BRIDGING RESOLUTIONS FOR DOMAIN-INVARIANT REPRESENTATION LEARNING IN TIME SERIES

Anonymous authors
Paper under double-blind review

ABSTRACT

Unsupervised time-series domain adaptation (TSDA) aims to mitigate distribution shifts between labeled source domains and unlabeled target domains. Existing methods typically rely on two extreme views of time series, namely fine-grained temporal representations and coarse global frequency representations, while overlooking crucial intermediate-resolution structures for progressive domain-invariant learning. To address this limitation, we investigate the previously unexplored intermediate-resolution hierarchy induced by the Stationary Wavelet Transform and reveal two key properties: different resolutions exhibit distinct transferability and discriminability characteristics under domain shift, while different resolutions exhibit distinct category-level decision strengths. Based on these properties, we propose **PRISM**, a TSDA framework that bridges progressive representations across resolutions for domain-invariant learning. PRISM introduces Cross-Resolution Knowledge Transfer to bridge multi-resolution representations by progressively exchanging complementary knowledge across resolutions, and Committee-based Consensus Training to bridge the cross-resolution decision hierarchy, enabling more consistent domain-invariant decisions. Extensive experiments on six datasets demonstrate that PRISM consistently outperforms existing methods, achieving an average accuracy improvement of 4.21% in cross-domain settings. Code is available at the anonymous link: `https://anonymous.4open.science/r/PRISM-E7BC/`.

1 INTRODUCTION

Time-series classification has achieved remarkable progress in the deep learning era, largely benefiting from strong discriminative representations learned from abundant labeled data (Fawaz et al., 2019; Wen et al., 2024; Zhang et al., 2023). However, when models are deployed in real-world environments, distribution shifts between training and test domains are often inevitable, which can substantially degrade the generalization ability of models (Ragab et al., 2023). Unsupervised domain adaptation (UDA) (Long et al., 2015; Ganin et al., 2016), which transfers knowledge from a labeled source domain to an unlabeled target domain, provides a promising solution to this challenge.

The core difficulty of time-series domain adaptation (TSDA) lies in the entanglement between domain shifts and complex temporal patterns. Existing methods mainly rely on two conventional views of time series: temporal representations in the original signal space and transformed frequency representations (He et al., 2023; Liu et al., 2024; Kim & Lee, 2025). Temporal signals preserve the full evolution of frequency components over time with fine-grained local detail, whereas frequency representations capture global spectral energy distribution, providing a more abstract and coarse-grained view. Although these views are complementary, relying solely on such extremes is inherently limited (Mallat, 1989). This overlooks crucial intermediate-resolution structures, leaving the properties of representations between these two extremes underexplored and hindering progressive domain-invariant learning.

This limitation motivates the need for an intermediate representation paradigm that bridges the temporal and frequency extremes. In this work, we employ the Stationary Wavelet Transform (SWT) to decompose original time series into a hierarchy of intermediate-resolution components (Nason & Silverman, 1995), thereby extending the conventional single-resolution view into a multi-resolution hierarchy. However, introducing a multi-resolution hierarchy alone does not answer how it should be

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

exploited for TSDA. Under domain shift, the interplay between transferability and discriminability is a central issue, naturally raising two fundamental questions: *how do transferability and discriminability vary across different resolutions?* And *how do these resolution-dependent differences affect the category-level decision strengths?* Motivated by these questions, we investigate the intrinsic relationships across resolutions and reveal two key insights for TSDA.

First, different resolutions exhibit distinct transferability and discriminability characteristics: lower-resolution representations tend to remain more domain-invariant and exhibit stronger transferability under domain shift, whereas higher-resolution representations vary more significantly across domains while preserving richer discriminative information. Second, different resolutions exhibit distinct category-level decision strengths. The original resolution provides the dominant decision information, while lower-resolution representations offer category-specific semantic advantages that can supplement the original resolution on particular categories.

Based on the above insights, we propose **PRISM** (**P**rogressive **R**esolution **I**nterplay for Unsupervised Time-**S**eries Do**M**ain Adaptation), the first TSDA framework that moves beyond the conventional temporal–frequency extremes and bridges resolutions for domain-invariant representation learning. To bridge the complementary properties of different resolutions, PRISM introduces **Cross-Resolution Knowledge Transfer**, which progressively exchanges complementary knowledge across resolutions through knowledge distillation-based mutual learning, enabling transferable and discriminative information to reinforce each other across the resolution hierarchy. To bridge the cross-resolution decision hierarchy, PRISM introduces **Committee-based Consensus Training**, which adopts anchor-guided committee consensus to achieve more consistent domain-invariant decisions through cross-resolution consensus.

The main contributions of this paper are summarized as follows:

- We move beyond conventional single-resolution TSDA paradigms and investigate the intrinsic relationships within previously unexplored intermediate-resolution representations, identifying two key insights: different resolutions exhibit distinct transferability and discriminability characteristics, while different resolutions exhibit distinct category-level decision strengths, motivating the design of effective domain-invariant learning mechanisms.
- Guided by these insights, we propose PRISM, which introduces two dedicated bridging modules to systematically exploit the identified cross-resolution properties: Cross-Resolution Knowledge Transfer to progressively exchange complementary knowledge across resolutions, and Committee-based Consensus Training to bridge the cross-resolution decision hierarchy through anchor-guided committee consensus.
- Extensive experiments verify that PRISM consistently outperforms existing methods, achieving average gains of 4.21% over state-of-the-art approaches across diverse datasets and application scenarios.

2 Related Work

General Unsupervised Domain Adaptation Unsupervised domain adaptation (UDA) transfers knowledge from a labeled source domain to an unlabeled but related target domain, and has been widely applied in practical scenarios (Zhang et al., 2018; Zou et al., 2018; Guo et al., 2021). Its goal is to reduce cross-domain discrepancy and improve target-domain generalization (Ben-David et al., 2010). Existing methods are mainly grouped into three categories: (1) Adversarial-based methods, which learn domain-invariant features via domain confusion, such as DANN (Ganin et al., 2016), CDAN (Long et al., 2018), and DIRT-T (Shu et al., 2018); (2) Discrepancy-based methods, which explicitly minimize feature-space distances, including DAN (Long et al., 2015), DeepCoral (Sun & Saenko, 2016), and HoMM (Chen et al., 2020); and (3) Self-training-based methods, which iteratively assign pseudo-labels to unlabeled target samples and select high-confidence predictions to augment training, such as PFAN (Chen et al., 2019), CST (Liu et al., 2021), and AdaMatch (Berthelot et al., 2021). Although effective in generic settings, these methods often underperform on time series due to the neglect of temporal structures and sequential dependencies.

Unsupervised Domain Adaptation for Time Series To date, relatively few methods are specifically designed for time-series UDA. Early studies mainly focus on temporal representation learning.

VRADA (Purushotham et al., 2017) and CoDATS (Wilson et al., 2020) design temporal feature extractors, SASA (Cai et al., 2021) employs LSTM (Shi et al., 2015) to model domain-invariant dependencies, AdvSKM (Liu & Xue, 2021) adapts MMD (Tzeng et al., 2014) for time-series alignment, and CLUDA (Ozyurt et al., 2022) learns contextual features via contrastive learning.

Recent works increasingly consider temporal–frequency modeling. RAINCOAT (He et al., 2023) introduces frequency-domain alignment via Sinkhorn divergence. ACON (Liu et al., 2024) combines mutual learning with adversarial training in the joint temporal–frequency space. TransPL (Kim & Lee, 2025) improves pseudo-labeling through vector-quantized transition matrices. SSSS-TSA (Ahad et al., 2025) and TDVR (Zhou et al., 2025) further exploit channel-wise characteristics of multivariate time series. However, existing methods mainly model aggregated temporal and frequency representations, while overlooking intermediate multi-resolution views and their complementary transferability and discriminability under domain shift.

3 Wavelet-Induced Multi-Resolution Properties for TSDA

3.1 Problem Setup

In this paper, we study unsupervised domain adaptation for time-series classification. Given a time-series sample $\mathbf{x}_i \in \mathbb{R}^{C\times T}$, where C and T denote the number of variables and time steps, respectively, our framework supports both univariate and multivariate settings. Under the standard UDA setting, we are given a labeled source dataset $\hat{P} = \{(\mathbf{x}_i^s, \mathbf{y}_i^s)\}_{i=1}^{n_s}$ and an unlabeled target dataset $\hat{Q} = \{(\mathbf{x}_i^t)\}_{i=1}^{n_t}$, sampled from two different distributions P and Q, respectively, where the superscripts s and t denote the source and target domains. The objective is to leverage the labeled source data together with the unlabeled target data to learn domain-invariant representations and a robust classifier that achieves high classification accuracy on the target domain.

3.2 Wavelet-Induced Multi-Resolution Hierarchy

Recent advances in unsupervised time-series domain adaptation have increasingly embraced joint temporal–frequency modeling paradigms, motivated by the complementary strengths of temporal and frequency representations (He et al., 2023; Liu et al., 2024). Specifically, temporal representations retain the most fine-grained view of the original sequence, while frequency representations provide a coarse-grained characterization (Zhou et al., 2022; Wu et al., 2022). However, such a design principle often relies on two extreme views of the same signal: temporal representations emphasize highly localized details, whereas frequency representations provide overly global abstractions. As a result, transitional patterns across different frequency bands are insufficiently characterized, limiting the ability to learn domain-invariant representations under distribution shift.

To address this limitation, we introduce the **Stationary Wavelet Transform (SWT)** (Nason & Silverman, 1995) to decompose the original sequence into a hierarchy of resolution-specific components, providing intermediate representations between fine-grained temporal details and coarse-grained spectral abstractions (Mallat, 1989). This hierarchical decomposition enables the model to explicitly capture how transferable and discriminative patterns evolve across different resolutions containing distinct frequency components, thereby supporting progressive domain-invariant learning. Specifically, given an input time series $\mathbf{x}_i \in \mathbb{R}^{C\times L}$, we apply the SWT and retain only the approximation coefficients at each level. The multi-resolution decomposition is defined as:

$$\{\mathbf{ca}_{i,j}\}_{j=0}^{m} = \mathcal{F}_{\text{SWT}}(\mathbf{x}_i, \psi, m), \mathbf{CA}_i = \{\mathbf{ca}_{i,j} \mid j = 0, \ldots, m\}, \tag{1}$$

where m denotes the number of decomposition levels and ψ represents the selected wavelet basis. $\mathbf{ca}_{i,0}$ corresponds to the original time series of the i-th sample, while $\mathbf{ca}_{i,j} \in \mathbb{R}^{C\times L}$ denotes the approximation coefficients at the j-th level. $\mathbf{CA}_i$ denotes the multi-resolution representation of the i-th sample, consisting of approximation coefficient sequences across all decomposition levels. Since the above recursive multi-resolution decomposition process preserves the underlying semantic category of the original sequence, all elements in $\mathbf{CA}_i$ share the same ground-truth label as $\mathbf{x}_i$.

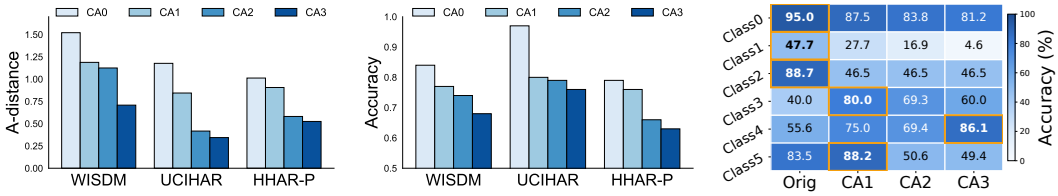
3.3 How to Exploit the Multi-Resolution Hierarchy for TSDA?

Constructing a multi-resolution hierarchy alone does not explain how it should be exploited for TSDA. To answer this question, we investigate the intrinsic relationships across resolutions from both representation and decision perspectives. Our analysis reveals two key insights that directly motivate the bridging mechanisms introduced in Section 4.

How do transferability and discriminability vary across different resolutions? Intuitively, lower-resolution representations primarily emphasize low-frequency components that capture global trends and coarse-grained variations, and thus are expected to exhibit stronger transferability across domains (Xu et al., 2019). In contrast, higher-resolution representations preserve a broader range of frequency information from the original sequence, retaining richer temporal details and thus providing stronger discriminability. To validate this intuition, we investigate the relationship between transferability and discriminability across the multi-resolution hierarchy.

According to domain adaptation theory (Ben-David et al., 2010), the target risk is jointly determined by the source risk, the domain discrepancy between the source and target domains, and the shared optimal error. Since the source risk is minimized through supervised learning, reducing domain discrepancy becomes a key objective for effective TSDA. We therefore adopt the A-distance (Ben-David et al., 2010) as a practical proxy for domain discrepancy, defined as $d_A = 2(1-2\epsilon)$, where ϵ is the error rate of a domain classifier distinguishing source and target samples. A smaller A-distance indicates stronger domain invariance and better transferability across domains.

As shown in Figure 1a, the A-distance decreases with increasing decomposition levels, indicating progressively improved domain alignment and stronger transferability of lower-resolution representations across domains. In contrast, Figure 1b shows that classification accuracy does not improve monotonically, suggesting that discriminative information is gradually weakened during decomposition and that higher-resolution representations generally retain stronger discriminability. These results validate our intuition and motivate us to bridge these complementary properties across resolutions through knowledge distillation-based mutual learning, as described in Section 4.2.



(a) d_A across CA levels (b) Accuracy in target domain (c) Class accuracy heatmap

Figure 1: Multi-resolution property analysis on the WISDM, UCIHAR, and HHAR-P datasets in DANN. (a) A-distance across the resolution hierarchy. (b) Target-domain accuracy across the resolution hierarchy. (c) Category-wise accuracy across the resolution hierarchy on the HHAR-P dataset.

How do these resolution-dependent differences affect the category-level decision strengths? Since different categories may be characterized by different dominant frequency components, different resolutions can naturally exhibit distinct category-level decision strengths. The original resolution preserves the complete frequency information of the input sequence and therefore has stronger overall discriminability, suggesting that it should play a more dominant role in the decision-making process. In contrast, lower-resolution representations focus on a subset of frequency components, which may nevertheless provide important discriminative cues for specific categories.

To verify this property, we analyze category-wise classification performance across resolutions. As shown in Figure 1b, the original resolution achieves the highest overall classification accuracy, indicating its dominant role in classification. However, Figure 1c shows that different resolutions exhibit distinct category-wise decision preferences, with lower-resolution representations outperforming the original resolution on specific categories, highlighting their complementary role in classification. This observation indicates that different resolutions provide distinct category-level decision strengths and motivate us to bridge the cross-resolution decision hierarchy through anchor-guided committee consensus, as described in Section 4.3.

4 APPROACH

Figure 2 illustrates the overall framework of the proposed model. Given an input time series $\mathbf{x}_i$, we first perform Multi-Resolution Decomposition via the stationary wavelet transform, and then obtain multi-resolution representations in Section 4.1. To further exploit the complementary transferability and discriminability across resolutions, we propose Cross-Resolution Knowledge Transfer (CRT) in Section 4.2. To achieve cross-resolution decision consistency, we propose Committee-based Consensus Training (CCT) in Section 4.3.

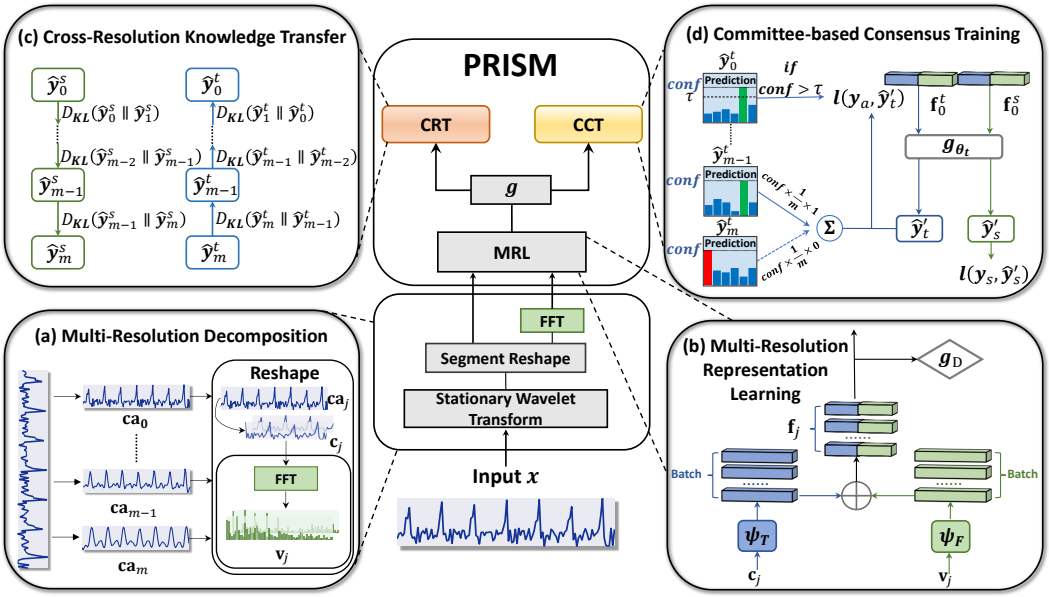


Figure 2: The overall architecture of PRISM. For clarity, the figure illustrates the processing procedure of a single sample within a mini-batch. (a) Multi-Resolution Decomposition first generates layer-wise approximation coefficients, (b) followed by Multi-Resolution Representation Learning to extract hybrid temporal–frequency features at each resolution level. (c) Cross-Resolution Knowledge Transfer serves as a key component to progressively propagate complementary knowledge across resolutions, achieving a better balance between transferability and discriminability. (d) Committee-based Consensus Training further organizes resolution-specific classifiers into a committee to provide committee-guided target supervision through cross-resolution consensus, leading to more consistent cross-domain decisions.

4.1 MULTI-RESOLUTION REPRESENTATION LEARNING

Given the multi-resolution sequence group $\mathbf{CA}_i$, we first construct resolution-wise temporal and frequency representations for subsequent adaptation. To improve efficiency on long time-series data, we adopt a periodic segmentation strategy (Wu et al., 2022). Specifically, the Fast Fourier Transform (FFT) is applied to identify the dominant frequency f_{peak} with the largest amplitude, and the corresponding period is computed as $p = \left\lceil \frac{T}{f_{\text{peak}}} \right\rceil$, where T denotes the original sequence length. Each resolution-level sequence is then independently reshaped into periodic segments. In parallel, FFT is further applied to the reshaped sequence to obtain frequency representations:

$$\mathbf{c}_{i,j} = \text{Reshape}_p(\mathbf{ca}_{i,j}), \quad \mathbf{v}_{i,j} = \text{Avg}\big(\text{FFT}(\mathbf{c}_{i,j})\big), \tag{2}$$

where $\mathbf{c}_{i,j} \in \mathbb{R}^{C \times f_{\text{peak}} \times p}$ and $\mathbf{v}_{i,j} \in \mathbb{C}^{C \times \lfloor \frac{p}{2} \rfloor}$ is obtained by averaging over the f_{peak} segments. Accordingly, the temporal–frequency representation set is denoted as $\{(\mathbf{c}_{i,j}, \mathbf{v}_{i,j}, \mathbf{y}_i)\}_{j=0}^{m}$. Following prior TSDA studies (Ragab et al., 2023; Liu et al., 2024), we jointly model temporal and frequency representations. Given the processed input, each resolution-level sequence is independently fed into a temporal feature extractor ψ_T and a frequency feature extractor ψ_F, avoiding interference across different resolutions.

To exploit complementary temporal and frequency information, the extracted features are concatenated into a hybrid representation $\mathbf{f}_{i,j}$. We further employ a domain discriminator g_D to align source and target hybrid features through adversarial learning. The hybrid features at each resolution are passed through a classifier g to obtain resolution-specific predictions, which are then supervised on labeled source samples:

$$\begin{aligned} \mathbf{f}_{i,j} &= \text{Concat}\big(\psi_T(\mathbf{c}_{i,j}), \psi_F(\mathbf{v}_{i,j})\big), \\ \mathcal{L}_D &= \sum_{j=0}^{m} \Big(-\mathbb{E}_{(\mathbf{c}^s_{i,j}, \mathbf{v}^s_{i,j}) \sim \hat{P}} \left[\log \left(g_D(\mathbf{f}^s_{i,j}) \right) \right] - \mathbb{E}_{(\mathbf{c}^t_{i,j}, \mathbf{v}^t_{i,j}) \sim \hat{Q}} \left[\log \left(1 - g_D(\mathbf{f}^t_{i,j}) \right) \right] \Big), \\ \mathcal{L}_C &= \sum_{j=0}^{m} \mathbb{E}_{(\mathbf{c}^s_{i,j}, \mathbf{v}^s_{i,j}, \mathbf{y}^s_i) \sim \hat{P}} \left[\ell \left(g(\mathbf{f}^s_{i,j}), \mathbf{y}^s_i \right) \right], \end{aligned} \tag{3}$$

where $\mathcal{L}_D$ is minimized with respect to g_D but maximized with respect to ψ_T and ψ_F.

4.2 Cross-Resolution Knowledge Transfer

Motivated by the finding that lower-resolution representations exhibit stronger transferability while higher-resolution representations retain stronger discriminability, we propose to explicitly transfer complementary knowledge across resolutions. Specifically, we employ the Kullback–Leibler (KL) divergence to enforce prediction consistency across adjacent resolutions, whose directional property naturally enables asymmetric guidance between representations. Since the wavelet transform performs hierarchical decomposition, with each coarser approximation coefficient generated by low-pass filtering the previous level, neighboring resolutions exhibit relatively small discrepancies. Such natural continuity provides a stable basis for progressive cross-resolution knowledge transfer.

Based on this principle, we further instantiate the mechanism in a domain-aware manner. In the source domain, finer-resolution predictions are treated as references to guide adjacent coarser levels, preserving discriminative semantic information during abstraction. In the target domain, coarser resolutions with stronger transferability are leveraged to regularize finer levels for more stable adaptation. Consequently, the model jointly enhances discriminability and transferability across resolutions. The resulting losses are formulated as:

$$\begin{aligned} \mathcal{L}_{M_s}(\psi_T, \psi_F, g) &= \sum_{j=0}^{m-1} \mathbb{E}_{(\mathbf{c}^s_{i,j}, \mathbf{v}^s_{i,j}) \sim \hat{P}} \left[D_{KL}\big(\hat{\mathbf{y}}^s_{i,j} \parallel \hat{\mathbf{y}}^s_{i,j+1}\big) \right], \\ \mathcal{L}_{M_t}(\psi_T, \psi_F, g) &= \sum_{j=0}^{m-1} \mathbb{E}_{(\mathbf{c}^t_{i,j}, \mathbf{v}^t_{i,j}) \sim \hat{Q}} \left[D_{KL}\big(\hat{\mathbf{y}}^t_{i,j+1} \parallel \hat{\mathbf{y}}^t_{i,j}\big) \right], \end{aligned} \tag{4}$$

where $\hat{\mathbf{y}}^s_{i,j}$ and $\hat{\mathbf{y}}^t_{i,j}$ denote the predictions at the j-th resolution level of the i-th sample in the source and target domains, respectively. The total multi-resolution knowledge transfer loss is:

$$\mathcal{L}_M = \mathcal{L}_{M_s} + \mathcal{L}_{M_t}. \tag{5}$$

4.3 Committee-based Consensus Training

Motivated by the finding that different resolutions exhibit distinct category-level decision strengths, we propose an anchor-guided decision strategy together with a committee consensus mechanism to exploit these complementary strengths across the resolution hierarchy. Considering the unsupervised nature of the target domain, directly using weighted aggregation or majority voting across multiple resolutions as the final prediction may become unreliable under large domain shifts and adversely affect training stability. Therefore, we introduce a dedicated target-domain classifier and leverage anchor-guided committee consensus to calibrate its decision learning across resolutions. Committee-guided target supervision is then coupled with labeled source supervision to construct a cross-domain calibration loss, enabling more stable domain-invariant decision learning. During training, the anchor decision for each target sample is obtained as

$$\mathbf{y}_a = \arg\max_c \ \hat{\mathbf{y}}^t_{i,0}[c], \quad \hat{p} = \max_c \ \hat{\mathbf{y}}^t_{i,0}[c], \tag{6}$$

where $\hat{\mathbf{y}}^t_{i,0}[c]$ denotes the predicted probability of class c, and $\hat{p}$ represents the confidence of the anchor committee member. To obtain reliable anchor decisions for the target-domain classifier under domain shift, we regularize the anchor prediction using the Tsallis entropy (Tsallis, 1988),

$$\mathcal{L}_{\text{Tsallis}}(\psi_T, \psi_F, g) = \mathbb{E}_{(\mathbf{c}^t_{i,0}, \mathbf{v}^t_{i,0})\sim\hat{Q}} \left[S_\alpha\big(g(\mathbf{f}^t_{i,0})\big) \right], \tag{7}$$

where $S_\alpha(\cdot)$ denotes the Tsallis entropy defined over the predicted probability distribution, and α controls its non-extensivity. We further retain only reliable anchor decisions using a predefined confidence threshold τ,

$$\mathbf{M} = \begin{cases} 1, & \hat{p} \geq \tau, \\ 0, & \text{otherwise}, \end{cases} \tag{8}$$

where $\mathbf{M}$ is a binary mask indicating reliable anchor decisions. To leverage the auxiliary decision strengths of the remaining resolutions, the resolution-specific classifiers participate as committee members to evaluate the anchor decision. Each committee member contributes only when its prediction agrees with the anchor committee member, and its contribution is weighted by its prediction confidence. The resulting committee consensus score is defined as

$$\mathbf{w} = \frac{1}{m} \sum_{j=1}^{m} \max_c \hat{\mathbf{y}}^t_{i,j}[c] \cdot \mathbb{I}\Big(\arg\max_c \hat{\mathbf{y}}^t_{i,0}[c] = \arg\max_c \hat{\mathbf{y}}^t_{i,j}[c] \Big), \tag{9}$$

where $\mathbf{w}$ denotes the committee consensus score, reflecting the confidence-weighted agreement between the anchor committee member and auxiliary committee members. The effective supervision weight is computed as $\tilde{\mathbf{w}} = \mathbf{M} \odot \mathbf{w}$, jointly considering the reliability of the anchor decision and the consensus support across resolutions. Guided by the committee consensus, the target-domain classifier g_{θ_t} is optimized as

$$\hat{\mathbf{y}}'_t = g_{\theta_t}(\mathbf{f}^t_{i,0}), \quad \mathcal{L}_g(g_{\theta_t}) = \mathbb{E}_{(\mathbf{c}^t_{i,0}, \mathbf{v}^t_{i,0})\sim\hat{Q}} \left[\tilde{\mathbf{w}}\, \ell(\hat{\mathbf{y}}'_t, \mathbf{y}_a) \right]. \tag{10}$$

To preserve consistent decision behavior across domains, we further couple the committee-guided target predictions with labeled source samples to construct the Committee-guided supervision loss:

$$\hat{\mathbf{y}}'_s = g_{\theta_t}(\mathbf{f}^s_{i,0}), \quad \mathcal{L}_{\text{sup}}(\psi_T, \psi_F, g) = \mathbb{E}_{(\mathbf{c}^s_{i,0}, \mathbf{v}^s_{i,0})\sim\hat{P}} \left[\ell(\hat{\mathbf{y}}'_s, \mathbf{y}_s) \right]. \tag{11}$$

where $\mathbf{y}_s$ denotes the ground-truth source-domain label. The Committee-guided supervision loss couples committee-guided target predictions with labeled source samples to obtain more stable domain-invariant decision learning.

4.4 Overall Training Objective

During training, we optimize the temporal feature extractor ψ_T, frequency feature extractor ψ_F, and classifier g by minimizing the supervised source-domain classification loss $\mathcal{L}_C$. Meanwhile, a Gradient Reversal Layer is employed to enable adversarial learning with the domain discriminator g_D through $\mathcal{L}_D$. We further introduce $\mathcal{L}_M$ to promote multi-resolution domain-invariant representation learning across resolutions and domains. For Committee-based Consensus Training, the Tsallis entropy loss $\mathcal{L}_{\text{Tsallis}}$ regularizes the confidence of the anchor committee member, while $\mathcal{L}_g$ optimizes the target-domain classifier under committee-guided target supervision. In addition, the committee-guided supervision loss $\mathcal{L}_{\text{sup}}$ couples the committee-guided target predictions with labeled source samples, encouraging consistent semantic decisions across source and target domains. For simplicity, we denote all parameters of ψ_T, ψ_F and g as θ. The overall objective is formulated as

$$\begin{aligned} &\min_{\theta} \ \mathcal{L}_C(\theta) + \mathcal{L}_M(\theta) + \mathcal{L}_{\text{Tsallis}}(\theta) + \mathcal{L}_{\text{sup}}(\theta) - \mathcal{L}_D(\theta, g_D), \\ &\min_{g_D} \ \mathcal{L}_D(\theta, g_D), \qquad \min_{g_{\theta_t}} \ \mathcal{L}_g(\theta, g_{\theta_t}). \end{aligned} \tag{12}$$

5 Experiments

5.1 Setup

Datasets Extensive experiments are conducted on six datasets: (1) UCIHAR (Anguita et al., 2013), HHAR-P (Ragab et al., 2023; Stisen et al., 2015), WISDM (Kwapisz et al., 2011), and

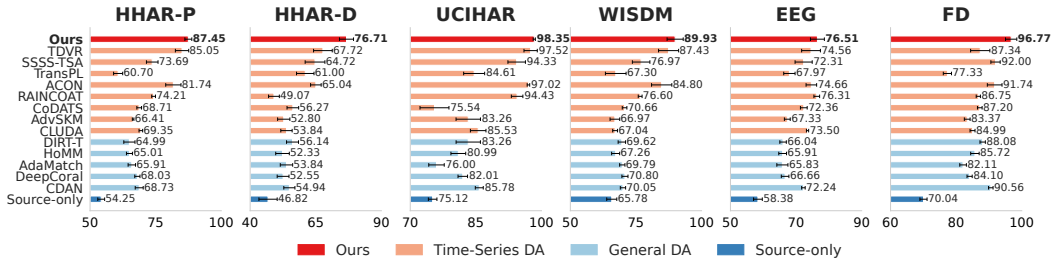


Figure 3: Average accuracy (%) across multiple datasets.

Table 1: Accuracy (%) on HHAR-D for unsupervised domain adaptation.

Method	0→1	0→2	0→3	0→4	1→0	1→3	1→4	2→1	3→4	4→1	Avg ± Std
Source-only	65.48	33.59	31.71	39.79	34.69	44.83	49.54	38.17	86.17	44.23	46.82 ± 3.48
CDAN	69.86	48.28	38.22	48.42	48.75	60.48	51.33	47.84	87.33	48.89	54.94 ± 2.12
DeepCoral	68.94	42.88	40.67	47.96	35.63	55.31	56.21	44.71	87.25	45.96	52.55 ± 2.25
AdaMatch	71.78	39.60	39.74	47.50	52.50	55.48	58.33	46.49	85.83	41.15	53.84 ± 2.36
HoMM	69.66	40.51	39.16	50.42	35.94	55.02	57.13	42.36	86.79	46.35	52.33 ± 2.59
DIRT-T	68.37	42.14	47.21	52.92	41.25	60.14	55.63	46.73	92.25	54.81	56.14 ± 2.28
CLUDA	71.78	39.60	39.74	47.50	52.50	55.48	58.33	46.49	85.83	41.15	53.84 ± 2.16
AdvSKM	67.93	40.71	40.19	47.33	37.19	55.65	59.54	42.69	87.46	49.33	52.80 ± 2.34
CoDATS	72.50	43.35	50.79	45.50	58.44	62.24	54.54	40.14	89.63	45.53	56.27 ± 2.22
RAINCOAT	74.47	36.52	48.82	35.29	51.25	41.49	41.50	34.28	88.58	38.46	49.07 ± 2.05
ACON	77.50	61.36	54.69	65.46	69.38	71.30	62.13	50.10	93.63	44.86	65.04 ± 2.02
TransPL	75.43	55.74	53.68	57.71	66.56	52.33	55.96	48.41	94.54	49.66	61.00 ± 3.33
SSSS-TSA	**84.66**	42.97	40.67	48.71	69.38	66.78	70.92	61.20	93.54	64.04	64.29 ± 3.76
TDVR	79.95	55.29	58.27	52.88	**80.63**	71.66	68.75	52.07	**99.29**	58.41	67.72 ± 3.56
Ours	82.31	**68.55**	**65.84**	**65.54**	78.44	**82.64**	**79.71**	**73.27**	94.50	**76.35**	**76.71 ± 2.68**

HHAR-D (Gagnon-Audet et al., 2023; Stisen et al., 2015) for sensor-based human activity recognition tasks; (2) FD (Ragab et al., 2023) for machine fault diagnosis; (3) EEG (Goldberger et al., 2000) for sleep stage classification. Consistent with existing domain adaptation methods in time series research (Berthelot et al., 2021; He et al., 2023; Liu et al., 2024), we sample 10 source-target domain pairs for each dataset to ensure comprehensive evaluation. Additional details regarding dataset processing and domain division are provided in Appendix A.

Baselines (1) We first report the source-only performance in the temporal domain to verify the effectiveness of UDA methods. (2) We compare with state-of-the-art TSDA methods, including CoDATS (Wilson et al., 2020), AdvSKM (Liu & Xue, 2021), CLUDA (Ozyurt et al., 2022), RAINCOAT (He et al., 2023), ACON (Liu et al., 2024), TransPL (Kim & Lee, 2025), SSSS-TSA (Ahad et al., 2025), and TDVR (Zhou et al., 2025). (3) We further include representative general UDA methods, such as CDAN (Long et al., 2018), DeepCoral (Sun & Saenko, 2016), AdaMatch (Berthelot et al., 2021), HoMM (Chen et al., 2020), and DIRT-T (Shu et al., 2018). For all baselines, we follow the original settings and ensure fair comparison by using consistent feature extractors, architectures, and hyperparameters as reported in the respective papers.

Implementation We adopt the implementation of AdaTime (Ragab et al., 2023) as our benchmark framework. We adopt CNN as the temporal feature extractor and a complex-valued linear layer as the frequency feature extractor. Evaluation is performed using classification accuracy and Macro-F1 score computed on the target test set. All experiments are repeated five times with different random seeds, and the average results are reported. Detailed model architectures and hyperparameters are provided in Appendix B.

5.2 Results

Figure 3 reports the average accuracy across 10 source-target pairs over six datasets and three scenarios. Consistent with our motivation, our proposed method consistently outperforms the best baseline, with notable gains on several datasets. Notably, PRISM outperforms all baselines across all six

evaluation tasks, achieving an average relative accuracy improvement of 4.21% over the strongest baseline while consistently exhibiting smaller standard-deviation error bars, demonstrating more reliable and stable performance. Specifically, our model enhances relative accuracy improvements by 2.82% on the HHAR-P dataset, 5.18% on the FD dataset, and 13.28% on the HHAR-D dataset, respectively, compared to the state-of-the-art baseline. Due to space limitations, we report the accuracy metric for selected source-target domain pairs, together with the average accuracy and the mean standard deviation on the HHAR-D dataset in Table 1. More accuracy results are given in Table 8–12 and the average Macro-F1 score is provided in Figure 10.

5.3 Analysis

Computational Cost Table 2 reports the training time and accuracy on HHAR-D. Although multi-resolution modeling introduces additional computational overhead, PRISM achieves the best accuracy of 76.71%, substantially outperforming both TDVR and ACON. Meanwhile, its training time remains competitive, being faster than both RAINCOAT and TDVR and only moderately slower than ACON. These results show that PRISM achieves a strong efficiency–performance trade-off.

Table 2: Training time and performance on HHAR-D.

HHAR-D	RAINCOAT	ACON	TDVR	Ours
Accuracy (%)	49.07	65.04	67.72	**76.71**
Training Time	1h7m47s	39m4s	1h2m32s	56m30s

Ablation Study We conduct ablation experiments on three datasets, UCIHAR, HHAR-P, and WISDM. For each dataset, we select the same 10 source-target domain pairs as described in Section 5.1. The results are reported in Table 3, while detailed experimental settings and additional analyses are provided in Appendix C.1. All proposed modules are effective. Each individual module improves performance (Rows 2–4), pairwise combinations yield further gains (Rows 5–7), and full integration (Row 8) achieves the best overall results.

Table 3: Ablation of the modules.

	MRL	CRT	CCT	UCIHAR	HHAR-P	WISDM	Average
1	-	-	-	95.13	76.72	76.48	82.78
2	✓	-	-	96.39	80.63	83.40	86.81
3	-	✓	-	97.04	83.58	86.01	88.88
4	-	-	✓	97.50	77.63	85.98	87.04
5	✓	✓	-	98.19	86.64	86.65	90.49
6	✓	-	✓	98.00	83.01	86.38	89.13
7	-	✓	✓	97.73	84.96	88.64	90.44
8	✓	✓	✓	**98.35**	**87.45**	**89.93**	**91.91**

Module Effectiveness Study To further investigate the proposed Cross-Resolution Knowledge Transfer module, we conduct additional studies on the directionality and symmetry of knowledge transfer across resolutions. Detailed results and analysis are provided in Appendix C.2. We further analyze whether Committee-based Consensus Training effectively exploits the hierarchical decision strengths across different resolutions and improves cross-resolution decision consistency. Detailed results and analysis are reported in Appendix C.3.

Sensitivity Analysis. To investigate the robustness of PRISM to hyperparameter choices, we extend the optimization objective with explicit trade-off coefficients and conduct sensitivity analyses on both the reformulated loss and key modeling hyperparameters. Detailed experimental settings and results are provided in Appendix C.4. The results demonstrate that PRISM is robust to hyperparameter variations while consistently maintaining superior performance across different settings.

6 Conclusion

In this paper, we propose PRISM, a TSDA framework that moves beyond conventional temporal–frequency modeling under two extreme views by introducing a wavelet-induced resolution hierarchy for domain-invariant representation learning. To effectively model the wavelet-induced resolution hierarchy for TSDA, we investigate the intrinsic relationships across resolutions and reveal two key insights for domain-invariant representation learning. Guided by these insights, PRISM introduces Cross-Resolution Knowledge Transfer to progressively bridge complementary properties across resolutions, and Committee-based Consensus Training to bridge cross-resolution decision consistency through anchor-guided committee calibration. Extensive experiments on diverse datasets demonstrate that PRISM consistently outperforms state-of-the-art methods, validating the effectiveness of bridging resolutions for domain-invariant representation learning in time-series domain adaptation.

AI USE STATEMENT

In this work, we used generative AI tools solely for language editing and writing refinement, including improving grammar, wording, sentence fluency, and overall readability of the manuscript. We did not use generative AI tools for the conception of research ideas, the design of the proposed methodology, experimental design, data collection, data analysis, result interpretation, or the generation of scientific claims and conclusions. These aspects of the work were developed and conducted entirely by the authors. All AI-assisted edits were carefully reviewed and revised by the authors to ensure that the technical content, terminology, and scientific meaning remained accurate and consistent with the original work. We take full responsibility for the final content of this paper, including all text, methods, experimental results, and conclusions.

REPRODUCIBILITY STATEMENT

We make every effort to ensure the reproducibility of this work. The complete methodology of PRISM is described in Section 4. The benchmark datasets, data sources, preprocessing procedures, and evaluation protocols are documented in Appendix A. Detailed experimental architectures, implementation settings, and training procedures are provided in Appendix B. The complete hyperparameter configurations and sensitivity analyses are reported in Appendix B and Appendix C.4. To facilitate independent verification and reproduction of all experiments, we provide an anonymous implementation, training scripts, and evaluation code in the supplementary anonymous repository: `https://anonymous.4open.science/r/PRISM-E7BC/`. All experiments are conducted using publicly available benchmark datasets under consistent preprocessing and evaluation protocols as described in the paper.

REFERENCES

Nauman Ahad, Mark A. Davenport, and Eva L. Dyer. Time series domain adaptation via channel-selective representation alignment. *TMLR*, 2025.

Davide Anguita, Alessandro Ghio, Luca Oneto, Xavier Parra, Jorge Luis Reyes-Ortiz, et al. A public domain dataset for human activity recognition using smartphones. In *ESANN*, 2013.

Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan. A theory of learning from different domains. *Machine Learning*, 2010.

David Berthelot, Rebecca Roelofs, Kihyuk Sohn, Nicholas Carlini, and Alexey Kurakin. Adamatch: A unified approach to semi-supervised learning and domain adaptation. In *ICLR*, 2021.

Ruichu Cai, Jiawei Chen, Zijian Li, Wei Chen, Keli Zhang, Junjian Ye, Zhuozhang Li, Xiaoyan Yang, and Zhenjie Zhang. Time series domain adaptation via sparse associative structure alignment. In *AAAI*, 2021.

Chao Chen, Zhihang Fu, Zhihong Chen, Sheng Jin, Zhaowei Cheng, Xinyu Jin, and Xian-Sheng Hua. Homm: Higher-order moment matching for unsupervised domain adaptation. In *AAAI*, 2020.

Chaoqi Chen, Weiping Xie, Wenbing Huang, Yu Rong, Xinghao Ding, Yue Huang, Tingyang Xu, and Junzhou Huang. Progressive feature alignment for unsupervised domain adaptation. In *CVPR*, 2019.

Hassan Ismail Fawaz, Germain Forestier, Jonathan Weber, Lhassane Idoumghar, and Pierre-Alain Muller. Deep learning for time series classification: A review. *Data Mining and Knowledge Discovery*, 2019.

Jean-Christophe Gagnon-Audet, Kartik Ahuja, Mohammad Javad Darvishi Bayazi, Pooneh Mousavi, Guillaume Dumas, and Irina Rish. Woods: Benchmarks for out-of-distribution generalization in time series. *TMLR*, 2023.

Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario March, and Victor Lempitsky. Domain-adversarial training of neural networks. *JMLR*, 2016.

Ary L. Goldberger, Luis A. Nunes Amaral, L. Glass, Jeffrey M. Hausdorff, Plamen Ch. Ivanov, Roger G. Mark, Joseph E. Mietus, George B. Moody, Chung-Kang Peng, and Harry Eugene Stanley. Physiobank, physiotoolkit, and physionet: Components of a new research resource for complex physiologic signals. *Circulation*, 2000.

Xiaoqing Guo, Chen Yang, Baopu Li, and Yixuan Yuan. Metacorrection: Domain-aware meta loss correction for unsupervised domain adaptation in semantic segmentation. In *CVPR*, 2021.

Huan He, Owen Queen, Teddy Koker, Consuelo Cuevas, Theodoros Tsiligkaridis, and Marinka Zitnik. Domain adaptation for time series under feature and label shifts. In *ICML*, 2023.

Jaeho Kim and Seulki Lee. Transpl: Vq-code transition matrices for pseudo-labeling of time series unsupervised domain adaptation. In *ICML*, 2025.

Jennifer R. Kwapisz, Gary M. Weiss, and Samuel A. Moore. Activity recognition using cell phone accelerometers. In *ACM SIGKDD*, 2011.

Hong Liu, Jianmin Wang, and Mingsheng Long. Cycle self-training for domain adaptation. In *NeurIPS*, 2021.

Mingyang Liu, Xinyang Chen, Yang Shu, Xiucheng Li, Weili Guan, and Liqiang Nie. Boosting transferability and discriminability for time series domain adaptation. In *NeurIPS*, 2024.

Qiao Liu and Hui Xue. Adversarial spectral kernel matching for unsupervised time series domain adaptation. In *IJCAI*, 2021.

Mingsheng Long, Yue Cao, Jianmin Wang, and Michael I. Jordan. Learning transferable features with deep adaptation networks. In *ICML*, 2015.

Mingsheng Long, Zhangjie Cao, Jianmin Wang, and Michael I. Jordan. Conditional adversarial domain adaptation. In *NeurIPS*, 2018.

Stephane G. Mallat. A theory for multiresolution signal decomposition: The wavelet representation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1989.

Guy P. Nason and Bernard Walter Silverman. The stationary wavelet transform and some statistical applications. *Lecture Notes in Statistics*, 1995.

Yilmazcan Ozyurt, Stefan Feuerriegel, and Ce Zhang. Contrastive learning for unsupervised domain adaptation of time series. In *ICLR*, 2022.

Sanjay Purushotham, Wilka Carvalho, Tanachat Nilanon, and Yan Liu. Variational recurrent adversarial deep domain adaptation. In *ICLR*, 2017.

Mohamed Ragab, Emadeldeen Eldele, Wee Ling Tan, Chuan-Sheng Foo, Zhenghua Chen, Min Wu, Chee-Keong Kwoh, and Xiaoli Li. Adatime: A benchmarking suite for domain adaptation on time series data. *ACM TKDD*, 2023.

Xingjian Shi, Zhourong Chen, Hao Wang, Dit-Yan Yeung, Wai-Kin Wong, and Wang-chun Woo. Convolutional lstm network: A machine learning approach to precipitation nowcasting. In *NeurIPS*, 2015.

Rui Shu, Hung Bui, Hirokazu Narui, and Stefano Ermon. A dirt-t approach to unsupervised domain adaptation. In *ICLR*, 2018.

Allan Stisen, Henrik Blunck, Sourav Bhattacharya, Thor Siiger Prentow, Mikkel Baun Kjærgaard, Anind Dey, Tobias Sonne, and Mads Møller Jensen. Smart devices are different: Assessing and mitigating mobile sensing heterogeneities for activity recognition. In *ACM SenSys*, 2015.

Baochen Sun and Kate Saenko. Deep coral: Correlation alignment for deep domain adaptation. In *ECCV*, 2016.

Constantino Tsallis. Possible generalization of boltzmann-gibbs statistics. *Journal of Statistical Physics*, 1988.

Eric Tzeng, Judy Hoffman, Ning Zhang, Kate Saenko, and Trevor Darrell. Deep domain confusion: Maximizing for domain invariance. *arXiv preprint arXiv:1412.3474*, 2014.

Dmitry Ulyanov, Andrea Vedaldi, and Victor S. Lempitsky. Instance normalization: The missing ingredient for fast stylization. *arXiv preprint arXiv:1607.08022*, 2016.

Shiyu Wang, Haixu Wu, Xiaoming Shi, Tengge Hu, Huakun Luo, Lintao Ma, James Y. Zhang, and Jun Zhou. Timemixer: Decomposable multiscale mixing for time series forecasting. In *ICLR*, 2024.

Shiyu Wang, Jiawei Li, Xiaoming Shi, Zhou Ye, Baichuan Mo, Wenze Lin, Shengtong Ju, Zhixuan Chu, and Ming Jin. Timemixer++: A general time series pattern machine for universal predictive analysis. In *ICLR*, 2025.

Yunshi Wen, Tengfei Ma, Tsui-Wei Weng, Lam M. Nguyen, and Anak Agung Julius. Abstracted shapes as tokens: A generalizable and interpretable model for time-series classification. In *NeurIPS*, 2024.

Garrett Wilson, Janardhan Rao Doppa, and Diane J. Cook. Multi-source deep domain adaptation with weak supervision for time-series sensor data. In *KDD*, 2020.

Haixu Wu, Tengge Hu, Yong Liu, Hang Zhou, Jianmin Wang, and Mingsheng Long. Timesnet: Temporal 2d-variation modeling for general time series analysis. In *ICML*, 2022.

Zhi-Qin John Xu, Yaoyu Zhang, Tao Luo, Yanyang Xiao, and Zheng Ma. Frequency principle: Fourier analysis sheds light on deep neural networks. *arXiv preprint arXiv:1901.06523*, 2019.

Yuchen Zhang, Mingsheng Long, Kaiyuan Chen, Lanxiang Xing, Ronghua Jin, Michael I. Jordan, and Jianmin Wang. Skilful nowcasting of extreme precipitation with nowcastnet. *Nature*, 2023.

Yue Zhang, Shun Miao, Tommaso Mansi, and Rui Liao. Task driven generative modeling for unsupervised domain adaptation: Application to x-ray image segmentation. In *MICCAI*, 2018.

Kexuan Zhou, Jiajing Geng, Boliang Hao, and Bailing Zhang. Transferability-driven variable recalibration for multivariate time series domain adaptation. In *ECAI*, 2025.

Tian Zhou, Ziqing Ma, Qingsong Wen, Xue Wang, Liang Sun, and Rong Jin. Fedformer: Frequency enhanced decomposed transformer for long-term series forecasting. In *ICML*, 2022.

Yang Zou, Zhiding Yu, B.V.K. Vijaya Kumar, and Jinsong Wang. Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In *ECCV*, 2018.

A DATASETS

A.1 DETAILED STATISTICS

We conduct extensive experiments using a wide range of time series datasets. The detailed statistics for each dataset are included in Table 4. For the HHAR-D dataset, we use the processed version released by WOODS (Gagnon-Audet et al., 2023). For the UCIHAR, HHAR-P, WISDM, EEG, and FD datasets, we use the processed versions released by AdaTime (Ragab et al., 2023).

Table 4: Summary of datasets.

Dataset	Subjects	Channels	Length	Class	Total	Task
EEG	20	1	3000	5	20590	SSC
FD	4	1	5120	3	10916	FD
UCIHAR	30	9	128	6	3290	HAR
HHAR-P	9	3	128	6	17934	HAR
WISDM	30	3	128	6	2070	HAR
HHAR-D	5	6	500	6	13674	HAR

A.2 DATASET PROCESSING

Each domain of each dataset is randomly divided into 80% training and 20% testing. Following AdaTime (Ragab et al., 2023), we apply Z-score normalization to both the training and testing splits using the statistics computed from the training split:

$$x_i^{\text{normalize}} = \frac{x_i - x_{\text{mean}}}{x_{\text{std}}}, \quad i = 1, 2, \ldots, N \tag{13}$$

where $N = N_s$ for the source domain data and $N = N_t$ for the target domain data. Note that both the training and testing splits are normalized based on the training set statistics only.

B EXPERIMENTAL DETAILS

B.1 MODEL ARCHITECTURE

Unlike existing methods, our approach employs a customized feature extractor to simultaneously capture multi-resolution time-series representations and facilitate knowledge transfer across different resolution levels. Multi-resolution sequences typically emphasize distinct frequency components, leading to different characteristics in terms of domain transferability and discriminability. Consequently, it is challenging to effectively model such heterogeneous multi-resolution features using a single CNN or MLP.

Inspired by recent models for multi-scale time-series analysis, which commonly employ scale-specific feed-forward layers followed by feature aggregation (Wang et al., 2024; 2025), and consistent with TSDA baselines (Ragab et al., 2023; Liu et al., 2024), we adopt resolution-specific feature extractors for each sequence level, consisting of an independent CNN for the temporal subspace and an independent MLP for the frequency subspace. To improve efficiency, only the magnitude information of frequency components is retained for modeling. Moreover, the coefficients at each resolution level are partitioned into periodic segments (Wu et al., 2022), and the original 1D-CNN is replaced with a lightweight 2D-CNN to process multiple segments simultaneously. This design substantially alleviates the additional overhead introduced by multi-resolution modeling. Finally, the extracted frequency features are projected to the same dimension as temporal features, enabling balanced hybrid representation learning across the two modalities.

B.2 MULTI-RESOLUTION IN TIME SERIES

To obtain multi-resolution representations of time series, we adopt stationary wavelet transform (Nason & Silverman, 1995) to decompose the original sequence. During training, the wavelet basis is selected from representative families, including Haar, Daubechies, and Symlets. Detailed selection

criteria are provided in Appendix C.4. To ensure stable optimization across different resolution levels, the approximation coefficients $\mathbf{ca}_{i,j}$ at each level are independently normalized using instance normalization (Ulyanov et al., 2016), which enforces consistent amplitude distributions across resolutions.

After preprocessing, features at each resolution level are independently extracted using separate feature extractors to avoid interference across resolutions, thereby enabling more reliable multi-resolution modeling. As shown in Figure 1b, the original-resolution sequence exhibits the strongest discriminability, while its transferability is substantially enhanced through multi-resolution learning, resulting in the best overall performance. Therefore, we adopt the prediction at the original resolution as the final output for model evaluation. To ensure a fair evaluation, hyperparameters are selected following an unsupervised protocol without accessing target-domain labels. All the experiments in this paper are conducted on a single NVIDIA A100-SXM4-40GB. The main configurations are summarized in Table 5, while the selection procedure and sensitivity analyses are detailed in Appendix C.4.

Table 5: Key hyperparameters

Hyperparameter	UCIHAR	HHAR-P	WISDM	HHAR-D	EEG	FD
Epoch	50	50	50	50	50	50
Batch Size	32	32	32	32	32	32
Decomposition levels m	3	1	4	1	1	1
Learning Rate	0.01	0.001	0.003	0.001	0.01	0.01

C Further Analysis

C.1 Ablation Study Analysis

We conduct ablation experiments on three datasets, UCIHAR, HHAR-P, and WISDM. For each dataset, we select the same 10 source-target domain pairs as described in Section 5.1. The ablation results are reported in Table 3.

To provide a fair evaluation of each proposed module, we adopt a simplified TSDA framework consisting of a single-resolution temporal feature extractor, a classifier, and a domain-adversarial discriminator as the baseline architecture. Each proposed module is then incorporated into this common framework individually or jointly, ensuring that the performance differences can be attributed solely to the introduced multi-resolution components.

Specifically, to evaluate the effectiveness of the proposed Multi-Resolution Representation Learning module, we replace the original single-resolution temporal feature extractor with the proposed multi-resolution temporal–frequency feature extractor. Comparing Rows 1, 2, and 8 demonstrates the contribution of multi-resolution representation learning. To assess the Cross-Resolution Knowledge Transfer module, we disable the cross-resolution transfer loss $\mathcal{L}_M$, thereby removing knowledge interaction among multi-resolution representations. The comparisons among Rows 1, 3, and 8 verify its contribution. To evaluate the Committee-based Consensus Training module, we remove the supervision terms that propagate committee-guided decision learning, namely $\mathcal{L}_{\text{Tsallis}}$ and $\mathcal{L}_{\text{sup}}$. The performance differences among Rows 1, 4, and 8 demonstrate the effectiveness of this module. Furthermore, Rows 5–7 consistently outperform their corresponding single-module variants (Rows 2–4), indicating that the proposed modules are complementary. When all components are jointly enabled (Row 8), PRISM achieves the best overall performance.

C.2 Further Analysis of Cross-Resolution Knowledge Transfer

To further evaluate the Cross-Resolution Knowledge Transfer module, we conduct controlled experiments from two perspectives: the effect of directional KL divergence and whether the performance gains arise from explicit cross-resolution transfer rather than generic distribution matching. For fair comparison, we use a simplified baseline consisting of the proposed multi-resolution temporal–frequency feature extractor, a classifier on fused features, and domain-adversarial alignment. The

Cross-Resolution Knowledge Transfer module is then incorporated without modifying the remaining components.

We first compare different KL divergence directions across resolutions. Table 6 shows that the proposed asymmetric strategy achieves the best performance. In the source domain, finer-resolution predictions supervise coarser levels to preserve discriminative information, whereas in the target domain, coarser-resolution predictions regularize finer levels to improve transferability. This observation is consistent with the analysis in Section 3.3. We further replace the directional KL divergence with several symmetric alternatives, including Jensen–Shannon divergence, symmetric KL divergence, and mean-squared error. Although these variants consistently outperform the baseline, they remain inferior to the proposed directional strategy, as shown in Table 6. This suggests that the improvement does not merely come from enforcing posterior consistency, but from modeling the distinct roles of different resolutions during adaptation.

Overall, the results support the effectiveness of cross-resolution knowledge transfer for balancing transferability and discriminability.

Table 6: Further analysis of Cross-Resolution Transfer on WISDM. We compare different transfer directions and symmetric alternatives under the same simplified baseline. "Fine→Coarse" means finer-resolution predictions guide coarser levels, while "Coarse→Fine" denotes the reverse direction.

Method	Source Direction	Target Direction	Divergence	Accuracy (%)
Baseline	–	–	–	83.40
Uniform F→C	Fine→Coarse	Fine→Coarse	KL	85.76
Uniform C→F	Coarse→Fine	Coarse→Fine	KL	85.23
Reversed Direction	Coarse→Fine	Fine→Coarse	KL	83.12
JS Matching	Bidirectional	Bidirectional	JS	85.14
Symmetric KL	Bidirectional	Bidirectional	SKL	83.93
L2 Consistency	Bidirectional	Bidirectional	MSE	85.23
Ours	Fine→Coarse	Coarse→Fine	KL	**86.01**

C.3 Further Analysis of Committee-based Consensus Training

To further understand the role of Committee-based Consensus Training (CCT), we examine whether CCT better exploits the hierarchical decision strengths across different resolutions than conventional consensus strategies. To isolate the effect of CCT from the other components in PRISM, we adopt a simplified baseline consisting only of temporal and frequency feature extractors, a classifier on fused features, and domain-adversarial alignment. Multi-resolution representations are constructed in the same way as PRISM, while all other modules are removed. All training settings and evaluation protocols remain identical to those in the main experiments.

Table 7: Further analysis of Committee-based Consensus Training on the WISDM, UCIHAR, and HHAR-P datasets.

Method	**WISDM**	**UCIHAR**	**HHAR-P**
Baseline	83.40	96.39	80.63
Confidence-Weighted Fusion	82.55	95.87	81.18
Anchor-only Calibration	84.91	97.57	82.54
CCT	86.38	98.00	83.01

We compare CCT with several representative multi-resolution consensus strategies, as summarized in Table 7. *Baseline* denotes the simplified TSDA model without cross-resolution decision learning. *Confidence-Weighted Fusion* directly aggregates predictions from all resolutions according to their class confidence to produce the final prediction. *Anchor-only Calibration* preserves the target-domain classifier together with the source–target calibration loss, but removes committee participation from the remaining resolutions, leaving only the original-resolution anchor for supervision. *CCT* denotes the full committee-based decision learning strategy proposed in this paper.

As shown in Table 7, Confidence-Weighted Fusion performs even worse than Baseline. Although Anchor-only Calibration provides moderate improvement over Baseline, it still falls considerably behind Committee-based Consensus Training.

These results are consistent with the insight in Section 3.3. Confidence-Weighted Fusion treats all resolutions equally and ignores the dominant semantic role of the original resolution, allowing less discriminative lower-resolution predictions to interfere with the anchor decision. In contrast, Anchor-only Calibration fails to exploit the category-specific decision strengths provided by the remaining resolutions, preventing the anchor classifier from benefiting from semantic cues preserved at lower resolutions. Committee-based Consensus Training explicitly models this cross-resolution decision hierarchy, preserving the dominant role of the original resolution while selectively incorporating category-specific semantic evidence from the committee, resulting in the best overall performance.

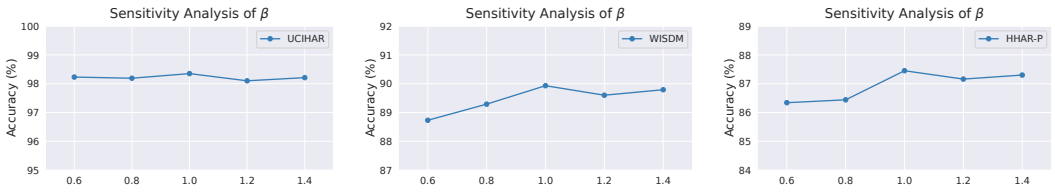
C.4 Sensitivity Analysis

To investigate the sensitivity of the proposed optimization objective to different trade-off hyperparameters, we reformulate the overall loss by explicitly introducing weighting coefficients for the Cross-Resolution Knowledge Transfer loss and the Committee-guided Supervision loss. The optimization objective in equation 12 is rewritten as follows:

$$\begin{aligned} &\min_{\theta} \ \mathcal{L}_C(\theta) + \beta\mathcal{L}_M(\theta) + \mathcal{L}_{\text{Tsallis}}(\theta) + \gamma\mathcal{L}_{\text{sup}}(\theta) - \mathcal{L}_D(\theta, g_D), \\ &\min_{g_D} \ \mathcal{L}_D(\theta, g_D), \qquad \min_{g_{\theta_t}} \ \mathcal{L}_g(\theta, g_{\theta_t}), \end{aligned} \tag{14}$$

Here, β and γ are trade-off hyperparameters controlling the contributions of the Cross-Resolution Knowledge Transfer loss and the Committee-guided supervision loss, respectively. We vary these two coefficients independently while keeping all other training settings unchanged to evaluate the robustness of PRISM. In addition, we conduct sensitivity analyses on several other key hyperparameters of PRISM, including the wavelet decomposition level m, the wavelet basis ψ, the Tsallis entropy parameter α, and the confidence threshold τ. Unless otherwise specified, all remaining training settings are kept identical to those used in the main experiments.

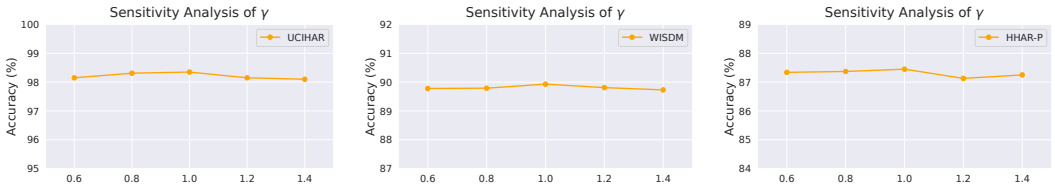
Sensitivity Analysis of the KL-Divergence Trade-off Coefficient β Following common practice in loss balancing, we initialize the trade-off coefficient β for $\mathcal{L}_M$ to 1 as a neutral setting. To validate the robustness of this choice, we further conduct a sensitivity analysis by varying β within the range of 0.6 to 1.4. As illustrated in Figure 4, PRISM exhibits consistently stable performance across different datasets with only minor variations. This indicates that the model is not overly sensitive to β within a reasonably wide interval. Based on both empirical stability and simplicity considerations, we fix $\beta = 1$ in all experiments.



(a) UCIHAR (b) WISDM (c) HHAR-P

Figure 4: Sensitivity analysis of the KL-divergence trade-off coefficient β.

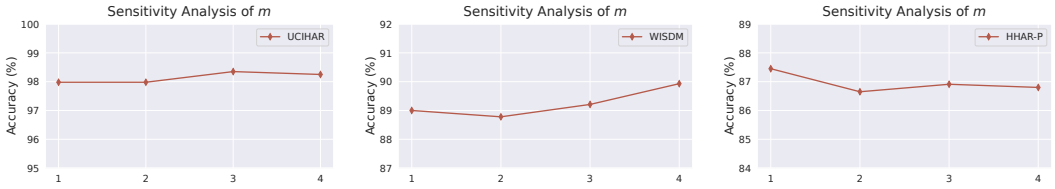
Sensitivity Analysis of the Committee-guided Supervision Loss Trade-off Coefficient γ Similarly, we initialize the trade-off coefficient γ for $\mathcal{L}_{\text{sup}}$ to 1, following the same principle of balanced weighting between loss components. We then perform a sensitivity study by varying γ from 0.6 to 1.4 to examine its effect on representation learning. As shown in Figure 5, the performance remains consistently stable across all evaluated datasets, with negligible fluctuations. This demonstrates that the proposed framework is robust to variations in γ within a broad range. Based on both empirical stability and simplicity considerations, we fix $\gamma = 1$ in all experiments.



(a) UCIHAR (b) WISDM (c) HHAR-P

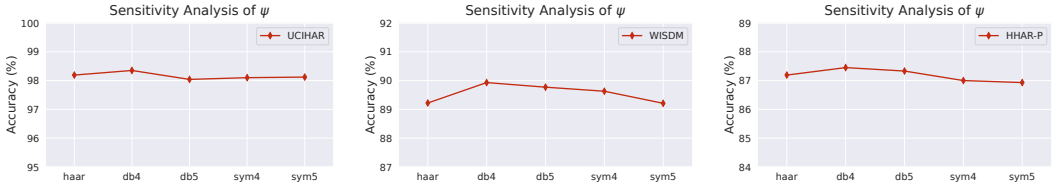
Figure 5: Sensitivity analysis of the committee-guided supervision loss trade-off coefficient γ.

Sensitivity Analysis of the Wavelet Decomposition Level m Considering that increasing the number of wavelet decomposition levels may introduce overly coarse representations and additional computational overhead, we conduct a sensitivity analysis on the decomposition level m. As shown in Figure 6, our method remains stable across different choices of m. Since the temporal characteristics vary across datasets, a fixed decomposition level may not be universally optimal. Therefore, to balance representation quality and computational complexity, we adopt dataset-specific decomposition levels, with detailed configurations reported in Table 5.



(a) UCIHAR (b) WISDM (c) HHAR-P

Figure 6: Sensitivity analysis of the wavelet decomposition Level m.



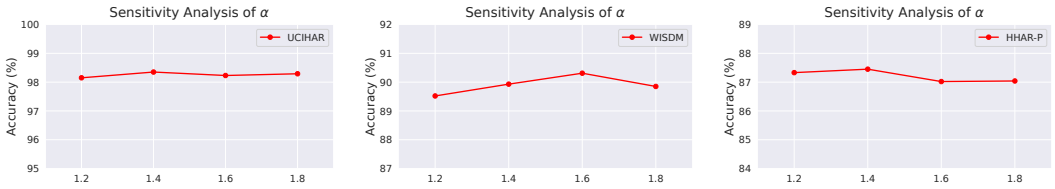
(a) UCIHAR (b) WISDM (c) HHAR-P

Figure 7: Sensitivity analysis of the wavelet basis ψ.

Sensitivity Analysis of the Wavelet Basis ψ The choice of wavelet basis ψ determines how the original time series is decomposed into multi-resolution components, affecting the trade-off between smoothness and fidelity in the extracted approximation coefficients. In our experiments, we consider representative wavelet families, including Haar, Daubechies, and Symlets. To balance smoothness and representation fidelity, we select `db4`, `db5`, `sym4`, and `sym5` as representative bases from the Daubechies and Symlets families, and conduct a sensitivity analysis over these choices. The results are presented in Figure 7. As observed, the performance of our model remains stable across different wavelet bases, indicating that the proposed framework is not sensitive to the specific choice of ψ. Considering both performance consistency and computational efficiency, we adopt `db4` as the default wavelet basis in all experiments.

Sensitivity Analysis of the Tsallis Entropy Parameter α We employ Tsallis entropy as a regularization term on the anchor prediction to encourage confident and consistent decision. It is defined as

$$S_\alpha(\mathbf{y}) = \frac{1}{\alpha - 1}\left(1 - \sum_c \mathbf{y}_{[c]}^\alpha\right), \tag{15}$$

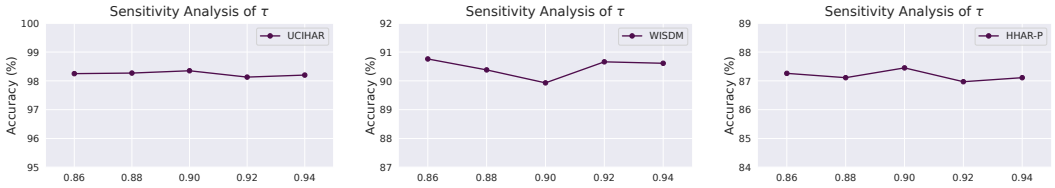


(a) UCIHAR (b) WISDM (c) HHAR-P

Figure 8: Sensitivity analysis of the Tsallis entropy parameter α.

where $\mathbf{y} \in \mathbb{R}^C$ denotes the predicted class probability distribution, C is the number of classes, and $\alpha > 1$ is a tunable parameter. The parameter α controls the sensitivity of the entropy regularization to the predicted probability distribution, thereby adjusting the degree of concentration encouraged in the anchor prediction. In our experiments, Tsallis entropy is used to stabilize anchor predictions by promoting confident and consistent semantic decisions. We discretize α within the range $(1.0, 2.0)$ with a step size of 0.2. As shown in Figure 8, the model remains stable across a wide range of α values. We therefore set $\alpha = 1.4$ in all experiments.

Sensitivity Analysis of the Confidence Threshold τ To determine an appropriate range for τ, we analyze the confidence of target-domain predictions at the original resolution during training. Specifically, for a given τ, we track the confidence-qualified sample ratio, defined as the proportion of samples within a mini-batch whose prediction confidence exceeds τ. A suitable threshold should yield a low ratio in the early stages—when predictions are unreliable—and gradually increase as the model becomes more confident over training. Based on this criterion, we observe that $\tau \approx 0.9$ provides a desirable balance, exhibiting the expected training dynamics. We therefore conduct a sensitivity analysis around this value. As shown in Figure 9, the model remains stable across a reasonable range of thresholds. Considering both robustness and effective noise filtering, we set $\tau = 0.9$ for all experiments.



(a) UCIHAR (b) WISDM (c) HHAR-P

Figure 9: Sensitivity analysis of the confidence threshold τ.

D Limitations

Although PRISM bridges resolutions for domain-invariant representation learning, its performance can still be unstable on certain datasets with poor stationarity and high volatility, similar to many existing DA methods. Enhancing robustness under such challenging data conditions remains an important direction for future work.

E Broader Impacts

We explore an unsupervised domain adaptation approach for time series classification via multi-resolution time series representation learning. This work intends to promote the development of time series analysis techniques and facilitate the practical deployment of time series applications, while ensuring no adverse social implications.

F Full Results

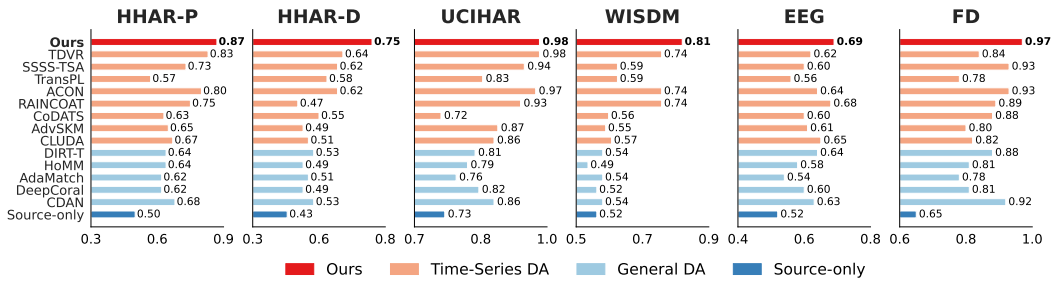


Figure 10: Average Macro-F1 score across multiple datasets.

Table 8: Accuracy (%) on WISDM for unsupervised domain adaptation.

Method	2→32	4→15	7→30	12→7	12→19	18→20	20→30	21→31	25→29	26→2	Average
Source-only	81.16	79.86	89.32	71.53	54.29	83.74	67.96	21.29	26.11	82.52	65.78
CDAN	89.37	65.97	84.79	70.48	51.01	88.62	77.02	46.58	44.33	83.33	70.05
DeepCoral	87.92	62.50	91.26	79.86	51.77	64.23	81.88	54.62	53.89	77.44	70.80
AdaMatch	74.39	78.47	89.64	73.26	55.30	75.20	74.76	31.32	57.78	87.20	69.79
HoMM	77.10	74.58	78.64	68.13	50.61	71.22	72.82	56.39	57.00	66.10	67.26
DIRT-T	77.78	70.83	90.61	70.20	51.51	85.36	71.84	54.41	60.04	66.46	69.62
CLUDA	73.91	67.36	86.40	65.97	49.24	83.74	72.49	49.97	35.00	86.47	67.04
AdvSKM	70.83	95.85	93.85	77.08	47.47	81.30	21.28	44.45	74.79	74.95	66.97
CoDATS	77.29	70.83	83.20	70.17	47.47	76.01	82.85	52.61	53.89	83.29	70.66
RAINCOAT	79.71	**97.91**	91.28	89.80	85.00	**92.23**	91.66	59.09	**82.97**	83.50	76.60
ACON	89.86	86.25	98.06	98.13	77.73	83.66	91.26	63.61	60.00	99.51	84.80
TransPL	72.19	66.25	91.88	77.08	81.41	70.63	77.71	55.00	14.38	66.56	67.30
SSSS-TSA	80.63	68.13	95.00	76.88	60.47	79.69	93.94	71.25	60.63	83.13	76.97
TDVR	**90.31**	88.75	98.96	99.17	70.47	90.00	93.96	**74.06**	68.63	**100.0**	87.43
Ours	87.19	92.50	**99.17**	**100.0**	**86.88**	89.38	**97.29**	73.75	74.28	98.75	**89.93**

Table 9: Accuracy (%) on UCIHAR for unsupervised domain adaptation.

Method	2→11	6→23	7→13	9→18	12→16	13→19	18→21	20→6	23→13	24→12	Average
Source-only	76.56	67.36	83.68	24.65	61.11	88.89	100.0	94.10	71.18	83.68	75.12
CDAN	85.42	87.50	92.01	58.86	66.67	96.52	100.0	95.13	82.64	93.40	85.78
DeepCoral	90.63	84.38	87.50	46.88	65.28	95.49	100.0	95.49	69.79	87.50	82.01
AdaMatch	75.00	80.20	85.76	56.59	49.65	94.79	100.0	84.37	68.75	70.83	76.07
HoMM	74.06	82.71	81.88	73.96	70.21	96.67	98.75	73.33	77.71	80.63	80.99
DIRT-T	80.21	74.31	82.99	59.03	67.01	99.30	98.61	92.36	74.72	94.27	83.26
CLUDA	81.77	92.01	99.31	67.71	65.28	94.44	98.96	97.22	72.92	99.31	85.53
AdvSKM	98.96	88.54	92.71	74.65	69.44	93.05	100.0	85.41	79.51	96.87	83.26
CoDATS	68.23	74.31	77.43	63.89	66.32	94.09	99.65	70.49	56.25	82.81	75.54
RAINCOAT	100.0	95.83	**100.0**	75.69	86.52	100.0	100.0	93.41	86.52	93.75	94.43
ACON	100.0	96.25	99.16	91.66	85.63	100.0	100.0	97.50	100.0	100.0	97.02
TransPL	85.94	87.29	89.58	70.63	85.41	86.67	94.38	71.76	82.50	92.08	84.61
SSSS-TSA	100.0	95.63	94.79	85.42	88.13	95.42	100.0	93.13	93.96	96.88	94.33
TDVR	100.0	100.0	98.96	86.46	**91.04**	100.0	100.0	98.75	100.0	**100.0**	97.52
Ours	**100.0**	**100.0**	99.38	**95.83**	88.54	**100.0**	**100.0**	**100.0**	**100.0**	99.79	**98.35**

Table 10: Accuracy (%) on HHAR-P for unsupervised domain adaptation.

Method	0→2	1→6	2→4	4→0	4→5	5→1	5→2	7→2	7→5	8→4	Average
Source-only	64.51	70.63	45.42	32.81	78.32	90.63	25.67	32.37	39.26	62.92	54.25
CDAN	76.19	92.57	52.57	29.09	97.27	96.16	35.04	37.05	75.26	96.11	68.73
DeepCoral	84.23	90.14	47.08	28.13	90.49	89.91	38.39	34.45	55.73	76.88	68.03
AdaMatch	84.78	92.31	54.50	36.45	78.45	94.20	41.96	37.65	63.80	64.69	65.91
HoMM	75.67	90.79	52.83	36.61	87.66	90.78	37.23	37.32	61.29	79.88	65.01
DIRT-T	77.83	88.54	50.69	32.22	93.16	91.86	38.62	38.10	72.46	65.83	64.99
CLUDA	79.84	93.40	45.90	38.84	94.08	95.57	33.93	37.80	77.57	96.52	69.35
AdvSKM	78.94	87.91	52.57	33.49	92.64	92.71	36.53	39.95	65.49	83.75	66.41
CoDATS	79.61	90.90	60.07	21.80	97.66	97.66	41.44	38.54	58.15	**97.01**	68.71
RAINCOAT	87.72	93.33	63.75	46.46	**98.05**	98.25	42.63	43.32	84.17	93.75	74.21
ACON	86.65	93.45	79.01	53.53	97.15	98.32	65.80	65.71	88.59	89.17	81.74
TransPL	60.49	60.17	60.46	43.04	74.96	90.82	50.40	47.50	58.44	60.75	60.70
SSSS-TSA	82.50	89.50	67.63	47.23	89.06	93.55	48.66	59.96	80.39	78.38	73.69
TDVR	**96.12**	93.79	81.83	**70.09**	98.01	**98.78**	**77.37**	66.21	90.80	73.05	85.05
Ours	90.71	**94.08**	**91.04**	69.01	97.85	97.70	74.96	**69.60**	**93.20**	96.38	**87.45**

Table 11: Accuracy (%) on FD for unsupervised domain adaptation.

Method	0→1	0→2	0→3	1→0	1→2	2→0	2→1	2→3	3→0	3→2	Average
Source-only	62.21	53.71	62.41	63.91	73.95	64.08	93.17	95.54	57.08	74.31	70.04
CDAN	91.29	71.83	90.13	96.50	90.09	83.10	99.38	99.98	95.40	87.95	90.56
DeepCoral	75.54	71.79	76.03	89.13	83.55	76.34	98.84	98.55	87.50	83.71	84.10
AdaMatch	67.81	55.38	62.88	92.21	98.57	79.08	89.96	90.40	87.23	97.57	82.11
HoMM	81.54	71.63	78.17	89.89	84.78	76.03	98.71	99.55	90.94	85.96	85.72
DIRT-T	75.94	70.85	76.36	98.10	90.27	81.92	100.0	99.98	97.06	90.29	88.08
CLUDA	90.47	82.63	88.68	89.06	92.23	61.92	93.91	90.80	82.01	78.17	84.99
AdvSKM	74.71	66.05	73.30	87.86	86.29	76.85	98.66	99.38	84.89	85.74	83.37
CoDATS	81.79	73.26	83.15	89.22	88.68	81.43	99.89	100.0	85.47	89.00	87.20
RAINCOAT	85.18	79.40	89.04	78.84	90.11	81.43	95.18	96.81	77.39	94.08	86.75
ACON	86.52	69.00	86.96	97.92	**99.80**	84.29	98.62	98.93	97.72	**97.66**	91.74
TransPL	65.94	68.50	67.77	62.88	89.26	66.14	98.55	97.90	69.35	87.05	77.33
SSSS-TSA	96.45	83.75	92.46	92.86	89.44	84.75	100.0	100.0	93.19	87.05	92.00
TDVR	79.58	71.58	76.92	95.89	85.47	83.97	100.0	99.42	93.64	86.96	87.34
Ours	**98.08**	**89.08**	**94.06**	**98.15**	97.68	**97.77**	**100.0**	**100.0**	**98.10**	94.75	**96.77**

Table 12: Accuracy (%) on EEG for unsupervised domain adaptation.

Method	0→11	2→5	12→5	7→18	16→1	9→14	4→12	10→7	6→3	8→10	Average
Source-only	55.60	60.03	72.01	53.91	40.21	75.00	48.76	67.86	75.20	35.21	58.38
CDAN	68.94	69.53	78.45	73.18	74.53	80.14	67.08	74.35	80.99	55.16	72.24
DeepCoral	57.22	60.41	75.00	65.82	69.53	82.22	64.97	76.05	78.38	36.79	66.66
AdaMatch	63.86	72.39	72.09	71.61	57.86	82.55	48.17	60.41	78.12	51.25	65.83
HoMM	55.08	66.60	72.66	59.38	72.66	76.95	58.79	73.28	76.37	47.34	65.91
DIRT-T	65.88	72.85	78.97	74.34	**81.82**	86.14	68.48	75.05	83.66	46.01	66.04
CLUDA	57.87	71.86	79.39	74.49	75.83	86.32	66.53	75.23	81.96	65.70	73.50
AdvSKM	56.51	65.62	76.49	60.93	72.96	76.75	66.14	74.31	78.90	44.76	67.33
CoDATS	69.53	71.83	79.28	73.19	75.32	81.64	**71.68**	73.31	83.59	44.22	72.36
RAINCOAT	**74.41**	**73.76**	79.81	75.32	78.64	87.17	69.86	77.23	84.58	62.35	76.31
ACON	61.91	71.24	79.46	77.63	75.99	87.09	66.17	77.38	84.19	65.57	74.66
TransPL	53.32	65.42	71.67	72.24	71.25	78.30	57.34	74.08	80.26	55.81	67.97
SSSS-TSA	47.23	68.58	79.79	75.99	79.32	**89.31**	68.49	70.30	83.20	60.94	72.31
TDVR	67.27	70.94	**80.24**	78.93	79.88	85.76	65.39	71.31	**84.63**	61.28	74.56
Ours	70.63	72.26	79.69	**80.96**	78.04	82.99	71.22	**78.63**	82.65	**68.03**	**76.51**