

ADARNN: Adaptive Learning and Forecasting for Time Series*

Yuntao Du¹, Jindong Wang², Wenjie Feng³, Sinno Pan⁴, Tao Qin², Renjun Xu⁵, Chongjun Wang¹

¹Nanjing University, Nanjing, China ²Microsoft Research Asia, Beijing, China

³Institute of Data Science, National University of Singapore ⁴Nanyang Technological University ⁵Zhejiang University
dz1833005@smail.nju.edu.cn, jindong.wang@microsoft.com

ABSTRACT

Though time series forecasting has a wide range of real-world applications, it is a very challenging task. This is because the statistical properties of a time series can vary with time, causing the distribution to change temporally, which is known as the distribution shift problem in the machine learning community. By far, it still remains unexplored to model time series from the distribution-shift perspective. In this paper, we formulate the Temporal Covariate Shift (TCS) problem for the time series forecasting. We propose Adaptive RNNs (ADARNN) to tackle the TCS problem. ADARNN is sequentially composed of two modules. The first module is referred to as Temporal Distribution Characterization, which aims to better characterize the distribution information in a time series. The second module is termed as Temporal Distribution Matching, which aims to reduce the distribution mismatch in the time series to learn an RNN-based adaptive time series prediction model. ADARNN is a general framework with flexible distribution distances integrated. Experiments on human activity recognition, air quality prediction, household power consumption, and financial analysis show that ADARNN outperforms some state-of-the-art methods by 2.6% in terms of accuracy on classification tasks and by 9.0% in terms of the mean squared error on regression tasks. We also show that the temporal distribution matching module can be extended to the Transformer architecture to further boost its performance.

CCS CONCEPTS

• **Computing methodologies** → **Transfer learning**: *Artificial intelligence*.

KEYWORDS

Time series, multi-task learning, transfer learning

ACM Reference Format:

Yuntao Du¹, Jindong Wang², Wenjie Feng³, Sinno Pan⁴, Tao Qin², Renjun Xu⁵, Chongjun Wang¹. 2021. ADARNN: Adaptive Learning and Forecasting for Time Series. In *Proceedings of the 30th ACM Int'l Conf. on Information and Knowledge Management (CIKM '21)*, November 1–5, 2021, Virtual Event, Australia. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/XXXXXX.XXXXXX>

*Corresponding author: J. Wang and W. Feng. This work is done when Y. Du was an intern at MSRA.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](https://permissions.acm.org).

CIKM '21, November 1–5, 2021, Virtual Event, Australia.

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8446-9/21/11...\$15.00

<https://doi.org/10.1145/XXXXXX.XXXXXX>




Figure 1: The temporal covariate shift problem in non-stationary time series. The raw time series data (x) is multivariate in reality. At time intervals A, B, C and the unseen test data, the distributions are different: $P_A(x) \neq P_B(x) \neq P_C(x) \neq P_{Test}(x)$. With the distribution changing over time, how to build an accurate and adaptive model?

1 INTRODUCTION

Time series (TS) data occur naturally in countless domains including financial analysis [55], medical analysis [29], weather condition prediction [46], and renewable energy production [5]. Forecasting is one of the most sought-after tasks on analyzing time series data (arguably the most difficult one as well) due to its importance in industrial, society and scientific applications. For instance, given the historical air quality data of a city for the last five days, how to predict the air quality in the future more accurately?

In real applications, it is natural that the statistical properties of TS are changing over time, i.e., the non-stationary TS. Over the years, various research efforts have been made for building reliable and accurate models for the non-stationary TS. Traditional approaches such as hidden Markov models (HMMs) [43], dynamic Bayesian networks [35], Kalman filters [9], and other statistical models (e.g. ARIMA [19]), have made great progress. Recently, better performance is achieved by the recurrent neural networks (RNNs) [37, 46]. RNNs make no assumptions on the temporal structure and can find highly non-linear and complex dependence relation in TS.

The non-stationary property of time series implies the *data distribution changes over time*. Given the example in Figure 1, data distributions $P(x)$ vary for different intervals A, B , and C where x, y are samples and predictions respectively; Especially for the test data which is unseen during training, its distribution is also different from the training data and makes the prediction more exacerbated. The conditional distribution $P(y|x)$, however, is usually considered to be unchanged for this scenario, which is reasonable in many real applications. For instance, in stock prediction, it is natural that the market is fluctuating which changes the financial factors

($P(x)$), while the economic laws remain unchanged ($P(y|x)$). It is an inevitable result that the afore-mentioned methods have inferior performance and poor generalization [22] since the distribution shift issue violates their basic I.I.D assumption.

Unfortunately, it remains unexplored to model the time series from the distribution perspective. The main challenges of the problem lies in two aspects. First, how to characterize the distribution in the data to maximally harness the common knowledge in these varied distributions? Second, how to invent an RNN-based distribution matching algorithm to maximally reduce their distribution divergence while capturing the temporal dependency?

In this paper, we formally define the *Temporal Covariate Shift (TCS)* in Figure 1, which is a more practical and challenging setting to model time series data. Based on our analysis on TCS, we propose ADARNN, a novel framework to learn an accurate and adaptive prediction model. ADARNN is composed of two modules. Firstly, to better characterize the distribution information in TS, we propose a *temporal distribution characterization (TDC)* algorithm to split the training data into K most diverse periods that are with large distribution gap inspired by the principle of maximum entropy. After that, we propose a *temporal distribution matching (TDM)* algorithm to dynamically reduce distribution divergence using a RNN-based model. Experiments on activity recognition, air quality prediction, household power consumption and stock price prediction show that our ADARNN outperforms the state-of-the-art baselines by 2.6% in terms of an accuracy on classification tasks and by 9.0% in terms of RMSE on regression tasks. ADARNN is agnostic to both RNN structures (i.e., RNNs, LSTMs, and GRUs) and distribution matching distances (e.g., cosine distance, MMD [3], or adversarial discrepancy [12]). ADARNN can also be extended to the Transformer architecture to further boost its performance.

To sum up, our main contributions are as follows,

- **Novel problem:** For the first time, we propose to model the time series from the distribution perspective, then we postulate and formulate the Temporal Covariate Shift (TCS) problem in non-stationary time series, which is more realistic and challenging.
- **General framework:** To solve TCS, we propose a general framework ADARNN that learns an accurate and adaptive model by proposing the temporal distribution characterization and Temporal Distribution matching algorithms.
- **Effectiveness:** We conduct extensive experiments on human activity recognition, air quality prediction, household power consumption, and stock price prediction datasets. Results show that ADARNN outperforms the state-of-the-art baselines on both classification and regression tasks.

2 RELATED WORK

2.1 Time series analysis

Traditional methods for time series classification or regression include distance-based [13, 32], feature-based [38], and ensemble methods [27]. The distance-based methods measure the distance and similarity of (segments of) raw series data with some metric, like the Euclidean distance or Dynamic time wrapping (DTW) [20]. The feature-based methods capture the global/local patterns in the time series rely on a set of manually extracted or learned features.

The ensemble methods combine multiple weak classifiers to boost the performance of the models.

However, involving heavy crafting on data pre-processing and labor-intensive feature engineering makes those methods struggling to cope with large-scale data and the performance is also limited for more complex patterns. Recurrent Neural Networks (RNNs), such as Gated Recurrent Unit (GRU) and Long Short Term Memory (LSTM) have been popular due to their ability to extract high quality features automatically and handle the long term dependencies. These deep methods tackle TS classification or forecasting by leveraging attention mechanisms [23, 33] or tensor factorization [40] for capturing shared information among series. Another trend is to model uncertainty by combining deep learning and State space models [37]. Furthermore, some methods [24, 37, 46] adopt seq2seq models for multi-step predictions. Unlike those methods based on statistical views, our ADARNN models TS from the perspective of the distributions. Recently, the Transformers structure [45] is proposed for sequence learning with the self-attention mechanism. The vanilla Transformer can be easily modified for time series prediction, while it still fails for distribution matching.

Time series segmentation and clustering are two similar topics to our temporal distribution characterization. Time series segmentation [6, 28] aim to separate time series into several pieces that can be used to discover the underlying patterns; it mainly uses change-point detection algorithms, include sliding windows, bottom-up, and top-down methods. The methods in segmentation almost don't utilize the distribution matching scheme and can't adapt to our problem. Time-series clustering [16, 34], on the other hand, aims at finding different groups consisting of similar (segments of) time series, which is clearly different from our situation.

2.2 Distribution matching

When the training and test data are coming from different distributions, it is common practice to adapt some domain adaptation (DA) algorithms to bridge their distribution gap, such that domain-invariant representations can be learned. DA often performs instance re-weighting or feature transfer to reduce the distribution divergence in training and test data [12, 44, 47, 48, 56]. Similar to DA in general methodology, domain generalization (DG) [49] also learns a domain-invariant model on multiple source domains in the hope that it can generalize well to the target domain [2, 26, 30]. The only difference between DA and DG is that DA assumes the test data is accessible while DG does not. It is clear that our problem settings are different from DA and DG. Firstly, DA and DG does not built for temporal distribution characterization since the domains are given a prior in their problems. Secondly, most DA and DG methods are for classification tasks using CNNs than RNNs. Thus, they might fail to capture the long-range temporal dependency in time series. Several methods applied pre-training to RNNs [4, 8, 11, 15, 51] which is similar to CNN-based pre-training.

3 PROBLEM FORMULATION

PROBLEM 1 (r -STEP AHEAD PREDICTION). *Given a time series of n segments $\mathcal{D} = \{\mathbf{x}_i, \mathbf{y}_i\}_{i=1}^n$, where $\mathbf{x}_i = \{x_i^1, \dots, x_i^{m_i}\} \in \mathbb{R}^{p \times m_i}$ is*

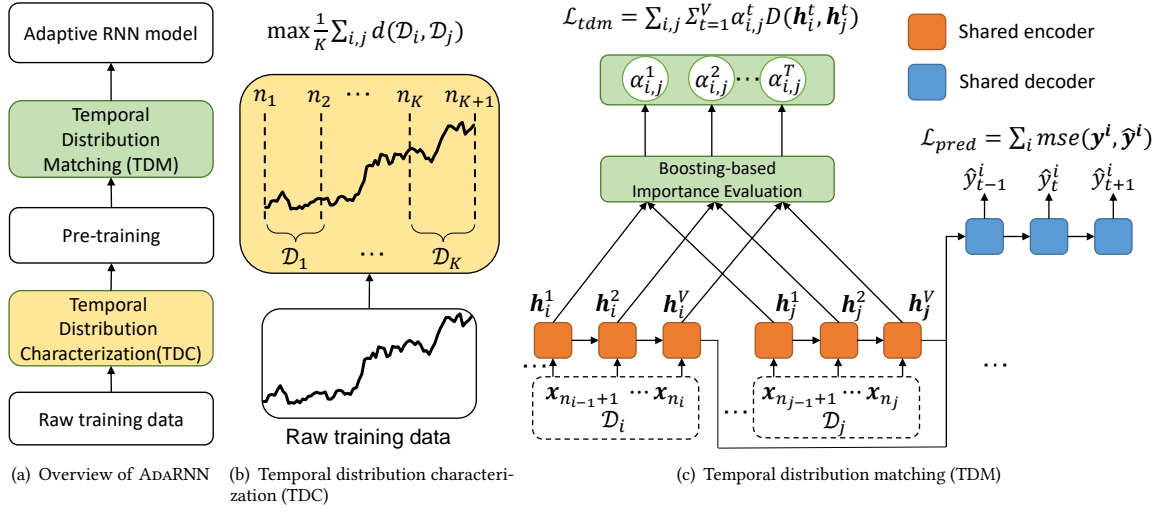


Figure 2: The framework of ADARNN.

a p -variate segment of length m_i^1 , and $\mathbf{y}_i = (y_i^1, \dots, y_i^c) \in \mathbb{R}^c$ is the corresponding label.

Learn a precise prediction model $M : \mathbf{x}_i \rightarrow \mathbf{y}_i$ on the future $r \in \mathbb{N}^+$ steps² for segments $\{\mathbf{x}_j\}_{j=n+1}^{n+r}$ in the same time series.

In most existing prediction approaches, all time series segments, $\{\mathbf{x}_i\}_{i=1}^{n+r}$, are assumed to follow the same data distribution; a prediction model learned from all labeled training segments are supposed to perform well on the unlabeled testing segments. However, this assumption is too strong to hold in many real-world applications. In this work, we consider a more practical problem setting: the distributions of the training segments and the test segments can be different, and the distributions of the training segments can also be changed over time as shown in Figure 1. In the following, we formally define our problem.

DEFINITION 1 (COVARIATE SHIFT [41]). Given the training and the test data from two distributions $P_{train}(\mathbf{x}, y), P_{test}(\mathbf{x}, y)$, covariate shift is referred to the case that the marginal probability distributions are different, and the conditional distributions are the same, i.e., $P_{train}(\mathbf{x}) \neq P_{test}(\mathbf{x})$, and $P_{train}(y|\mathbf{x}) = P_{test}(y|\mathbf{x})$.

Note that the definition of covariate shift is for non-time series data. We now extend its definition for time series data as follows, which is referred to as *temporal covariate shift* (TCS) in this work.

DEFINITION 2 (TEMPORAL COVARIATE SHIFT). Given a time series data \mathcal{D} with n labeled segments. Suppose it can be split into K periods or intervals, i.e., $\mathcal{D} = \{\mathcal{D}_1, \dots, \mathcal{D}_K\}$, where $\mathcal{D}_k = \{\mathbf{x}_i, \mathbf{y}_i\}_{i=n_k+1}^{n_{k+1}}$, $n_1 = 0$ and $n_{k+1} = n$. Temporal Covariate Shift (TCS) is referred to the case that all the segments in the same period i follow the same data distribution $P_{\mathcal{D}_i}(\mathbf{x}, y)$, while for different time periods $1 \leq i \neq j \leq K$, $P_{\mathcal{D}_i}(\mathbf{x}) \neq P_{\mathcal{D}_j}(\mathbf{x})$ and $P_{\mathcal{D}_i}(y|\mathbf{x}) = P_{\mathcal{D}_j}(y|\mathbf{x})$.

To learn a prediction model with good generalization performance under TCS, a crucial research issue is to capture the *common knowledge* shared among different periods of \mathcal{D} [22]. Take the stock price prediction as an example. Though the financial factors ($P(x)$)

vary with the change of market (i.e., different distributions in different periods of time), some common knowledge, e.g., the economic laws and patterns ($P(y|x)$) in different periods, can still be utilized to make precise predictions over time.

However, the number of periods K and the boundaries of each period under TCS are usually unknown in practice. Therefore, before designing a model to capture commonality among different periods in training, we need to first discover the periods by comparing their underlying data distributions such that segments in the same period follow the same data distributions. The formal definition of the problem studied in this work is described as follows.

PROBLEM 2 (r -STEP TS PREDICTION UNDER TCS). Given a time series of n segments with corresponding labels, $\mathcal{D} = \{\mathbf{x}_i, \mathbf{y}_i\}_{i=1}^n$, for training. Suppose there exist an unknown number of K underlying periods in the time series, so that segments in the same period i follow the same data distribution $P_{\mathcal{D}_i}(\mathbf{x}, y)$, while for different time periods $1 \leq i \neq j \leq K$, $P_{\mathcal{D}_i}(\mathbf{x}) \neq P_{\mathcal{D}_j}(\mathbf{x})$ and $P_{\mathcal{D}_i}(y|\mathbf{x}) = P_{\mathcal{D}_j}(y|\mathbf{x})$. Our goal is to automatically discover the K periods in the training time series data, and learn a prediction model $M : \mathbf{x}_i \rightarrow \mathbf{y}_i$ by exploiting the commonality among different time periods, such that it makes precise predictions on the future r segments, $\mathcal{D}_{test} = \{\mathbf{x}_j\}_{j=n+1}^{n+r}$. Here we assume that the test segments are in the same time period, $P_{\mathcal{D}_{test}}(\mathbf{x}) \neq P_{\mathcal{D}_i}(\mathbf{x})$ and $P_{\mathcal{D}_{test}}(y|\mathbf{x}) = P_{\mathcal{D}_i}(y|\mathbf{x})$ for any $1 \leq i \leq K$.

Note that learning the periods underlying the training time series data is challenging as both the number of periods and the corresponding boundaries are unknown, and the search space is huge. Afterwards, with the identified periods, the next challenge is how to learn a good generalizable prediction model M to be used in the future periods. If the periods are not discovered accurately, it may affect the generalization performance of the final prediction model.

4 OUR PROPOSED MODEL: ADARNN

In this section, we present our proposed framework, termed Adaptive RNNs (ADARNN), to solve the temporal covariate shift problem. Figure 2(a) depicts the overview of ADARNN. As we can see, it mainly consists of two novel algorithms:

¹In general, m_i may be different for different time segments.

²When $r = 1$, it becomes the one-step prediction problem.

- *Temporal Distribution Characterization (TDC)*: to characterize the distribution information in TS;
- *Temporal Distribution Matching (TDM)*: to match the distributions of the discovered periods to build a time series prediction model \mathcal{M} .

Given the training data, ADARNN first utilizes TDC to split it into the periods which fully characterize its distribution information. It then applies the TDM module to perform distribution matching among the periods to build a generalized prediction model \mathcal{M} . The learned \mathcal{M} is finally used to make r -step predictions on new data.

The rationale behind ADARNN is as follows. In TDC, the model \mathcal{M} is expected to work under the *worst* distribution scenario where the distribution gaps among the different periods are large, thus the optimal split of periods can be determined by *maximizing their dissimilarity*. In TDM, \mathcal{M} utilizes the common knowledge of the learned time periods by matching their distributions via RNNs with a regularization term to make precise future predictions.

4.1 Temporal distribution characterization

Based on *the principle of maximum entropy* [17], maximizing the utilization of shared knowledge underlying a times series under temporal covariate shift can be done by finding periods which are most dissimilar to each other, which is also considered as the worst case of temporal covariate shift since the cross-period distributions are the most diverse. Therefore, as Figure 2(b) shows, TDC achieves this goal for splitting the TS by solving an optimization problem whose objective can be formulated as:

$$\begin{aligned} \max_{0 < K \leq K_0} \max_{n_1, \dots, n_K} \frac{1}{K} \sum_{1 \leq i \neq j \leq K} d(\mathcal{D}_i, \mathcal{D}_j) \\ \text{s.t. } \forall i, \Delta_1 < |\mathcal{D}_i| < \Delta_2; \sum_i |\mathcal{D}_i| = n \end{aligned} \quad (1)$$

where d is a distance metric, Δ_1 and Δ_2 are predefined parameters to avoid trivial solutions (e.g., very small values or very large values may fail to capture the distribution information), and K_0 is the hyperparameter to avoid over-splitting. The metric $d(\cdot, \cdot)$ in Eq. (1) can be any distance function, e.g., Euclidean or Editing distance, or some distribution-based distance / divergence, like MMD [14] and KL-divergence. We will introduce how to select $d(\cdot, \cdot)$ in Section 4.3.

The learning goal of the optimization problem (1) is to maximize the averaged period-wise distribution distances by searching K and the corresponding periods so that the distributions of each period are as diverse as possible and the learned prediction model has better a more generalization ability.

We explain in detail for the splitting scheme in (1) with the principle of maximum entropy (ME). Why we need the most dissimilar periods instead of the most similar ones? According to ME, when no prior assumptions are made on the splitting of the time series data, it is reasonable that the distributions of each period become as diverse as possible to maximize the entropy of the total distributions. This allows more general and flexible modeling of the future data. Therefore, since we have no prior information on the test data which is unseen during training, it is more reasonable to train a model at the worst case, which can be simulated using diverse period distributions. If the model is able to learn from the worst

case, it would have better generalization ability on unseen test data. This assumption has also been validated in theoretical analysis of time series models [22, 31] that diversity matters in TS modeling.

In general, the time series splitting optimization problem in (1) is computationally intractable and may not have a closed-form solution. By selecting a proper distance metric, however, the optimization problem (1) can be solved with dynamic programming (DP) [36]. In this work, by considering the scalability and the efficiency issues on large-scale data, we resort to a greedy algorithm to solve (1). Specifically, for efficient computation and avoiding trivial solutions, we evenly split the time series into $N = 10$ parts, where each part is the minimal-unit period that cannot be split anymore. We then randomly search the value of K in $\{2, 3, 5, 7, 10\}$. Given K , we choose each period of length n_j based on a greedy strategy. Denote the start and the end points of the time series by A and B , respectively. We first consider $K = 2$ by choosing 1 splitting point (denote it as C) from the 9 candidate splitting points via maximizing the distribution distance $d(S_{AC}, S_{CB})$. After C is determined, we consider $K = 3$ and use the same strategy to select another point D . A similar strategy is applied to different values of K . Our experiments show that this algorithm can select the more appropriate periods than random splits. Also, it shows that when K is very large or very small, the final performance of the prediction model becomes poor (Section 5.6.1).

4.2 Temporal distribution matching

Given the learned time periods, the TDM module is designed to learn the common knowledge shared by different periods via matching their distributions. Thus, the learned model \mathcal{M} is expected to generalize well on unseen test data compared with the methods which only rely on local or statistical information.

The loss of TDM for prediction, \mathcal{L}_{pred} , can be formulated as:

$$\mathcal{L}_{pred}(\theta) = \frac{1}{K} \sum_{j=1}^K \frac{1}{|\mathcal{D}_j|} \sum_{i=1}^{|\mathcal{D}_j|} \ell(y_i^j, \mathcal{M}(x_i^j; \theta)), \quad (2)$$

where (x_i^j, y_i^j) denotes the i -th labeled segment from period \mathcal{D}_j , $\ell(\cdot, \cdot)$ is a loss function, e.g., the MSE loss, and θ denotes the learnable model parameters.

However, minimizing (2) can only learn the predictive knowledge from each period, which cannot reduce the distribution diversities among different periods to harness the common knowledge. A naïve solution is to match the distributions in each period-pair \mathcal{D}_i and \mathcal{D}_j by adopting some popular distribution matching distance ($d(\cdot, \cdot)$) as the regularizer. Following existing work on domain adaptation [12, 42] that distribution matching is usually performed on high-level representations, we apply the distribution matching regularization term on the final outputs of the cell of RNN models. Formally, we use $\mathbf{H} = \{\mathbf{h}^t\}_{t=1}^V \in \mathbb{R}^{V \times q}$ to denote the V hidden states of an RNN with feature dimension q . Then, the period-wise distribution matching on the final hidden states for a pair $(\mathcal{D}_i, \mathcal{D}_j)$ can be represented as:

$$\mathcal{L}_{dm}(\mathcal{D}_i, \mathcal{D}_j; \theta) = d(\mathbf{h}_i^V, \mathbf{h}_j^V; \theta).$$

Unfortunately, the above regularization term fails to capture the *temporal dependency* of each hidden state in the RNN. As each hidden state only contains *partial* distribution information of an input sequence, *each* hidden state of the RNN should be considered

while constructing a distribution matching regularizer as described in the following section.

4.2.1 Temporal distribution matching. As shown in Figure 2(c), we propose the *Temporal Distribution Matching (TDM)* module to adaptively match the distributions between the RNN cells of two periods while capturing the temporal dependencies. TDM introduces the *importance vector* $\alpha \in \mathbb{R}^V$ to learn the relative importance of V hidden states inside the RNN, where all the hidden states are weighted with a normalized α . Note that for each pair of periods, there is an α , and we omit the subscript if there is no confusion. In this way, we can *dynamically* reduce the distribution divergence of cross-periods.

Given a period-pair $(\mathcal{D}_i, \mathcal{D}_j)$, the loss of temporal distribution matching is formulated as:

$$\mathcal{L}_{tdm}(\mathcal{D}_i, \mathcal{D}_j; \theta) = \sum_{t=1}^V \alpha_{i,j}^t d(\mathbf{h}_i^t, \mathbf{h}_j^t; \theta), \quad (3)$$

where $\alpha_{i,j}^t$ denotes the distribution importance between the periods \mathcal{D}_i and \mathcal{D}_j at state t .

All the hidden states of the RNN can be easily computed by following the standard RNN computation. Denote by $\delta(\cdot)$ the computation of a next hidden state based on a previous state. The state computation can be formulated as

$$\mathbf{h}_i^t = \delta(\mathbf{x}_i^t, \mathbf{h}_i^{t-1}). \quad (4)$$

By integrating (2) and (3), the final objective of temporal distribution matching (one RNN layer) is:

$$\mathcal{L}(\theta, \alpha) = \mathcal{L}_{pred}(\theta) + \lambda \frac{2}{K(K-1)} \sum_{i,j}^{i \neq j} \mathcal{L}_{tdm}(\mathcal{D}_i, \mathcal{D}_j; \theta, \alpha), \quad (5)$$

where λ is a trade-off hyper-parameter. Note that in the second term, we compute the average of the distribution distances of all pairwise periods. For computation, we take a mini-batch of \mathcal{D}_i and \mathcal{D}_j to perform forward operation in RNN layers and concatenate all hidden features. Then, we can perform TDM using (5).

4.2.2 Boosting-based importance evaluation. We propose a *Boosting-based importance evaluation* algorithm to learn α . Before that, we first perform pre-training on the network parameter θ using the fully-labeled data across all periods, i.e., use (2). This is to learn better hidden state representations to facilitate the learning of α . We denote the pre-trained parameter as θ_0 .

With θ_0 , TDM uses a boosting-based [39] procedure to learn the hidden state importance. Initially, for each RNN layer, all the weights are initialized as the same value in the same layer, i.e., $\alpha_{i,j} = \{1/V\}^V$. We choose the cross-domain distribution distance as the indicator of boosting. If the distribution distance of epoch $n+1$ is larger than that of epoch n , we increase the value of $\alpha_{i,j}^{t,(n+1)}$ to enlarge its effect for reducing the distribution divergence. Otherwise, we keep it unchanged. The boosting is represented as:

$$\alpha_{i,j}^{t,(n+1)} = \begin{cases} \alpha_{i,j}^{t,(n)} \times G(d_{i,j}^{t,(n)}, d_{i,j}^{t,(n-1)}) & d_{i,j}^{t,(n)} \geq d_{i,j}^{t,(n-1)} \\ \alpha_{i,j}^{t,(n)} & \text{otherwise} \end{cases} \quad (6)$$

where

$$G(d_{i,j}^{t,(n)}, d_{i,j}^{t,(n-1)}) = (1 + \sigma(d_{i,j}^{t,(n)} - d_{i,j}^{t,(n-1)})) \quad (7)$$

Algorithm 1 ADARNN

Require: A time series dataset \mathcal{D} , λ , max iteration number $Iter$

- 1: Perform temporal distribution characterization by solving (1) to get K periods $\{\mathcal{D}_i\}$'s.
- 2: End-to-end pre-training by minimizing (2) with $\lambda = 0$
- 3: **For** iteration $i = 1$ to $Iter$:
- 4: Learn the value of α and θ by solving (5) and using (6)
- 5: **End For**
- 6: **return** Optimal parameters θ^*, α^* .

is the updating function computed by distribution matching loss at different learning stages. And $d_{i,j}^{t,(n)} = D(\mathbf{h}_i^t, \mathbf{h}_j^t; \alpha_{i,j}^{t,(n)})$ is the distribution distance at time step t in epoch n . $\sigma(\cdot)$ is the sigmoid function. It is easy to see that $G(\cdot, \cdot) > 1$, ensuring the increment of the importance. Similar to naive method, we normalize it as $\alpha_{i,j}^{t,(n+1)} = \frac{\alpha_{i,j}^{t,(n+1)}}{\sum_{j=1}^V \alpha_{i,j}^{t,(n+1)}}$. By (3) and (6), we can learn the value of α .

REMARK 1. *The objective (3) is defined on one RNN layer, it can be easily implemented for multiple layers to achieve multi-layer distribution matching. As for inference, we simply use the optimal network parameter θ^* . The computation of domain-wise distribution distances can also be made efficient following existing work [26].*

REMARK 2. *In fact, there is also a naive approach to acquire α by using another network with weight $\mathbf{W}_{i,j} \in \mathbb{R}^{2q \times V}$ that takes $(\mathbf{H}_i, \mathbf{H}_j)$ as inputs and output α . For instance, α can be learned as $\alpha_{i,j} = g(\mathbf{W}_{i,j} \odot [\mathbf{H}_i, \mathbf{H}_j]^\top)$ where g is an activation and normalization function and \odot is element-wise production. After computing $\alpha_{i,j}$, we normalize it by a softmax function. Unfortunately, this naive approach does not work in reality for two reasons. First, since $\alpha_{i,j}$ and θ are highly correlated and at early training stage, the hidden state representations learned by θ are less meaningful, which will result in insufficient learning of $\alpha_{i,j}$. Second, it is expensive to learn such $\mathbf{W}_{i,j}$ for each domain pair in RNN models.*

4.3 Computation of the distribution distance

ADARNN is agnostic to RNN structures (vanilla RNNs, LSTMs, and GRUs) and the distribution distances, i.e., the function $d(\cdot, \cdot)$. Technically, $d(\cdot, \cdot)$ in TDC and TDM can be the same choice. We adopt several popular distances: cosine distance, MMD [14] and adversarial distance [12]. Also note that the TDC and TDM are currently two separate stages for controlling computation complexity. It shall be a future work to end-to-end optimize these two stages. The complete process of ADARNN is presented in Algorithm 1. The detailed implementations of these distances are presented in Appendix.

5 EXPERIMENTS

5.1 Setup

We mainly use GRUs (Gated Recurrent Units) [7] as the base RNN cells for its high performance in real applications³. We test ADARNN in four real-world datasets as summarized in Table 1.

We compare ADARNN with four categories of methods:

³It is natural to use simple RNNs or LSTMs.

Table 1: Information of the datasets used in this paper.

Dataset	#Train	#Valid	#Test	dim. (p)	Task
UCI activity	6,000	1,352	2,947	128	Classification
Air quality	29,232	2,904	2,832	6	Regression
Electric power	1,235,964	413,280	413,202	7	Regression
Stock price	778,200	271,690	504,870	360	Regression

- Traditional TS modeling methods, including ARIMA [53], FBProphet [10], LightGBM, and GRU [7];
- Latest TS models, including LSTNet [23] and STRIPE [24]⁴;
- Transformer [45] by removing its mask operation;
- Variants of popular domain adaptation / generalization methods and we denote them as “MMD-RNN” and “DANN-RNN”.

Note that the comparison is mainly for the final prediction performance using temporal distribution matching after obtaining the time periods using TDC, which the evaluation of TDC is in Section 5.6.1. Transformer [45] was originally proposed for machine translation. Therefore, we alter it by removing the masks for TS data by following implementation⁵. We use 4 encoder layers of Transformer with 8 heads for self-attention tuned by cross-validation to ensure that it achieves the best performance on the datasets. Existing DA / DG methods cannot be directly applied to our problem, since most of them are for CNN-based classification tasks. Therefore, we use variants of them. Specifically, we use “MMD-RNN” to denote popular MMD-based methods such as MEDA [47, 48] and MMD-AAE [26]. We use “DANN-RNN” to denote popular domain-adversarial training methods such as RevGrad [12] and DAAN [52]. These two distances are widely adopted in recent DA and DG literature. We change their backbone network to be RNNs which are the same as other comparison methods and then add the MMD or adversarial distances to the final output of the RNN layers.

All methods use the same train-valid-test split and periods after TDC and the final results are averaged on five random runs. Optimal hyperparameters are tuned on the validation set. Note that to balance the efficiency and performance in ADARNN, we limit the range of K in $\{2, 3, 5, 7, 10\}$ to acquire the best splitting of periods. We set the pre-training epochs to be 10 in TDM. We use different distance function d and the main results are based on the MMD [14]. Other distances are discussed later. Code of ADARNN is at <https://github.com/jindongwang/transferlearning/tree/master/code/deep/adarnn>.

5.2 Human activity recognition

The UCI smartphone activity [1] is a time-series dataset for classification task. It is collected by thirty volunteers wearing a smartphone to record the accelerometer, gyroscope, and magnetometer signals while performing six different activities (walking, upstairs, downstairs, sitting, standing, and lying). We use 7,352 instances for training and 2,947 instances as the test following [1]. It is clear that training and testing sets have different distributions and TCS exists in the training set due to different activities and users. We firstly clip the data using sliding window with the size of 128. As there are 9 channels for each instance, we regard it a sequence with length of 128 and feature of 9 dimension. For this task, we use accuracy

⁴STRIPE is an improved and stronger version of DILATE [46].

⁵<https://timeseriestransformer.readthedocs.io/en/latest/README.html>

Table 2: Results on UCI time series classification dataset.

Method	ACC	P	R	F1	AUC
LightGBM	84.11	83.73	83.63	84.91	90.23
GRU	85.68	85.62	85.51	85.46	91.33
MMD-RNN	86.39	86.80	86.26	86.38	91.77
DANN-RNN	85.88	85.59	85.62	85.56	91.41
ADARNN	88.44	88.71	88.59	88.63	93.19

(ACC), precision (P), recall (R), F1, and area under curve (AUC) as the evaluation metrics.

The network is composed of two-layer GRUs with the hidden state dimension 32 and a FC layer with output to be $r = 6$. We optimize by Adam optimizer with the learning rate of 0.005 with a batch size 64. We set $\lambda = 0.1$. On this dataset, we did not run the latest methods such as LSTNet and STRIPE as well as FBProphet and ARIMA since they are built for regression tasks. Transformer is not compared since we did not get good performance.

Table 2 shows the results on this classification task. It is obvious that ADARNN achieves the best performance on all evaluation metrics that outperforms the adversarial-based DANN method by 2.56% in accuracy and 3.07% in F1 score. This means that ADARNN is effective in non-stationary time series classification tasks.

5.3 Air quality prediction

The air-quality dataset [54] contains hourly air quality information collected from 12 stations in Beijing from 03/2013 to 02/2017. We randomly chose four stations (Dongsi, Tiantan, Nongzhanguan, and Dingling) and select six features (PM2.5, PM10, SO2, NO2, CO, and O3). Since there are some missing data, we simply fill the empty slots using averaged values. Then, the dataset is normalized before feeding into the network to scale all features into the same range. This process is accomplished by max-min normalization and ranges data between 0 and 1.

The network is composed of two-layer GRUs with the dimension of the hidden state to be 64 and two bottleneck layers which contains two FC layers, and an FC layer with output to be $r = 1$. We set $\lambda = 0.5$. We optimize all the models by Adam optimizer with the learning rate of 0.005, we set the batch size to be 36.

Table 3 presents the results where the last column “ $\Delta(\%)$ ” denotes the average increment over the baseline GRU. It can be seen that the proposed method achieves the best results in all four stations. Our ADARNN significantly outperforms the strong baseline STRIPE [24] by 8.97% in RMSE reduction.

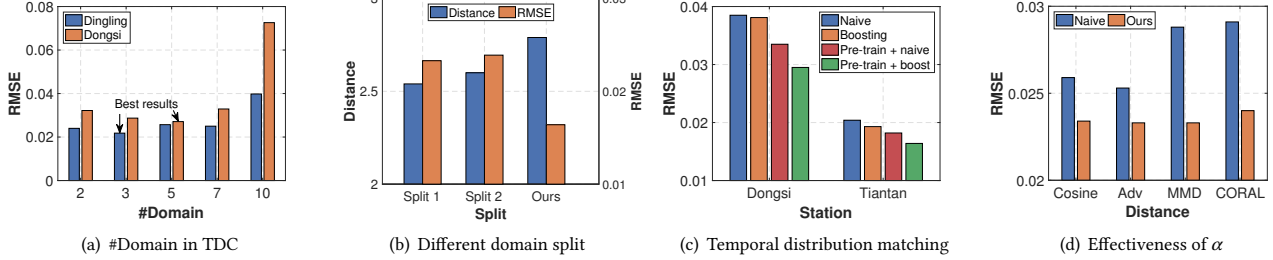
5.4 Household electric power consumption

This dataset [18] contains 2,075,260 measurements gathered between 16/12/2006 and 26/11/2010 with a one-minute sampling rate. It contains several attributes (global active power, global reactive power, voltage, global intensity, and three substrings). We get 2,062,346 measurements for the whole dataset after removing all the null data. We construct the train/valid/test sets with a ratio of 6 : 2 : 2, i.e., train on data from 16/12/2006 to 24/04/2009 and test on data from 25/02/2010 to 26/11/2010 and the rest is valid set. We use the data of 5 consecutive days to predict the value of next day.

The network is composed of two-layer GRUs with the dimension set to be 64 and two bottleneck layers which contains two FC

Table 3: Results on air quality prediction (left) and household electric power consumption prediction (right).

	Dongsi		Tiantan		Nongzhanguan		Dingling		$\Delta(\%)$	Electric Power
	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE		
FBProphet [10]	0.1866	0.1403	0.1434	0.1119	0.1551	0.1221	0.0932	0.0736	-	0.080
ARIMA	0.1811	0.1356	0.1414	0.1082	0.1557	0.1156	0.0922	0.0709	-	-
GRU	0.0510	0.0380	0.0475	0.0348	0.0459	0.0330	0.0347	0.0244	0.00	0.093
MMD-RNN	0.0360	0.0267	0.0183	0.0133	0.0267	0.0197	0.0288	0.0168	-61.31	0.082
DANN-RNN	0.0356	0.0255	0.0214	0.0157	0.0274	0.0203	0.0291	0.0211	-59.97	0.080
LightGBM	0.0587	0.0390	0.0412	0.0289	0.0436	0.0319	0.0322	0.0210	-11.08	0.080
LSTNet [23]	0.0544	0.0651	0.0519	0.0651	0.0548	0.0696	0.0599	0.0705	-	0.080
Transformer [45]	0.0339	0.0220	0.0233	0.0164	0.0226	0.0181	0.0263	0.0163	-61.20	0.079
STRIFE [24]	0.0365	0.0216	0.0204	0.0148	0.0248	0.0154	0.0304	0.0139	-64.60	0.086
ADARNN	0.0295	0.0185	0.0164	0.0112	0.0196	0.0122	0.0233	0.0150	-73.57	0.077

**Figure 3: Ablation study. (a) Analysis of K for temporal distribution characterization. (b) Different domain split strategies. (c) Comparison of w. or w/o. adaptive weights in distribution matching. (d) Effectiveness of temporal distribution matching.**

layers, and a FC layer with output to be $r = 1$. We set $\lambda = 0.001$. We optimize all the models by Adam optimizer with the learning rate of 0.0005 and batch size of 36.

The right part of Table 3 shows the RMSE of predicting the household power consumption using different methods. Again, our ADARNN achieves the best performance.

5.5 Stock price prediction

We further evaluate the performance of ADARNN on a large private finance dataset where the task is to predict the stock price given the historical data from 01/2007 to 12/2019. This dataset contains over 2.8M samples with 360 financial factors as the features. We use the data from 01/2017 to 12/2019 as the test set, and the data from 01/2007 to 12/2014 as the training set with the rest as the validation set. We use information coefficient (IC) and information ratio (IR) as the evaluation metrics following existing work [21]. The detailed description and computation of these metrics are in appendix.

The network is composed of two-layer GRUs with the dimension of hidden state to be 64 and two bottleneck layers which contains one FC layers, and a FC layer with output to be $r = 1$. We optimize all the models by Adam optimizer with the learning rate of 0.0002, we set the batch size to be 800 and $\lambda = 0.5$.

Table 4 shows that our proposed ADARNN outperforms other comparison methods in all metrics. While achieving the best IC values, its RankIC can still be higher than other methods. This means that ADARNN is relatively steady in real applications.

5.6 Ablation study

5.6.1 Temporal distribution characterization. To evaluate TDC, we design two experiments. Firstly, we compare the results of $K = 2, 3, 5, 7, 10$ to analyze the importance of K in Figure 3(a). We see

Table 4: Results on stock price prediction.

Method	IC	ICIR	RankIC	RankICIR	RawIC	RawICIR
LightGBM	0.063	0.551	0.056	0.502	0.064	0.568
GRU	0.106	0.965	0.098	0.925	0.109	0.986
MMD-RNN	0.107	0.962	0.101	0.926	0.108	0.967
DANN-RNN	0.106	0.964	0.103	0.924	0.107	0.966
Transformer [45]	0.104	0.935	0.099	0.920	0.101	0.946
STRIFE [24]	0.108	0.987	0.101	0.946	0.109	0.993
ADARNN	0.115	1.071	0.110	1.035	0.116	1.077

that the performance is changing with the increment of K and the best result is achieved when $K = 3, 5$ for two stations on air quality prediction, respectively. Then, when K continues to increase, the performance drops. It means that K is important in characterizing the distribution information in non-stationary TS.

Secondly, Figure 3(b) shows the comparison of domain split algorithms using the same distribution matching method when $K = 2$. Here, “Split 1” denotes a random split and “Split 2” denotes the split where all periods are with similar distributions (i.e., just the opposite of TDC). The “distance” denotes the distribution distance. The results indicate that our TDC achieves the best performance in RMSE with the largest distribution distance. The results of other values of K are also the same. Therefore, it is important that we split the periods to the *worst* case when the distributions are the most diverse. Both experiments show that TDC could effectively characterize the distribution information in time series. Evaluation of TDC using other distances are in appendix.

5.6.2 Temporal distribution matching. We analyze the TDM module by evaluating the performance of pre-training and the boosting-based method. We compare four variants of TDM: (1) naive approach (i.e., learning α by a feed-forward layer as discussed in remark 2) w/o pre-training, (2) boosting w/o pre-training (3) pre-training + naive, (4) pre-training + boosting (Ours).




Figure 4: (a) Multi-step prediction of air quality. (b) TDM is agnostic to distribution distances in different stations. (c) The change of distribution distance. (d) Convergence analysis. (e) Training time.

Figure 3(c) shows the performance on two stations. We observe that the pre-training technique helps both naïve and boosting-based methods to get better performance. More importantly, our boosting-based method can achieve the best performance with pre-training. Compared with the STRIPE method from Table 3, we see that almost all four variants of TDM can achieve better performance than STRIPE. We also observe that by manually setting the same importance to each hidden state, the results are not comparable. To summarize, it is necessary to consider the importance of each hidden state when trying to build an adaptive RNN model.

To evaluate the importance vector α , we compare the results of with or without importance vector in Figure 3(d). The results indicate that by taking the same distribution distance, our ADARNN achieve the best results. This demonstrates that the different distributions of each hidden state is important to model in RNN-based time series modeling, to which our AdaRNN is capable of dynamically evaluating their importance to achieve the best performance.

5.7 Further analysis

5.7.1 Distribution distance. ADARNN is agnostic to the distribution distance. Here we evaluate this property by running TDM with four different kinds of distribution distances: MMD [3], CORAL [42], Cosine, and domain-adversarial distance [12]. Figure 4(a) presents the results on three stations for air quality prediction. Combining the results from Table 3, we observe that ADARNN is robust to the choice of distribution distance while achieving the best performance. Moreover, Figure 4(b) shows that by adding the importance weights α , our ADARNN achieves lowest distribution distances compared to MMD and adversarial training.

5.7.2 Multi-step prediction. We analyze the performance of ADARNN for multi-step ($r > 1$) air quality prediction and compare with STRIPE [24] by predicting air quality index in next 6, 12, 18, 24 hours (the output sequence length is 6, 12, 18, 24 accordingly). Figure 4(c) shows the results on one station. We see that ADARNN achieves the best performance for multi-step prediction. In addition, as the prediction steps increase (6 \rightarrow 12 \rightarrow 18 \rightarrow 24), the performances of all methods are becoming worse since it is more difficult to model the longer sequences and also affected by the cumulative errors.

5.7.3 Convergence and time complexity analysis. We analyze the convergence of ADARNN by recording its training loss for prediction and distribution matching in Figure 4(d). We see that ADARNN can be easily trained with quick convergence. On the other hand, ADARNN is based on RNNs which are autoregressive models. So it is reasonable that ADARNN takes more time for training than

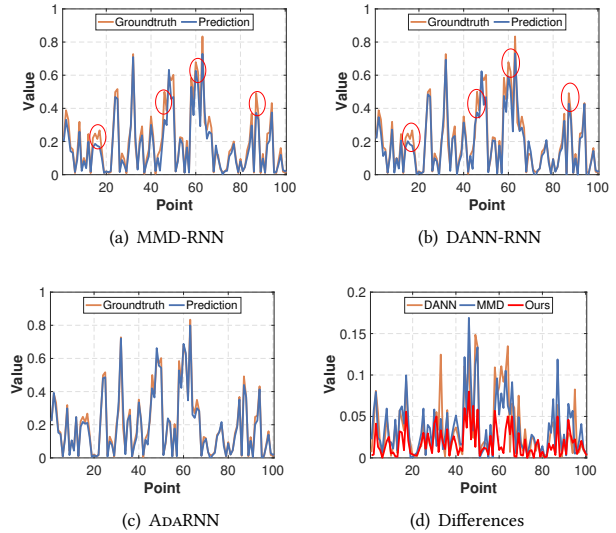


Figure 5: Prediction vs. ground truth.

traditional RNNs. Figure 4(e) shows the training time comparison between ADARNN, GRU, and STRIPE [24]. We see that other than the best performance, the training time of ADARNN is still comparable to existing methods. The inference time of ADARNN is the same as existing RNN-based models since we only need to compute Eq. (2).

5.8 Case study: prediction vs. ground truth

Taking the station Dongsu as an example, we show the prediction vs. ground truth values of different methods in Figure 5, where the red circles denote the bad cases. We see that DANN and MMD achieve worse results than our ADARNN since they have relatively worse predictions in many points. Especially in Figure 5(d), we show the absolute difference between the prediction of each method and the ground-truth. We can clearly see that our ADARNN outperforms all methods with much less mistakes.

6 EXTENSIONS

Apart from RNN-based models, in this section, we add the adaptive importance weights to each hidden representations in the Transformer [45] model, which is claimed more powerful than RNN-based models in the latest studies. We add the importance weights α to the representations of each self-attention block and term this extension as *AdaTransformer*. As shown in Table 5, this pilot study

shows the power of the ADARNN in Transformers. We believe more improvements can be made in the future.

Table 5: Results of vanilla and adaptive transformer

Method	Station 1	Station 2
Vanilla Transformer	0.028	0.035
AdaTransformer	0.025	0.030

7 CONCLUSIONS AND FUTURE WORK

This paper proposed to study the more realistic and challenging temporal covariate shift (TCS) problem for non-stationary time series. We proposed the ADARNN framework to learn an adaptive RNN model for better generalization. ADARNN consists of two novel algorithms: temporal distribution characterization to characterize the distribution information in TS and temporal distribution matching to build the generalized RNN model by distribution matching. Extensive experiments show its effectiveness in three datasets.

In the future, we plan to extend this work from both algorithm and application level. First, explore deeper extension of ADARNN to Transformer for better performance. Second, integrate the two stages of ADARNN in an end-to-end network for easier training.

A DISTRIBUTION DISTANCE

We adopt four widely used distribution distance functions following existing work on DA/DG, i.e., cosine, MMD [3], CORAL [42], and domain adversarial discrepancy [12].

Cosine. Cosine distance is widely used in time series data [25, 50] to measure the similarity between two distributions. Given the data \mathbf{h}_s and \mathbf{h}_t , cosine distance is defined as:

$$d_{\text{cosine}}(\mathbf{h}_s, \mathbf{h}_t) = 1 - \frac{\langle \mathbf{h}_s, \mathbf{h}_t \rangle}{\|\mathbf{h}_s\| \cdot \|\mathbf{h}_t\|}. \quad (8)$$

Maximum Mean Discrepancy (MMD). The MMD [14] is a non-parametric distance between two probability distributions. MMD measures the distance of the mean embeddings of the samples from two distributions in a Reproducing Kernel Hilbert Space (RKHS). MMD can be empirically estimated by

$$d_{\text{mmd}}(\mathbf{h}_s, \mathbf{h}_t) = \frac{1}{n_s^2} \sum_{i,j=1}^{n_s} k(h_{s_i}, h_{s_j}) + \frac{1}{n_t^2} \sum_{i,j=n_s+1}^{n_s+n_t} k(h_{s_i}, h_{s_j}) - \frac{2}{n_s n_t} \sum_{j=n_s+1}^{n_s+n_t} k(h_{s_i}, h_{s_j}), \quad (9)$$

where $k(\cdot, \cdot)$ is the kernel function such as RBF kernel and linear kernel, and $n_s = |\mathbf{h}_s|$, $n_t = |\mathbf{h}_t|$ are the number of the data from two distributions, respectively.

Deep Correlation Alignment (CORAL). The CORAL distance [42] is defined as the distance of the second-order statistic (covariance) of the samples from two distributions:

$$d_{\text{coral}}(\mathbf{h}_s, \mathbf{h}_t) = \frac{1}{4q^2} \|\mathbf{C}_s - \mathbf{C}_t\|_F^2, \quad (10)$$

where q is the dimension of the features, and \mathbf{C}_s and \mathbf{C}_t are the covariance matrices of the two distributions, respectively.

Domain Adversarial Discrepancy. The domain discrepancy can be parameterized as a neural network and it is referred to as the domain adversarial discrepancy [12]. An additional network named

Table 6: Results of different distances in TDC

Split	$ \mathcal{D}_1 / \mathcal{D}_2 $	$d(\text{MMD})$	$d(\text{DTW})$	$d(\text{CORAL})$	$d(\text{Cosine})$	RMSE	MAE
Split1	0.75/0.25	0.898	2.54	9.9999884E-1	0.9748	0.0233	0.0179
Split2	0.25/0.75	0.910	2.60	9.9999902E-1	0.9966	0.0239	0.0197
Ours	0.66/0.33	0.932	2.79	9.9999919E-1	0.9861	0.0164	0.0112

domain discriminator denoted by D is introduced. The domain adversarial objective is defined as:

$$l_{\text{adv}}(\mathbf{h}_s, \mathbf{h}_t) = \mathbb{E}[\log[D(\mathbf{h}_s)]] + \mathbb{E}[\log[1 - D(\mathbf{h}_t)]], \quad (11)$$

where \mathbb{E} denotes expectation. Hence, the adversarial discrepancy is

$$d_{\text{adv}}(\mathbf{h}_s, \mathbf{h}_t) = -l_{\text{adv}}(\mathbf{h}_s, \mathbf{h}_t). \quad (12)$$

B DIFFERENT DISTANCES IN TDC

We analyze temporal distribution characterization (TDC) by using different distance functions. We set $K = 2$ for TDC and compare its performance with one random domain split and the opposite domain split (i.e., with the smaller distribution distance).

The results in Table 6 indicate that our TDC algorithm achieves the best performance in RMSE and MAE with the largest distribution distance under most distance functions. When using cosine distance, our split achieves the second largest distribution distance but it gets the best performance in RMSE and MAE. The results show that TDC could effectively characterize the distribution information and is agnostic to distance functions.

C STOCK PRICE PREDICTION: METRICS

The evaluation metrics for stock price prediction have been widely adopted in previous work [21]. The information coefficient (IC) describes the correlation between predicted and actual stock returns, sometimes used to measure the contribution of a financial analyst. The IC can range from 1.0 to -1.0 . An IC of $+1.0$ indicates a perfect linear relationship between predicted and actual returns, while an IC of 0.0 indicates no linear relationship. An IC of -1.0 indicates that the analyst always fails at making a correct prediction. IC is defined as:

$$IC = \text{corr}(f_{t-1}, r_t), \quad (13)$$

where f_{t-1} is the factor value of the stock in period $t-1$ and r_t is the stock return rate of period t . $\text{corr}(\cdot, \cdot)$ is the correlation functions. Besides, there are also some variants of IC, such as rankIC, which is

$$\text{rankIC} = \text{corr}(\text{order}_{t-1}^f, \text{order}_t^r), \quad (14)$$

where order_{t-1}^f is the factor rank of each stock in period $t-1$ and order_t^r is the ranking of each stock's return rate in period t .

The information ratio (IR) of an actively managed portfolio is defined as the portfolio's mean return over a benchmark portfolio return divided by the tracking error (standard deviation of the excess return), which is defined as

$$IR = IC \times \sqrt{BR}, \quad (15)$$

where BR is breadth or number of independent information sources.

We also call IR the ICIR. Based on IC and ICIR, we further use RankIC and RankICIR to evaluate the robustness of models against permutations, and use RawIC and RawICIR to denote the performance of unnormalized data. To sum up, the higher they are, the better the model is.

REFERENCES

- [1] Bandar Almaslukh, Jalal AlMuhtadi, and Abdelmonim Artoli. 2017. An effective deep autoencoder approach for online smartphone-based human activity recognition. *Int. J. Comput. Sci. Netw. Secur.* 17, 4 (2017), 160–165.
- [2] Yogesh Balaji, Swami Sankaranarayanan, and Rama Chellappa. 2018. Metareg: Towards domain generalization using meta-regularization. In *NeurIPS*. 998–1008.
- [3] Karsten M Borgwardt, Arthur Gretton, Malte J Rasch, Hans-Peter Kriegel, Bernhard Schölkopf, and Alex J Smola. 2006. Integrating structured biological data by kernel maximum mean discrepancy. *Bioinformatics* 22, 14 (2006), e49–e57.
- [4] Lingzhen Chen and Alessandro Moschitti. 2019. Transfer learning for sequence labeling using source model and target data. In *AAAI*, Vol. 33. 6260–6267.
- [5] Edward Choi, Mohammad Taha Bahadori, Jimeng Sun, Joshua Kulas, Andy Schuetz, and Walter Stewart. 2016. Retain: An interpretable predictive model for healthcare using reverse time attention mechanism. In *NeurIPS*. 3504–3512.
- [6] Fu-Lai Chung, Tak-Chung Fu, Vincent Ng, and Robert WP Luk. 2004. An evolutionary approach to pattern-based time series segmentation. *IEEE transactions on evolutionary computation* 8, 5 (2004), 471–489.
- [7] Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. 2014. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555* (2014).
- [8] Wanyun Cui, Guangyu Zheng, Zhiqiang Shen, S. Jiang, and Wei Wang. 2019. Transfer Learning for Sequences via Learning to Collocate. In *ICLR*.
- [9] Emmanuel de Bézenac, Syama Sundar Rangapuram, Konstantinos Benidis, Michael Bohlke-Schneider, Richard Kurl, Lorenzo Stella, Hilaf Hasson, Patrick Gallinari, and Tim Januschowski. 2020. Normalizing Kalman Filters for Multivariate Time Series Analysis. *NeurIPS* 33 (2020).
- [10] Facebook. Accessed in Jan 2021. FBProphet. <https://facebook.github.io/prophet/>.
- [11] Hassan Ismail Fawaz, G. Forestier, Jonathan Weber, L. Idoumghar, and Pierre-Alain Muller. 2018. Transfer learning for time series classification. *Big Data* (2018), 1367–1376.
- [12] Yaroslav Ganin, E. Ustinova, Hana Ajakan, P. Germain, H. Larochelle, F. Laviolette, M. Marchand, and V. Lempitsky. 2016. Domain-Adversarial Training of Neural Networks. *JMLR* 17 (2016), 59:1–59:35.
- [13] Tomasz Górecki and Maciej Luczak. 2015. Multivariate time series classification with parametric derivative dynamic time warping. *Expert Syst. Appl.* 42 (2015), 2305–2312.
- [14] A. Gretton, Bharath K. Sriperumbudur, D. Sejdinovic, Heiko Strathmann, S. Balakrishnan, M. Pontil, and K. Fukumizu. 2012. Optimal kernel choice for large-scale two-sample tests. In *NeurIPS*.
- [15] Priyanka Gupta, P. Malhotra, L. Vig, and G. Shroff. 2018. Transfer Learning for Clinical Time Series Analysis using Recurrent Neural Networks. In *ML for Medicine and Healthcare Workshop at KDD*.
- [16] David Hallac, Sagar Vare, Stephen Boyd, and Jure Leskovec. 2017. Toeplitz inverse covariance-based clustering of multivariate time series data. In *SIGKDD*. 215–223.
- [17] Edwin T Jaynes. 1982. On the rationale of maximum-entropy methods. *Proc. IEEE* 70, 9 (1982), 939–952.
- [18] Kaggle. Accessed in Jan 2021. Household electric power consumption dataset. <https://www.kaggle.com/uciml/electric-power-consumption-data-set>.
- [19] Konstantinos Kalpakis, Dhiral Gada, and Vasundhara Puttagunta. 2001. Distance measures for effective clustering of ARIMA time-series. In *ICDM*. IEEE, 273–280.
- [20] Eamonn Keogh and Chotirat Ann Ratanamahatana. 2005. Exact indexing of dynamic time warping. *Knowl. Inf. Syst.* 7, 3 (2005), 358–386.
- [21] Kazuhiro Kohara, Tsutomu Ishikawa, Yoshimi Fukuhara, and Yukihiro Nakamura. 1997. Stock price prediction using prior knowledge and neural networks. *Intell. Syst. Account. Finance Manag.* 6, 1 (1997), 11–22.
- [22] Vitaly Kuznetsov and Mehryar Mohri. 2014. Generalization bounds for time series prediction with non-stationary processes. In *ALT*. Springer, 260–274.
- [23] Guokun Lai, Wei-Cheng Chang, Yiming Yang, and Hanxiao Liu. 2018. Modeling long-and short-term temporal patterns with deep neural networks. In *SIGIR*. 95–104.
- [24] Vincent Le Guen and Nicolas Thome. 2020. Probabilistic Time Series Forecasting with Structured Shape and Temporal Diversity. In *Advances in Neural Information Processing Systems*.
- [25] Qi Lei, Jinfeng Yi, Roman Vaculin, Lingfei Wu, and Inderjit S. Dhillon. 2019. Similarity Preserving Representation Learning for Time Series Clustering. In *IJCAI*.
- [26] Haoliang Li, Sinno Jialin Pan, Shiqi Wang, and Alex C Kot. 2018. Domain generalization with adversarial feature learning. In *CVPR*. 5400–5409.
- [27] J. Lines and Anthony J. Bagnall. 2014. Time series classification with ensembles of elastic distance measures. *Data Min. Knowl. Discov.* 29 (2014), 565–592.
- [28] Xiaoyan Liu, Zhenjiang Lin, and Huaqing Wang. 2008. Novel online methods for time series segmentation. *IEEE TKDE* 20, 12 (2008), 1616–1626.
- [29] Yasuko Matsubara, Yasushi Sakurai, Willem G Van Panhuis, and Christos Faloutsos. 2014. FUNNEL: automatic mining of spatially coevolving epidemics. In *KDD*. 105–114.
- [30] Krikamol Muandet, David Balduzzi, and Bernhard Schölkopf. 2013. Domain generalization via invariant feature representation. In *ICML*. 10–18.
- [31] Boris N Oreshkin, Dmitri Carpov, Nicolas Chapados, and Yoshua Bengio. 2021. Meta-learning framework with applications to zero-shot time-series forecasting. In *AAAI*.
- [32] C. Orsenigo and C. Vercellis. 2010. Combining discrete SVM and fixed cardinality warping distances for multivariate time series classification. *Pattern Recognit.* 43 (2010), 3787–3794.
- [33] Yao Qin, Dongjin Song, Haifeng Chen, Wei Cheng, Guofei Jiang, and Garrison Cottrell. 2017. A dual-stage attention-based recurrent neural network for time series prediction. In *AAAI*.
- [34] Goce Ristanoski, Wei Liu, and James Bailey. 2013. Time series forecasting using distribution enhanced linear regression. In *PAKDD*. 484–495.
- [35] Joshua W Robinson, Alexander J Hartemink, and Zoubin Ghahramani. 2010. Learning Non-Stationary Dynamic Bayesian Networks. *Journal of Machine Learning Research* 11, 12 (2010).
- [36] Sheldon M Ross. 2014. *Introduction to stochastic dynamic programming*. Academic press.
- [37] David Salinas, Valentin Flunkert, Jan Gasthaus, and Tim Januschowski. 2020. DeepAR: Probabilistic forecasting with autoregressive recurrent networks. *Int. J. Forecast.* 36, 3 (2020), 1181–1191.
- [38] P. Schäfer. 2015. Scalable time series classification. *Data Min. Knowl. Discov.* 30 (2015), 1273–1298.
- [39] Robert E Schapire. 2003. The boosting approach to machine learning: An overview. *Nonlinear estimation and classification* (2003), 149–171.
- [40] Rajat Sen, Hsiang-Fu Yu, and Inderjit S Dhillon. 2019. Think globally, act locally: A deep neural network approach to high-dimensional time series forecasting. In *NeurIPS*. 4837–4846.
- [41] Hidetoshi Shimodaira. 2000. Improving predictive inference under covariate shift by weighting the log-likelihood function. *Journal of statistical planning and inference* 90, 2 (2000), 227–244.
- [42] Baochen Sun and Kate Saenko. 2016. Deep CORAL: Correlation Alignment for Deep Domain Adaptation. In *ECCV*.
- [43] Kerem Sinan Tuncel and Mustafa Gokce Baydogan. 2018. Autoregressive forests for multivariate time series modeling. *Pattern recognition* 73 (2018), 202–215.
- [44] E. Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. 2017. Adversarial Discriminative Domain Adaptation. In *CVPR*. 2962–2971.
- [45] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*. 5998–6008.
- [46] LE Vincent and Nicolas Thome. 2019. Shape and Time Distortion Loss for Training Deep Time Series Forecasting Models. In *NeurIPS*. 4189–4201.
- [47] Jindong Wang, Yiqiang Chen, Wenjie Feng, Han Yu, Meiyu Huang, and Qiang Yang. 2020. Transfer learning with dynamic distribution adaptation. *ACM Transactions on Intelligent Systems and Technology (TIST)* 11, 1 (2020), 1–25.
- [48] Jindong Wang, Wenjie Feng, Yiqiang Chen, Han Yu, Meiyu Huang, and Philip S Yu. 2018. Visual domain adaptation with manifold embedded distribution alignment. In *Proceedings of the 26th ACM international conference on Multimedia*. 402–410.
- [49] Jindong Wang, Cuiling Lan, Chang Liu, Yidong Ouyang, Wenjun Zeng, and Tao Qin. 2021. Generalizing to Unseen Domains: A Survey on Domain Generalization. In *International Joint Conference on Artificial Intelligence (IJCAI)*.
- [50] Zhiguang Wang and Tim Oates. 2015. Imaging Time-Series to Improve Classification and Imputation. In *IJCAI*.
- [51] Z. Yang, R. Salakhutdinov, and William W. Cohen. 2017. Transfer Learning for Sequence Tagging with Hierarchical Recurrent Networks. In *ICLR*.
- [52] Chaohui Yu, Jindong Wang, Yiqiang Chen, and Meiyu Huang. 2019. Transfer learning with dynamic adversarial adaptation network. In *2019 IEEE International Conference on Data Mining (ICDM)*. IEEE, 778–786.
- [53] G Peter Zhang. 2003. Time series forecasting using a hybrid ARIMA and neural network model. *Neurocomputing* 50 (2003), 159–175.
- [54] Shuyi Zhang, Bin Guo, Anlan Dong, Jing He, Ziping Xu, and Song Xi Chen. 2017. Cautionary tales on air-quality improvement in Beijing. *Proc. Math. Phys. Eng. Sci.* 473, 2205 (2017), 20170457.
- [55] Yunyue Zhu and Dennis Shasha. 2002. Statstream: Statistical monitoring of thousands of data streams in real time. In *VLDB*. Elsevier, 358–369.
- [56] Yongchun Zhu, Fuzhen Zhuang, Jindong Wang, Guolin Ke, Jingwu Chen, Jiang Bian, Hui Xiong, and Qing He. 2020. Deep subdomain adaptation network for image classification. *IEEE transactions on neural networks and learning systems* 32, 4 (2020), 1713–1722.