Transfer Learning in Financial Time Series with Gramian Angular Field

1st Hou-Wan Long
The Chinese University of Hong Kong, Hong Kong
houwanlong@link.cuhk.edu.hk

2nd On-In Ho
University of Macau, Macao
MC15410@um.edu.mo

3rd Qi-Qiao He
Foshan University, China
hesammul@gmail.com

4th Yain-Whar Si
University of Macau, Macao
fstasp@um.edu.mo

Abstract—In financial analysis, time series modeling is often hampered by data scarcity, limiting neural network models' ability to generalize. Transfer learning mitigates this by leveraging data from similar domains, but selecting appropriate source domains is crucial to avoid negative transfer. This study enhances source domain selection in transfer learning by introducing Gramian Angular Field (GAF) transformations to improve time series similarity functions. We evaluate a comprehensive range of baseline similarity functions, including both basic and state-of-the-art (SOTA) functions, and perform extensive experiments with Deep Neural Networks (DNN) and Long Short-Term Memory (LSTM) networks. The results demonstrate that GAF-based similarity functions significantly reduce prediction errors. Notably, Coral (GAF) for DNN and CMD (GAF) for LSTM consistently deliver superior performance, highlighting their effectiveness in complex financial environments.

Index Terms—Time Series, Transfer Learning, Similarity Functions, Gramian Angular Field, Deep Learning

I. Introduction

Transfer learning is a powerful and increasingly popular machine learning technique designed to enhance the performance of models in a target domain by leveraging knowledge acquired from a related source domain. This technique is particularly valuable in scenarios where the target domain suffers from data scarcity, which often hinders the development of robust models and increases the risk of overfitting. By effectively transferring information from a source domain that has a more abundant and rich dataset, transfer learning enhances model accuracy and generalization capabilities. This makes it an essential tool across various fields, especially in situations where labeled data is limited or expensive to obtain. The ability of transfer learning is evident in its widespread application across diverse areas, particularly in time series analysis. It has proven to be effective in domains such as fault diagnosis, human activity recognition (HAR), and brain-computer interfaces (BCI), where time-dependent patterns are crucial for accurate predictions. A systematic study by Weber et al. [1] reviewed 223 publications and highlighted the extensive application of transfer learning across more than 40 different time series domains, underscoring its broad utility and impact.

However, the process of transferring knowledge is not without its challenges. One of the most significant risks associated with transfer learning is the phenomenon known as negative transfer. Negative transfer occurs when the inclusion of data from a source domain, instead of enhancing, actually deteriorates the model's performance in the target domain. This issue, first identified by Rosenstein et al. [2], has been a focal point of transfer learning research, prompting the development of various strategies to mitigate its effects. For instance, in time series classification, Fawaz et al. [3] observed instances of negative transfer when using Convolutional Neural Networks (CNNs). To address this, researchers have proposed various sophisticated domain selection techniques to better align the source and target domains and reduce the likelihood of negative transfer [4]. These techniques aim to ensure that the source domain's data is sufficiently relevant and beneficial to the target domain, thereby maximizing the effectiveness of the transfer learning process.

A crucial aspect of mitigating negative transfer and ensuring effective knowledge transfer is the selection of an appropriate source domain, which relies heavily on similarity functions. Similarity functions play a key role in measuring how closely related the source and target domains are. They do so by quantifying the distance between distributions or feature representations of the two domains. This measurement is often referred to as a distance measure. Distance measures, such as Euclidean distance, Pearson correlation, Maximum Mean Discrepancy (MMD) [9], Central Moment Discrepancy (CMD) [13], Correlation Alignment (CORAL) [6], Dynamic Time Warping (DTW), Time Warp Edit Distance (TWED) [23], and Wasserstein distance [11], serve as the mathematical foundation for similarity functions. These measures calculate the degree of difference or similarity between the domains, guiding the transfer learning process by indicating the most relevant and beneficial source domain for the target task. For instance, Euclidean distance calculates the straight-line distance between two points in a feature space, while DTW aligns sequences that may vary in time or speed, making it particularly useful for time series data. Similarly, MMD and CMD measure discrepancies in the statistical distributions of the source and target domains.

Despite the wide variety of similarity functions and distance measures currently available, their effectiveness in the context

of transfer learning—especially for time series data—remains underexplored. The selection of an appropriate similarity function is crucial for minimizing negative transfer and enhancing model performance, yet there has been insufficient research into how these functions perform across different transfer learning scenarios. Specifically, in time series analysis, where data is inherently complex and temporally dependent, there is a notable gap in comprehensive evaluations of these methods. Given that time series data often contain rich temporal and angular relationships, we are motivated to explore whether enhancing these similarity functions with Gramian Angular Field (GAF) [8] transformations could significantly improve their accuracy and robustness. To address these gaps, we not only evaluate the existing baseline similarity functions, including both basic and SOTA functions, but also perform extensive experiments to assess how these functions perform when enhanced with GAF transformations.

Gramian Angular Field (GAF) transformation is a technique that transforms time series data into two-dimensional images (see Fig.1), capturing both the temporal structure and angular relationships within the data. GAF has shown potential in preserving the intrinsic patterns of time series data, making it a valuable tool for enhancing similarity functions. By converting time series data into a form that can be processed similarly to images, GAF allows for the extraction of richer, more informative features that can improve the accuracy of domain distance measurements.

Our contributions are threefold:

(i) We evaluate baseline similarity functions for source dataset selection in both single-source and multi-source transfer learning scenarios.

(ii) We introduce a GAF-based method for measuring domain distance and demonstrate that GAF scaling significantly improves performance compared to unscaled similarity functions.

(iii) We provide a detailed analysis of the reasons behind the superior performance of GAF-based methods, offering deeper insights into their comparative advantages over traditional similarity functions.

The paper is organized as follows: Section 2 reviews related literature, Section 3 introduces the GAF-based similarity functions, Section 4 evaluates the baseline similarity functions for transfer learning, Section 5 is the experiment section, and Section 6 concludes the paper.

II. RELATED WORK

This section reviews the use of deep learning models and advanced similarity functions to enhance transfer learning in time series.

A. Time Series Problems and Transfer Learning Approaches

Time series problems supported by transfer learning include classification, regression, and clustering, as noted by Weber et al. [1]. Anomaly detection and time forecasting fall under classification and regression. In their study, 223 publications were included, with 164 focused on time series classification.

Transfer learning in time series primarily uses deep learning models, with four main approaches: model-based, feature-based, instance-based, and hybrid. Model-based transfer learning is the most common, utilizing CNN, RNN, and LSTM models, with combinations like CNN-RNN and CNN-LSTM also prevalent. Model-based techniques include pre-training and fine-tuning, partial freezing, and architecture modification. Domain-adversarial learning, dedicated model objectives, and ensemble-based transfer are also used. Transforming time series into GAF, proposed by Oates and Wang [8], is notable in transfer learning.
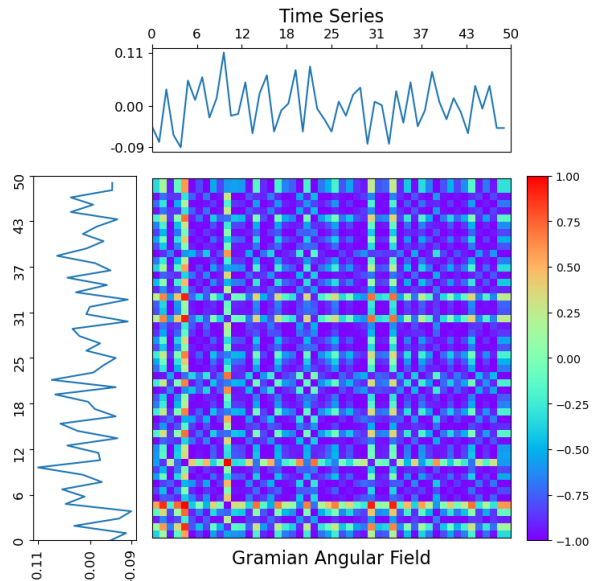

Fig. 1. GAF Transformation on Time Series

B. Distance Measures and Transfer Learning Techniques

Ben-David et al. [20] utilized A-distance, a type of distance measurement, to quantify the difference between domains in transfer learning. This approach laid the groundwork for the development of more advanced similarity functions. Chen et al. [9] built upon this by introducing a transfer learning method that uses fine-grained Maximum Mean Discrepancy (MMD) as a distance measure to assess domain discrepancies, serving as a similarity function that guides the transfer process. Wang et al. [18] and Wang et al. [19] further refined this approach by estimating a balance factor μ, which adjusts the balance between marginal and conditional distribution adaptation, with A-distance being a key component in defining the similarity function.

In geometrical feature transformation methods, Sun et al. [6] proposed CORrelation Alignment (CORAL), which employs second-order statistics as a similarity function to align the source and target domains. The core of CORAL is learning a transformation matrix that minimizes the distance between the second-order statistics of the source and target domains, thereby reducing distribution divergence. This method was
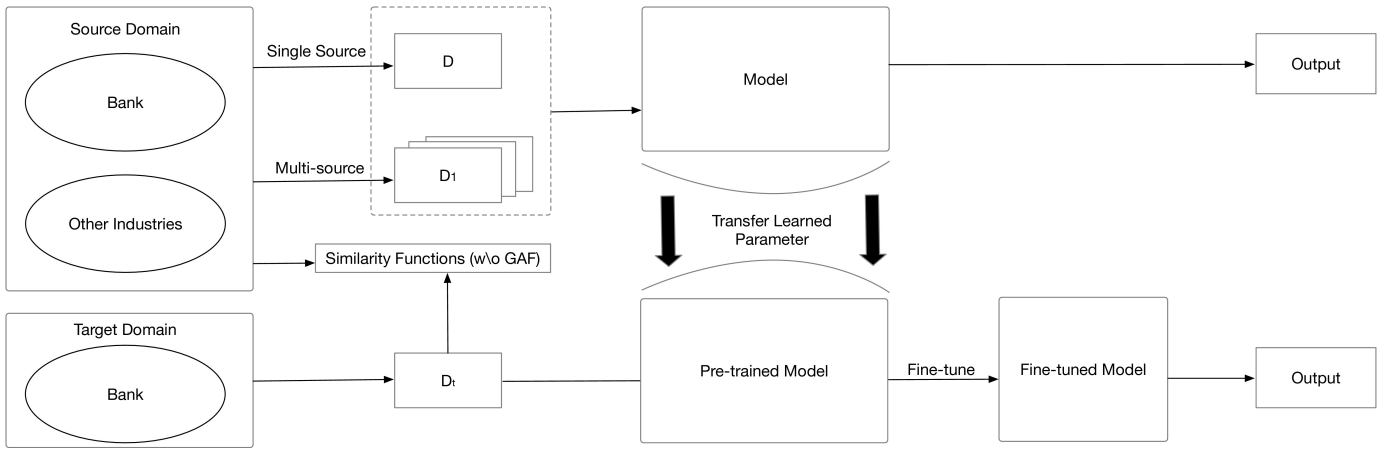
Fig. 2. The Procedure of GAF-based Transfer Learning

extended to Deep CORAL [7], where CORAL loss replaced MMD loss, effectively acting as a new similarity function for aligning feature distributions. Wasserstein distance, another critical distance measure rooted in optimal transport theory, was applied in transfer learning by Courty et al. [11]. They used it as a similarity function to learn feature transformations that reduce the marginal distribution distance, further extending this approach with Joint Distribution Optimal Transport (JDOT) to account for both marginal and conditional distributions [17].

In deep transfer learning, Tzeng et al. [12] introduced Deep Domain Confusion (DDC), where MMD was used as a distance measure to define a similarity function for aligning feature distributions in deep networks. Long et al. [21] advanced this with Deep Adaptation Networks (DAN), which replaced single-kernel MMD with multiple-kernel MMD, offering a more flexible similarity function for distribution adaptation. Zhu et al. [22] proposed Deep Subdomain Adaptation Networks (DSAN), introducing a local MMD distance that computes weighted MMD distances as a refined similarity function. Zellinger et al. [13] proposed a new regularization method that directly minimizes discrepancies between domain-specific latent feature representations in the hidden activation space, introducing Central Moment Discrepancy (CMD) as a novel distance-based similarity function. This approach was tested on benchmark datasets for object recognition and sentiment analysis.

In the context of time series classification, Serrà and Arcos [5] evaluated various distance measures used as similarity functions, including Euclidean distance, Fourier Coefficients (FC), auto-regression model-based (AR) measures, DTW, Edit Distance on Real Sequences (EDR), Time Warp Edit Distance (TWED), and Minimum Jump Costs dissimilarity (MJC). Their study concluded that TWED, as a similarity function, consistently outperformed others. No significant differences were found among DTW, EDR, and MJC, while Euclidean distance, another basic similarity function, performed worse than these methods. FC and AR measures, also acting as similarity functions, were ranked below Euclidean distance in performance.

III. Transfer Learning Based on Gramian Angular Field (GAF)

There are different variations of transfer learning, including single source transfer learning and multi-source transfer learning.

In the case of single source transfer learning, the procedure starts with the training of a model on a single source dataset, denoted as D. During this stage, the model learns the patterns, features, and knowledge inherent to this specific domain. Once the model has been trained, the learned parameters are transferred to the next stage, where the model undergoes fine-tuning using data from the target domain D_t, which, in this scenario, belongs to the bank sector. The fine-tuning process is essential as it allows the model to adjust and adapt the pre-learned knowledge to better fit the unique characteristics of the target domain's data. After fine-tuning, the model is tested on the test set of D_t to evaluate its performance and to ensure that the model generalizes well to the new domain, ultimately producing the final output.

To ensure the selection of the most relevant source domain for transfer learning, a set of similarity functions is applied. The selection algorithm iterates over each similarity function F in the set of similarity functions. For each function, the algorithm calculates the similarity between each source dataset D_s in the set of possible sources S and the target dataset D_t. If the GAF transformation is being used, both D_s and D_t are transformed before calculating the similarity. These calculated similarity values are stored in a list. If the similarity function F is Pearson correlation, the algorithm selects the source dataset with the maximum similarity value. For other similarity functions, the source dataset with the minimum similarity value is selected. The chosen source dataset D is then used for the initial training phase of the transfer learning model. The pseudo-code for this process is shown in Algorithm 1.

For multi-source transfer learning, the process involves training the model on multiple datasets, denoted as D_1, D_2,

Algorithm 1 Selecting Source Domains for Single-source Transfer Learning

---
```
for F in similarityfunctions do
    values = [ ]
    for D_s in S do
        if use_GAF_transform then
            D_s ← GAF(D_s)
            D_t ← GAF(D_t)
        end if
        value = F(D_s, D_t)
        values.append(value)
    end for
    if F == Pearson then
        i = argmax(values)
    else
        i = argmin(values)
    end if
    D = S[i]
end for
```
---

and potentially others, which may come from different source domains like bank and other industries. This approach is designed to integrate diverse knowledge from various domains, with the aim of enhancing the model's robustness and generalization capabilities. The model, having been trained on these multiple sources, is then fine-tuned on the target domain D_t, similar to the single-source approach. The fine-tuning process refines the model's performance specifically for the target data, leveraging the combined knowledge from the multiple sources. Subsequently, the fine-tuned model is tested on the test set of D_t, and the final output is generated, ensuring that the model performs well on the new domain.

The selection of source domains in the multi-source scenario follows a more complex algorithm. Initially, the similarity function is applied to identify the first most relevant source dataset, denoted as D_1, from the set of possible sources S. If the GAF transformation is being used, both D_s and D_t are transformed before calculating the similarity. After selecting D_1, it is removed from S, and the similarity function is then reapplied to the remaining datasets in the modified set S'. This process identifies the second most relevant source dataset, denoted as D_2. The algorithm ensures that the selected datasets D_1 and D_2 are the most appropriate for training the model in a multi-source transfer learning setting. The pseudo-code for this multi-source selection process is shown in Algorithm 2.

A critical component in our proposed transfer learning process is the use of GAF-based similarity functions to measure the distance between the source domains and the target domain. This approach addresses a key limitation of the baseline similarity functions discussed in Section IV, which operate on 1-dimensional time series data and may not fully capture the complexity of the data. By transforming the data into a more representative form, these similarity functions can more effectively compare the source domain datasets with the target domain dataset D_t. The primary purpose of these functions is to aid in selecting the most appropriate source domains, thereby optimizing the transfer learning process and minimizing the risk of negative transfer. By accurately measuring the similarity between domains, these functions ensure that the knowledge being transferred is both beneficial

Algorithm 2 Selecting Source Domains for Multi-source Transfer Learning

---
```
for F in similarityfunctions do
    values = [ ]
    for D_s in S do
        if use_GAF_transform then
            D_s ← GAF(D_s)
            D_t ← GAF(D_t)
        end if
        value = F(D_s, D_t)
        values.append(value)
    end for
    if F == Pearson then
        i = argmax(values)
    else
        i = argmin(values)
    end if
    D_1 = S[i]
    S' = S.remove(D_1)
    values = [ ]
    for D_s in S' do
        if use_GAF_transform then
            D_s ← GAF(D_s)
            D_t ← GAF(D_t)
        end if
        value = F(D_s, D_t)
        values.append(value)
    end for
    if F == Pearson then
        i = argmax(values)
    else
        i = argmin(values)
    end if
    D_2 = S'[i]
end for
```
---

and relevant to the target task.

A. Turning 1-d Time Series to 2-d Gramian Angular Field

There are three steps in transforming time series to GAF. The first step is to scale the time series to [-1,1] using:

$$\tilde{x}_i = \frac{(x_i - \max(X)) + (x_i - \min(X))}{\max(X) - \min(X)} \quad (1)$$

x_i represents an individual data point in the time series. X denotes the entire time series dataset. The equation scales each data point to a range of [-1,1] by adjusting it relative to the maximum and minimum values in the entire dataset X.

Next, the scaled time series is transformed to polar coordinates:

$$\begin{cases} \phi = \arccos(\tilde{x}_i), & -1 \le \tilde{x}_i \le 1, \tilde{x}_i \in \tilde{X} \\ r = \frac{t_i}{N}, & t_i \in \mathbb{N} \end{cases} \quad (2)$$

where ϕ is the angle and r is the radius.

Finally, the polar coordinates are converted to Gramian Angular Summation Field (GASF) image as shown in Fig.1:

$$GASF = [\cos(\phi_i + \phi_j)] \quad (3)$$

where GASF represents the pixel value at the (i, j) position in the GASF image.

According to Oates and Wang [10], GAFs preserve temporal dependencies. As we move from the top-left to the bottom of GAFs, time increases, capturing temporal correlations effectively. GASF is a specific type of GAF that focuses on the summation of angular values. In this study, we generally refer

to these transformations as GAF to encompass both GASF and other potential variations.

B. Similarity Functions Applied on GAF

Pearson, DTW, and TWED are designed for 1-dimensional series and are excluded for being applied on GAF. MMD is also excluded due to high memory consumption. Thus, the similarity functions applied on GAF include (i) Euclidean distance, (ii) Coral, (iii) CMD, and (iv) Wasserstein distance.

Additionally, three image similarity functions are applied on GAF. Before applying these, GAFs are scaled from [-1,1] to [0,255]:

$$GAF_{scaled} = GAF \times 127.5 + 127.5 \tag{4}$$

The image similarity functions include Adapted Rand error (ARE), Peak Signal Noise Ratio (PSNR), and Structural Similarity (SSIM). Skimage metrics module [14] is used for these calculations.

Adapted Rand error is calculated as:

$$ARE = 1 - \frac{\sum_{ij} p_{ij}^2}{\alpha \sum_k s_k^2 + (1 - \alpha) \sum_k t_k^2} \tag{5}$$

where p_{ij} is the probability that the label of a pixel in the test image matches the true image, t_k and s_k are the probabilities of pixel labels in the true and test images, respectively, and α is set to 0.5.

Structural similarity (SSIM) [14] is defined as:

$$SSIM(x, y) = \frac{(2\mu_x \mu_y + c_1)(2\sigma_{xy} + c_2)}{(\mu_x^2 + \mu_y^2 + c_1)(\sigma_x^2 + \sigma_y^2 + c_2)} \tag{6}$$

where $c_1 = (k_1 L)^2$ and $c_2 = (k_2 L)^2$. L is the dynamic range of pixel values, set to 255, and k_1 and k_2 are small constants. x and y represent the two images (or signals) being compared for structural similarity. μ_x represents the mean of the variable x, while μ_y represents the mean of the variable y. Similarly, σ_x^2 denotes the variance of x, σ_y^2 denotes the variance of y, and σ_{xy} represents the covariance between x and y.

Peak Signal Noise Ratio (PSNR) measures the ratio between the maximum possible power of a signal and the power of noise affecting image quality:

$$PSNR = 10 \cdot \log_{10} \left(\frac{MAX_I^2}{MSE} \right) \tag{7}$$

where MSE is the mean squared error between two images, and MAX_I is the maximum pixel value, set to 255.

IV. EVALUATION OF BASELINE SIMILARITY FUNCTIONS FOR TRANSFER LEARNING

We conduct baseline evaluations to assess the performance of (i) Euclidean Distance, (ii) Pearson Correlation, (iii) Maximum Mean Discrepancy (MMD), (iv) Central Moment Discrepancy (CMD), (v) Correlation Alignment (CORAL), (vi) Dynamic Time Warping (DTW), (vii) Time Warp Edit Distance (TWED) and (viii) Wasserstein Distance for transfer learning. We perform two sets of baseline analyses: single-source and multi-source transfer learning. In single-source

transfer learning, the model is trained on a single source dataset before being fine-tuned on the target dataset [15]. We compare the effectiveness of similarity functions when the source dataset is from the same industry (banking) versus different industries in predicting bank stock prices. In multi-source transfer learning, the model is trained on multiple source datasets before fine-tuning on the target dataset [16]. Three scenarios are considered: sources are selected either solely from banking, from other industries, or a combination of both. In each evaluation, the similarity functions are used to select the most similar source datasets for transfer learning. The setting of baseline evaluation can be seen in Table III.

A. Dataset Description

We use Hong Kong stock price series data, comprising daily closing prices from the beginning of 2022 to the end of 2023, retrieved via the yfinance Python module. The study forecasts the next day's closing price based on the previous 10 days' prices. The dataset is split into training and test sets in a 7:3 ratio. Before training, the time series are scaled using MinMaxScaler, with the feature range set from -1 to 1. As shown in Table I, the stock datasets include bank stocks and stocks from other industries. The bank stocks represent the top 11 market capitalizations on HKEX, while the non-bank stocks are the top 10 constituents of the Hang Seng Index by market capitalization.

B. Network Architecture and Evaluation Metrics

This paper use the network in [15]. We focus on forecasting financial time series using two network architectures: Deep Neural Network (DNN) and Long Short-Term Memory (LSTM). The DNN used in our study consists of four layers, including two hidden layers with 256 nodes each and an output layer with one node. LSTM, suitable for sequence prediction tasks, is composed of multiple LSTM cells arranged in layers. Our LSTM network comprises two layers with 256 cells each, followed by a single-node output layer. The hyper-parameters of both networks are shown Table II.

Model performance is evaluated using Mean Absolute Percentage Error (MAPE[1]), Root Mean Squared Error (RMSE), and R-squared (R2), where lower MAPE and RMSE and higher R2 indicate better performance.

MAPE measures the accuracy of the model by calculating the percentage difference between the predicted values and the actual values. It provides an intuitive understanding of the prediction error in percentage terms, making it easier to interpret the model's accuracy.

$$MAPE = \frac{100\%}{n} \sum_{i=1}^{n} \left| \frac{y_i - \hat{y}_i}{y_i} \right| \tag{8}$$

RMSE is used to measure the average magnitude of the prediction errors, with errors being squared before averaging.

---

[1]We use nnMAPE and lstmMAPE to distinguish between the MAPE values obtained from the deep neural network (nn) and the LSTM models, respectively; similar distinctions are applied for other metrics like RMSE and R-squared (nnRMSE, lstmRMSE, nnR^2, and lstmR^2)

TABLE I
STOCK DATASETS

Bank Stocks	0939. HK	3988. HK	3328. HK	0011. HK	1288. HK	1398. HK	2388. HK	2888. HK	3968. HK	1658. HK	0005. HK
Other Industry	0700. HK	0941. HK	1299. HK	3690. HK	1810. HK	9988. HK	0883. HK	9999. HK	9618. HK	0388. HK	

TABLE II
HYPER-PARAMETERS OF NEURAL NETWORK

Hyper-parameters	Pre-train	Fine-tune
Epochs	100	100
Batch size	200	200
Optimizer	Adam	Adam
Learning rate	0.001	0.00001
First moment	0.9	0.9
Second moment	0.999	0.999
Loss function	MSE	MSE

This metric gives more weight to large errors, which is useful when large deviations are particularly undesirable.

$$\text{RMSE} = \sqrt{\frac{\sum (y_i - \hat{y}_i)^2}{n}} \quad (9)$$

R-squared (R2) is a statistical measure that represents the proportion of the variance in the dependent variable that is predictable from the independent variables. It indicates how well the model's predictions approximate the actual data points, with values closer to 1 suggesting a better fit.

$$R^2 = 1 - \frac{\sum (y_i - \hat{y}_i)^2}{\sum (y_i - \bar{y}_i)^2} \quad (10)$$

where y_i is the actual value and $\hat{y}_i$ is the predicted value. n is the prediction period.

TABLE III
SINGLE SOURCE AND MULTI-SOURCE TRANSFER LEARNING BASELINE
EVALUATIONS AND EXPERIMENTS

Evaluation&Experiment	Source (S_A)	Source (S_B)	Target
Evaluation&Experiment 1	Bank	—	Bank
Evaluation&Experiment 2	Other Industries	—	Bank
Evaluation&Experiment 3	Bank	Bank	Bank
Evaluation&Experiment 4	Other Industries	Other Industries	Bank
Evaluation&Experiment 5	Bank	Other Industries	Bank

C. Summary Analysis for Baseline Similarity Functions

Table IV to VIII conclude the effectiveness of different similarity functions in transfer learning across various scenarios. In single-source transfer learning, CMD stands out as the best similarity function when the source is from other industries, while multiple functions, including Euclidean and Pearson, perform well for banking datasets. In multi-source transfer learning, CMD consistently excels, especially when sources are from the banking sector or a mix of industries. MMD is most effective for multi-source scenarios involving non-banking industries. This analysis underscores CMD's robustness and adaptability across different transfer learning contexts.

The similarity functions correspond to the lowest errors are considered the baseline best similarity functions. Table IX presents the baseline best similarity functions for single

source transfer learning and Table X presents the baseline best similarity functions for multi-source transfer learning, providing a standard against which the GAF-based similarity functions will be compared in subsequent sections. 0005.HK is used as the target dataset, i.e $D_t = 0005.HK$, and also for the experiments in Section V.

V. EXPERIMENT

To evaluate the GAF-based similarity functions, we repeat the baseline evaluations from Section IV using these GAF similarity functions. The experiments are denoted as experiments 1-5 as shown in Table III, corresponding to the previous baseline evaluations.

A. Experiment Results

The results of these experiments are shown in Table XI to Table XV. The errors of the GAF-based similarity functions are compared with the baseline best similarity functions' results in Tables IX and X. If the error obtained by GAF-based similarity functions is lower than that of similarity functions without GAF scaling, then GAF enhances the model's performance. The similarity functions corresponding to the lowest errors are considered the best similarity functions.

B. Summary Analysis for GAF-based Similarity Functions

Table XVI combines the results of Section IV and the experiment results, showcasing the best similarity functions under different conditions. The analysis highlights that scaling the time series by GAF and applying similarity functions can significantly improve transfer learning outcomes in several scenarios. Here, we discuss the findings in more detail:

Single Source Transfer Learning (Banking): When both the source and target datasets consist of banking stocks, applying GAF-based similarity functions leads to notable improvements. Specifically, Coral (GAF) for DNN and CMD (GAF) for LSTM yielded the best performance. This indicates that the temporal dependencies and correlations captured by GAF transformations enhance the model's ability to generalize from the source to the target domain within the same industry. This is particularly beneficial in financial sectors where capturing intricate temporal patterns is crucial.

Single Source Transfer Learning (Other Industries): In single source transfer learning scenarios with source data from other industries, both GAF-based and traditional similarity functions showed advantages for DNN and LSTM models. Coral (GAF) and CMD were the best performers for both models, indicating that GAF transformations improved temporal dependency capture, but CMD without GAF also yielded strong results. This suggests that while GAF-based

TABLE IV
RESULT OF BASELINE EVALUATION 1

Similarity Functions	D	nnMAPE	nnRMSE	nnR^2	lstmMAPE	lstmRMSE	lstmR^2
Euclidean, Pearson, MMD, Coral, DTW, TWED, Wasserstein	2888.HK	*0.9090*	*0.6784*	*0.8862*	*1.4436*	*1.1158*	*0.6923*
CMD	2388.HK	0.9093	0.6792	0.8860	2.0817	1.6230	0.3489

[a] D stands for source dataset for single source transfer learning.
[b] Euclidean, Pearson, MMD, Coral, DTW, TWED, Wasserstein being in the same cell means that the source dataset selected by these similarity functions are the same, i.e 2888.HK.

TABLE V
RESULT OF BASELINE EVALUATION 2

Similarity Functions	D	nnMAPE	nnRMSE	nnR^2	lstmMAPE	lstmRMSE	lstmR^2
Euclidean, Pearson, Coral, DTW, TWED, Wasserstein	0941.HK	0.9104	0.6808	0.8854	2.2173	1.5995	0.3677
CMD	1299.HK	*0.8985*	*0.6679*	*0.8898*	*1.0438*	*0.7731*	*0.8523*

TABLE VI
RESULT OF BASELINE EVALUATION 3

Similarity Functions	D_1	D_2	nnMAPE	nnRMSE	nnR^2	lstmMAPE	lstmRMSE	lstmR^2
Euclidean, MMD, Coral, DTW, Wasserstein	2888.HK	3968.HK	*0.8910*	0.6667	0.8901	0.9477	0.7025	0.8780
Pearson	2888.HK	1288.HK	0.9122	0.6781	0.8863	1.0180	0.7493	0.8612
CMD	2388.HK	1658.HK	0.8980	*0.6650*	*0.8907*	*0.9274*	*0.6939*	*0.8810*
TWED	2888.HK	3988.HK	0.9240	0.6869	0.8834	0.9624	0.7196	0.8720

[a] D_1 and D_2 stand for the two source datasets for multi-source transfer learning.
[b] Euclidean, MMD, Coral, DTW, Wasserstein being in the same cell means that the source datasets selected by these similarity functions are the same, i.e 2888.HK and 3968.HK.

TABLE VII
RESULT OF BASELINE EVALUATION 4

Similarity Functions	D_1	D_2	nnMAPE	nnRMSE	nnR^2	lstmMAPE	lstmRMSE	lstmR^2
Euclidean, DTW, Wasserstein	0941.HK	1299.HK	0.9017	0.6704	0.8889	1.1433	0.8624	0.8162
Pearson, TWED	0941.HK	0883.HK	0.9157	0.6807	0.8855	1.1643	0.8597	0.8173
MMD	0941.HK	9618.HK	*0.8928*	*0.6654*	*0.8906*	*0.9520*	*0.7058*	*0.8769*
CMD	1229.HK	0883.HK	0.9101	0.6769	0.8867	0.9840	0.7237	0.8706
Coral	0941.HK	1810.HK	0.9110	0.6741	0.8877	1.0257	0.7616	0.8566

TABLE VIII
RESULT OF BASELINE EVALUATION 5

Similarity Functions	D_1	D_2	nnMAPE	nnRMSE	nnR^2	lstmMAPE	lstmRMSE	lstmR^2
Euclidean, Pearson, MMD, Coral, DTW, TWED, Wasserstein	2888.HK	0941.HK	0.9017	0.6704	0.8889	1.1433	*0.8624*	*0.8162*
CMD	2388.HK	1299.HK	*0.9157*	*0.6807*	*0.8855*	*1.1643*	0.8597	0.8173

TABLE IX
BASELINE BEST SIMILARITY FUNCTIONS FOR SINGLE SOURCE TRANSFER LEARNING

	DNN	LSTM
(S_A: Bank, T: Bank)	Euclidean, Pearson, MMD, Coral, DTW, TWED, Wasserstein	Euclidean, Pearson, MMD, Coral, DTW, TWED, Wasserstein
(S_A: Other Industries, T: Bank)	CMD	CMD

TABLE X
BASELINE BEST SIMILARITY FUNCTIONS FOR MULTISOURCE TRANSFER LEARNING

	DNN	LSTM
(S_A: Bank, S_B: Bank, T: Bank)	CMD	CMD
(S_A: Other Industries, S_B: Other Industries, T: Bank)	MMD	MMD
(S_A: Bank, S_B: Other Industries, T: Bank)	CMD	Euclidean, Pearson, MMD, Coral, DTW, TWED, Wasserstein

methods enhance model generalization, traditional CMD remains equally effective. This highlights the utility of GAF

TABLE XI
RESULT OF EXPERIMENT 1

Similarity Functions with GAF	D	nnMAPE	nnRMSE	nnR^2	lstmMAPE	lstmRMSE	lstmR^2
Euclidean, ARE, PSNR, SSIM	2888.HK	0.9090	0.6784	0.8862	1.4436	1.1158	0.6923
CMD	3988.HK	0.9166	0.6824	0.8849	*1.1058*	*0.8169*	*0.8350*
Coral	0011.HK	*0.8984*	*0.6708*	*0.8888*	1.7527	1.3123	0.5744
Wasserstein	1288.HK	0.9152	0.6800	0.8857	1.1467	0.8497	0.8215
Best result without GAF	2888.HK	0.9090	0.6784	0.8862	1.4436	1.1158	0.6923

TABLE XII
RESULT OF EXPERIMENT 2

Similarity Functions with GAF	D	nnMAPE	nnRMSE	nnR^2	lstmMAPE	lstmRMSE	lstmR^2
Euclidean, CMD, Wasserstein	9999.HK	0.9234	0.6815	0.8852	1.1029	0.8153	0.8357
Coral	1299.HK	*0.8985*	*0.6679*	*0.8898*	*1.0438*	*0.7731*	*0.8523*
ARE, PSNR, SSIM	0883.HK	0.9128	0.6787	0.8862	1.2799	0.9424	0.7805
Best result without GAF	1299.HK	*0.8985*	*0.6679*	*0.8898*	*1.0438*	*0.7731*	*0.8523*

TABLE XIII
RESULT OF EXPERIMENT 3

Similarity Functions with GAF	D_1	D_2	nnMAPE	nnRMSE	nnR^2	lstmMAPE	lstmRMSE	lstmR^2
Euclidean, Wasserstein, ARE, PSNR, SSIM	2888.HK	1288.HK	0.9122	0.6781	0.8863	1.0180	0.7493	0.8612
CMD	3988.HK	3328.HK	0.9096	0.6752	0.8873	1.0194	0.7554	0.8590
Coral	0011.HK	3988.HK	0.9096	0.6758	0.8871	0.9972	0.7388	0.8651
Wasserstein	1288.HK	2888.HK	0.9146	0.6825	0.8849	1.3271	1.0161	0.7448
Best Result without GAF	2888.HK	3968.HK	*0.8910*	-	-	-	-	-
	2388.HK	1658.HK	-	*0.6650*	*0.8907*	*0.9274*	*0.6939*	*0.8810*

TABLE XIV
RESULT OF EXPERIMENT 4

Similarity Functions with GAF	D_1	D_2	nnMAPE	nnRMSE	nnR^2	lstmMAPE	lstmRMSE	lstmR^2
Euclidean, Wasserstein, PSNR	9999.HK	0883.HK	0.9201	0.6838	0.8844	0.9760	0.7193	0.8721
CMD	9999.HK	1299.HK	0.9071	0.6741	0.8877	1.0747	0.8001	0.8418
Coral	1299.HK	9999.HK	0.9203	0.6792	0.8860	1.1314	0.8416	0.8249
ARE	0883.HK	3690.HK	0.8947	0.6655	0.8905	*0.9184*	*0.6841*	*0.8843*
SSIM	0883.HK	9999.HK	0.9216	0.6803	0.8856	1.0749	0.7952	0.8437
Best Result without GAF	0941.HK	0700.HK	*0.8928*	*0.6654*	*0.8906*	0.9520	0.7058	0.8769

TABLE XV
RESULT OF EXPERIMENT 5

Similarity Functions with GAF	D_1	D_2	nnMAPE	nnRMSE	nnR^2	lstmMAPE	lstmRMSE	lstmR^2
Euclidean, PSNR	2888.HK	9999.HK	0.9289	0.6848	0.8841	1.2318	0.9340	0.7844
CMD	3988.HK	9999.HK	0.9294	0.6859	0.8837	1.0670	*0.7849*	*0.8477*
Coral	0011.HK	1299.HK	*0.8930*	*0.6644*	*0.8909*	1.2119	0.9197	0.7910
Wasserstein	1288.HK	9999.HK	0.9294	0.6858	0.8837	1.1039	0.8224	0.8329
ARE, SSIM	2888.HK	0883.HK	0.9164	0.6807	0.8855	1.0774	0.7950	0.8438
Best Result without GAF	2388.HK	1299.HK	0.8942	0.6648	0.8908	*1.0497*	-	-
	2388.HK	0941.HK	-	-	-	-	0.7858	0.8474

transformations in preserving temporal structures and validates the robustness of traditional CMD in such scenarios.

Multi-source Transfer Learning (Banking): In multi-source transfer learning scenarios where the source data is from banking industries, the results demonstrated that baseline similarity functions without GAF transformations provided the best performance for both DNN and LSTM models. This suggests that while GAF-based methods can enhance performance in other contexts, baseline similarity functions are more robust and effective in multi-source transfer learning when dealing with banking data.

Multi-source Transfer Learning (Other Industries): For multi-source transfer learning using DNN and LSTM models, when the sources are from other industries, the GAF-based similarity functions demonstrated mixed results. Specifically, ARE (GAF) proved to be particularly effective for the LSTM model, showing significant improvement in capturing temporal structures and enhancing the model's performance. However, the DNN models did not benefit as much from the GAF transformations and did not exhibit a notable improvement compared to traditional methods. This suggests that while GAF transformations can effectively enhance LSTM models by preserving important temporal structures, their benefits may not extend to DNN models in multi-source transfer learning scenarios involving diverse industries.

Multi-source Transfer Learning (Banking and Other Industries): In multi-source transfer learning scenarios where the first source is from banking industries and the second source is from other industries, both Coral (GAF) for DNN and CMD (GAF) for LSTM showed the best results. This finding is significant because it indicates that GAF transformations facilitate effective knowledge transfer even when the source domains are heterogeneous. The ability to integrate information from both similar and dissimilar sources can lead to more robust models, particularly in complex financial environments where different sectors may exhibit interconnected dynamics.

TABLE XVI
SUMMARY OF RESULTS

	DNN	LSTM
(S_A: Bank, T: Bank)	**Coral (GAF)**	**CMD (GAF)**
(S_A: Other Industries, T: Bank)	Coral (GAF), CMD	Coral (GAF), CMD
(S_A: Bank, S_B: Bank, T: Bank)	CMD	CMD
(S_A: Other Industries, S_B: Other Industries, T: Bank)	MMD	**ARE (GAF)**
(S_A: Bank, S_B: Other Industries, T: Bank)	**Coral (GAF)**	**CMD (GAF)**

C. Discussion

In this subsection, we provide a deeper analysis of why GAF performs better. Discussing the underlying reasons for the performance improvements would provide more insights into the effectiveness of GAF in transfer learning, particularly in the context of financial time series data.

1) Enhanced Temporal Feature Preservation: One of the key advantages of GAF-based similarity functions is their ability to capture and preserve temporal dependencies in time series data, which traditional methods like Euclidean distance or Pearson correlation often overlook. Traditional methods treat time points as isolated entities, missing the overall context and flow of the series. GAF transforms time series into 2D images, where each pixel encodes relationships between time points, maintaining the temporal structure and rhythm of the data. This richer representation is particularly valuable in financial markets, where the timing of price movements is critical to predicting future trends and patterns.

2) Robustness to Noise and Nonlinearities: Financial time series data are typically noisy and nonlinear, and traditional similarity functions often struggle to handle these characteristics effectively. GAF-based methods, however, excel in these scenarios by smoothing out high-frequency noise and focusing on underlying trends. The angular transformation used in GAF highlights the direction and strength of changes, making it easier to detect and work with nonlinear patterns in the data. This leads to more robust models that can generalize across different market conditions, a crucial advantage for financial forecasting, where volatility and unexpected shifts are common.

3) Improved Interpretability and Visual Analysis: Another advantage of GAF-based similarity functions is the improved interpretability they offer by transforming time series into images. This visual representation allows for a more intuitive understanding and comparison of different time series, making it easier to identify patterns, anomalies, or correlations between datasets. This interpretability is particularly useful in multi-source transfer learning, where understanding relationships between domains is crucial. By visually inspecting GAF images, researchers can quickly assess the similarity of different time series, aiding in the selection of the most appropriate source domains for transfer learning.

4) Synergistic Integration with Deep Learning Models: GAF-based methods integrate seamlessly with deep learning models, especially those designed for image analysis like convolutional neural networks (CNNs). When time series are transformed into GAF images, they can be processed similarly to other image data, allowing the deep learning model to leverage spatial patterns and dependencies that traditional similarity functions might miss. This synergy between GAF and deep learning architectures enables the extraction of both temporal and spatial features, leading to more accurate predictions and better overall performance. This is particularly beneficial in complex financial environments where subtle patterns are crucial for accurate forecasting.

5) Flexibility in Handling Heterogeneous Data Sources: GAF-based methods also provide significant flexibility in handling data from diverse sources, an area where traditional similarity functions can struggle. In multi-source transfer learning, datasets often come from different industries or exhibit varying behaviors. GAF standardizes these datasets by transforming them into a common 2D representation, allowing for a more consistent comparison across domains. This flexibility is invaluable in financial applications, where datasets can vary widely in structure and behavior. By enabling consistent

comparisons through GAF, transfer learning models can more effectively integrate heterogeneous data, leading to improved generalization and performance across different domains.

VI. Conclusion

In this study, we evaluated and compared the performance of various baseline similarity functions and Gramian Angular Field (GAF)-based similarity functions in the context of transfer learning for financial time series. The results from the baseline evaluations indicated that the Central Moment Discrepancy (CMD) function consistently performed well across different scenarios, especially in minimizing prediction errors and enhancing the transfer learning process. This robustness was particularly evident in multi-source transfer learning scenarios, where CMD effectively integrated knowledge from multiple sources.

When introducing GAF-based similarity functions, the findings revealed that these functions significantly improved the performance of transfer learning in certain contexts, particularly for single-source transfer learning within the banking industry. Specifically, Coral (GAF) for Deep Neural Networks (DNN) and CMD (GAF) for Long Short-Term Memory (LSTM) networks consistently delivered superior results, highlighting the effectiveness of GAF transformations in capturing temporal dependencies and correlations. However, the effectiveness of GAF-based functions was less consistent in multi-source transfer learning scenarios, particularly when the source domains involved a mix of banking and other industries. While some GAF-based functions like Adapted Rand Error (ARE) showed significant improvements for LSTM models, the benefits were not as pronounced for DNN models. This suggests that while GAF transformations can enhance the transfer learning process, particularly in capturing complex temporal structures, their utility may vary depending on the specific model and transfer learning scenario.

Overall, the study underscores the importance of selecting appropriate similarity functions for transfer learning, particularly when dealing with complex financial time series data. The findings suggest that while traditional similarity functions like CMD remain robust across various scenarios, GAF-based functions offer significant advantages in certain contexts, particularly in enhancing model performance in single-source transfer learning within the same industry. These insights provide a foundation for further research into optimizing transfer learning methodologies in financial time series analysis.

References

[1] Weber, M., Auch, M., Doblander, C., Mandl, P., Jacobsen, H.-A.: Transfer Learning With Time Series Data: A Systematic Mapping Study. IEEE Access, vol. 9, pp. 165409–165432 (2021).

[2] Rosenstein, M., Marx, Z., Pack, L., Dietterich, T.: To Transfer or Not To Transfer. In: Neural Information Processing Systems (2005)

[3] Fawaz, H. I., Forestier, G., Weber, J., Idoumghar, L., Muller, P.-A.: Transfer Learning for Time Series Classification. In: 2018 IEEE International Conference on Big Data (Big Data), pp. 1367–1376 (2018)

[4] Yao, Y., Doretto, G.: Boosting for transfer learning with multiple sources. In: Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition, pp. 1855–1862, CA (2010).

[5] Serrà, J., Arcos, J. Ll.: An empirical evaluation of similarity measures for time series classification. Knowledge-Based Systems, vol. 67, pp. 305–314 (2014).

[6] Sun, B., Feng, J., Saenko, K.: Correlation Alignment for Unsupervised Domain Adaptation. arXiv (Cornell University) (2016).

[7] B. Sun and K. Saenko, "Deep CORAL: Correlation alignment for deep domain adaptation," *arXiv*, Jul. 2016. [Online]. Available: https://doi.org/10.48550/arxiv.1607.01719

[8] Wang, Z., Oates, T.: Imaging Time-Series to Improve Classification and Imputation. arXiv.org (2015).

[9] Chen, Y., Wang, J., Huang, M., Yu, H.: Cross-position Activity Recognition with Stratified Transfer Learning. arXiv (2018).

[10] Wang, J., Chen, Y., Feng, W., Yu, H., Huang, M., Yang, Q.: Transfer Learning with Dynamic Distribution Adaptation. ACM Transactions on Intelligent Systems and Technology, vol. 11, no. 1, pp. 1–25 (2020).

[11] Courty, N., Flamary, R., Tuia, D., Rakotomamonjy, A.: Optimal Transport for Domain Adaptation. IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 39, no. 9, pp. 1853–1865 (2017).

[12] Tzeng, E., Hoffman, J., Zhang, N., Saenko, K., Darrell, T.: Deep Domain Confusion: Maximizing for Domain Invariance. arXiv (2014)

[13] Zellinger, W., Grubinger, T., Lughofer, E., Natschläger, T., Saminger-Platz, S.: Central Moment Discrepancy (CMD) for Domain-Invariant Representation Learning. arXiv (2017). https://arxiv.org/abs/1702.08811

[14] Wang, Z., Bovik, A. C., Sheikh, H. R., Simoncelli, E. P.: Image Quality Assessment: From Error Visibility to Structural Similarity. IEEE Transactions on Image Processing, vol. 13, no. 4, pp. 600–612 (2004).

[15] He, Qi-Qiao, Siu, Shirley Weng In, Si, Yain-Whar: Instance-based deep transfer learning with attention for stock movement prediction. Applied Intelligence, vol. 53, no. 6, pp. 6887–6908 (2023). Springer

[16] Lai, Yue-Hui, Cheong, Se-Hang, Zhang, Hui, Si, Yain-Whar: Coverage hole detection in WSN with force-directed algorithm and transfer learning. Applied Intelligence, vol. 52, no. 5, pp. 5435–5456 (2022). Springer

[17] N. Courty, R. Flamary, A. Habrard, and A. Rakotomamonjy, "Joint distribution optimal transportation for domain adaptation," in *Neural Information Processing Systems*, 2017.

[18] J. Wang, W. Feng, Y. Chen, H. Yu, M. Huang, and P. S. Yu, "Visual domain adaptation with manifold embedded distribution alignment," in *Proceedings of the 26th ACM International Conference on Multimedia*, Oct. 2018. doi: https://doi.org/10.1145/3240508.324051210.1145/3240508.3240512.

[19] J. Wang, Y. Chen, W. Feng, H. Yu, M. Huang, and Q. Yang, "Transfer learning with dynamic distribution adaptation," *ACM Transactions on Intelligent Systems and Technology*, vol. 11, no. 1, pp. 1–25, Feb. 2020. doi: https://doi.org/10.1145/336030910.1145/3360309.

[20] S. Ben-David, J. Blitzer, K. Crammer, and F. Lobo Pereira, "Analysis of representations for domain adaptation," in *Advances in Neural Information Processing Systems 19 (NIPS 2006)*, vol. 9, Jan. 2007. doi: https://doi.org/10.7551/mitpress/7503.003.002210.7551/mitpress/7503.003.0022.

[21] M. Long, Y. Cao, J. Wang, and M. Jordan, "Learning transferable features with deep adaptation networks," in *Proceedings of the 32nd International Conference on Machine Learning*, pp. 97–105, 2015.

[22] Y. Zhu *et al.*, "Deep subdomain adaptation network for image classification," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 32, no. 4, pp. 1713–1722, Apr. 2021. doi: https://doi.org/10.1109/TNNLS.2020.298892810.1109/TNNLS.2020.2988928.

[23] P.-F. Marteau, "Time Warp Edit Distance with Stiffness Adjustment for Time Series Matching," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 31, no. 2, pp. 306-318, Feb. 2009, doi: 10.1109/TPAMI.2008.76.

[24] Wu, Wangyu, Qiu, Xianglin, Song, Siqi, Chen, Zhenhong, Huang, Xiaowei, Ma, Fei, and Xiao, Jimin. "Image Augmentation Agent for Weakly Supervised Semantic Segmentation." arXiv preprint arXiv:2412.20439, 2024.

[25] Wu, Wangyu, Dai, Tianhong, Chen, Zhenhong, Huang, Xiaowei, Xiao, Jimin, Ma, Fei, and Ouyang, Renrong. "Adaptive Patch Contrast for Weakly Supervised Semantic Segmentation." Engineering Applications of Artificial Intelligence, vol. 139, pp. 109626, 2025. Elsevier.

[26] Wu, Wangyu, Dai, Tianhong, Huang, Xiaowei, Ma, Fei, and Xiao, Jimin. "Image Augmentation with Controlled Diffusion for Weakly-Supervised Semantic Segmentation." In: ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 6175–6179, 2024. IEEE.