

Aligning Non-Causal Factors for Transformer-Based Source-Free Domain Adaptation

Sunandini Sanyal* Ashish Ramayee Asokan* Suvaansh Bhambri Pradyumna YM
 Akshay Kulkarni Jogendra Nath Kundu R Venkatesh Babu
 Vision and AI Lab, Indian Institute of Science, Bengaluru

Abstract

Conventional domain adaptation algorithms aim to achieve better generalization by aligning only the task-discriminative causal factors between a source and target domain. However, we find that retaining the spurious correlation between causal and non-causal factors plays a vital role in bridging the domain gap and improving target adaptation. Therefore, we propose to build a framework that disentangles and supports causal factor alignment by aligning the non-causal factors first. We also investigate and find that the strong shape bias of vision transformers, coupled with its multi-head attentions, make it a suitable architecture for realizing our proposed disentanglement. Hence, we propose to build a Causality-enforcing Source-Free Transformer framework (C-SFTrans¹) to achieve disentanglement via a novel two-stage alignment approach: a) non-causal factor alignment: non-causal factors are aligned using a style classification task which leads to an overall global alignment, b) task-discriminative causal factor alignment: causal factors are aligned via target adaptation. We are the first to investigate the role of vision transformers (ViTs) in a privacy-preserving source-free setting. Our approach achieves state-of-the-art results in several DA benchmarks.

1. Introduction

Machine learning models often fail to generalize well in scenarios where the test data distribution (*source domain*) differs a lot from the training data distribution (*target domain*). In practice, a model often encounters data from unseen domains *i.e.* *domain shift*. This leads to a poor deployment performance, which critically impacts many real-world applications such as autonomous driving [6], surveillance systems [42], etc.

Unsupervised domain adaptation (UDA) methods [9]

*Equal Contribution

¹Project Page: <https://val.cds.iisc.ac.in/C-SFTrans/>




Figure 1. **A. Proposed method.** We incorporate non-causal factors to learn domain-invariant representations using a subsidiary non-causal factor classification task. **B. w/o Non-Causal Factor Alignment.** Conventional domain-invariance methods aim to align only causal factors, leading to sub-optimal alignment between the source and target domain. **C. w/ Non-Causal Factor Alignment.** Non-causal factor alignment improves global alignment, leading to better task-discriminative causal factor alignment.

aim to address the challenges of domain shift by learning the task knowledge of the labeled source domain and adapting to an unlabeled target domain. However, these works [10] require joint access to the source and target data. Such a constraint is highly impractical as data sharing is usually restricted in most real-life applications due to privacy concerns. Hence, in this work, we focus on the practical problem setting of source-free domain-adaptation (SFDA) [19] where a *vendor* trains a source model and shares only the source model with a *client* for target adaptation.

Conventional DA works [10] aim to learn domain-invariant representations by aligning only the task-related features between the source and target domain. We refer to these features as *causal factors* that heavily influence the goal task. We also denote factors that capture contextual information as *non-causal factors*. Causal factor alignment leads to a low target error, thereby improving the adaptation performance, as shown theoretically by Ben-David et al.

[1]. But these methods require concurrent access to the source and target domain data, which is impractical in restricted data sharing scenarios of SFDA. Further, causal and non-causal factors are spuriously correlated [47] and this correlation may break when a domain shift occurs. Hence, in our work, we propose to retain this spurious correlation through disentanglement and learning of both causal and non-causal factors.

Motivated by the remarkable success of vision transformer architectures [14], we propose to explore the possibility of disentanglement and alignment of causal/non-causal factors using transformers. Recent domain-invariance based SFDA works [17, 50] have been found to be highly effective on convolution-based architectures. However, in our analysis, we find that a simple vision transformer (ViT) baseline outperforms the state-of-the-art CNN-based methods, implying that transformers are highly robust to domain shifts [28]. Secondly, we observe that the domain-invariance methods do not significantly impact the performance of vision transformers due to their inherent shape bias [30]. Based on these observations, we propose to leverage ViTs for a realizable disentanglement of causal and non-causal factors. Hence, in our work, we seek an answer to an important question, “*How do we develop a framework for disentanglement using vision transformers to retain the spurious correlation between causal and non-causal factors, in the challenging source-free setting?*”.

Conventional approaches [10, 26, 50] completely ignore the non-causal factors and align only the causal factors, which leads to sub-optimal alignment of the class-specific clusters as shown in Fig. 1B. This results in poor performance, especially in cases of large domain gap between the source and target. Hence, we propose to first align the non-causal factors, which leads to an overall global alignment between the source and target domain. We then align the goal task-discriminative causal factors that optimally aligns the local class-specific clusters (Fig. 1C). Hence, in our work, we pinpoint that “*aligning the non-causal factors is crucial for improving target adaptation performance*”.

We next seek to devise a framework that explicitly guides the process of disentanglement and alignment. We develop a novel framework - *Causality-enforcing Source-Free Transformers (C-SFTrans)* that comprises of two stages, **a)** Non-causal factor alignment, and **b)** Task-discriminative causal factor alignment. To enable non-causal factor alignment, we propose a subsidiary task of non-causal factor classification for a source-free setting. Prior works [5] show that multi-head self-attentions in transformers focus on redundant features. To facilitate the separate learning of diverse non-causal factors and causal factors, we utilize this inherent potential of transformers to designate non-causal and causal attention heads for training. We propose a novel *Causal Influence Score Criterion* to perform the head se-

lection. The two stages of non-causal factor alignment and task-discriminative causal factor alignment are performed alternatively to achieve domain-invariance. We outline the major contributions of our work:

- To the best of our knowledge, we are the first to explore vision transformers (ViTs) for a practical source-free DA setting. We investigate domain-invariance in vision transformers and provide insights to improve target adaptation via non-causal factor alignment.
- We propose a novel two-stage disentanglement and alignment framework *Causality-enforcing Source-Free Transformers (C-SFTrans)* to preserve the spurious correlation between causal and non-causal factors.
- We define a novel attention head selection criterion, *Causal Influence Score Criterion*, to select attention heads for non-causal factor alignment. We also introduce a novel Style Characterizing Input (SCI) to further aid the head selection.
- We achieve state-of-the-art results on several source-free benchmarks of single-source and multi-target DA.

2. Related works

Unsupervised Domain Adaptation. Unsupervised Domain Adaptation (DA) aims to adapt a source-trained model to a given target domain. DA approaches can be broadly classified into two categories: 1) methods using generative models [23, 39, 44] to create synthetic target-like images for adaptation. 2) methods focusing on aligning the source and target feature distributions using statistical distance measures on the source/target features [32, 40, 43, 56], and adversarial training [12, 13, 25]. Recent works [16, 20, 35, 38] address a more restrictive and privacy-preserving setting of *Source-Free Domain Adaptation (SFDA)*, where the source data is inaccessible during target adaptation. SFDA works of SHOT [21] and SHOT++ [22] use pseudo-labeling and information maximization to align the source and target domains. Our work also addresses the challenging SFDA setting, intending to improve the adaptation performance.

Transformers for Domain Adaptation. Despite the success of Vision Transformers (ViTs) [7] in several vision tasks, their application to DA has been relatively less explored. TransDA [54] improves the model generalizability by incorporating the transformer’s attention module in a convolutional network. CDTrans [53] proposes a transformer framework comprising three weight-sharing branches for cross-attention and self-attention using the source and target samples, while SSRT [46] introduces a self-training strategy that uses perturbed versions of the target samples to refine the ViT backbone. TVT [55] improves the transferability of ViTs through adversarial training.




Figure 2. **A.** SOTA domain-invariance-based DA works, Feature-Mixup [17] and DIPE [50], do not improve over the baseline for vision transformers despite large gains for CNNs. **B.** We observe that correlation S and Z in preserved after target adaptation. (Office-Home) **C.** With our proposed style task training, the overall source-target domain gap (pink) is lower indicating better domain-invariance. Further, we observe a lower domain gap considering only causal factors (yellow) indicating that non-causal alignment helps causal alignment. **D.** Causal Graph representing causal factors S and non-causal factors Z , which are spuriously correlated via confounder C .

Domain Invariance for Domain Adaptation. These methods aim to learn domain-invariant feature representations between the source and target domains. SHOT [21] and SHOT++ [22] prevent updates to the source hypothesis, which enables the feature extractor to learn domain-invariant representations. Feature-Mixup [17] constructs an intermediate domain whose representations preserve the task discriminative information while being domain-invariant. In contrast, DIPE [50] trains the domain-invariant parameters of the source-trained model rather than learning domain-invariant representations between the domains.

Causal representation learning. Causality mechanisms [45] focus on learning invariant representations and recovering causal features [26, 51] that improve the model’s generalizability. Some works [4, 15, 58] attempt this via texture-invariant representation learning. However, such works are less effective as they align only the causal factors towards improving generalization. In contrast, we propose a novel way of learning domain-invariant representations by taking into account both the non-causal and causal factors in the target domain to achieve the best adaptation performance.

3. Approach

3.1. Preliminaries

Problem Setting We consider the problem setting of closed-set DA, with a labeled source domain dataset $\mathcal{D}_s = \{(x_s, y_s) : x_s \in \mathcal{X}, y_s \in \mathcal{C}_g\}$ where \mathcal{X} is the input space and \mathcal{C}_g is the class label set. The unlabeled target dataset is denoted as $\mathcal{D}_t = \{x_t : x_t \in \mathcal{X}\}$. The task of DA is to predict the label for each target sample x_t from the label space \mathcal{C}_g . Following Xu et al. [53], we use a vision transformer backbone ViT-B [7] as the feature extractor, denoted as $f : \mathcal{X} \rightarrow \{\mathcal{Z}_c, \mathcal{Z}_1, \mathcal{Z}_2, \dots, \mathcal{Z}_{N_P}\}$. \mathcal{Z}_c denotes the class token feature-space and $\mathcal{Z}_1, \mathcal{Z}_2, \dots, \mathcal{Z}_{N_P}$ are patch token feature-spaces. N_P denotes the number

of patches. We train a goal task classifier f_g on the class-token as $f_g : \mathcal{Z}_c \rightarrow \mathcal{C}_g$. In this work, we operate under the vendor-client paradigm of source-free domain adaptation [19] where a vendor trains a source model on the labeled source domain dataset and shares the model with the client. The client, on the other hand, trains the model with the unlabeled target domain data for target adaptation.

Causal Model. Let X denote the input variable and Y denote the output variable or label. We represent the structural causal graph G as shown in Fig. 2D where we introduce latent variables S and Z to capture the generative concepts (e.g. object shape, backgrounds, textures) that lead to the observed variables X and Y . We explicitly separate causal factors S that causally influence the class-label Y (e.g. shape) and non-causal factors Z which refers to contextual information (like background, texture, etc.). We also assume that S is spuriously correlated with Z shown through an unobserved confounder C . This spurious correlation may vary across domains. For instance, in a source-free DA setting, the source domain might have a “ball” in a “football field”, while the target domain may have a “ball” in “table tennis”. For simplicity, we use Fig. 2(b) of Sun *et al.* [47] where domain shifts are represented as changes in the probability distributions $P(C)$ or $P(S, Z|C)$.

Analyzing Causality in Source-Free DA. Since S and Z are correlated spuriously, a source model trained on source domain data X inherits this spurious correlation in the form of bias (e.g. contextual information of objects and scenes occurring together in the source domain). In the target domain, this spurious correlation changes, leading to worsened performance of the source model. In our work, we aim to retain the spurious correlation between S and Z in the source-free DA setting when domain shift occurs. We construct an explicitly disentangled network to model the causal factors S and the non-causal factors Z through separate learning objectives. When the domain changes, we

leverage the disentanglement and our designed training objectives to retain the $P(S, Z|C)$ correlation.

We begin by first investigating the impact of existing CNN-based DA methods on vision transformers (ViTs) and draw insights towards achieving the desired disentanglement.

3.2. Exploring domain-invariance for ViTs

Motivated by the impressive performance of vision transformers (ViT) across several vision tasks [14], we propose to investigate the impact of domain-invariance methods on vision transformers in the highly practical and privacy-oriented source-free setting. A few recent works [17, 50] propose to learn domain-invariant representations without access to source data. However, we find that these approaches mainly work only with CNN architectures. We first extend two state-of-the-art domain-invariance works, Feature Mixup [17] and DIPE [50], for transformers (see Fig. 2A). We observe that these works exhibit marginal drops compared to a baseline SHOT [21] for transformers, although they have significant gains for CNNs.

But why does this happen? Recent works [36] show that ViTs incorporate more global information than CNNs. This is because the multi-head self-attention captures causal factors via shape bias [30], while convolutions capture non-causal factors via texture bias. As convolutions inherently capture texture bias, domain-invariance methods perform well as they align causal factors and improve the shape bias in CNNs [17]. Since transformers implicitly possess a stronger shape-bias, it is more robust to domain shifts [28] and existing methods do not improve the adaptation performance significantly, unlike in CNNs. Based on these observations, we hypothesize that ViTs can better model the causal factors and can thus accommodate disentanglement of non-causal factors, which are easier to learn. However, the question remains, “*How can we leverage such a disentanglement in ViTs to retain the spurious correlation between causal and non-causal factors, while operating under the challenging yet practical source-free DA constraints?*”

Next, we propose a method for disentanglement in ViTs for modeling the causal and non-causal factors.

3.3. Non-causal factor alignment for source-free DA

Conventional approaches ignore non-causal factors while aligning only the causal factors. We know that the spurious correlation between causal and non-causal factors can heavily influence the classification performance [47], especially in scenarios with a large source-target domain gap. Hence, disentangling and aligning non-causal factors can help bridge the domain gap to a large extent, especially for ViT architectures with a strong shape bias. Therefore, we come up with the following insight of effectively utilizing the non-causal factors to enable better alignment of causal factors in the challenging source-free DA setting.

Insight 1. (Non-causal factor alignment positively influences causal factor alignment) *Aligning non-causal factors leads to an overall global alignment between the source and target domain for source-free DA settings. In other words, it implicitly improves the alignment of causal factors as well, thereby improving the adaptation performance.*

Remarks. Non-causal factor alignment forces the model to focus on the non-causal factors. Causal factors modeled through the inherent shape bias of transformers lead to substantial gains over CNNs (as shown in Fig. 2A). However, aligning the residual non-causal factors can further improve the overall global alignment which helps to preserve the correlation between causal and non-causal factors in the target domain. We demonstrate this phenomenon in Fig. 2C (pink bars), where our non-causal factor alignment improves the overall global alignment between the two domains, leading to a significant reduction in the domain-gap between the source and target domain. This shows that non-causal factor alignment positively influences causal factor alignment as the spurious correlation between causal and non-causal factors is preserved after target adaptation (Fig. 2B).

As Insight 1 motivates that aligning non-causal factors is extremely crucial for preserving the spurious correlation between causal and non-causal factors, a natural question arises, “*How do we enable alignment of non-causal factors between a source and target domain in a SFDA setting?*”

Insight 2. (Style clsf. for non-causal factor extraction) *Stylization can facilitate controllable access to the local non-causal factors. Thus, non-causal factors can be aligned using a subsidiary task of style classification, while respecting the source-free constraints.*

Remarks. To enable non-causal factor alignment, we propose a subsidiary task of style classification on both the vendor and client-side in a source-free setting. We make use of label-preserving augmentations [34] (see Suppl. for details) to construct novel styles and train a style classifier f_n on a novel style token $z_n \in \mathcal{Z}_n$ in the transformer backbone f . Through the subsidiary task of style classification, we propose to extract the non-causal factors and project the local features of the source and target domain into a common feature space. This implicitly aligns the two domains, even without concurrent access to source and target data. Further, it can be easily enabled in a practical source-free setting by sharing only the augmentation information between the vendor and the client, without sharing the data.

Analysis Experiments. To analyze the effectiveness of our proposed insights, we examine the effect of non-causal factor alignment on the causal factors. In Fig. 2B, we observe that the \mathcal{A} -distance [2] between the causal and non-causal factors remain almost same in both source and target domains, indicating that the spurious correlation between S and Z is preserved with our approach. For (Fig. 2C), we construct a domain-invariant feature (as a proxy

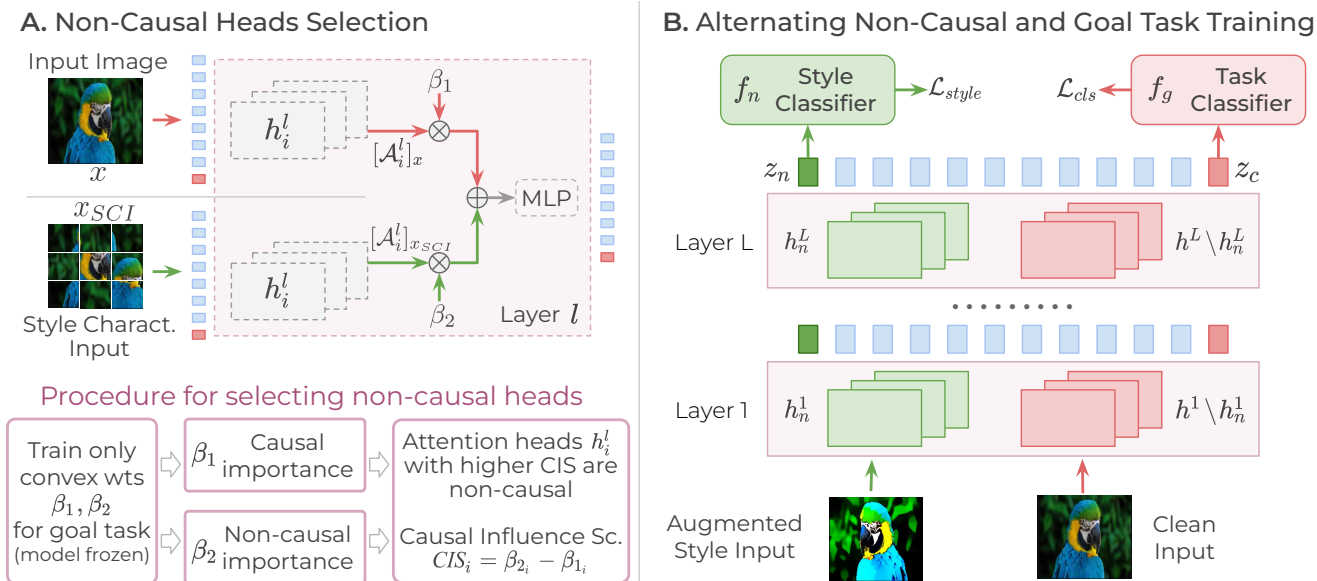


Figure 3. **A.** We estimate the causal importance of each attention head in a block by training convex weights β_1 and β_2 that combine the output of the attention heads for normal input x and style characterizing input x_{SCI} . A higher β_2 indicates that the attention head gave more importance to non-causal style factors and can be chosen as a non-causal head. **B.** The selected non-causal heads h_n are updated to train the style classifier f_n with the style classification loss \mathcal{L}_{style} via a style token z_n . The task classifier is trained by updating only causal heads $\mathcal{H} \setminus h_n$ with the task classification loss \mathcal{L}_{cls} via the class token z_c . The two steps are executed alternately.

for the causal factors) by taking the mean of class tokens across augmentations. Between the mean class tokens from source and target domains, we compute the \mathcal{A} -distance [2], which measures the separation between the two distributions/domains. We observe a lower \mathcal{A} -distance at the class token for our method as compared to the baseline SHOT, indicating that non-causal factor alignment leads to improved causal factor alignment (Insight 1) and global alignment between the source and target domains (Insight 2).

3.4. Training Algorithm

We propose a Causality-enforcing Source-Free Transformer (C-SFTrans) framework which involves a two-stage feature alignment for learning causal representations using vision transformers.

3.4.1 Vendor-side source training

The vendor trains C-SFTrans on the source dataset in two steps: (a) *Non-causal factor alignment* using a style classification task, and (b) *Goal task-discriminative feature alignment*, each of which are discussed in detail below.

a) Non-causal factor alignment.

To align the non-causal factors, we introduce a style classification task that is trained with a novel style token $z_n \in \mathcal{Z}_n$ and updates only non-causal attention heads $h_n^l \in \mathcal{H}$. Here, $\mathcal{H} = \{h_i^l\}_{i=1, l=1}^{N_h, L}$ denotes the set of self-attention heads across all blocks, N_h denotes the number of self-attention

heads in each block of the backbone f , and L denotes the number of blocks. Each head h_i^l computes self-attention as $h_i^l = \text{Softmax}(\alpha K Q^T) V$. Here, $K = z_j^l W_K$, $Q = z_j^l \cdot W_Q$ and $V = z_j^l W_V$, $\alpha = \frac{1}{\sqrt{d_k}}$, and d_k is the dimension of z_j^l . Let $\mathcal{A}_{j_x}^l = h_i^l(\{z_c^l, z_1^l, \dots, z_{N_P}^l\}_x)$ where x is the input sample that gets divided into N_P patch tokens and $\mathcal{A}_{j_x}^l$ denotes the output of the self-attention head. For the non-causal factor alignment which will update only non-causal attention heads, we first need to select these non-causal attention heads. We choose the heads that give higher importance to the non-causal style features than the causal task-discriminative features. Next, we outline the procedure for non-causal attention heads selection.

Non-causal attention heads selection. We propose a novel selection criteria to select attention heads based on its contribution towards the causal goal task features and the non-causal style-related features. For this, we first construct a novel *Style Characterizing Input (SCI)* to preserve only style-related features of the input samples. To construct an SCI, we apply a task-destructive transformation (patch shuffling) [27], which keeps the style information intact. Intuitively, we preserve the higher order statistics of style [3] by shuffling patches while the class information is destroyed. We pass both the clean input x and the SCI x_{SCI} as input to each head (Fig. 3A). Let $\mathcal{A}_{i_x}^l$ be the output of each head computed using the clean input x computed as follows:

$$\mathcal{A}_{i_x}^l = h_i^l(\{z_c^l, z_1^l, \dots, z_{N_P}^l\}_x) \quad (1)$$

Table 1. Single-Source Domain Adaptation (SSDA) on Office-Home benchmark. SF denotes *source-free* adaptation. ResNet-based methods (top) and Transformer-based methods (bottom). * indicates results taken from [53].

Method	SF	Office-Home												
		Ar→Cl	Ar→Pr	Ar→Rw	Cl→Ar	Cl→Pr	Cl→Rw	Pr→Ar	Pr→Cl	Pr→Rw	Rw→Ar	Rw→Cl	Rw→Pr	Avg
ResNet-50 [11]	✗	34.9	50.0	58.0	37.4	41.9	46.2	38.5	31.2	60.4	53.9	41.2	59.9	46.1
A ² Net [52]	✓	58.4	79.0	82.4	67.5	79.3	78.9	68.0	56.2	82.9	74.1	60.5	85.0	72.8
GSFDA [57]	✓	57.9	78.6	81.0	66.7	77.2	77.2	65.6	56.0	82.2	72.0	57.8	83.4	71.3
NRC [56]	✓	57.7	80.3	82.0	68.1	79.8	78.6	65.3	56.4	83.0	71.0	58.6	85.6	72.2
SHOT [21]	✓	57.1	78.1	81.5	68.0	78.2	78.1	67.4	54.9	82.2	73.3	58.8	84.3	71.8
SHOT++ [22]	✓	57.9	79.7	82.5	68.5	79.6	79.3	68.5	57.0	83.0	73.7	60.7	84.9	73.0
TVT [55]	✗	74.8	86.8	89.4	82.7	87.9	88.2	79.8	71.9	90.1	85.4	74.6	90.5	83.5
SSRT-B [46]	✗	75.1	88.9	91.0	85.1	88.2	89.9	85.0	74.2	91.2	85.7	78.5	91.7	85.4
CDTrans [53]	✗	68.8	85.0	86.9	81.5	87.1	87.3	79.6	63.3	88.2	82.0	66.0	90.6	80.5
SHOT-B*	✓	67.1	83.5	85.5	76.6	83.4	83.7	76.3	65.3	85.3	80.4	66.7	83.4	78.1 (+2.5)
DIPE [50]	✓	66.0	80.6	85.6	77.1	83.5	83.4	75.3	63.3	85.1	81.6	67.7	89.6	78.2 (+2.4)
Mixup [17]	✓	65.3	82.1	86.5	77.3	81.7	82.4	77.1	65.7	84.6	81.2	70.1	88.3	78.5 (+2.1)
C-SFTrans (Ours)	✓	70.3	83.9	87.3	80.2	86.9	86.1	78.9	65.0	87.7	82.6	67.9	90.2	80.6

Let $\mathcal{A}_{i_{x_{SCI}}}^l$ be the output of the heads computed using the input x_{SCI} as follows,

$$\mathcal{A}_{i_{x_{SCI}}}^l = h_i^l(\{\mathcal{Z}_c^l, \mathcal{Z}_1^l, \dots, \mathcal{Z}_{N_P}^l\}_{x_{SCI}}) \quad (2)$$

Let $\beta_1, \beta_2 \in \mathbb{R}^{N_h}$ be the importance weights for the domain and task feature outputs respectively. We constrain these using $\beta_2 = 1 - \beta_1$, which simplifies the optimization. We compute the weighted output as follows,

$$\mathcal{A}_i^l = \beta_{1_i} \times \mathcal{A}_{i_x}^l + \beta_{2_i} \times \mathcal{A}_{i_{x_{SCI}}}^l \quad (3)$$

This weighted output \mathcal{A}_i^l is propagated further across the layers. We keep the entire model frozen, training only the two parameters β_1 and β_2 for each attention head (Fig. 3A). The following objective is used for attention heads selection,

$$\min_{\beta_1, \beta_2} \mathbb{E}_{(x_s, y_s) \in \mathcal{D}_s} [\mathcal{L}_{cls}] \text{ where } \mathcal{L}_{cls} = \mathcal{L}_{ce}(f_g(z_c), y_c) \quad (4)$$

See Suppl. for more details. Next, we define the criterion for selecting non-causal heads using the optimized β_1, β_2 .

Definition 1. (Causal Influence Score Criterion) *Causal Influence Score (CIS) is computed as $CIS_i = \beta_{2_i} - \beta_{1_i}$ for an attention head h_i . We choose each attention head h_i as non-causal for which the condition $CIS_i > \tau$ is satisfied. The remaining heads are designated as causal heads.*

Remarks. Since we train β_1 and β_2 with the task classification objective, a higher value of β_2 indicates that the attention head gives more weightage to the style information and is inherently more suitable for the style classification task. Empirically, we set 30% heads as non-causal for our experiments, keeping the remaining heads fixed as causal heads for the goal task classification task. While a more sophisticated block-wise strategy could be used, we find this to be insensitive over a wide range of values (see Suppl).

Style classification task. Once the causal and non-causal attention heads are chosen, the vendor prepares the augmented datasets $\mathcal{D}_s^{(i)} = \{(x_s^{[i]}, y_s, y_n)\} \forall i \in [N_a]$ by augmenting each source sample x_s (where N_a is the number of augmentations). Here, an augmentation $a_i : \mathcal{X} \rightarrow \mathcal{X}$ is applied to get $x_s^{[i]} = a_i(x_s)$. Each input is assigned a style label $y_n = i$ where i denotes the augmentation label. We use five label-preserving augmentations that simulate novel styles [18]. Refer to Suppl. for more details.

For the non-causal factor learning task (Fig. 3B), we train only the non-causal heads for style classification as follows,

$$\min_{\theta_{h_n}, \theta_{f_n}} \mathbb{E}_{(x, y_s) \in \cup_i \mathcal{D}_s^{(i)}} [\mathcal{L}_{style}] \text{ where } \mathcal{L}_{style} = \mathcal{L}_{ce}(f_n(z_n), y_n) \quad (5)$$

where θ_{h_n} contains non-causal heads selected using Def. 1.

b) Task-discriminative causal factor alignment.

After one round of style classifier training, we perform the goal task training, where we update only the causal heads (Fig. 3B). The vendor trains the source model consisting of the backbone f and the task classifier f_g with the source labeled dataset \mathcal{D}_s and the task classification loss as follows,

$$\min_{\theta_f \setminus \theta_{h_n}, \theta_{f_g}} \mathbb{E}_{(x_s, y_s) \in \mathcal{D}_s} [\mathcal{L}_{cls}] \text{ where } \mathcal{L}_{cls} = \mathcal{L}_{ce}(f_g(z_c), y_c) \quad (6)$$

where z_c is the class token, $\theta_f \setminus \theta_{h_n}$ are the parameters of the backbone excluding the parameters of non-causal heads and θ_{f_g} are the parameters of task classifier f_g .

The two steps of non-causal factor alignment and goal task-discriminative feature alignment are performed in alternate iterations, one after the other (see Suppl. for details).

Table 2. Multi-Target Domain Adaptation (MTDA) on Office-Home. SF denotes *source-free* adaptation.

Method	SF	Office-Home				
		Ar→	Cl→	Pr→	Rw→	Avg.
MT-MTDA [29]	✗	64.6	66.4	59.2	67.1	64.3
CDAN+DCL [24]	✗	63.0	66.3	60.0	67.0	64.1
D-CGCT [37]	✗	70.5	71.6	66.0	71.2	69.8
D-CGCT-B [37]	✗	77.0	78.5	77.9	80.9	78.6
SHOT-B*	✓	75.4	79.3	73.6	77.1	76.4
C-SFTrans-B (Ours)	✓	77.3	82.9	74.4	76.9	77.8 (+1.4)

Table 3. Single-Source Domain Adaptation (SSDA) on Office-31 and VisDA benchmarks. SF denotes *source-free* adaptation. ResNet-based methods (top) and Transformer-based methods (bottom). * indicates results taken from [53].

Method	SF	Office-31	VisDA
ResNet-50 [11]	✗	76.1	52.4
NRC [56]	✓	89.4	85.9
SHOT [21]	✓	88.6	82.9
SHOT++ [22]	✓	89.2	87.3
TVT [55]	✗	93.8	83.9
CGDM-B* [8]	✗	91.2	82.3
CD-Trans [53]	✗	92.6	88.4
SSRT-B [46]	✗	93.5	88.7
SHOT-B*	✓	91.4 (+0.9)	85.9 (+2.4)
DIPE [50]	✓	90.5 (+1.8)	82.8 (+5.3)
Mixup [17]	✓	91.7 (+0.6)	86.3 (+2.0)
C-SFTrans (Ours)	✓	92.3	88.3

3.4.2 Client-side target adaptation

The vendor shares the trained C-SFTrans model with the client for target adaptation. The client performs the non-causal factor alignment in the same way as described earlier by augmenting the target data for style classification. Note that vendor and client may share training and augmentation strategies without sharing the training data [17]. This step leads to the non-causal factor alignment between source and target domains. For the goal task-discriminative training, the client uses the standard information maximization loss [21] as follows,

$$\min_{\theta_f \setminus \theta_{h_n}, \theta_{f_g}} \mathbb{E}_{\mathcal{D}_t} [\mathcal{L}_{ent} + \mathcal{L}_{div}] + \min_{\theta_{h_n}, \theta_{f_n} \cup_i \mathcal{D}_t^{(i)}} \mathbb{E} [\mathcal{L}_{style}] \quad (7)$$

where \mathcal{L}_{ent} , \mathcal{L}_{div} denote entropy and diversity losses, respectively and \mathcal{L}_{style} is supervised CE loss. Note that only original unlabeled target data is used to optimize \mathcal{L}_{ent} , \mathcal{L}_{div} . The two steps of non-causal factor alignment and task-discriminative causal factor alignment are done one after the other on the client-side as well.

4. Experiments

In this section, we evaluate our proposed approach by comparing with existing works on several benchmarks and analyze the significance of each component of the approach.

Datasets. We evaluate our approach on four existing standard object classification benchmarks for Domain Adaptation: OfficeHome, Office-31, VisDA, and DomainNet. The **Office-Home** dataset [49] contains images from 65 categories of objects found in everyday home and office environments. The images are grouped into four domains - Art (**Ar**), Clipart (**Cl**), Product (**Pr**) and Real World (**Rw**). The **Office-31 (Office)** dataset [41] consists of images from three domains - Amazon (**A**), DSLR (**D**), and Webcam (**W**). The three domains contain images from 31 classes of objects that are found in a typical office setting. The **VisDA** [33] dataset is a large-scale synthetic-to-real benchmark with images from 12 categories. **DomainNet** [31] is the largest and the most challenging dataset among the four standard benchmarks due to severe class imbalance and diversity of domains. It contains 345 categories of objects from six domains - Clipart (**clp**), Infograph (**inf**), Painting (**pnt**), Quickdraw (**qdr**), Real (**rel**), Sketch (**skt**).

Implementation details. To ensure fair comparisons, we make use of DeiT-Base [48] with patch size 16×16 and follow the experimental setup outlined in CDTrans [53]. We use Stochastic Gradient Descent (SGD) with a weight decay ratio of 1×10^{-4} , and a momentum of 0.9 for the training process. Refer to Suppl. for more implementation details.

Table 5. Ablation study for various stages of training on VisDA Single-Source Domain Adaptation (SSDA) benchmark.

Training Phase	Method	Goal Task	Style Task	Avg.
Source-Side	Source-Only	✓	✗	65.1
	Ours	✗	✓	66.2 (+1.1)
		✓	✓	68.8 (+3.7)
Target-Side	SHOT-B	✓	✗	85.9
	Ours	✗	✓	87.5 (+1.6)
		✓	✓	88.3 (+2.4)

4.1. Comparison with prior arts

a) Single-Source Domain Adaptation (SSDA). We provide comparisons between our proposed method, C-SFTrans, and earlier SSDA works in Tables 1 and 3. Our method provides the best performance among source-free works for the three standard DA benchmarks. On Office-Home, C-SFTrans outperforms the transformer based source-free prior work SHOT-B* by 2.5% and shows competitive performance w.r.t. the non-source-free method CD-Trans [53]. On the Office-31 benchmark (Table 3), our technique outperforms the source-free SHOT-B* by 0.9% and achieves competitive performance when compared to non-source-free works. Table 3 also demonstrates that our method shows 2.4% improvement over SHOT-B* and is on par with the non-source-free methods CDTrans and SSRT [46] on the larger and more challenging VisDA benchmark. We also achieve a significant improvement of 3.6% on the DomainNet benchmark (Table 4) over SHOT-B baseline.

Table 4. Single-Source Domain Adaptation (SSDA) on the DomainNet benchmark. * indicates results taken from [46].

MDD+SCDA [59]	clp	inf	pnt	qdr	rel	skt	Avg.	DeiT-B [48]	clp	inf	pnt	qdr	rel	skt	Avg.	SHOT-B [21]	clp	inf	pnt	qdr	rel	skt	Avg.
clp	-	20.4	43.3	15.2	59.3	46.5	36.9	clp	-	24.3	49.6	15.8	65.3	52.1	41.4	clp	-	45.9	49.7	16.5	65.4	53.2	46.1
inf	32.7	-	34.5	6.3	47.6	29.2	30.1	inf	45.9	-	45.9	6.7	61.4	39.5	39.9	inf	46.4	-	45.9	7.4	60.6	40.1	40.1
pnt	46.4	19.9	-	8.1	58.8	42.9	35.2	pnt	53.2	23.8	-	6.5	66.4	44.7	38.9	pnt	54.6	25.7	-	8.1	66.3	49.0	40.7
qdr	31.1	6.6	18.0	-	28.8	22.0	21.3	qdr	31.9	6.8	15.4	-	23.4	20.6	19.6	qdr	33.3	6.8	15.5	-	23.8	24.0	20.7
rel	55.5	23.7	52.9	9.5	-	45.2	37.4	rel	59.0	25.8	56.3	9.16	-	44.8	39.0	rel	59.3	28.1	57.4	9.0	-	47.3	40.2
skt	55.8	20.1	46.5	15.0	56.7	-	38.8	skt	60.6	20.6	48.4	16.5	61.2	-	41.5	skt	64.0	26.5	55.0	18.2	63.8	-	45.5
Avg.	44.3	18.1	39.0	10.8	50.2	37.2	33.3	Avg.	50.1	20.3	43.1	10.9	55.5	40.3	36.7	Avg.	51.5	26.6	44.7	11.8	56.0	42.7	38.9

CDTrans* [53]	clp	inf	pnt	qdr	rel	skt	Avg.	SSRT-B* [46]	clp	inf	pnt	qdr	rel	skt	Avg.	C-SFTrans (Ours)	clp	inf	pnt	qdr	rel	skt	Avg.
clp	-	27.9	57.6	27.9	73.0	58.8	49.0	clp	-	33.8	60.2	19.4	75.8	59.8	49.8	clp	-	26.6	53.6	23.6	71.4	54.6	46.0
inf	58.6	-	53.4	9.6	71.1	47.6	48.1	inf	55.5	-	54.0	9.0	68.2	44.7	46.3	inf	55.9	-	51.7	11.4	69.6	46.0	46.9
pnt	60.7	24.0	-	13.0	69.8	49.6	43.4	pnt	61.7	28.5	-	8.4	71.4	55.2	45.0	pnt	60.0	25.2	-	14.3	71.2	51.1	44.4
qdr	2.9	0.4	0.3	-	0.7	4.7	1.8	qdr	42.5	8.8	24.2	-	37.6	33.6	29.3	qdr	43.2	8.2	17.4	-	40.2	28.8	27.5
rel	49.3	18.7	47.8	9.4	-	33.5	31.7	rel	69.9	37.1	66.0	10.1	-	58.9	48.4	rel	60.4	28.1	56.5	12.2	-	49.8	41.4
skt	66.8	23.7	54.6	27.5	68.0	-	48.1	skt	70.6	32.8	62.2	21.7	73.2	-	52.1	skt	66.7	26.5	56.2	25.1	71.0	-	49.1
Avg.	47.7	18.9	42.7	17.5	56.5	38.8	37.0	Avg.	60.0	28.2	53.3	13.7	65.3	50.4	45.2	Avg.	57.2	22.9	47.1	17.3	64.7	46.1	42.5

b) Multi Target Domain Adaptation (MTDA). In Table 2, we compare our proposed framework, C-SFTrans, with existing works on multi-target domain adaption on the Office-Home dataset. Our method achieves a 1.4% improvement over the source-free baseline (SHOT-B) and is comparable to the non-source-free method D-CGCT [37] despite using a pure transformer backbone while the latter uses a hybrid convolution-transformer feature extractor.

4.2. Analysis

We perform a thorough ablation study of our proposed approach and analyze the contribution of each component of our approach in Table 5.

a) Effect of non-causal factor classification. In Table 5, we study the effect of non-causal factor classification on the goal task performance using a subsidiary style classification task. We first train only the style task while keeping the goal task classifier f_g and causal parameters $\theta_f \setminus \theta_{h_n}$ fixed. Here, we observe a significant improvement of 1.1% on the source-side and 1.6% on the target-side. This validates Insight 2 since the non-causal style classification task improves the causal factor alignment, thereby improving the goal task performance.

b) Effect of both goal and non-causal tasks. Next, in Table 5, we use both goal task and non-causal style task alternately as proposed in Sec. 3. As per Insight 1, this should improve the alignment between source and target causal factors, and result in optimal clustering of task-related features. The alternate training of style and goal task yields an overall improvement of 3.7% on source-side and 2.4% on target-side, which validates Insight 1.

c) Comparisons with different backbones. In Table 6, we provide results for our approach with the ViT-Base backbone pre-trained on the ImageNet-21K dataset, and the DeiT-S backbone pre-trained on the ImageNet-1K dataset. We observe that our method achieves a 1.4% improvement

Table 6. Single-Source Domain Adaptation (SSDA) on Office-Home with ViT-B Backbone pretrained on the ImageNet-21K dataset (*bottom*), and the DeiT-S backbone pretrained on ImageNet-1K dataset (*top*). SF indicates *source-free* adaptation.

Method	SF	Office-Home				
		Ar→Cl	Cl→Pr	Pr→Rw	Rw→Ar	Avg.
CDTrans-S [53]	✗	60.7	75.6	84.4	77.0	74.4
SHOT-S	✓	56.3	73.7	81.3	76.7	72.0
C-SFTrans-S	✓	63.3	79.7	83.0	76.8	75.7 (+3.7)
SSRT-B [46]	✗	75.2	88.3	91.3	85.7	85.1
SHOT-B	✓	69.1	85.3	88.1	83.9	81.6
C-SFTrans-B (Ours)	✓	73.1	87.3	83.5	88.1	83.0 (+1.4)

over SHOT-B with the ViT-B backbone, and a 3.7% improvement over SHOT-S with the DeiT-S backbone.

5. Conclusion

In this work, we study the concepts of source free domain-adaptation from the perspective of causality. We conjecture that the spurious correlation among causal and non-causal factors are crucial to preserve in the target domain to improve the adaptation performance. Hence, we provide insights showing that the disentangling and aligning non-causal factors positively influence the alignment of causal factors in SFDA. Further, we first investigate the behaviour of vision transformers in SFDA and propose a novel Causality-enforcing Source-free Transformer (C-SFTrans) architecture for non-causal factor alignment. Based on our insights, we introduce a non-causal factor classification task to align non-causal factors. We also propose a novel Causal Influence Score criterion to improve the training. The proposed approach leads to improved task-discriminative causal factor alignment and outperforms the prior works on DA benchmarks of single-source and multi-target SFDA. **Acknowledgements.** Sunandini Sanyal was supported by the Prime Minister’s Research Fellowship, Govt of India.

References

- [1] Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan. A theory of learning from different domains. *Machine learning*, 79(1-2):151–175, 2010. [2](#)
- [2] Shai Ben-David, John Blitzer, Koby Crammer, and Fernando Pereira. Analysis of representations for domain adaptation. 2006. [4](#), [5](#)
- [3] Chao Chen, Zhihang Fu, Zhihong Chen, Sheng Jin, Zhaowei Cheng, Xinyu Jin, and Xian-Sheng Hua. Homm: Higher-order moment matching for unsupervised domain adaptation. In *AAAI*, 2020. [5](#)
- [4] Hongrui Chen, Chen Wu, Yonghao Xu, and Bo Du. Unsupervised domain adaptation for semantic segmentation via low-level edge information transfer. *arXiv preprint arXiv:2109.08912*, 2021. [3](#)
- [5] Tianlong Chen, Zhenyu Zhang, Yu Cheng, Ahmed Awadallah, and Zhangyang Wang. The principle of diversity: Training stronger vision transformers calls for reducing all levels of redundancy. In *CVPR*, 2022. [2](#)
- [6] 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 *CVPR*, 2016. [1](#)
- [7] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021. [2](#), [3](#)
- [8] Zhekai Du, Jingjing Li, Hongzu Su, Lei Zhu, and Ke Lu. Cross-domain gradient discrepancy minimization for unsupervised domain adaptation. In *CVPR*, 2021. [7](#)
- [9] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *ICML*, 2015. [1](#)
- [10] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *The Journal of Machine Learning Research*, 17(1):2096–2030, 2016. [1](#), [2](#)
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. [6](#), [7](#)
- [12] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei Efros, and Trevor Darrell. CyCADA: Cycle-consistent adversarial domain adaptation. In *ICML*, 2018. [2](#)
- [13] Lanqing Hu, Meina Kan, Shiguang Shan, and Xilin Chen. Duplex generative adversarial network for unsupervised domain adaptation. In *CVPR*, 2018. [2](#)
- [14] Salman Khan, Muzammal Naseer, Munawar Hayat, Syed Waqas Zamir, Fahad Shahbaz Khan, and Mubarak Shah. Transformers in vision: A survey. *ACM Computing Surveys (CSUR)*, 2021. [2](#), [4](#)
- [15] Myeongjin Kim and Hyeran Byun. Learning texture invariant representation for domain adaptation of semantic segmentation. In *CVPR*, 2020. [3](#)
- [16] Jogendra Nath Kundu, Suvaansh Bhambri, Akshay Kulkarni, Hiran Sarkar, Varun Jampani, and R Venkatesh Babu. Concurrent subsidiary supervision for unsupervised source-free domain adaptation. In *ECCV*, 2022. [2](#)
- [17] Jogendra Nath Kundu, Akshay Kulkarni, Suvaansh Bhambri, Deepesh Mehta, Shreyas Kulkarni, Varun Jampani, and R. Venkatesh Babu. Balancing discriminability and transferability for source-free domain adaptation. In *ICML*, 2022. [2](#), [3](#), [4](#), [6](#), [7](#)
- [18] Jogendra Nath Kundu, Akshay Kulkarni, Amit Singh, Varun Jampani, and R. Venkatesh Babu. Generalize then adapt: Source-free domain adaptive semantic segmentation. In *ICCV*, 2021. [6](#)
- [19] Jogendra Nath Kundu, Naveen Venkat, M V Rahul, and R. Venkatesh Babu. Universal source-free domain adaptation. In *CVPR*, 2020. [1](#), [3](#)
- [20] Jonghyun Lee, Dahuin Jung, Junho Yim, and Sungroh Yoon. Confidence score for source-free unsupervised domain adaptation. In *ICML*, 2022. [2](#)
- [21] Jian Liang, Dapeng Hu, and Jiashi Feng. Do we really need to access the source data? source hypothesis transfer for unsupervised domain adaptation. In *ICML*, 2020. [2](#), [3](#), [4](#), [6](#), [7](#), [8](#)
- [22] Jian Liang, Dapeng Hu, Yunbo Wang, Ran He, and Jiashi Feng. Source data-absent unsupervised domain adaptation through hypothesis transfer and labeling transfer. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. [2](#), [3](#), [6](#), [7](#)
- [23] Ming-Yu Liu and Oncel Tuzel. Coupled generative adversarial networks. In *NeurIPS*, 2016. [2](#)
- [24] Mingsheng Long, Zhangjie Cao, Jianmin Wang, and Michael I Jordan. Conditional adversarial domain adaptation. In *NeurIPS*, 2017. [7](#)
- [25] Mingsheng Long, Han Zhu, Jianmin Wang, and Michael I Jordan. Deep transfer learning with joint adaptation networks. In *ICML*, 2017. [2](#)
- [26] Fangrui Lv, Jian Liang, Shuang Li, Bin Zang, Chi Harold Liu, Ziteng Wang, and Di Liu. Causality inspired representation learning for domain generalization. In *CVPR*, 2022. [2](#), [3](#)
- [27] Yu Mitsuzumi, Go Irie, Daiki Ikami, and Takashi Shibata. Generalized domain adaptation. In *CVPR*, 2021. [5](#)
- [28] Muhammad Muzammal Naseer, Kanchana Ranasinghe, Salman H Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Intriguing properties of vision transformers. In *NeurIPS*, 2021. [2](#), [4](#)
- [29] Le Thanh Nguyen-Meidine, Atif Belal, Madhu Kiran, Jose Dolz, Louis-Antoine Blais-Morin, and Eric Granger. Unsupervised multi-target domain adaptation through knowledge distillation. In *WACV*, 2021. [7](#)
- [30] Namuk Park and Songkuk Kim. How do vision transformers work? In *ICLR*, 2022. [2](#), [4](#)
- [31] Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. In *ICCV*, 2019. [7](#)
- [32] Xingchao Peng, Zijun Huang, Ximeng Sun, and Kate Saenko. Domain agnostic learning with disentangled representations. In *ICML*, 2019. [2](#)
- [33] Xingchao Peng, Ben Usman, Neela Kaushik, Judy Hoffman, Dequan Wang, and Kate Saenko. VisDA: The visual domain

- adaptation challenge. *arXiv preprint arXiv:1710.06924*, 2017. 7
- [34] Vihari Piratla, Praneeth Netrapalli, and Sunita Sarawagi. Efficient domain generalization via common-specific low-rank decomposition. In *ICML*, 2020. 4
- [35] Sanqing Qu, Guang Chen, Jing Zhang, Zhijun Li, Wei He, and Dacheng Tao. BMD: A general class-balanced multi-centric dynamic prototype strategy for source-free domain adaptation. In *ECCV*, 2022. 2
- [36] Maithra Raghu, Thomas Unterthiner, Simon Kornblith, Chiyuan Zhang, and Alexey Dosovitskiy. Do vision transformers see like convolutional neural networks? In *NeurIPS*, 2021. 4
- [37] Subhankar Roy, Evgeny Krivosheev, Zhun Zhong, Nicu Sebe, and Elisa Ricci. Curriculum graph co-teaching for multi-target domain adaptation. In *CVPR*, 2021. 7, 8
- [38] Subhankar Roy, Martin Trapp, Andrea Pilzer, Juho Kannala, Nicu Sebe, Elisa Ricci, and Arno Solin. Uncertainty-guided source-free domain adaptation. In *ECCV*, 2022. 2
- [39] Paolo Russo, Fabio M Carlucci, Tatiana Tommasi, and Barbara Caputo. From source to target and back: symmetric bi-directional adaptive gan. In *CVPR*, 2018. 2
- [40] Kate Saenko, Brian Kulis, Mario Fritz, and Trevor Darrell. Adapting visual category models to new domains. In *ECCV*, 2010. 2
- [41] Kate Saenko, Brian Kulis, Mario Fritz, and Trevor Darrell. Adapting visual category models to new domains. In *ECCV*, 2010. 7
- [42] K. Sage and S. Young. Computer vision for security applications. In *Proceedings IEEE 32nd Annual 1998 International Carnahan Conference on Security Technology (Cat. No.98CH36209)*, pages 210–215, 1998. 1
- [43] Kuniaki Saito, Kohei Watanabe, Yoshitaka Ushiku, and Tatsuya Harada. Maximum classifier discrepancy for unsupervised domain adaptation. In *CVPR*, 2018. 2
- [44] Swami Sankaranarayanan, Yogesh Balaji, Carlos D Castillo, and Rama Chellappa. Generate to adapt: Aligning domains using generative adversarial networks. In *CVPR*, 2018. 2
- [45] Bernhard Schölkopf, Francesco Locatello, Stefan Bauer, Nan Rosemary Ke, Nal Kalchbrenner, Anirudh Goyal, and Yoshua Bengio. Toward causal representation learning. *Proceedings of the IEEE*, 109(5):612–634, 2021. 3
- [46] Tao Sun, Cheng Lu, Tianshuo Zhang, and Haibin Ling. Safe self-refinement for transformer-based domain adaptation. In *CVPR*, 2022. 2, 6, 7, 8
- [47] Xinwei Sun, Botong Wu, Xiangyu Zheng, Chang Liu, Wei Chen, Tao Qin, and Tie-Yan Liu. Recovering latent causal factor for generalization to distributional shifts. In *NeurIPS*, 2021. 2, 3, 4
- [48] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *ICML*, 2021. 7, 8
- [49] Hemanth Venkateswara, Jose Eusebio, Shayok Chakraborty, and Sethuraman Panchanathan. Deep hashing network for unsupervised domain adaptation. In *CVPR*, 2017. 7
- [50] Fan Wang, Zhongyi Han, Yongshun Gong, and Yilong Yin. Exploring domain-invariant parameters for source free domain adaptation. In *CVPR*, 2022. 2, 3, 4, 6, 7
- [51] Yunqi Wang, Furui Liu, Zhitang Chen, Yik-Chung Wu, Jianye Hao, Guangyong Chen, and Pheng-Ann Heng. Contrastive-ACE: Domain generalization through alignment of causal mechanisms. *IEEE Transactions on Image Processing*, 32:235–250, 2022. 3
- [52] Haifeng Xia, Handong Zhao, and Zhengming Ding. Adaptive adversarial network for source-free domain adaptation. In *ICCV*, 2021. 6
- [53] Tongkun Xu, Weihua Chen, Pichao Wang, Fan Wang, Hao Li, and Rong Jin. CDTrans: Cross-domain transformer for unsupervised domain adaptation. In *ICLR*, 2022. 2, 3, 6, 7, 8
- [54] Guanglei Yang, Hao Tang, Zhun Zhong, Mingli Ding, Ling Shao, Nicu Sebe, and Elisa Ricci. Transformer-based source-free domain adaptation. In *APIN*, 2022. 2
- [55] Jinyu Yang, Jingjing Liu, Ning Xu, and Junzhou Huang. TVT: Transferable vision transformer for unsupervised domain adaptation. In *WACV*, 2023. 2, 6, 7
- [56] Shiqi Yang, Yaxing Wang, Joost van de Weijer, Luis Herranz, and Shangling Jui. Exploiting the intrinsic neighborhood structure for source-free domain adaptation. In *NeurIPS*, 2021. 2, 6, 7
- [57] Shiqi Yang, Yaxing Wang, Joost van de Weijer, Luis Herranz, and Shangling Jui. Generalized source-free domain adaptation. In *ICCV*, 2021. 6
- [58] Yalan Ye, Ziqi Liu, Yangwuyong Zhang, Jingjing Li, and Hengtao Shen. Alleviating style sensitivity then adapting: Source-free domain adaptation for medical image segmentation. In *ACMMM*, 2022. 3
- [59] Yuchen Zhang, Tianle Liu, Mingsheng Long, and Michael Jordan. Bridging theory and algorithm for domain adaptation. In *ICML*, 2019. 8