment Inference for Invariant Learning. In *International Conference on Machine Learning*.
- Dau, H. A.; Bagnall, A.; Kamgar, K.; Yeh, C.-C. M.; Zhu, Y.; Gharghabi, S.; Ratanamahatana, C. A.; and Keogh, E. 2019. The UCR time series archive. *IEEE/CAA Journal of Automatica Sinica*.
- Dou, Q.; de Castro, D. C.; Kamnitsas, K.; and Glocker, B. 2019. Domain Generalization via Model-Agnostic Learning of Semantic Features. In *Conference on Neural Information Processing Systems*.
- Gagnon-Audet, J.-C.; Ahuja, K.; Darvishi-Bayazi, M.-J.; Dumas, G.; and Rish, I. 2022. WOODS: Benchmarks for Out-of-Distribution Generalization in Time Series Tasks. *arXiv preprint arXiv:2203.09978*.
- Ganin, Y.; Ustinova, E.; Ajakan, H.; Germain, P.; Larochelle, H.; Laviolette, F.; Marchand, M.; and Lempitsky, V. S. 2016. Domain-Adversarial Training of Neural Networks. *Journal of Machine Learning Research*.
- 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*.
- Ginossar, 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. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*.
- Han, J.; Min, M. R.; Han, L.; Li, L. E.; and Zhang, X. 2021. Disentangled Recurrent Wasserstein Autoencoder. In *International Conference on Learning Representations*.
- Higgins, I.; Matthey, L.; Pal, A.; Burgess, C.; Glorot, X.; Botvinick, M.; Mohamed, S.; and Lerchner, A. 2016. beta-vaes: Learning basic visual concepts with a constrained variational framework. In *International Conference on Learning Representations*.
- Hoffman, J.; Darrell, T.; and Saenko, K. 2014. Continuous manifold based adaptation for evolving visual domains. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Hu, W.; Niu, G.; Sato, I.; and Sugiyama, M. 2018. Does Distributionally Robust Supervised Learning Give Robust Classifiers? In *International Conference on Machine Learning*.
- Ilse, M.; Tomczak, J. M.; Louizos, C.; and Welling, M. 2020. Diva: Domain invariant variational autoencoders. In *Medical Imaging with Deep Learning*.
- Kirkpatrick, J.; Pascanu, R.; Rabinowitz, N.; Veness, J.; Desjardins, G.; Rusu, A. A.; Milan, K.; Quan, J.; Ramalho, T.; Grabska-Barwinska, A.; et al. 2017. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*.
- Koyama, M.; and Yamaguchi, S. 2020. Out-of-Distribution Generalization with Maximal Invariant Predictor. *arXiv preprint arXiv:2008.01883*.

- Krueger, D.; Caballero, E.; Jacobsen, J.; Zhang, A.; Binas, J.; Zhang, D.; Priol, R. L.; and Courville, A. C. 2021. Out-of-Distribution Generalization via Risk Extrapolation (REx). In *International Conference on Machine Learning*.
- Kumar, A.; Ma, T.; and Liang, P. 2020. Understanding self-training for gradual domain adaptation. In *International Conference on Machine Learning*.
- Lao, Q.; Jiang, X.; Havaei, M.; and Bengio, Y. 2020. Continuous domain adaptation with variational domain-agnostic feature replay. *arXiv preprint arXiv:2003.04382*.
- Li, Y.; Tian, X.; Gong, M.; Liu, Y.; Liu, T.; Zhang, K.; and Tao, D. 2018. Deep Domain Generalization via Conditional Invariant Adversarial Networks. In *European Conference on Computer Vision*.
- Liu, E. Z.; Haghighi, B.; Chen, A. S.; Raghunathan, A.; Koh, P. W.; Sagawa, S.; Liang, P.; and Finn, C. 2021. Just Train Twice: Improving Group Robustness without Training Group Information. In *International Conference on Machine Learning*.
- Liu, H.; Long, M.; Wang, J.; and Wang, Y. 2020. Learning to Adapt to Evolving Domains. In *Conference on Neural Information Processing Systems*.
- Namkoong, H.; and Duchi, J. C. 2016. Stochastic Gradient Methods for Distributionally Robust Optimization with f-divergences. In *Conference on Neural Information Processing Systems*.
- 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. In *Conference on Neural Information Processing Systems*.
- Parascandolo, G.; Neitz, A.; Orvieto, A.; Gresele, L.; and Schölkopf, B. 2021. Learning explanations that are hard to vary. In *International Conference on Learning Representations*.
- Pearl, J. 2009. *Causality*. Cambridge University Press, 2 edition.
- Pesaranghader, A.; and Viktor, H. L. 2016. Fast Hoeffding Drift Detection Method for Evolving Data Streams. In *European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases*.
- Peters, J.; Bühlmann, P.; and Meinshausen, N. 2016. Causal inference by using invariant prediction: identification and confidence intervals. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*.
- Qin, T.; Wang, S.; and Li, H. 2022. Generalizing to Evolving Domains with Latent Structure-Aware Sequential Autoencoder. In *International Conference on Machine Learning*.
- Rame, A.; Dancette, C.; and Cord, M. 2021. Fishr: Invariant Gradient Variances for Out-of-distribution Generalization. *arXiv preprint arXiv:2109.02934*.
- Rosenfeld, E.; Ravikumar, P.; and Risteski, A. 2020. The risks of invariant risk minimization. *arXiv preprint arXiv:2010.05761*.
- Sagawa, S.; Koh, P. W.; Hashimoto, T. B.; and Liang, P. 2020. Distributionally Robust Neural Networks. In *International Conference on Learning Representations*.
- Schölkopf, B.; Locatello, F.; Bauer, S.; Ke, N. R.; Kalchbrenner, N.; Goyal, A.; and Bengio, Y. 2021. Toward Causal Representation Learning. *Proceedings of the Institute of Electrical and Electronics Engineers*.
- Shi, Y.; Seely, J.; Torr, P.; N, S.; Hannun, A.; Usunier, N.; and Synnaeve, G. 2022. Gradient Matching for Domain Generalization. In *International Conference on Learning Representations*.
- Sun, B.; and Saenko, K. 2016. Deep CORAL: Correlation Alignment for Deep Domain Adaptation. In *European Conference on Computer Vision*.
- Vapnik, V. 1991. Principles of Risk Minimization for Learning Theory. In *Conference on Neural Information Processing Systems*.
- Wald, Y.; Feder, A.; Greenfeld, D.; and Shalit, U. 2021. On Calibration and Out-of-Domain Generalization. In *Conference on Neural Information Processing Systems*.
- Wang, H.; He, H.; and Katabi, D. 2020. Continuously Indexed Domain Adaptation. In *Proceedings of the 37th International Conference on Machine Learning*.
- Wang, H.; Li, B.; and Zhao, H. 2022. Understanding gradual domain adaptation: Improved analysis, optimal path and beyond. In *International Conference on Machine Learning*.
- Wang, W. W.; Xu, G.; Pu, R.; Li, J.; Zhou, F.; Shui, C.; Ling, C.; Gagné, C.; and Wang, B. 2022. Evolving Domain Generalization. *arXiv preprint arXiv:2206.00047*.
- Yan, S.; Song, H.; Li, N.; Zou, L.; and Ren, L. 2020. Improve unsupervised domain adaptation with mixup training. *arXiv preprint arXiv:2001.00677*.
- Yao, H.; Choi, C.; Cao, B.; Lee, Y.; Koh, P. W.; and Finn, C. 2022a. Wild-time: A benchmark of in-the-wild distribution shift over time. *arXiv preprint arXiv:2211.14238*.
- Yao, H.; Wang, Y.; Li, S.; Zhang, L.; Liang, W.; Zou, J.; and Finn, C. 2022b. Improving Out-of-Distribution Robustness via Selective Augmentation. In *International Conference on Machine Learning*.
- Zenke, F.; Poole, B.; and Ganguli, S. 2017. Continual learning through synaptic intelligence. In *International conference on machine learning*.
- Zhang, D.; Chen, M.; Li, C.; Huang, L.; and Zhang, J. 2023. Aggregation of disentanglement: Reconsidering domain variations in domain generalization. *arXiv preprint arXiv:2302.02350*.
- Zhang, M.; Marklund, H.; Dhawan, N.; Gupta, A.; Levine, S.; and Finn, C. 2021. Adaptive risk minimization: Learning to adapt to domain shift. *Advances in Neural Information Processing Systems*, 34: 23664–23678.
- Zhang, M.; Sohoni, N. S.; Zhang, H. R.; Finn, C.; and Ré, C. 2022a. Correct-N-Contrast: a Contrastive Approach for Improving Robustness to Spurious Correlations. In *International Conference on Machine Learning*.
- Zhang, Y.; Gong, M.; Liu, T.; Niu, G.; Tian, X.; Han, B.; Schölkopf, B.; and Zhang, K. 2022b. CausalAdv: Adversarial Robustness through the Lens of Causality. In *International Conference on Learning Representations*.