SOURCE-FREE OBJECT DETECTION VIA MULTI-SOURCE KNOWLEDGE FUSION

Anonymous authors
Paper under double-blind review

ABSTRACT

Source-free domain adaptation (SFDA) adapts a pretrained source model to an unlabeled target domain without accessing source data, which is useful when source data cannot be retained or transferred. Multi-source source-free domain adaptation (MSFDA) further exploits complementary knowledge from multiple pretrained source models, but heterogeneous source-target shifts and noisy pseudo labels make multi-source object detection particularly challenging. We propose a multi-source knowledge-fusion framework that addresses these issues through three complementary mechanisms. First, Text-Driven Feature Augmentation (TFA) transforms unlabeled target images toward each source-domain appearance using only coarse semantic descriptions of the source domains, reducing the mismatch seen by the corresponding source experts. Second, the source experts are locally adapted on their source-stylized target views, while an aggregator integrates their knowledge using adaptively learned contribution weights. Third, mutual confidence selection between the aggregator and domain experts reduces the influence of noisy pseudo labels during self-training. Extensive experiments across multiple real-world domain shifts show consistent improvements over existing SFDA and MSFDA baselines under the standard MSFDA protocol.

1 INTRODUCTION

Source-Free Domain Adaptation (SFDA) has recently gained prominence as a solution to practical challenges such as data privacy, distributed data storage, and inconvenient data transmission Huang et al. (2021); Li et al. (2021b; 2022). SFDA aims to adapt pre-trained model on the source domain to unlabeled target domain without accessing actual source data. Multi-source free domain adaptation extends SFDA to incorporate pre-trained models from multiple source domains, allows knowledge aggregation from a broader variety of data, which is particularly advantageous when the target domain is diverse and spans a wide range of possible scenarios Peng et al. (2019); Dong et al. (2021); Ahmed et al. (2021).

However, properly aggregating knowledge from multiple source domains without access to actual data poses a set of unique challenges. First, the source domains may not comprehensively cover all aspects of the target domain, leading to missing or poorly represented features. Even with multiple source models, certain characteristics of the target domain might remain unaccounted for, hindering generalization and degrading adaptation performance. Second, without access to source data, determining the relative importance of each source model is non-trivial. Some source models may be more relevant than others for a given target domain, but misaligned weighting or reliance on less relevant sources can negatively impact adaptation Li et al. (2024). Last, source models trained on different domains may encode conflicting feature representations, leading to inconsistencies in target predictions Wang et al. (2019); Ding et al. (2016). Disagreements among source models can introduce noise into pseudo-labeling, propagating errors across training stages and ultimately compromising adaptation performance.

While few existing methods have made progress in addressing MSFDA by quantifying the contributions of multiple source models and finding optimal combinations through techniques like joint feature alignment Dong et al. (2021); Ahmed et al. (2021); Peng et al. (2019), or attention mechanisms Li et al. (2024), these approaches have some inherent limitations. For example, joint feature alignment aims to learn a shared common space from multiple source models for the generalization to target domain by quantifying contribution of each domain or find the optimal combination of source

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.

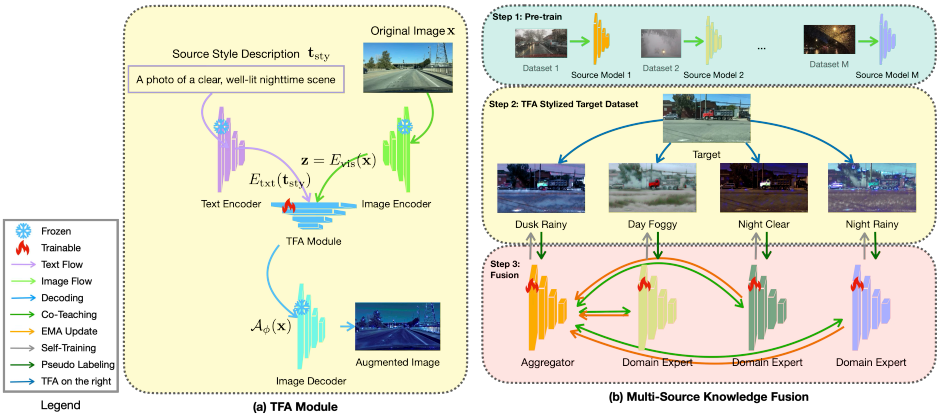

Figure 1: (a) **TFA module:** Given a target image and a style description, TFA uses CLIP features to transform the image representation to match the source domain style. (b) **Overall pipeline:** (1) Pre-train source models on their respective source datasets; (2) Generate stylized target datasets with TFA; (3) Train the multi-source knowledge fusion framework, where domain experts are adapted locally and the aggregator is updated via EMA-based knowledge integration.

domains Dong et al. (2021); Ahmed et al. (2021); Shen et al. (2023). However, source domains can exhibit significant divergence in feature distributions due to differences in data characteristics, such as lighting conditions, object categories, or scene variations. These differences are often not easily captured through a single alignment process, and aligning features can lead to loss of important domain-specific information. Attention mechanisms have been used to focus on the most relevant parts of the source features during adaptation Li et al. (2024). However, attention alone does not fully address the semantic gaps that arise due to differences in how source models perceive the target data. Attention mechanisms can focus on certain features but may not be able to comprehensively reduce the domain shift between the source and target domains.

Our central idea is to make multi-source fusion reliable by first reducing the mismatch seen by each source expert and then aggregating only reliable cross-domain knowledge. As shown in Fig. 1, TFA creates source-aligned views of the unlabeled target data from coarse source-domain semantics; local expert adaptation uses these aligned views to retain domain-specific expertise; adaptive contribution weighting integrates complementary experts into a single aggregator; and mutual confidence selection suppresses noisy pseudo labels during self-training. Rather than treating these modules as independent improvements, they form a single pipeline that addresses the two coupled challenges of MSFDA: source-target distribution mismatch and unreliable cross-source pseudo supervision.

Contributions. Our key contributions are three-folded: (1) We introduce TFA, a source-free augmentation mechanism that uses coarse source-domain text descriptions to construct source-stylized target views while regularizing target content. (2) We propose a multi-source knowledge-fusion strategy that combines local expert adaptation with adaptive contribution weighting, allowing complementary source knowledge to be integrated without accessing source data. (3) We introduce mutual confidence selection between the aggregator and domain experts to reduce pseudo-label noise during self-training, and provide extensive evaluation across diverse detection benchmarks, source-model architectures, and domain shifts.

2 Related Works

Multi-source-free domain adaptation. Multi-Source-Free Domain Adaptation (MSFDA) aims to distill knowledge from multiple pre-trained models and adapts to an unlabeled target domain without access to the actual source data Dong et al. (2021); Shen et al. (2023); Yeh et al. (2023); Li et al. (2023; 2024); Peng et al. (2019). Dong et al., propose to quantify the contributions of multiple source models with a source-specific transferable perception module. It then improve the quality of the pseudo label with a confident-anchor-induced pseudo label generator Dong et al. (2021). Aiming to find the optimal combination of source models, Ahmed et al., learn a set weights by minimizing the conditional entropy of transferring each source model to the unlabeled target. They also provide

intuitive theoretical insights to justify their methodology Ahmed et al. (2021). Shen et al., balance domain aggregation, pseudo-labeling, and joint feature alignment with information-theoretic bound on the generalization error Shen et al. (2023). Focusing on balancing between instance specificity and domain consistency, Li et al., propose a parameter-tuning free method for MSFDA with an attention module that learns both intra-domain weights and inter-domain ensemble weights Li et al. (2024). In contrast to these existing methods, our method focuses on explicitly reducing the semantic gaps among different source models and the unlabeled target data distribution. We propose a text-driven feature augmentation technique to achieve style transfer given only images from the target domain, and a simple description of the source domain style.

Source-free domain adaptation for object detection. Source-free domain adaptation for object detection (SFOD) operates under the assumption that only the pre-trained model on the source domain is accessible, while the actual source data is not available, presenting itself as a promising area of research. Conventional SFOD methods commonly employ the pseudo labeling paradigm, involving a cyclic process of model adaptation that oscillates between predicting pseudo labels and fine-tuning the model Huang et al. (2021); Li et al. (2021b); Xiong et al. (2021). Some recent efforts attempt to address these problems by using self-entropy descent as a confidence threshold to select high-quality pseudo labels Li et al. (2021b). Other efforts directly learn domain-invariant features through devising domain perturbation Xiong et al. (2021), graph alignment constraint Li et al. (2022), adversarial alignment of the target images Chu et al. (2023), instance relation graph network Vibashan et al. (2023), or teacher-student models Lin et al. (2023); Liu et al. (2023). Although existing SFOD methods Li et al. (2021b); Lin et al. (2023); Liu et al. (2023); Li et al. (2022) have shown promise, MSFDA is less explored for object detection tasks. In this work, we propose to utilize multiple pre-trained source models to address source-free domain adaptation for object detection.

Text-based style transfer. Style transfer aims to transform a content image by transferring the semantic texture of a style image. Traditional style transfer approaches require a reference style image for learning the style texture the texture to change the style of the content image, which might not be always available. Under this condition, using text information to convey the desired style has emerged as a solution. Current text-based style transfer methods can be categorized into two parts: (1) The generative-based methods; StackGAN Zhang et al. (2017a) integrated text conditions to multi-scale generative model for high-quality image synthesis. AttnGAN Xu et al. (2018) further improved the performance with attention mechanism on text and image features. ManiGAN Li et al. (2020) proposed a modules for simultaneously embedding the text and image features. StyleCLIP Patashnik et al. (2021) performed attribute manipulation with exploring learned latent space of StyleGAN Karras et al. (2019). StyleGAN-NADA Gal et al. (2022) proposed a model modification method with using text conditions only, and modulates the trained model into a novel domain without additional training images. (2) The non-generative-based methods; CLIPstyler Kwon & Ye (2022) design a modulation of the style of content images only with a single text condition using the pre-trained text-image embedding model of CLIP, and propose a patch-wise text-image matching loss with multiview augmentations for realistic texture transfer. PODA Fahes et al. (2023) propose a prompt-driven instance normalization (PIN) layer to do style transfer, where affine transformations of low-level features are optimized such that the representation in CLIP latent space matches the one of text-based prompt. PromptStyler Cho et al. (2023) simulates various distribution shifts in the joint space by synthesizing diverse styles via prompts without using any images to deal with source-free domain generalization. In this work, we propose a novel non-generative text-based style transfer method, TFA, that focuses on aligning the target image with the source style text in both high-level semantics and low-level visual textures.

3 METHODOLOGY

The proposed approach aims to reduce the domain shift between multiple source domains and a target domain. To achieve this, we introduce text-driven feature augmentation (TFA), which uses simple textual descriptions of the source domains to augment target images with corresponding styles. This allows for the direct application of pre-trained models to the augmented target images. Additionally, we propose a multi-source knowledge fusion framework to integrate knowledge from multiple source models. To enhance the quality of pseudo-labels, we use confidence selection in a mutual refinement manner, offering a more robust approach to handling noisy labels.

Problem formulation. We consider M labeled source domains denoted as $\{D^{s_i}\}_{i=1}^{M}$, where each domain $D^{s_i} = \{(\mathbf{x}_k^{s_i}, y_k^{s_i})\}_{k=1}^{N_i}$ consists of N_i images. Each image $\mathbf{x}_k^{s_i}$ is paired with an annotation $y_k^{s_i} = (\mathbf{b}_k^{s_i}, c_k^{s_i})$, where $\mathbf{b}_k^{s_i}$ denotes the bounding box coordinates and $c_k^{s_i}$ represents the class label of the object in the k-th image of the i-th source domain. In addition, we are given an unlabeled target domain $D^t = \{\mathbf{x}_k^t\}_{k=1}^{N_t}$, which contains N_t images without annotations. During adaptation, we do not have access to the raw source data. Instead, we are provided with: (1) the M pre-trained source models $\{\theta^{s_i}\}_{i=1}^{M}$, each trained on a different source domain, and (2) the text descriptions $\{\mathbf{t}_{\text{sty}}^{s_i}\}_{i=1}^{M}$ that characterize the style of each source domain. The goal is to leverage the pre-trained source models and their textual style descriptions to adapt effectively to the unlabeled target domain.

CLIP. We leverage the vision–language (V-L) space of pre-trained CLIP models Radford et al. (2021) for TFA. CLIP consists of a visual encoder E_{vis} and a text encoder E_{txt}, which map images and text into a shared embedding space. Given an image $\mathbf{x}$, the visual encoder produces $\mathbf{v} = E_{\text{vis}}(\mathbf{x})$. For the i-th category, we construct a prompt "A photo of a [class-i]", tokenize it as $\mathbf{p}_i$, and encode it into $\mathbf{t}_i = E_{\text{txt}}(\mathbf{p}_i)$. The prediction for class y is obtained via a softmax over cosine similarities:

$$p(\hat{y} = y|\mathbf{x}) = \frac{\exp(\cos(\mathbf{v}, \mathbf{t}_y)/\tau)}{\sum_{k=1}^{K} \exp(\cos(\mathbf{v}, \mathbf{t}_k)/\tau)}. \tag{1}$$

This aligns visual and textual features, enabling cross-modal representation learning.

3.1 Text-driven Feature Augmentation

To bridge the semantic gap between pre-trained source models and target images, we propose to use Text-Driven Feature Augmentation (TFA), which transforms target image features to match the source domain style guided by textual descriptions. The style transformation is performed by aligning the target image's features with textual representations of the source domain style in a shared vision–language space, where image and text features are well-aligned. By minimizing the gap between the target image feature and the textual source style feature, the target image is effectively transformed to reflect the source domain style. Specifically, we introduce a learnable augmentation module $\mathcal{A}_\phi(\cdot)$ that adapts target features in the vision–language (V-L) space. Given a target image $\mathbf{x}$, its CLIP visual feature is extracted as $\mathbf{z} = E_{\text{vis}}(\mathbf{x})$, and a style description of a source domain is encoded as $E_{\text{txt}}(\mathbf{t}_{\text{sty}})$. The augmentation module generates $\mathcal{A}_\phi(\mathbf{z})$, which is optimized to align with the textual style while preserving image content. The learning objective is:

$$\phi^* = \min_{\phi} \lambda_1 \mathcal{L}_{\text{style}} + \lambda_2 \mathcal{L}_{\text{content}}, \tag{2}$$

$$\mathcal{L}_{\text{style}} = 1 - \cos(\mathcal{A}_\phi(\mathbf{z}), E_{\text{txt}}(\mathbf{t}_{\text{sty}})), \tag{3}$$

$$\mathcal{L}_{\text{content}} = \|\mathcal{A}_\phi(\mathbf{z}) - \mathbf{z}\|_1, \tag{4}$$

We implement $\mathcal{A}_\phi$ as a learnable feature augmentation network parameterized by ϕ. As shown in Figure 1 (a), $\mathcal{A}_\phi$ is a lightweight multilayer perceptron (MLP) that takes the CLIP visual feature $\mathbf{z} = E_{\text{vis}}(\mathbf{x})$ and textual feature $E_{\text{txt}}(\mathbf{t}_{\text{sty}})$ as input and outputs a transformed feature $\mathcal{A}_\phi(\mathbf{z})$ aligned with the source domain style. The parameters ϕ are optimized with the objective in Eq. 2, while the CLIP encoders E_{vis}, E_{txt}, and the image decoder remain frozen. Finally, a pretrained ClipStyler image decoder Kwon & Ye (2022) is used to reconstruct stylized images from $\mathcal{A}_\phi(\mathbf{z})$. This design ensures that TFA operates efficiently in the joint V-L feature space without requiring image-space generation.

In Eq. 3, $\cos(\cdot)$ computes the cosine similarity between augmented feature embedding and style text embedding: $\cos(\mathcal{A}_\phi(\mathbf{z}), E_{\text{txt}}(\mathbf{t}_{\text{sty}})) = \frac{\mathcal{A}_\phi(\mathbf{z}) \cdot E_{\text{txt}}(\mathbf{t}_{\text{sty}})}{\|\mathcal{A}_\phi(\mathbf{z})\| \|E_{\text{txt}}(\mathbf{t}_{\text{sty}})\|}$ measures the semantic alignment between the augmented target image features and the text description of the source domain style. By maximizing the cosine similarity, we ensure that the style of the augmented image is consistent with the desired source style, which is represented in the CLIP text embedding space. Eq. 4 acts as an $\mathcal{L}_1$ regularizer to match high-level image content between the augmented image and the original target image.

3.2 Multi-Source Knowledge Fusion Framework

Using TFA, we construct multiple stylized target datasets $\{D^{t\to s_1}, D^{t\to s_2}, ..., D^{t\to s_M}\}$, where $D^{t\to s_i} = \{\mathcal{A}_{\phi_i}^{s_i}(\mathbf{x}) | \mathbf{x} \in D^t\}$. In this section, we introduce the multi-source knowledge fusion framework, which effectively integrates knowledge from multiple source models while minimizing

the negative impact of noisy pseudo labels during adaptation. Specifically, given pre-trained source models $\{\theta^{s_1}, \theta^{s_2}, \ldots, \theta^{s_M}\}$ as domain experts, we designate one model as the aggregator θ^{agg}, and denote the remaining models as domain experts θ_i^{DE}. The aggregator simultaneously acts as a domain expert for its own source domain while also serving as the central model for integrating knowledge across all domains. During adaptation, both the aggregator and domain experts are locally updated on their corresponding augmented target datasets, while the aggregator further aggregates cross-domain knowledge through an additional EMA update from the domain experts.

Domain expert update via self-training. Each domain expert θ_i^{DE} (including the aggregator θ^{agg} of its own domain) is adapted to its corresponding augmented dataset $D^{t\to s_i} = \{\mathcal{A}_{\phi_i}^{s_i}(\mathbf{x}) \mid \mathbf{x} \in D^t\}$ via self-training. Specifically, the expert generates pseudo labels $\tilde{y}_i$ for the augmented images $\mathbf{z}_i = \mathcal{A}_{\phi_i}^{s_i}(\mathbf{x})$, which are then used to update itself with the local detection loss:

$$\theta_i^{\text{DE}} \leftarrow \theta_i^{\text{DE}} - \gamma_i \nabla_{\theta_i^{\text{DE}}} \mathcal{L}_i^{\text{local}}, \tag{5}$$

where γ_i is the learning rate. Local loss $\mathcal{L}_i^{\text{local}}$ follows the standard Faster R-CNN detection loss Ren et al. (2015):

$$\mathcal{L}_i^{\text{local}} = \mathcal{L}_{cls}^{rpn}(\mathbf{z}_i, \tilde{y}_i) + \mathcal{L}_{reg}^{rpn}(\mathbf{z}_i, \tilde{y}_i) + \mathcal{L}_{cls}^{roi}(\mathbf{z}_i, \tilde{y}_i) + \mathcal{L}_{reg}^{roi}(\mathbf{z}_i, \tilde{y}_i). \tag{6}$$

Here, $\mathcal{L}_{cls}^{rpn}$ and $\mathcal{L}_{reg}^{rpn}$ are the classification and regression losses of the region proposal network (RPN), while $\mathcal{L}_{cls}^{roi}$ and $\mathcal{L}_{reg}^{roi}$ are those of the region-of-interest (ROI) head. The classification terms are standard cross-entropy losses for category prediction, and the regression terms measure localization error. In this way, each domain expert iteratively improves itself by using its own high-confidence predictions as supervision, following the principle of self-training.

Aggregation via adaptive contribution weighting. To integrate knowledge across source domains, we update the aggregator using an EMA-style weighted combination of the locally adapted models:

$$\theta^{\text{agg}} \leftarrow \alpha^{\text{agg}} \theta^{\text{agg}} + \sum_{i=1}^{M-1} \alpha_i^{\text{DE}} \theta_i^{\text{DE}}, \tag{7}$$

where the non-negative contribution weights satisfy $\alpha^{\text{agg}} + \sum_i \alpha_i^{\text{DE}} = 1$.

Rather than using fixed weights, we learn the contribution weights with a lightweight network $F_\omega(\cdot)$. Given the prediction statistics $s_i(x)$ of each model on a target batch, the contribution vector is

$$\boldsymbol{\alpha} = \text{softmax}(F_\omega([s_1(x), \ldots, s_M(x)])) \,. \tag{8}$$

We form a weighted ensemble prediction

$$\tilde{p}(y_k|x;\omega) = \sum_{i=1}^{M} \alpha_i \, p(y_k|x, \theta^{s_i}), \tag{9}$$

and optimize F_ω by minimizing its prediction entropy:

$$\mathcal{L}_{\text{ent}} = -\sum_{k=1}^{K} \tilde{p}(y_k|x;\omega) \log \tilde{p}(y_k|x;\omega). \tag{10}$$

This objective assigns larger contribution weights to experts whose predictions are confident and consistent with the multi-source ensemble, while reducing the influence of less reliable experts. The learned weights are then used in Eq. 7 for cross-domain knowledge integration.

Self-training with mutual confidence selection. Confidence selection is usually coupled with self-training to filter out noisy pseudo labels in SFDA. Usually, it filters out less confident predictions with high entropy via a pre-defined threshold Li et al. (2021b); Kim et al. (2021). This approach suffers from a critical limitation: as training progresses, the model tends to overfit noisy pseudo-labels, becoming increasingly confident in incorrect predictions. This overfitting amplifies label noise, making it progressively harder to distinguish clean from noisy samples. To address this issue, we introduce an external validation mechanism by enabling the aggregator and domain experts to teach each other Han et al. (2018).

Table 1: Per-class results on Day Foggy with different components enabled. SS stands for single-source model, MS stands for multi-source knowledge fusion, and CS stands for confidence selection. As usual, for single-source, Day Clear is used as the source model, and for multi-source, Day Clear and the rest three targets except for Day Foggy are used as source models. The last row is our proposed method.

					AP							mAP
Source	SS	MS	TFA	CS	Bus	Bike	Car	Motor	Person	Rider	Truck	All
✗	✓	✗	✗	✗	30.8	29.3	28.5	32.7	30.8	32.4	35.8	31.4
✗	✗	✓	✗	✗	38.5	30.6	29.9	35.5	34.1	33.9	38.0	33.5
✗	✗	✓	✓	✗	**38.8**	31.2	30.5	**36.4**	34.8	34.5	**38.4**	34.2
✗	✗	✓	✓	✓	37.4	**36.5**	**45.8**	35.9	**40.8**	**36.0**	36.0	**38.4**

For example, given a domain expert θ_i^{DE} and its corresponding augmented target dataset $D^{t \to s_i}$, the aggregator θ^{agg} computes the loss of each example and ranks them in ascending order, where examples with loss higher than a predefined threshold will be considered as noisy labels and be discarded. The rationale behind this mechanism is twofold: (1) The aggregator and domain experts are trained with different augmentations of the same target data distribution. This correlation allows each model to identify and filter out noisy labels that might have been incorrectly assigned in the early stages of training. (2) If the domain expert were used for confidence selection, errors from noisy pseudo-labels generated during early training would propagate and accumulate, leading to the model becoming increasingly overfitted to these erroneous labels. As the model continues to train, it would gradually lose the ability to distinguish between correct and noisy labels. However, by using the aggregator, which is less affected by overfitting to the domain expert's noisy data, we can mitigate this issue and ensure that the pseudo-labels used for domain expert training remain more accurate and reliable. Similarly, we randomly choose a domain expert to filter out the noisy pseudo labels of the aggregator to avoid error accumulation of a single domain expert.

4 Experiments

In this section, we first introduce the datasets, evaluation metrics, comparative baselines and the implementation details. Then the quantitative and qualitative results are presented to evaluate the effectiveness of our method. Some additional results are presented in the Appendix.

4.1 Experiment Details

Datasets and evaluation metric. For evaluation, we adopt three cross-domain datasets: (1) Diverse Weather Dataset (DWD) Wu & Deng (2022), which includes driving scenes under various weather and time conditions; (2) Cityscapes Cordts et al. (2016), FoggyCityscapes Sakaridis et al. (2018), and KITTI Geiger et al. (2013a), a combination that covers both real-world urban driving scenes and synthetic car images; and (3) Art Inoue et al. (2018), which consists of images rendered in diverse artistic styles. More detailed dataset descriptions are provided in the Appendix. In all our experiments, we use the Mean Average Precision (mAP) as our metric. Specifically, we report the mAP@0.5, which considers a prediction as a true positive if it matches the ground-truth label and has an intersection over union (IOU) score of more than 0.5 with the ground-truth bounding box. All reported experimental results are based on a single run.

Comparison baselines We compare our method with (1) *Source-free domain adaptation (SFDA)* methods including SED Li et al. (2021b), HCL Huang et al. (2021), and IRG Vibashan et al. (2023), LODS Li et al. (2022), LPLD Yoon et al. (2024), SF-UT Hao et al. (2024), DRU Khanh et al. (2024), and (2) *Multi-source-free domain adaptation (MSFDA)* methods including Mean-Teacher Tarvainen & Valpola (2017), MixUp Zhang et al. (2017b), MSFDAOD Zhao et al. (2024), CAiDA Dong et al. (2021), Selective Self-Training (SST) Shen et al. (2023), Bi-ATEN Li et al. (2024). For those SFDA and MSFDA methods that are not designed for object detection task, the reported results are based on our re-implementation. We additionally report the performance of Faster R-CNN (FR) Ren et al. (2015) initialized with ImageNet pre-trained weights.

Implementation details. For pre-trained source models, we used the implementation of Faster R-CNN Ren et al. (2015) from the MMDetection library Chen et al. (2019). We use ResNet50 as the backbone, with the learning rate equal to 0.01 and the max epoch set to 8. And we use the last

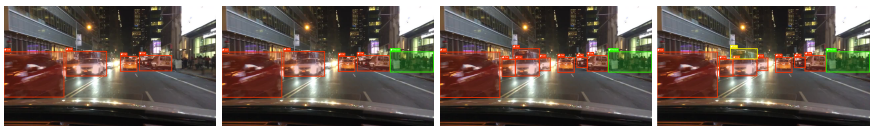

(a) SS (b) MS (c) MS+TFA (d) MS+TFA+CS

Figure 2: Qualitative visualization: Bounding box predictions from different settings, where SS stands for single source model, MS stands for multi-source knowledge fusion, and CS stands for confidence selection.

checkpoint as our source model. The training is conducted using 4 P4. The contribution network $\mathcal{F}(\cdot)$ is implemented as a lightweight two-layer MLP. Its input is the concatenated per-model prediction statistics on the current target batch, including class probability/confidence information from all source models. The network outputs normalized EMA weights for adaptive aggregation. Specifically, $\mathcal{F}(\cdot)$ consists of two fully connected layers with 64 hidden units, each followed by ReLU, and a final linear layer with softmax normalization to produce the aggregation weights. For TFA, we use the Layer 1 target feature maps of the pre-trained CLIP-ResNet-50 model Radford et al. (2021). The text description of each source domain are shown in Table 6. Optimization was done for 1000 iterations using SGD with a learning rate of 0.01, momentum of 0.9, and weight decay of 1e-4. The training is conducted using 1 A100. During adaptation, we first initialize the target model with the chosen closest pre-trained source model. Then with the backbone frozen, we fine-tune the classifier and bounding box head with the augmented feature and pseudo labels. Non-maximum Suppression (NMS) Hosang et al. (2017) is utilized to eliminate duplicate detections and select the most relevant bounding boxes.

4.2 Main Results

Table 2 presents a comprehensive comparison between our method and state-of-the-art SFDA and MSFDA approaches for object detection on the DWD dataset, which includes five domains: Day Clear (DC), Night Clear (NC), Day Foggy (DF), Dusk Rainy (DR) and Night Rainy (NR). We first observe that MSFDA methods consistently outperform SFDA approaches, reaffirming the advantage of leveraging diverse domain knowledge. The presence of multiple source domains provides a richer representation space, enabling better generalization to the target domain. Secondly, our proposed method outperforms all existing MSFDA methods across all four domains. This demonstrates the effectiveness of our approach in mitigating domain gaps by incorporating both local and global updates in a structured multi-source knowledge fusion paradigm. Additionally, by learning domain-specific contributions dynamically, the multi-source knowledge fusion framework avoids the pitfalls of naïve model aggregation, leading to more stable adaptation. Night Rainy remains the most challenging DWD target because severe illumination degradation, rain, and reduced visibility substantially weaken visual cues for all methods. Our method obtains 20.3 mAP, improving over the strongest compared MSFDA baseline (19.8 mAP), but the absolute performance remains considerably below that on the other target domains. We therefore view TFA and multi-source fusion as mitigating rather than completely eliminating extreme domain shifts. This limitation is discussed further in Appendix D. A detailed per-class analysis for each domain is provided in the Appendix.

Additional results on the Art dataset The results in Table 3 show consistent improvements of our method across all three domains when tested under the "leave-one-domain-out" setting. Compared with prior approaches such as Mean-Teacher, MixUp, and MSFDAOD, our method achieves the highest performance on Clipart1k (48.45%), Watercolor2k (50.33%), and Comic2k (46.39%). The gains over the strongest baseline Bi-ATEN are 1.62%, 1.57%, and 1.15% respectively. These results highlight that our method is not only effective in the weather/lighting scenarios but also generalizes well to domain shifts caused by stylistic variations, such as differences between clipart, watercolor, and comic images. The improvements across all target domains indicate that the combination of domain-specific self-training and adaptive aggregation contributes to robust cross-domain transfer even under large appearance gaps.

Table 2: Multi-source domain adaptation results (mAP). For each target domain, Day Clear and the rest three domains are used as the source domains for the multi-source setting. For the single-source UDA and SFDA, Day Clear is used as the source following the typical setting Wu & Deng (2022); Vidit et al. (2023); Fahes et al. (2023).

		mAP			
Method	Multi-Source	NC	DR	NR	DF
FR	✗	34.4	26.0	12.4	32.0
SED	✗	33.4	21.1	15.1	29.4
HCL	✗	33.8	21.9	16.3	30.2
IRG	✗	42.7	30.5	18.4	35.2
LODS	✗	33.5	25.7	13.5	31.2
LPLD	✗	34.7	28.5	14.2	32.8
SF-UT	✗	36.8	30.0	16.9	34.2
DRU	✗	35.7	28.5	15.8	33.4
Mean-Teacher	✓	44.1	32.0	19.1	36.8
MixUp	✓	36.0	30.0	16.7	31.5
MSFDAOD	✓	42.1	30.8	18.7	34.4
CAiDA	✓	43.4	31.7	19.5	35.2
SST	✓	43.8	32.0	19.6	35.4
ATEN	✓	43.9	32.2	19.8	35.5
Bi-ATEN	✓	43.9	32.3	19.5	35.4
Ours	✓	**44.5**	**32.5**	**20.3**	**38.4**

Table 3: Comparison of detection performance (mAP %) on the Art dataset across three domains: Clipart1k, Watercolor2k, and Comic2k.

Method	Clipart1k	Watercolor2k	Comic2k
Mean-Teacher	41.98	42.55	40.32
MixUp	41.26	41.98	39.85
MSFDAOD	42.43	44.29	41.37
CAiDA	43.99	45.42	42.87
SST	46.45	48.33	44.92
Bi-ATEN	46.83	48.76	45.24
Ours	**48.45**	**50.33**	**46.39**

4.3 Ablation Study

Effectiveness of different components In this ablative study, we investigate the impact of each individual component in our framework. The detailed setting and results are shown in Table 1. We observe that, compared to a single source, the utilization of the multi-source knowledge fusion framework largely improves the mAP by 2.1 and the improvement is observed over all categories. By using the TFA, we managed to minimize the domain gap between all the source models and target data, which increased the mAP from 33.5 to 34.2. The confidence selection further improve the performance by 4.2 by filtering out noisy labels.

Text-driven feature augmentation In Figure 3, we show some stylized image rendered with TFA and the other two text-based style transfer methods. Take two images from the Day Clear domain, we augmented them with Day Foggy and Night Clear styles using TFA, CLIPstyler Kwon & Ye (2022) and PODA Fahes et al. (2023), respectively. From the augmented images, we observe that TFA renders augmented images with desirable style effects, including low-level texture and high-level concepts (like fog) while keeping the original content from being distorted. However, for CLIPstyler and PODA, the rendered style effects shroud the original content, making the objects in the original images hard to recognize. We present more qualitative results for TFA in the Appendix.

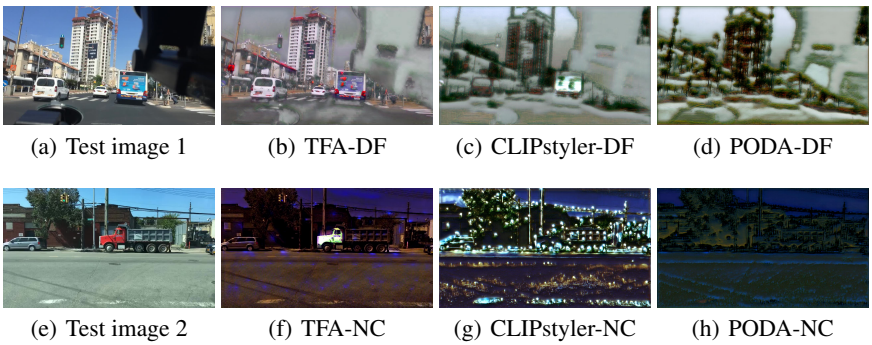

(a) Test image 1 (b) TFA-DF (c) CLIPstyler-DF (d) PODA-DF

(e) Test image 2 (f) TFA-NC (g) CLIPstyler-NC (h) PODA-NC

Figure 3: Two images (a)/(e) from Day Clear are stylized with different domain styles: Day Foggy (DF) and Night Clear (NC), using different text-based style transfer methods: TFA, CLIPstyler and PODA.

Robustness to source-domain descriptions. We evaluate whether TFA is sensitive to the choice of textual descriptions for the source domains. Descriptions 1 and 2 in Table 4 provide two independently phrased, coarse descriptions of the same source-domain appearance, while Description 3 contains intentionally unrelated text and serves as a negative control. Descriptions 1 and 2 achieve comparable performance despite substantial differences in wording, obtaining 38.4 and 36.5 mAP on Day Foggy, respectively. This suggests that TFA is not dependent on a specific prompt template. In contrast, the unrelated description substantially degrades performance to 12.0 mAP, as it provides no meaningful information about the source-domain appearance. Overall, these results indicate that TFA primarily benefits from coarse semantic relevance (e.g., 'rainy night" or 'foggy day") rather than carefully engineered prompts or precise low-level descriptions.

Table 4: Source domain text descriptions

Source Domain	Text Descriptions 1	Text Descriptions 2	Text 3: Unrelated negative controls
Day Clear	["Driving at clear day time"]	["Sunny day street view"]	["Lost in enchanted forest"]
Night Clear	["Driving at clear night time"]	["Clear night time street views"]	["Time traveler meets past self."]
Night Rainy	["Driving at rainy night time"]	["Rainy night street view"]	["Robot falls in love"]
Dusk Rainy	["Driving at rainy dusk"]	["Rainy dusk street view"]	["Magic potion gone wrong"]
Day Foggy	["Driving at foggy day time"]	["Foggy day street view"]	["Alien invasion thwarted heroically."]

4.4 Qualitative Study

We show some qualitative performance in Figure 2 with an example from Night Clear target domain. We observe that SS falsely detects multiple cars as one, misclassifies the Bus as Car, and fails to detect Person. By using the multi-source knowledge fusion, less mistakes were made, showing that the utilization of multiple sources help improve the generalization to some extend. While couple with TFA, MS+TFA successfully detects multiple cars, since the application of TFA reduce the domain gap between the multiple source model and the target images. Along with confidence selection, the noisy labels are filtered out, and we are able to detect all the objects.

5 Conclusion

In this study, we introduce a novel approach to perform multi-source-free domain adaptation, addressing the challenge of integrating multiple sources to harness information effectively. Our method first mitigates domain shift between multiple source domains and the target domain by transforming target images to match the styles of the source domains, utilizing text-based style transfer with textual descriptions of the source domains. Furthermore, we aggregate the information from multiple source models with a novel knowledge fusion framework, where the aggregator and domain experts are updated globally and locally, and the pseudo label quality are mutually enhanced. Experimental evaluations on diverse weather datasets demonstrate the efficacy of our proposed model across different domains. In future research, we plan to explore dynamic target domains where the target data distribution evolves over time.

AI USE STATEMENT

In this work, we have not used generative AI tools for generate synthetic data sets, help develop theoretical models or conceptual frameworks, formulate mathematical claims, provide critical ingredients for proving mathematical claims, design or provide feedback on research methodology or experiments, implement methods, assist with translation, clean and reformat dataset, support qualitative and thematic data analysis, interpret results, and assist in the writing of proofs, propose or refine hypotheses, are not applicable to this work. Additionally, we used generative AI tools for brainstorming, edit the paper to improve readability. We have reviewed all AI-assisted work. We checked LLM-generated research ideas for potential plagiarism through a manual literature survey. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

REFERENCES

Sk Miraj Ahmed, Dripta S Raychaudhuri, Sujoy Paul, Samet Oymak, and Amit K Roy-Chowdhury. Unsupervised multi-source domain adaptation without access to source data. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 10103–10112, 2021.

Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European conference on computer vision*, pp. 213–229. Springer, 2020.

Kai Chen, Jiaqi Wang, Jiangmiao Pang, Yuhang Cao, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jiarui Xu, Zheng Zhang, Dazhi Cheng, Chenchen Zhu, Tianheng Cheng, Qijie Zhao, Buyu Li, Xin Lu, Rui Zhu, Yue Wu, Jifeng Dai, Jingdong Wang, Jianping Shi, Wanli Ouyang, Chen Change Loy, and Dahua Lin. MMDetection: Open mmlab detection toolbox and benchmark. *arXiv preprint arXiv:1906.07155*, 2019.

Junhyeong Cho, Gilhyun Nam, Sungyeon Kim, Hunmin Yang, and Suha Kwak. Promptstyler: Prompt-driven style generation for source-free domain generalization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 15702–15712, 2023.

Sungha Choi, Sanghun Jung, Huiwon Yun, Joanne T Kim, Seungryong Kim, and Jaegul Choo. Robustnet: Improving domain generalization in urban-scene segmentation via instance selective whitening. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11580–11590, 2021.

Qiaosong Chu, Shuyan Li, Guangyi Chen, Kai Li, and Xiu Li. Adversarial alignment for source free object detection. *arXiv preprint arXiv:2301.04265*, 2023.

Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.

Zhengming Ding, Ming Shao, and Yun Fu. Incomplete multisource transfer learning. *IEEE transactions on neural networks and learning systems*, 29(2):310–323, 2016.

Jiahua Dong, Zhen Fang, Anjin Liu, Gan Sun, and Tongliang Liu. Confident anchor-induced multi-source free domain adaptation. *Advances in Neural Information Processing Systems*, 34: 2848–2860, 2021.

Mohammad Fahes, Tuan-Hung Vu, Andrei Bursuc, Patrick Pérez, and Raoul de Charette. Poda: Prompt-driven zero-shot domain adaptation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 18623–18633, 2023.

Rinon Gal, Or Patashnik, Haggai Maron, Amit H Bermano, Gal Chechik, and Daniel Cohen-Or. Stylegan-nada: Clip-guided domain adaptation of image generators. *ACM Transactions on Graphics (TOG)*, 41(4):1–13, 2022.

Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The kitti dataset. *International Journal of Robotics Research (IJRR)*, 2013a.

Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research*, 32(11):1231–1237, 2013b.

Bo Han, Quanming Yao, Xingrui Yu, Gang Niu, Miao Xu, Weihua Hu, Ivor Tsang, and Masashi Sugiyama. Co-teaching: Robust training of deep neural networks with extremely noisy labels. *Advances in neural information processing systems*, 31, 2018.

Yan Hao, Florent Forest, and Olga Fink. Simplifying source-free domain adaptation for object detection: Effective self-training strategies and performance insights. In *European Conference on Computer Vision*, pp. 196–213. Springer, 2024.

Mahmoud Hassaballah, Mourad A Kenk, Khan Muhammad, and Shervin Minaee. Vehicle detection and tracking in adverse weather using a deep learning framework. *IEEE transactions on intelligent transportation systems*, 22(7):4230–4242, 2020.

Jan Hosang, Rodrigo Benenson, and Bernt Schiele. Learning non-maximum suppression. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4507–4515, 2017.

Jiaxing Huang, Dayan Guan, Aoran Xiao, and Shijian Lu. Model adaptation: Historical contrastive learning for unsupervised domain adaptation without source data. *Advances in Neural Information Processing Systems*, 34:3635–3649, 2021.

Lei Huang, Yi Zhou, Fan Zhu, Li Liu, and Ling Shao. Iterative normalization: Beyond standardization towards efficient whitening. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 4874–4883, 2019.

Naoto Inoue, Ryosuke Furuta, Toshihiko Yamasaki, and Kiyoharu Aizawa. Cross-domain weakly-supervised object detection through progressive domain adaptation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 5001–5009, 2018.

Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 4401–4410, 2019.

Trinh Le Ba Khanh, Huy-Hung Nguyen, Long Hoang Pham, Duong Nguyen-Ngoc Tran, and Jae Wook Jeon. Dynamic retraining-updating mean teacher for source-free object detection. In *European Conference on Computer Vision*, pp. 328–344. Springer, 2024.

Youngeun Kim, Donghyeon Cho, Kyeongtak Han, Priyadarshini Panda, and Sungeun Hong. Domain adaptation without source data. *IEEE Transactions on Artificial Intelligence*, 2(6):508–518, 2021.

Gihyun Kwon and Jong Chul Ye. Clipstyler: Image style transfer with a single text condition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 18062–18071, 2022.

Aristotelis Leventidis, Laura Di Rocco, Wolfgang Gatterbauer, Renée J Miller, and Mirek Riedewald. Domainnet: Homograph detection for data lake disambiguation. *arXiv preprint arXiv:2103.09940*, 2021.

Bowen Li, Xiaojuan Qi, Thomas Lukasiewicz, and Philip HS Torr. Manigan: Text-guided image manipulation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7880–7889, 2020.

Shuai Li, Jianqiang Huang, Xian-Sheng Hua, and Lei Zhang. Category dictionary guided unsupervised domain adaptation for object detection. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, pp. 1949–1957, 2021a.

Shuaifeng Li, Mao Ye, Xiatian Zhu, Lihua Zhou, and Lin Xiong. Source-free object detection by learning to overlook domain style. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 8014–8023, 2022.

Xianfeng Li, Weijie Chen, Di Xie, Shicai Yang, Peng Yuan, Shiliang Pu, and Yueting Zhuang. A free lunch for unsupervised domain adaptive object detection without source data. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pp. 8474–8481, 2021b.

Xinyao Li, Jingjing Li, Fengling Li, Lei Zhu, and Ke Lu. Agile multi-source-free domain adaptation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 13673–13681, 2024.

Zijian Li, Ruichu Cai, Guangyi Chen, Boyang Sun, Zhifeng Hao, and Kun Zhang. Subspace identification for multi-source domain adaptation. In *Adv. Neural Inform. Process. Syst.*, 2023.

Luojun Lin, Zhifeng Yang, Qipeng Liu, Yuanlong Yu, and Qifeng Lin. Run and chase: Towards accurate source-free domain adaptive object detection. In *2023 IEEE International Conference on Multimedia and Expo (ICME)*, pp. 2453–2458. IEEE, 2023.

Qipeng Liu, Luojun Lin, Zhifeng Shen, and Zhifeng Yang. Periodically exchange teacher-student for source-free object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 6414–6424, 2023.

Xingang Pan, Ping Luo, Jianping Shi, and Xiaoou Tang. Two at once: Enhancing learning and generalization capacities via ibn-net. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 464–479, 2018.

Xingang Pan, Xiaohang Zhan, Jianping Shi, Xiaoou Tang, and Ping Luo. Switchable whitening for deep representation learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 1863–1871, 2019.

Or Patashnik, Zongze Wu, Eli Shechtman, Daniel Cohen-Or, and Dani Lischinski. Styleclip: Text-driven manipulation of stylegan imagery. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 2085–2094, 2021.

Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 1406–1415, 2019.

Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pp. 8748–8763. PMLR, 2021.

Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28, 2015.

Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Semantic foggy scene understanding with synthetic data. *International Journal of Computer Vision*, 126:973–992, 2018.

Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Acdc: The adverse conditions dataset with correspondences for semantic driving scene understanding. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 10765–10775, 2021.

Maohao Shen, Yuheng Bu, and Gregory W Wornell. On balancing bias and variance in unsupervised multi-source-free domain adaptation. In *International Conference on Machine Learning*, pp. 30976–30991. PMLR, 2023.

Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *Advances in neural information processing systems*, 30, 2017.

VS Vibashan, Poojan Oza, and Vishal M Patel. Instance relation graph guided source-free domain adaptive object detection. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3520–3530. IEEE, 2023.

Vidit Vidit, Martin Engilberge, and Mathieu Salzmann. Clip the gap: A single domain generalization approach for object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3219–3229, 2023.

Zirui Wang, Zihang Dai, Barnabás Póczos, and Jaime Carbonell. Characterizing and avoiding negative transfer. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 11293–11302, 2019.

Aming Wu and Cheng Deng. Single-domain generalized object detection in urban scene via cyclic-disentangled self-distillation. In *Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition*, pp. 847–856, 2022.

Aming Wu, Rui Liu, Yahong Han, Linchao Zhu, and Yi Yang. Vector-decomposed disentanglement for domain-invariant object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 9342–9351, 2021.

Lin Xiong, Mao Ye, Dan Zhang, Yan Gan, Xue Li, and Yingying Zhu. Source data-free domain adaptation of object detector through domain-specific perturbation. *International Journal of Intelligent Systems*, 36(8):3746–3766, 2021.

Minghao Xu, Hang Wang, Bingbing Ni, Qi Tian, and Wenjun Zhang. Cross-domain detection via graph-induced prototype alignment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 12355–12364, 2020.

Tao Xu, Pengchuan Zhang, Qiuyuan Huang, Han Zhang, Zhe Gan, Xiaolei Huang, and Xiaodong He. Attngan: Fine-grained text to image generation with attentional generative adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1316–1324, 2018.

Hao-Wei Yeh, Qier Meng, and Tatsuya Harada. Misalignment-free relation aggregation for multi-source-free domain adaptation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 4313–4322, 2023.

Ilhoon Yoon, Hyeongjun Kwon, Jin Kim, Junyoung Park, Hyunsung Jang, and Kwanghoon Sohn. Enhancing source-free domain adaptive object detection with low-confidence pseudo label distillation. In *European Conference on Computer Vision*, pp. 337–353. Springer, 2024.

Fisher Yu, Haofeng Chen, Xin Wang, Wenqi Xian, Yingying Chen, Fangchen Liu, Vashisht Madhavan, and Trevor Darrell. Bdd100k: A diverse driving dataset for heterogeneous multitask learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 2636–2645, 2020.

Han Zhang, Tao Xu, Hongsheng Li, Shaoting Zhang, Xiaogang Wang, Xiaolei Huang, and Dimitris N Metaxas. Stackgan: Text to photo-realistic image synthesis with stacked generative adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pp. 5907–5915, 2017a.

Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017b.

Sicheng Zhao, Huizai Yao, Chuang Lin, Yue Gao, and Guiguang Ding. Multi-source-free domain adaptive object detection. *International Journal of Computer Vision*, 132(12):5950–5982, 2024.

Appendix

Appendix Contents

A Notations and Algorithms **16**

B Experimental Details **17**
B.1 Implementation Details 17
Implementation details of the multi-source knowledge fusion framework. 17
B.2 Details of the Datasets 17

C Additional Experimental Results **18**
C.1 Per-class Analysis of the DWD Dataset 18
Multi-source→Day Foggy. 19
Multi-source→Dusk Rainy. 19
Multi-source→Night Clear. 19
Multi-source→Night Rainy. 19
C.2 Per-class analysis on Foggy-Cityscapes 19
C.3 Additional Ablation Studies 19
TFA hyperparameter λ_1, λ_2. 19
Heterogeneous source model architectures. 20
Impact of aggregator selection. 22
Ablation on backbone architectures. 22
Adaptation and inference costs. 23
Impact of different text-based style transfer techniques. 23
Impact of fixed EMA rate. 23
Evaluation on segmentation and classification tasks. 23
C.4 Qualitative Results 23
C.5 Illustration of failed examples 24
C.6 Comparison of TFA with Existing Text-Based Style Transfer Methods 27
C.7 Comparison of Domain Adaptation Methods with Access to the Source Data . . 27

D Potential Social Impact and Limitations **28**

E Source Code **29**

A NOTATIONS AND ALGORITHMS

In this section, we present the table of notations used in the main paper, and the pseudo code for our algorithms. Table 5 summarizes the major notations used in the paper. Our algorithm consists of two main components. In Algorithm 1, we introduce the pseudo code for TFA. In Algorithm 2, we outline the multi-source free domain adaptation with the mean-teacher framework and the confidence selection via co-teaching.

Table 5: Notations

Notation	Description
D^{s_i}	i-th source domain
D^t	target domain
θ^{s_i}	pre-trained i-th source domain model
$x_k^{s_i}$	k-th feature from i-th source domain
$y_k^{s_i}$	k-th label from i-th source domain
$b_k^{s_i}$	k-th bounding box from i-th source domain
$c_k^{s_i}$	k-th class label from i-th source domain
x_k^t	k-th feature from target domain
M	number of source models
N_i	number of data samples in i-th source domain
N_t	number of data samples in target domain
ϕ	TFA augmentation parameters
E_{txt}	CLIP text encoder
E_{vis}	CLIP image encoder
$\cos(\cdot)$	cosine similarity
$\mathbf{t}_{\text{sty}}$	domain style description
$\mathcal{L}_{\text{style}}$	text image style consistent loss
$\mathcal{L}_{\text{content}}$	content preservation loss
$D^{t\to s_i}$	augmented target domain datasets with i-th source domain style
θ^{agg}	aggregator
θ_i^{DE}	i-th domain expert
$\mathcal{L}_i^{\text{local}}$	i-th local loss
$\mathcal{L}_{cls}^{rpn}$	region classification loss
$\mathcal{L}_{reg}^{rpn}$	region proposal loss
$\mathcal{L}_{cls}^{roi}$	bounding box classification loss
$\mathcal{L}_{reg}^{roi}$	bounding box regression loss
$\boldsymbol{\alpha}$	EMA learning rates

Algorithm 1 Text-driven Feature Augmentation (TFA)

INPUT Target dataset D^t, text description of style $T^{src} = \{\mathbf{t}_{\text{sty}}^{s_i}\}_{i=1}^M$ of source domain s_i, ϕ
OUTPUT Multiple augmented datasets $\{D^{t\to s_i}\}_{i=1}^M$
for each target image $\mathbf{x} \in D^t$ **do**
  Using CLIP image encoder to extract target image feature: $\mathbf{z} = E_{\text{vis}}(\mathbf{x})$
  for each source style text $\mathbf{t}_{\text{sty}}^{s_i} \in T^{src}$ **do**
    Using CLIP text encoder to extract source text feature: $E_{\text{txt}}(\mathbf{t}_{\text{sty}}^{s_i})$
    while not converged **do**
      Obtain augmented image feature with s-th source style using $\mathcal{A}_\phi(\cdot)$: $\mathcal{A}_\phi(\mathbf{z})$ (i.e., $E_{\text{vis}}(\mathcal{A}_\phi(\mathbf{x}))$)
      Update ϕ with Equation (2)
    end while
  end for
end for

Algorithm 2 Multi-Source Domain Adaptation (MSDA)

INPUT Multiple pre-trained source models $\{\theta^{s_i}\}_{i=1}^{M}$, target dataset D^t, and multiple stylized target datasets $\{D^{t\to s_i}\}_{i=1}^{M}$, EMA parameters $\boldsymbol{\alpha}$, target dataset D^t
Choose a pre-trained source model as aggregator θ^{agg}
Choose the rest pre-trained source models as domain expert $\{\theta_1^{\text{DE}}, ..., \theta_{M-1}^{\text{stu}}\}$
while not converged **do**
  // Update aggregator and domain experts
  for each model θ^{s_i} in $\{\theta^{s_i}\}_{i=1}^{M}$ **do**
    Sample a batch of data from $D^{t\to s_i}$
    Generate pseudo label $\tilde{y}_i$ with θ^{s_i}
    // Confidence selection in a co-teaching manner
    if the model is a domain expert **then**
      Use the aggregator to filter out noisy labels
    else
      Randomly select a domain expert to filter out noisy labels
    end if
    Update θ^{s_i} with Equation (6)
  end for
  Update aggregator
  // Update EMA parameters $\boldsymbol{\alpha}$
  Sample a batch of data from target dataset D^t
  for each model θ^{s_i} in $\{\theta^{s_i}\}_{i=1}^{M}$ **do**
    Compute the prediction probability of each model on the data
  end for
  Update $\boldsymbol{\alpha}$
end while

B Experimental Details

B.1 Implementation Details

Implementation details of the multi-source knowledge fusion framework. To complement the pseudo-code, we outline the training and evaluation procedure of our proposed method:

- **Step 1: Pretraining source models.** We first pretrain a set of source models on their respective labeled source-domain datasets. The backbone architecture of each model follows the experimental setting (e.g., Faster R-CNN, ATSS, or YOLOv7).
- **Step 2: Pretraining the feature augmentation module.** The feature augmentation module is trained following the procedure in Algorithm 1.
- **Step 3: Generating target-augmented datasets.** Using the trained augmentation module, we translate target-domain samples into multiple source-domain styles, producing augmented datasets for subsequent training.
- **Step 4: Training the knowledge fusion framework.** With both pretrained source models and target-augmented datasets, we train the proposed multi-source fusion framework as described in Section 3.2 and detailed in Algorithm 2. In this stage, domain experts are updated via self-training on the target-style data, while the aggregator progressively integrates knowledge across domains through EMA updates.
- **Step 5: Evaluation.** After training, the aggregated model is evaluated on the held-out target domain. We report standard detection metrics (e.g., mAP, mAP@0.5) to assess the performance and cross-domain generalization ability of the framework.

B.2 Details of the Datasets

The images of the DWD dataset was selected from three primary datasets, Berkeley Deep Drive 100K (BBD-100K) Yu et al. (2020), Cityscapes Cordts et al. (2016) and Adverse-Weather Hassaballah et al. (2020). Additionally, rainy images are rendered by Wu et al. (2021), and some of the foggy images

Table 6: Source style description for DWD.

Domain	Style Descriptions
Day Clear	“A photo of a clear, sunny daytime scene."
Night Clear	“A photo of a clear, well-lit nighttime scene."
Night Rainy	“A photo of a rainy, dimly-lit nighttime scene."
Dusk Rainy	“A photo of a rainy scene at dusk."
Day Foggy	“A photo of a foggy scene during the day."

Table 7: Source style description for Art.

Domain	Style Descriptions
Clipart	“An illustration in a flat, clipart style."
Comic	“A drawing in bold, comic-book style."
Watercolor	“A painting in soft, watercolor brush strokes."

are synthetically generated from Sakaridis et al. (2018). The daytime clear dataset consists of 27708 images, the night clear dataset contains 26158 images, the dusk rainy dataset has 3501 images, the night rainy dataset has 2494 images, and the daytime foggy dataset has 3775 images. All the datasets contain bounding box annotations for the 7 classes objects: *bus, bike, car, motorbike, person, rider*, and *truck*. For text augmentation, we utilize the domain description in Table 6.

For the combined dataset, **Cityscapes**[1] Cordts et al. (2016), **Foggy-Cityscapes**[2] Sakaridis et al. (2018), and **KITTI**[3] Geiger et al. (2013b) for further evaluation. Cityscapes consist of 2975 training images and 500 testing images, have a total of 8 categories captured under normal weather. Foggy-Cityscapes applys images of Cityscapes to simulate foggy as well as inherits the annotations of Cityscapes. KITTI contains 7,481 urban images of the same classes which are different from Cityscapes. For the comparative baselines, the training set of Cityscapes are utlized to pre-trained the source model, and test on the test set of Foggy-Cityscapes following the general setting Xu et al. (2020); Li et al. (2021a). In addition, we incorporate the validation set of KITTI dataset as an additional source model, which includes 1870 images, and then test our model on the test set of Foggy-Cityscapes. For text augmentation, we utilize the domain description in Table 13.

The Art dataset contains different artistic styles including Clipart1k, Comic2k, and Watercolor2k Inoue et al. (2018). Clipart1k contains 1000 clipart images across 20 classes, Watercolor2k and Comic2k contains 2000 watercolor/comic images across 6 classes. For text augmentation, we utilize the domain description in Table 7.

C Additional Experimental Results

In this section, we first provide the per-class analysis of each domain of DWD dataset in Appendix C.1. We than provide experimental results on the Art dataset in Table 3. In Appendix C.3, we conduct ablation studies including the impact of source domain style descriptions, the impact of TFA hyperparameters λ_1, λ_2, the impact of heterogeneous architecture of source models, the impact of aggregator, and some qualitative results. In Appendix C.4, we show some qualitative results of using TFA with different source descriptions, and their detection results. In Appendix C.6, we compare the qualitative results of TFA with other two text-based style transfer methods CLIPstyler Kwon & Ye (2022) and PODA Fahes et al. (2023).

C.1 Per-class Analysis of the DWD Dataset

In this section, we present per-class analysis for each domain of the DWD Dataset, including Day Foggy, Dusk Rainy, Night Clear and Night Rainy.

[1] https://github.com/tiancity-NJU/da-faster-rcnn-PyTorch
[2] https://www.cityscapes-dataset.com/downloads/
[3] http://www.cvlibs.net/datasets/kitti/

Multi-source→Day Foggy. We show the detailed results of each class of Day Foggy in Table 9. In general MSFDA outperform SFDA in most categories, except for Car, FR ranks the top. Our method achieves the best performance on Bus and Truck, and performs consistently better on other classes, resulting the top mAP performance.

Multi-source→Dusk Rainy. In Table 10, our method outperforms SFDA and MSFDA in most of the categories except for Car, Motor and Truck, achieving the best average mAP over all categories. Compared to SFDA, our method either ranks the top or the second except for Motor, resulting in the top 1 average mAP over all categories. Mean-Teacher ranks the top for Car and Motor, but our method excel it by outperforming on other classes consistently.

Multi-source→Night Clear. In Table 11, our method performs consistently better in most categories, ranking the top at four of them, achieving the highest mAP. For SFDA, FR ranks the top in Car. And for MSFDA, Mean-Teacher ranks the top at Bus and Motor.

Multi-source→Night Rainy. As shown in Table 12, in the more challenging Night Rainy dataset, our method achieve the best average mAP by performing consistently well through in two categories, and consistently better in other categories.

C.2 Per-class analysis on Foggy-Cityscapes

In Table 8, we present a per-class analysis on the Foggy-Cityscapes dataset using Cityscapes and KITTI as source domains. The results demonstrate that our method achieves the highest mAP by consistently ranking at the top across multiple classes. Similarly, MSFDA methods outperform SFDA methods, highlighting the advantages of leveraging multiple source domains.

Table 8: Per-class results on Foggy-Cityscapes

		AP								mAP
Method	Multi-Source	Bus	Bike	Car	Motor	Person	Rider	Truck	Train	All
SED	✗	11.8	34.3	40.4	34.5	21.7	44.0	32.6	25.3	30.6
SED(Moisac)	✗	22.2	39.0	40.7	34.1	25.5	44.5	33.2	28.4	33.5
HCL	✗	25.0	**46.0**	41.3	**35.9**	26.9	40.7	33.0	28.1	34.6
SOAP	✗	37.2	37.9	48.4	31.8	35.9	45.0	23.9	24.3	35.5
LODS	✗	39.7	37.8	48.8	33.2	34.0	**45.7**	27.3	19.6	35.8
IRG	✗	39.6	41.6	**51.9**	31.5	37.4	45.2	24.4	25.2	37.1
LPLD	✗	37.4	37.8	48.7	32.0	36.1	45.3	24.0	24.5	35.6
SF-UT	✗	38.5	40.5	50.7	30.8	36.5	44.6	23.5	24.1	36.1
DRU	✗	39.2	40.9	50.9	31.2	36.7	44.9	23.9	24.3	36.4
Mean-Teacher	✓	39.2	40.2	47.0	27.6	35.9	45.0	31.2	27.1	36.3
MixUp	✓	38.5	39.8	44.0	26.1	30.4	42.8	29.1	26.4	34.2
MSFDAOD	✓	39.9	33.2	47.3	29.5	33.8	45.0	32.4	29.8	37.6
CAiDA	✓	41.8	44.5	47.3	29.8	34.0	45.8	33.5	32.0	37.8
SST	✓	42.2	44.2	44.9	30.3	33.5	47.2	36.0	33.2	37.9
ATEN	✓	42.6	44.4	48.8	30.0	34.1	**47.5**	36.4	**33.3**	38.1
Bi-ATEN	✓	**42.6**	44.5	47.9	29.7	34.0	47.4	**36.5**	33.1	38.0
Ours	✓	40.1	39.8	49.9	35.4	40.4	**46.1**	35.5	27.5	**39.2**

C.3 Additional Ablation Studies

In this section, we present the ablation study of using different source domain text descriptions, the impact of TFA hyperparameters λ_1, and λ_2 in Equation (2), and the impact of heterogeneous source model architectures.

TFA hyperparameter λ_1, λ_2. In this ablation, we test the hyperparameters λ_1, λ_2 for TFA in Eq. (2). In the paper, the default weights of λ_1, λ_2 are 2, 1, respectively. In this ablation, we show some different combination of the weights and the per-class results are shown in Table 14. As observed, by

Table 9: Per-class results on multi-source to Day Foggy (the setting of source models is the same as in Table 2).

		AP							mAP
Method	Multi-Source	Bus	Bike	Car	Motor	Person	Rider	Truck	All
FR	✗	28.1	29.7	**49.7**	26.3	33.2	35.5	21.5	32.0
SED	✗	28.4	29.1	28.5	24.1	33.9	30.4	32.7	29.4
HCL	✗	32.5	31.3	32.1	25.9	28.0	34.2	31.8	30.2
IRG	✗	33.8	33.9	34.2	36.8	37.5	38.9	34.8	35.2
LODS	✗	28.5	29.4	33.8	29.7	34.5	34.9	21.2	31.2
LPLD	✗	28.9	30.2	34.2	30.5	34.8	35.8	23.0	32.8
SF-UT	✗	30.0	30.8	35.4	31.7	35.8	36.9	24.2	34.2
DRU	✗	29.8	30.2	35.0	31.0	35.2	36.1	22.8	33.4
Mean-Teacher	✓	35.4	**37.9**	40.2	**39.2**	31.5	33.4	32.9	36.8
MixUp	✓	33.2	32.4	33.5	26.8	29.1	35.5	33.2	31.5
MSFDAOD	✓	31.5	32.8	36.0	33.4	38.2	38.9	29.5	34.4
CAiDA	✓	32.4	33.6	37.2	34.2	39.5	39.9	30.4	35.2
SST	✓	32.8	33.9	37.4	34.5	40.3	**40.8**	30.5	35.4
ATEN	✓	33.0	34.1	37.9	34.8	**40.6**	41.1	30.9	35.5
Bi-ATEN	✓	32.7	34.2	37.5	34.2	40.4	40.7	30.6	35.4
Ours	✓	**36.4**	35.5	45.8	34.9	39.8	35.0	**35.0**	**38.4**

Table 10: Per-class results on multi-source to Dusk Rainy (the setting of source models is the same as in Table 2).

		AP							mAP
Method	Multi-Source	Bus	Bike	Car	Motor	Person	Rider	Truck	All
FR	✗	28.5	20.3	**58.2**	6.5	23.4	11.3	33.9	26.0
SED	✗	20.4	21.5	20.6	20.8	27.5	18.3	24.5	21.1
HCL	✗	26.7	14.2	22.4	14.2	22.9	14.3	30.5	21.9
IRG	✗	35.2	21.4	29.8	15.9	26.5	22.4	38.7	30.5
LODS	✗	27.1	30.9	23.4	19.7	16.3	30.7	28.2	25.7
LPLD	✗	27.9	31.4	24.2	22.5	19.9	31.5	29.2	28.5
SF-UT	✗	28.2	31.6	24.5	22.9	20.9	31.8	29.4	30.0
DRU	✗	28.0	31.2	24.5	22.4	20.4	31.6	29.4	28.5
Mean-Teacher	✓	29.0	31.2	33.8	**35.6**	30.5	27.8	26.2	32.0
MixUp	✓	30.5	30.1	28.8	34.7	**32.5**	28.4	28.2	30.8
MSFDAOD	✓	30.4	30.3	29.2	34.9	32.3	28.1	28.7	30.8
CAiDA	✓	31.2	31.1	30.5	35.4	33.5	28.9	29.7	31.7
SST	✓	31.4	31.3	30.8	35.5	33.8	29.1	30.2	32.0
ATEN	✓	31.5	31.4	31.1	35.8	34.0	29.0	30.4	32.2
Bi-ATEN	✓	31.6	31.2	31.5	35.9	34.2	29.3	30.4	32.3
Ours	✓	**34.5**	**33.1**	30.4	29.5	31.8	**34.6**	**34.7**	**32.5**

gradually increasing the weight of $\mathcal{L}_{\text{style}}$ from 0.1 to 10, the per-class performances improve steadily, showing the effectiveness of style transfer via text. When we increase the weight of $\mathcal{L}_{\text{content}}$ from 1 to 5, the performance improves from 35.5 to 38.4, showing that preserving the content information is very important during style transfer.

Heterogeneous source model architectures. In our experiment, we use faster R-CNN as both pre-trained source and target models to demonstrate the effectiveness of our proposed multi-source free domain adaptation algorithm. To further validate the effectiveness, we include the additional architecture experiments with ATSS and YOLO-v7 as model architecture choices, and compare with the latest source-free baselines. The results are presented in Table 15. Despite the choices of model architectures, our methods shows consistently better performance. The algorithm can be further

Table 11: Per-class results on multi-source to Night Clear (the setting of source models is the same as in Table 2).

Method	Multi-Source	AP							mAP
		Bus	Bike	Car	Motor	Person	Rider	Truck	All
FR	✗	34.7	32.0	**56.6**	13.6	37.4	27.6	38.6	34.4
SED	✗	31.9	34.5	33.8	31.2	32.5	34.9	33.7	33.4
HCL	✗	33.4	32.9	33.4	33.1	34.7	35.1	34.5	33.8
IRG	✗	43.2	41.8	42.4	42.4	43.5	44.1	42.4	42.7
LODS	✗	32.0	34.6	33.9	31.4	32.4	35.1	33.8	33.5
LPLD	✗	33.2	32.7	33.5	33.2	34.5	35.0	34.6	34.7
SF-UT	✗	35.1	34.2	35.9	35.3	35.8	36.2	35.9	36.8
DRU	✗	34.0	33.5	34.8	34.8	34.4	35.8	33.9	35.7
Mean-Teacher	✓	**44.5**	42.3	45.6	**45.2**	43.2	44.5	43.4	44.1
MixUp	✓	37.2	35.4	35.8	35.6	36.4	36.8	35.6	36.0
MSFDAOD	✓	42.8	41.2	41.5	41.6	43.0	43.2	41.5	42.1
CAiDA	✓	43.3	41.4	41.8	41.9	43.5	43.6	41.9	43.4
SST	✓	43.9	41.8	42.4	42.6	44.0	44.1	42.4	43.8
ATEN	✓	43.9	41.9	42.6	42.8	44.1	44.3	42.5	43.9
Bi-ATEN	✓	43.8	41.9	42.6	43.0	44.2	44.2	42.3	43.9
Ours	✓	44.0	**42.9**	43.5	44.8	**44.9**	**45.5**	**45.0**	**44.5**

Table 12: Per-class results on multi-source to Night Rainy (the setting of source models is the same as in Table 2).

Method	Multi-Source	AP							mAP
		Bus	Bike	Car	Motor	Person	Rider	Truck	All
FR	✗	16.8	6.9	**26.3**	0.6	11.6	9.4	15.4	12.4
SED	✗	15.8	14.5	14.2	18.6	6.9	16.5	18.8	15.1
HCL	✗	16.6	13.8	14.5	17.0	17.1	15.6	16.2	16.3
IRG	✗	18.5	24.8	15.4	18.2	18.3	16.7	16.4	18.4
LODS	✗	16.7	7.2	25.4	0.9	12.4	10.8	16.5	13.5
LPLD	✗	16.9	9.8	25.5	3.2	15.2	12.4	18.4	14.2
SF-UT	✗	17.4	10.5	25.2	5.2	16.5	12.8	18.2	16.9
DRU	✗	17.1	10.0	25.3	3.4	15.5	12.8	18.6	15.8
Mean-Teacher	✓	14.9	18.8	19.4	17.6	17.2	25.5	**20.7**	19.1
MixUp	✓	16.9	14.2	14.8	17.5	17.7	16.0	16.5	16.7
MSFDAOD	✓	19.2	16.4	17.0	19.1	18.0	17.2	17.4	18.7
CAiDA	✓	19.9	17.4	18.2	20.0	19.2	18.3	18.5	19.5
SST	✓	19.9	17.7	18.0	19.8	19.1	18.0	18.6	19.6
ATEN	✓	**20.2**	18.0	18.4	**20.2**	**19.3**	18.2	18.5	19.8
Bi-ATEN	✓	19.8	17.5	17.8	19.6	18.8	17.8	18.5	19.5
Ours	✓	16.5	**19.6**	20.2	19.4	18.4	**26.2**	20.5	**20.3**

Table 13: Source Domain Text Description

Source Domain	Descriptions
Cityscapes	["Driving at city streets"]
KITTI	["Driving at urban streets"]

extended to multiple architecture setting with minor modifications. Specifically, instead of having only one aggregator, we can initialize multiple aggregator-domain expert pairs with the different architectures, where each pair share the same architecture for one dataset. For the EMA update of

Table 14: Per-class results on multi-source to Day Foggy with Ablations on TFA hyperparameters in Equation (2).

Hyperparameters	AP							mAP
(λ_1, λ_2)	Bus	Bike	Car	Motor	Person	Rider	Truck	All
$(0.1, 1)$	31.0	29.6	28.4	32.8	30.8	32.1	35.6	34.3
$(1, 1)$	31.4	29.8	29.1	33.4	31.5	32.8	36.2	34.9
$(10, 1)$	32.1	30.7	29.6	33.8	32.0	33.1	**36.5**	35.5
$(2, 1)$ (default)	**36.4**	**35.5**	**45.8**	**34.9**	**39.8**	**35.0**	35.0	**38.4**

the aggregator, we update it with the corresponding domain expert, while the domain expert update remains unchanged.

During inference, the ensemble of domain expert is utilized. We present the result of this extended setting as follows, where we set four model pairs, two with faster-RCNN for Night Clear and Dusk Rainy datasets and two ATSS and YOLO-v7 for Night Rainy and Day Foggy datasets, respectively, denoted as FR+ATSS and YOLO-v7 in the table below.

Table 15: Multi-source domain adaptation results (mAP) using different model architectures (ATSS, YOLO-v7). For each target domain, Day Clear and the rest three other domains are used as the source domains for the multi-source setting.

Method	Multi-Source	Night Clear	Dusk Rainy	Night Rainy	Day Foggy
SED (ATSS)	✗	32.9	20.5	24.7	28.9
Mean-Teacher (ATSS)	✓	43.2	31.4	28.6	36.1
MixUp (ATSS)	✓	41.5	30.4	26.1	31.0
Ours (ATSS)	✓	**44.1**	**32.0**	**28.7**	**37.8**
Ours (FR+ATSS)	✓	43.8	31.6	28.3	37.4
SED (YOLO-v7)	✗	34.2	22.0	26.2	30.6
Mean-Teacher (YOLO-v7)	✓	45.2	32.8	30.0	38.1
MixUp (YOLO-v7)	✓	43.0	31.7	27.5	32.8
Ours (YOLO-v7)	✗	**45.6**	**33.5**	**30.2**	**39.6**
Ours (FR+YOLO-v7)	✓	45.3	33.3	30.0	39.2

Impact of aggregator selection. The final aggregator is used directly for target-domain inference without test- time stylization. We initialize it from a source model whose coarse domain semantics are close to the target (e.g., weather and time attributes). This is a simple initialization heuristic rather than an oracle selection rule. Table 16 shows that performance varies with the initialization, while semantically closer choices generally perform better. For Night Rainy, Dusk Rainy yields the best result (20.3 mAP), followed by Night Clear. These results motivate the adaptive fusion stage and also indicate that automatic aggregator initialization is a useful direction for future work.

Table 16: mAP on Night Rainy using different aggregator.

Aggregator	Night Clear	Day Foggy	Dusk Rainy	Day Clear
mAP	18.9	17.5	**20.3**	18.4

Ablation on backbone architectures. To verify the robustness and generalizability of our proposed method beyond the initial Faster R-CNN + ResNet-50 setup, we evaluate our framework using a more modern detector, DETR with a ViT backbone Carion et al. (2020). This ablation addresses two key questions: (1) whether the proposed framework improvements are architecture-agnostic, and (2) whether the method provides consistent gains even when applied to stronger baselines. The results in Table 17 show that our method consistently outperforms all baselines across every target domain with DETR + ViT. These findings indicate that the method's effectiveness is not tied to a specific backbone; it adapts well to both CNN-based and transformer-based detectors. Moreover, the improvements demonstrate that multi-source knowledge fusion scales with stronger architectures,

highlighting its flexibility and applicability across different detection paradigms. This confirms that the framework can robustly handle domain shifts while leveraging the representational power of modern detection backbones.

Table 17: Comparison of methods on the NC → DR → NR → NC adaptation setting.

Method	NC	DR	NR	NC (cycle)
SED	32.4	29.0	16.3	30.2
Mean-Teacher	43.1	33.0	21.4	37.2
MixUp	39.6	30.0	18.3	33.3
Ours	**44.8**	**34.3**	**22.6**	**38.7**

Adaptation and inference costs. We distinguish the one-time target adaptation cost from deployment-time inference cost. Table 18 reports wall-clock adaptation time, GPU-hours, inference throughput, and model size under matched hardware and evaluation settings. For our method, the adaptation time includes training the TFA module, generating and caching all source-stylized target images, local updates of the domain experts, and aggregator updates. TFA generation is performed once and the resulting stylized images are reused during adaptation. Our method requires 1.075 hours (2.15 GPU-hours) of target adaptation and maintains 17 FPS at inference with 126M parameters. We note that, as in other multi-source methods, adaptation cost grows with the number of source experts; this cost is incurred only during adaptation and does not increase the deployment-time cost of the final detector.

Impact of different text-based style transfer techniques. In this ablation, we investigate the effect of different text-based style transfer methods, including our proposed TFA, PODA Fahes et al. (2023), and ClipStyler Kwon & Ye (2022). As noted earlier, TFA achieves superior style transfer by effectively altering the image style while preserving content integrity. We further evaluate the performance of our multi-source knowledge fusion framework using target images augmented by PODA and ClipStyler on the Day Foggy adaptation. With PODA-augmented images, the framework achieves 37.5, while ClipStyler achieves 38.2, both of which are lower than TFA's 38.4. This demonstrates that the choice of style transfer technique has a measurable impact on adaptation performance, and that TFA provides the most effective domain-stylized augmentations for our framework.

Impact of fixed EMA rate. In this ablation, we investigate the effect of using fixed EMA rates instead of dynamically learned ones. Specifically, for domain experts (excluding the aggregator), we assign equal weights $\alpha^{\text{DE}} = \frac{1-\alpha^{\text{agg}}}{M-1}$ and vary the aggregator's EMA rate α^{agg} from 0.50 to 0.99. The resulting mAP values are shown in table 19. Although performance slightly improves with higher α^{agg}, the best result (38.2 mAP) still falls short of our full method with adaptive weighting (38.4 mAP). This demonstrates that simply assigning equal importance to non-aggregator models is suboptimal, and that our proposed learned weighting strategy meaningfully contributes to the final performance.

Evaluation on segmentation and classification tasks. Our framework can also be applied to classification and segmentation tasks. To further support the generality of the framework, we additionally conduct experiments on classification (DomainNet Leventidis et al. (2021), Table 20) and semantic segmentation (ACDC Sakaridis et al. (2021), Table 21), confirming that the same framework transfers well to other tasks without architectural changes. These experiments indicate that the proposed fusion mechanism is not restricted to the Faster R-CNN + ResNet-50 configuration and can be instantiated with the evaluated alternative architectures. They provide evidence of architectural robustness within the tested adaptation settings.

C.4 Qualitative Results

From Figure 4-Figure 8, we show some qualitative results by making use of the pre-trained image decoder from CLIPstyler Kwon & Ye (2022) of some stylized target images generated with TFA by sampling sampled from different domains with different text-based description styles described in Table 22. In general, the text-based style transfer successfully transfer the domain style from the

Table 18: Comparison of target-domain adaptation and inference efficiency. Adaptation time is measured from the pretrained source models to the final adapted detector and includes all method-specific preprocessing. GPU-hours account for the number of GPUs used. Inference FPS is measured under the same hardware, batch size, and image resolution.

Method	Adapt. Time (h)	GPU-hours	FPS	Params (M)
MSFDAOD	1.320	2.64	17	133
CAiDA	1.420	2.84	15	180
Mean Teacher	1.185	2.37	17	128
Ours	1.075	2.15	17	126

Table 19: Performance of using different EMA rate α.

α^{agg}	0.50	0.80	0.95	0.99
mAP	37.5	37.7	38.2	38.1

source domain to the target domain with the text descriptions. For example, in Figure 4, taking an image from Day Clear domain, the Night Clear and Night Rainy styles transfer the image to a dark night environment; the Dusk Rainy style imparts a pink dusk ambiance to the image; the Day Foggy style introduces fog into the image. In conclusion, the text-based style transfer technique is able to change the weather and time conditions given an image. In other instances, the City View and Urban View styles largely maintain their similarity in style descriptions, as they are closely related. Conversely, the Video Game style transforms a realistic image into a simulated one. Additionally, the random texts effectively incorporate relevant elements corresponding to the text descriptions.

In addition, we present some qualitative results on the evaluation of the aforementioned augmented images in Figure 9, with the target image sampled from Day Clear and the augmented images from Figure 4. As observed in the first row, when directly predicting the original image with different pre-trained source models, the domain gap between them tends to lead to the erroneous predictions. While in the second row, the pre-trained source models perform well on the corresponding augmented stylized images even they are not sampled from the corresponding domains. They all give perfect predictions except for Night Rainy, which has the largest domain gap with Day Clear. This show that the text based style transfer has reasonably reduced the domain gap between the source models and the target image.

C.5 Illustration of failed examples

In this section, we present some examples where the model struggles to detect objects in the target images, as shown in Figure 10. The results show that under extreme conditions, where the domain gap is large, or the classes are unseen in the source, the model's performance significantly degrades.

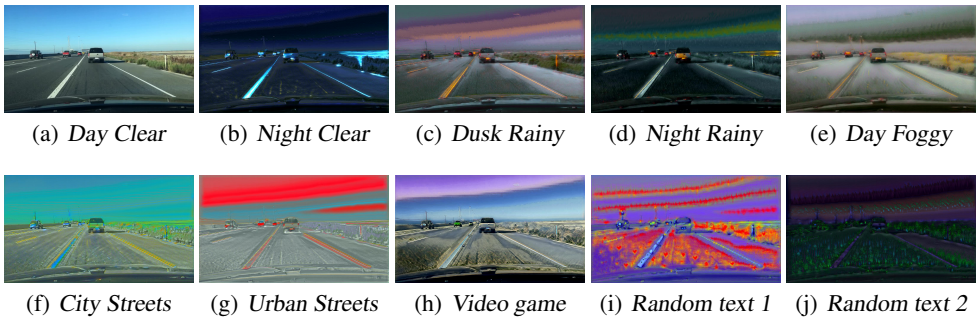

(a) *Day Clear* (b) *Night Clear* (c) *Dusk Rainy* (d) *Night Rainy* (e) *Day Foggy*

(f) *City Streets* (g) *Urban Streets* (h) *Video game* (i) *Random text 1* (j) *Random text 2*

Figure 4: Stylized images of different styles with image sampled from Day Clear domain of DWD dataset. (a) is the original target image, and (b)-(j) are text-based stylized images with corresponding description in Table 22.

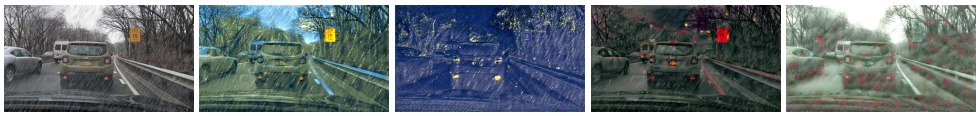

(a) *Dusk Rainy* (b) *Day Clear* (c) *Night Clear* (d) *Night Rainy* (e) *Day Foggy*



(f) *City Streets* (g) *Urban Streets* (h) *Video game* (i) *Random text 1* (j) *Random text 2*

Figure 5: Stylized images of different styles with an image sampled from Dusk Rainy domain of the DWD dataset. (a) is the original target image, and (b)-(j) are text-based stylized images with corresponding description in Table 22.



(a) *Night Rainy* (b) *Day Clear* (c) *Night Clear* (d) *Dusk Rainy* (e) *Day Foggy*



(f) *City Streets* (g) *Urban Streets* (h) *Video game* (i) *Random text 1* (j) *Random text 2*

Figure 6: Stylized images of different styles with an image sampled from Night Rainy domain of the DWD dataset. (a) is the original target image, and (b)-(j) are text-based stylized images with corresponding description in Table 22.

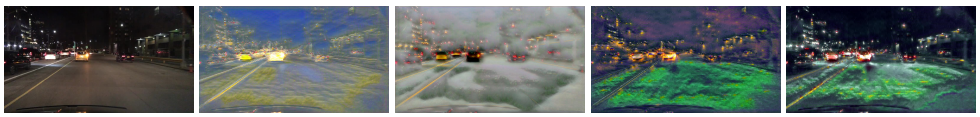

(a) *Night Clear* (b) *Day Clear* (c) *Day Foggy* (d) *Dusk Rainy* (e) *Night Rainy*

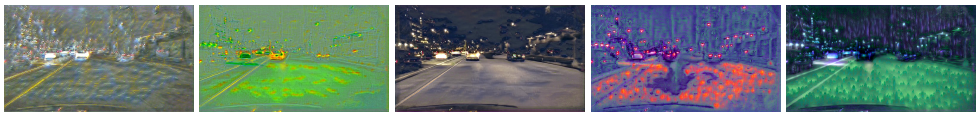

(f) *City Streets* (g) *Urban Streets* (h) *Video game* (i) *Random text 1* (j) *Random text 2*

Figure 7: Stylized images of different styles with an image sampled from Night Clear domain of the DWD dataset. (a) is the original target image, and (b)-(j) are text-based stylized images with corresponding description in Table 22.

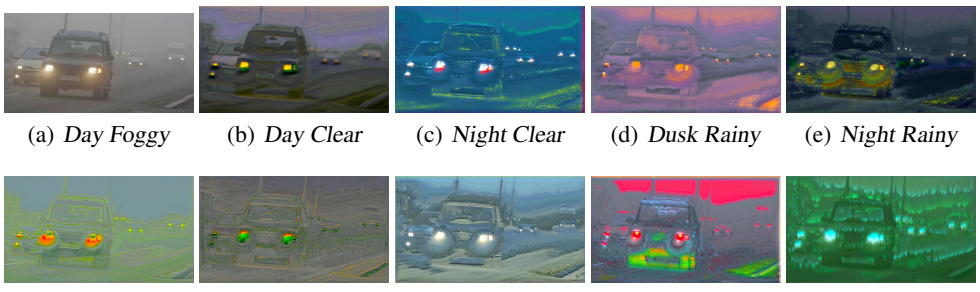

(a) *Day Foggy* (b) *Day Clear* (c) *Night Clear* (d) *Dusk Rainy* (e) *Night Rainy*

(f) *City Streets* (g) *Urban Streets* (h) *Video game* (i) *Random text 1* (j) *Random text 2*

Figure 8: Stylized images of different styles with an image sampled from Day Foggy domain of the DWD dataset. (a) is the original target image, and (b)-(j) are text-based stylized images with corresponding description in Table 22.

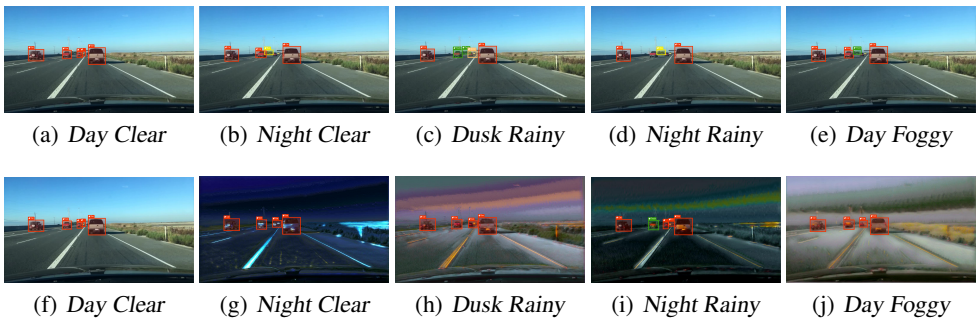

(a) *Day Clear* (b) *Night Clear* (c) *Dusk Rainy* (d) *Night Rainy* (e) *Day Foggy*

(f) *Day Clear* (g) *Night Clear* (h) *Dusk Rainy* (i) *Night Rainy* (j) *Day Foggy*

Figure 9: Prediction comparison on original and stylized images of different styles for each source model trained on Day Clear (a, f), Night Clear (b, g), Dusk Rainy (c, h), Night Rainy (d, i) and Day Foggy (e, j). (a)-(e) is the original target image, and (f)-(j) are text-based stylized images with corresponding description in Table 22. The colors mean different class labels (red: car, yellow: bus, green: truck).

Table 20: Performance across DomainNet domains.

Domain	AP	AP_{base}	AP_{novel}
Clipart	88.65	90.87	87.24
Infograph	86.84	87.57	86.82
Painting	84.27	85.53	84.75
Quickdraw	85.43	84.96	86.44
Real	83.62	84.58	82.66
Sketch	87.34	87.52	86.43

Table 21: Performance across BDD100K weather conditions.

Domain	AP	AP_{base}	AP_{novel}
Foggy	32.5	33.4	29.8
Nighttime	32.2	32.8	29.7
Rainy	29.3	30.7	26.5
Snowy	29.1	30.2	26.6

Table 22: Text Descriptions

Source Domain	Descriptions
Day Clear	["Driving at clear day time"]
Night Clear	["Driving at clear night time"]
Night Rainy	["Driving at rainy night time"]
Dusk Rainy	["Driving at rainy dusk time"]
Day Foggy	["Driving at foggy day time"]
City streets	["Driving at city streets"]
Urban street	["Driving at urban streets"]
Video game	["Driving in a video game"]
Random Text 1	["Robot falls in love"]
Random Text 2	["Lost in enchanted forest"]

C.6 Comparison of TFA with Existing Text-Based Style Transfer Methods

In this section, we present the augmented images generated with different text-based style transfer methods including CLIPstyler Kwon & Ye (2022), PODA Fahes et al. (2023) and our TFA using the same pre-trained CLIPstyler image decoder. In Figure 11, we present the detection results of different stylized images using the same model trained on Day Clear domain. The detection results of the two original images sampled from Day Clear domain serve as oracle results. Compared to CLIPstyler and PODA, TFA manage to maintain the details of the original images while rendering corresponding styles, such as foggy style and night style. The detection performance of TFA aligns closely with the oracle performance. While CLIPstyler misclassifies car as bus, bus as truck with Day Foggy style augmentation and fails to detection truck with Night Clear style augmentation. PODA renders a very strong style for the augmentation that even lose most of the details in the original content, which is somewhat reasonable since PODA only updates its style statistics with CLIP style loss without a content preservation regularization.

C.7 Comparison of Domain Adaptation Methods with Access to the Source Data

We compare our method with *Unsupervised Domain Adaptation (UDA)* methods which have access to the source data including SW Pan et al. (2019), IBN-Net Pan et al. (2018), IterNorm Huang et al. (2019), ISW Choi et al. (2021), PODA Fahes et al. (2023), CLIP-Aug Vidit et al. (2023), and S-DGOD Wu & Deng (2022). From Table 23, we observe that even access to the source data, our model outperforms the UDA methods with access to the source data in the Night Clear and Night Rainy domain. This is attributed to the incorporation of multi-source pre-trained models. We effectively utilize the knowledge from different pre-trained source models with our proposed mean-teacher framework and benefit the generalization to the target domain without needing the access to the source data. As for the Dusk Rainy and Day Foggy domains, our method achieves comparable performance

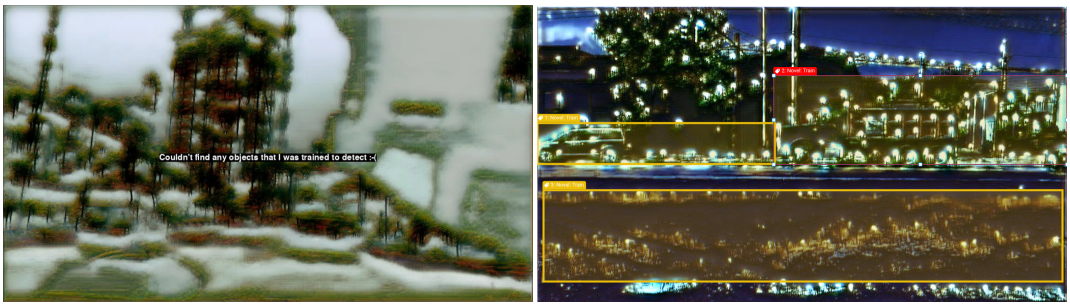

(a) *Failed to detect due to weather condition.* (b) *Fail to detect due to bad lighting.*

Figure 10: Example of failures.



(a) *Test image 1* (b) *TFA - day foggy* (c) *CLIPstyler - day foggy* (d) *PODA - day foggy*



(e) *Test image 2* (f) *TFA - day foggy* (g) *CLIPstyler - day foggy* (h) *PODA - day foggy*

Figure 11: Two images (a)/(e) from the Day Clear domain are stylized with different domain styles: Day Foggy and Night Clear, using different text-based style transfer methods: TFA (ours), CLIPstyler and PODA.

as PODA and CLIP-Aug, and outperforms the rest UDA methods, which helps demonstrate the effectiveness of our approach. This indicates that leveraging multi-source pre-trained models with our mean-teacher framework provides a strong advantage in domain adaptation, even without direct access to the source data. Our proposed text-based augmentation successfully reduces the domain gap between the target images and the pre-trained models. Our method generalizes well across diverse target domains, outperforming traditional UDA methods in challenging conditions such as Night Clear and Night Rainy while maintaining competitive results in other domains like Dusk Rainy and Day Foggy.

D Potential Social Impact and Limitations

Source-free adaptation can reduce the need to retain or transmit source-domain training data, which is useful when data sharing is restricted. However, pretrained models may still encode information from their training data, and source-free adaptation should therefore not be interpreted as providing a formal privacy guarantee.

Our method has four main limitations. First, it assumes that a coarse semantic description of each source domain is available. The description need not contain precise low-level details: our prompt ablation shows that different relevant phrasings yield comparable performance, whereas intentionally unrelated text serves as a negative control and degrades adaptation. In practice, short domain-level descriptions such as “rainy night”, “foggy day”, or “urban street” are sufficient for the evaluated settings.

Table 23: Multi-source domain adaptation results (mAP). For each target domain, Day Clear and the rest three domains are used as the source domains for the multi-source setting. For the single-source UDA and SFDA, Day Clear is used as the source following the typical setting Wu & Deng (2022); Vidit et al. (2023); Fahes et al. (2023).

		mAP			
Method	Source	NC	DR	NR	DF
SW	✓	33.4	26.3	13.7	30.8
IBN-Net	✓	32.1	26.1	14.3	29.6
IterNorm	✓	29.6	22.8	12.6	28.4
ISW	✓	33.2	25.9	14.1	31.8
S-DGOD	✓	36.6	28.2	16.6	33.5
CLIP-Aug	✓	36.9	18.7	18.7	38.5
PODA	✓	43.4	**40.2**	19.5	**44.4**
Ours	✗	**44.5**	32.5	**20.3**	38.4

Second, the method incurs a one-time target-adaptation cost because a source-stylized target set is constructed for each source expert and the experts are locally updated. Consequently, adaptation cost increases with the number of source domains. The stylized images are generated once and cached, and these operations are not required during deployment, so the final inference cost does not grow with the number of source experts used during adaptation.

Third, very large source-target gaps remain challenging. On the Night Rainy target, for example, our method achieves the best result among the compared MSFDA methods but the absolute detection accuracy remains substantially lower than on the other DWD domains. This suggests that text-driven style alignment can mitigate, but cannot fully remove, information loss caused by severe illumination and adverse weather.

Finally, we evaluate under the standard MSFDA assumption that unlabeled samples from the target domain are available during adaptation. We do not evaluate generalization to a completely unseen/OOD target distribution without target adaptation; that setting corresponds more closely to domain generalization or open-world adaptation and is outside the scope of the present formulation. Future work could combine uncertainty-aware source selection, localization-aware pseudo-label quality estimation, and more scalable expert aggregation to address these limitations.

E Source Code

For source code, please refer to `https://anonymous.4open.science/r/sfda_aug-ADB0/Readme.txt`.