Received: 7 November 2022 | Revised: 26 December 2022 | Accepted: 26 December 2022 | Published online: 5 January 2023

Artificial Intelligence and Applications
2023, Vol. 1(1) 43–51
DOI: 10.47852/bonviewAIA2202524

RESEARCH ARTICLE

Distribution Alignment Using Complement Entropy Objective and Adaptive Consensus-Based Label Refinement for Partial Domain Adaptation

Sandipan Choudhuri[1,*], Suli Adeniye[1] and Arunabha Sen[1]

[1] *Arizona State University, USA*

Abstract: In this work, we address a realistic case of unsupervised domain adaptation, where the source label set subsumes that of the target. This relaxation in the requirement of an identical label set assumption, as witnessed in the standard closed-set variant, poses a challenging obstacle of negative transfer that potentially misleads the learning process from the intended target classification objective. To counteract this issue, we propose a novel framework for a partial domain adaptation setup that enforces domain and category-level alignments through optimization of intra- and inter-class distances, uncertainty suppression on classifier predictions, and target supervision with an adaptive consensus-based sample filtering. In this work, we aim to modify the latent space arrangement where samples from identical classes are forced to reside in close proximity while that from distinct classes are well separated in a domain-agnostic fashion. In addition, the proposed model addresses a challenging issue of uncertainty propagation by employing a complement entropy objective that requires the incorrect classes to have uniformly distributed low-prediction probabilities. Target supervision is ensured by employing a robust technique for adaptive pseudo-label generation using a nonparametric classifier. The methodology employs a strategy that permits supervision from target samples with prediction probabilities higher than an adaptive threshold. We conduct experiments involving a range of partial domain adaptation tasks on two benchmark datasets to thoroughly assess the proposed model's performance against the state-of-the-art methods. In addition, we performed an ablation study to validate the necessity of the incorporated modules and highlight their contribution to the proposed framework. The experimental findings obtained manifest the superior performance of the proposed model when compared to the benchmarks.

Keywords: partial domain adaptation, domain adaptation, conditional distribution alignment, complement entropy objective, object recognition, pseudo-label-based supervision

1. Introduction

Deep neural networks have remarkably improved the performance of a wide range of frameworks designed to address complex machine learning tasks (Choudhuri et al., 2018; Dang et al., 2019; Guo et al., 2021; Liu et al., 2021; Wang et al., 2019). The availability of extensive annotated data is a prerequisite for the generalizability of such models. In some real-world situations where data collection and subsequent annotation involve high costs, this supervision utilizing richly annotated data is often challenging. By transferring useful information from a large-scale dataset that has previously been labeled in a relevant area, domain adaptation (*da*) approaches (Ganin et al., 2016; Li et al., 2020) can mitigate this annotation demand. An implicit assumption by the majority of the currently employed *da* approaches (Ganin & Lempitsky, 2015; Ganin et al., 2016; Li et al., 2020) is that labeled and unlabeled domains share an identical label set. However, acquiring a source with a suitable label set under such a rigorous label space constraint is highly challenging in practice. This restriction is relaxed in a partial domain adaptation (*pda*) setup (Cao et al., 2018b), which caters to a more realistic scenario where the source label set subsumes the target label set (Figure 1 highlights the difference in the label set relationships between the two domains in closed-set and partial domain adaptation).

Although the constraint relaxation increases the likelihood of large-scale labeled dataset availability to serve as the source domain, it introduces the issue of unwanted information transfer from classes private to the source (*negative transfer*) (Cao et al., 2018b, 2019), impairing the effectiveness of classification framework. Existing approaches (Cao et al., 2018a, 2018b, 2019; Zhang et al., 2018) have sought to down-weight these samples either by reweighting them or by aggregating all target sample predictions at the category level to determine the common class information existing between the domains. However, such an estimation is noise-prone, especially at the initial stages of classifier training, and could adversely affect the learning process by misleading it from the intended objective. In this work, we train the classifier on the entire source data while ensuring the category distributions are well separated (see Figure 2). This is achieved using objectives that yield more distinct class distributions in a domain-agnostic setup.

*Corresponding author: Sandipan Choudhuri, Arizona State University, USA. Email: schoud13@asu.edu

Figure 1
Closed-set and partial domain adaptation scenarios

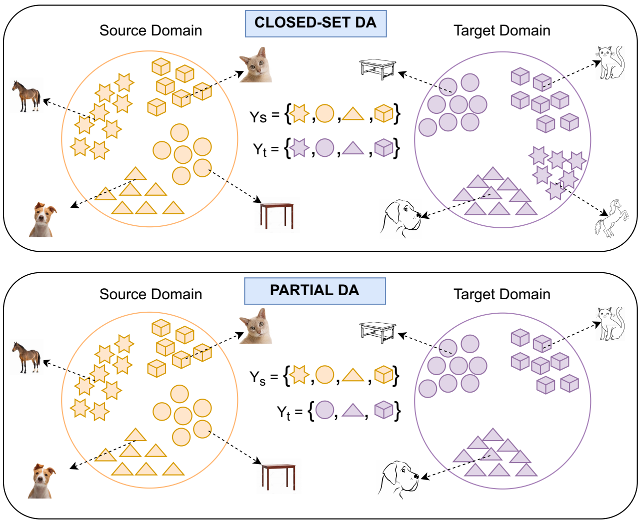

Figure 2
Significance of domain + class distribution alignment

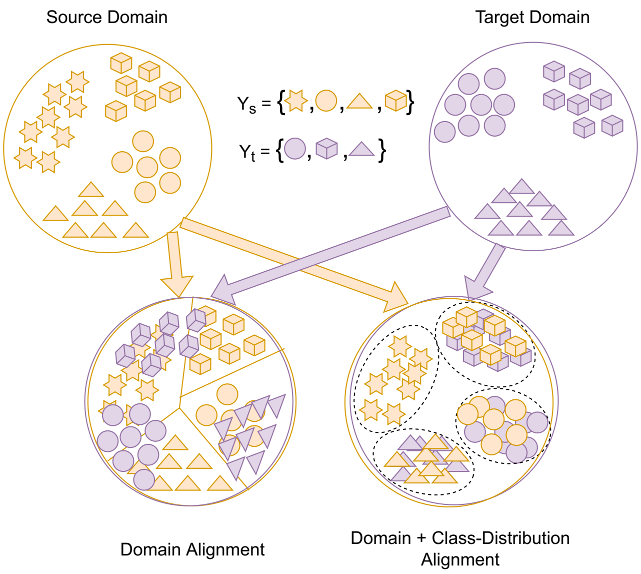

The standard domain adaptation approaches often trade off feature discriminability for feature transferability when using cross-entropy loss in the labeled source domain. Despite witnessing a reduction in the domain shift, the classifier may thus perform worse on the target data. This is because the source classes are not evenly spaced from one another and may cause confusion (or uncertainty) to spread to the target predictions (an issue coined as uncertainty propagation). Some unlabeled target data that might be readily misclassified are pushed to match these source data during domain alignment. The classification heads of the existing domain adaptation models (Cao et al., 2018a; Ganin et al., 2016; Zhang et al., 2018) often overlook such a significant issue and apply the standard cross-entropy loss in the labeled source domain to learn transferable features across domains relevant to the classification task. This loss exclusively supervises the ground-truth class and disregards the scores of the incorrect classes. Consequently, samples lying at the decision boundaries of two nearby classes might confuse the classifier, eventually leading to uncertainty propagation. Prior approaches (Kumar et al., 2018; Shu et al., 2018) have addressed this issue by focusing on increasing the gap between two distinct classes using adversarial training (Miyato et al., 2018). However, the increased set of parameters that require tuning in these models elevates the model complexity. In this work, we follow an objective training approach presented by Chen et al. (2018) and incorporate a complement entropy objective module in the proposed network that enforces uniformly distributed low-probability values among the incorrect classes. This objective is defined as the average of sample-wise entropy over complement categories in a mini-batch. The predicted probability of complement classes is neutralized as the number of classes increases by maximizing the complement entropy (since entropy is maximized when events are equally likely to occur). In other words, increasing complement entropy evenly distributes the complement classes' expected probability. This reduces the likelihood of an incorrect category (a) having an adequately high-prediction probability and (b) challenging the ground-truth category.

Prior works (Cao et al., 2018a, 2018b, 2019; Choudhuri et al., 2020, 2022; Zhang et al., 2018) achieve domain invariance by aligning the distributions of two domains via an adversarial objective. This, however, is not a sufficient condition for improving the target classifier performance (Jing et al., 2020). In this approach, we leverage the pseudo-labeling technique for target supervision, which is essential for domain and class-level alignment. Since most of the pseudo-labels generated at the initial stages of model training are unreliable and could cause the learning process to deviate from its intended objective, we utilize a nonparametric classifier that measures a sample's likelihood of being assigned to a source cluster. Subsequent selection is conducted to obtain a subset of the target samples that yield correct class prediction probabilities higher than an adaptive threshold parameter. These predictions, aggregated over a fixed number of iterations for increased robustness, are utilized for target supervision.

To sum up, we have made the following contributions to this work:

- A strategy to modify the latent space arrangement by maximizing the intra-category compactness and the inter-category separation in conjunction with achieving domain alignment.
- A technique to reduce uncertainty propagation in the classifier from hard samples, using a complement entropy objective.
- A robust and adaptive method for target supervision using high-confidence samples in the target domain.

2. Related Work

The effectiveness of the modern transfer learning approaches for mitigating domain discrepancy and transferring underlying information between domains has been thoroughly examined in several existing works (Hoffman et al., 2014; Oquab et al., 2014; Yosinski et al., 2014). This encourages remarkable research efforts in many applications and minimizes a load of manual labeling by analyzing the information from other sources. A majority of the present works address the domain discrepancy problem through learning domain-invariant representations or applying instance reweighting schemes (Pan & Yang, 2010). The feature learning strategies primarily focus on generating latent subspaces that capture the properties of both domains. For instance, the authors in Ghifary et al. (2016) present a strategy for minimizing the domain distribution differences while preserving class discriminability. The Transfer Component Analysis framework, proposed by Pan et al. (2010) utilizes the Maximum Mean Discrepancy (MMD) metric to learn transferable components in the Reproducing Kernel Hilbert Space (RKHS) to reduce the marginal distribution gap between two domains. The work proposed in Long et al. (2013) employs a pseudo-labeling technique to jointly align the source and target domain's marginal and conditional distributions by minimizing the distance between the representative centers of the two domains and the proximity of instances from the same category. The classical instance reweighting schemes (Chen et al., 2011; Huang et al., 2006) address the adaptation task by reweighting the source samples. However, these approaches yield unsatisfactory performance on complex adaptation tasks since they operate by learning shallow features for both domains.

To eliminate this issue, recent adaptation works have leveraged deep learning frameworks to disentangle different factors of variations behind data and learn complex representations which exhibit satisfactory transferability across similar tasks (Tzeng et al., 2014; Yosinski et al., 2014). These works estimate the distribution means and match them in the adaptation layers. For instance, Tzeng et al. work on Deep Domain Confusion (Tzeng et al., 2014) utilizes a domain confusion loss by leveraging the MMD metric to learn the domain-invariant representations. A similar line of work (Zhang et al., 2018) effectively aligns the distribution across domains and eliminates domain discrepancy by using high-order statistical properties centering around MMD. A different line of approach couples adversarial loss with a domain classifier to induce confusion and modify sample data in a domain-agnostic fashion. Adversarial learning is utilized by the authors of Ganin et al. (2016), Li et al. (2019), and Tzeng et al. (2017) to create a mini-max game for extracting domain-invariant features. Compared to the MMD-based approaches, it is thus often difficult to establish reliable solutions. Unfortunately, they only work in a limited closed-set domain adaptation context (source label set identical to the target label set) and do not scale well in a partial domain adaptation setting.

In the real world, it is more plausible to drop the identical label set constraint of a closed-set scenario and acquire a large-scale source dataset that subsumes the smaller target dataset. The adaptation process can thus be improved by necessitating the transfer of some significant information from the source to the target. Selective Adversarial Network (SAN) (Cao et al., 2018b) addresses such a partial domain adaptation scenario by down-weighing private source category samples, using multiple adversarial networks, to ensure effective knowledge transfer. By extending this concept, the authors of Cao et al. (2018a) develop a class importance weight estimation framework by aggregating target sample prediction scores. A similar line of concept is proposed by Zhang et al. (2018) in their work on Importance Weighted Adversarial Nets (IWAN), which utilizes an auxiliary domain discriminator to determine the degree of closeness a source sample shares with the target domain. The Example Transfer Network (ETN) (Cao et al., 2019) uses discriminative information to quantify the transferability of source domain samples, thereby yielding a soft metric for separating the common categories from the private source classes. Despite outperforming the closed-set domain adaption techniques, poor classification performance during the first training phases in these models may result in significant inaccuracies when determining private source categories. In this work, we have attempted to address the shortcomings highlighted above.

3. Methodology

Reweighing the degree of source sample contribution in the training process has been adopted by a majority of the current *partial domain adaptation* approaches that aim to alleviate the negative transfer induced by classes private to the source domain. However, assigning sample/class importance weights to source samples at the initial stages of classifier training could potentially derail the learning process from the intended objective. Instead, we focus on training a classifier on the entire source data while ensuring the optimization of intra- and inter-category distances. Alongside this, we utilize a complement entropy objective to counteract the negative effects of the uncertain source samples residing near the decision boundaries of classes. The following sections give a mathematical formulation of the problem statement and describe the proposed domain adaptation framework.

3.1. Problem settings

In this work, we consider a typical *domain adaptation* setup with *two* datasets, representing the *source (s)* and the *target (t)* domains. The *source* dataset $\mathbb{D}_s = \{\langle x_s^i, y_s^i \rangle\}_{i=1}^{|\mathbb{D}_s|}$, consisting of $|\mathbb{D}_s|$ labeled points, is sampled from a distribution $\mathbb{P}_s$. The *target* dataset $\mathbb{D}_t = \{x_t^j\}_{j=1}^{|\mathbb{D}_t|}$ contains $|\mathbb{D}_t|$ unlabeled samples, drawn from distribution $\mathbb{P}_t$, where $x_{s/t} \in \mathbb{R}^d$ and $\mathbb{P}_s \neq \mathbb{P}_t$. Since target class information (Y_t) is *absent* during adaptation, the *closed-set* variation assumes that samples in $\mathbb{D}_s$ and $\mathbb{D}_t$ are classified into categories from the known source label set Y_s $(Y_t = Y_s)$. Obtaining such a source dataset that exhibits complete alignment with the target label set is challenging. Citing this, we relax the equality constraint by addressing a more realistic *partial domain adaptation (pda)* scenario where the label set of s subsumes that of t (i.e., $Y_t \subseteq Y_s$). To summarize, we aim to design a classifier $H : \mathcal{X} \to Y$ $(H \in \mathcal{H}$, where $\mathcal{H}$ denotes the hypothesis space) that operates under a *pda* setting, by leveraging source domain supervision under an adaptation setup, to *reduce the target classification risk.*

This relaxation, however, introduces the *negative transfer* problem where samples x_s private to the source domain $(\langle x_s, y_s \rangle \in \mathbb{D}_s, \quad y_s \in Y_s - Y_t)$ promotes superfluous knowledge transfer, thereby misleading the classification process. Mitigating this requires strategic estimation of categories common to both domains for enhancing the target classification accuracy.

3.2. Proposed approach

Unlike recent approaches, which center around class/sample reweighting schemes to localize outlier categories $(Y_s - Y_t)$, we propose to take a different route through strategic selection of confident target samples for *domain* and *category-level* alignment. The proposed network (Figure 3) achieves both by maximizing intra-class compactness and inter-class separation between samples from the source and target domains. To counteract the adverse effects of uncertainty propagation brought on by the prediction probabilities of inaccurate classes, we apply a complement entropy objective in conjunction with the regular cross-entropy loss. A detailed description of the proposed modules is presented in the following sections.

3.2.1. Sample classification

As highlighted in Section 3.1, the objective is to design a hypothesis classifier $H : \mathcal{X} \to Y$ that represents the label space as a function of the input feature space. In our approach, h consists of two components, E and C, with $H = C \circ E$. Here, $E : \mathcal{X} \to \mathcal{Z}$ denotes the feature encoder that maps samples in the input feature space $\mathcal{X}$ to the latent space $\mathcal{Z}$, and $C : \mathcal{Z} \to Y$ denotes a classifier network that maps points in $\mathcal{Z}$ to the label space Y. Source supervision is obtained by training E and C with the categorical cross-entropy loss $l_{cr\_ent}(\cdot, \cdot)$ between the ground-truth labels of source samples and their predicted classification scores. Alongside, the classifier is trained on a subset $\mathbb{D}_\tau \in \mathbb{D}_t$ that consists of target samples with confidently predicted pseudo-labels $\bar{y}_\tau$ (label generation strategy described in

Figure 3
Architectural diagram of the proposed domain adaptation model (model training phase)

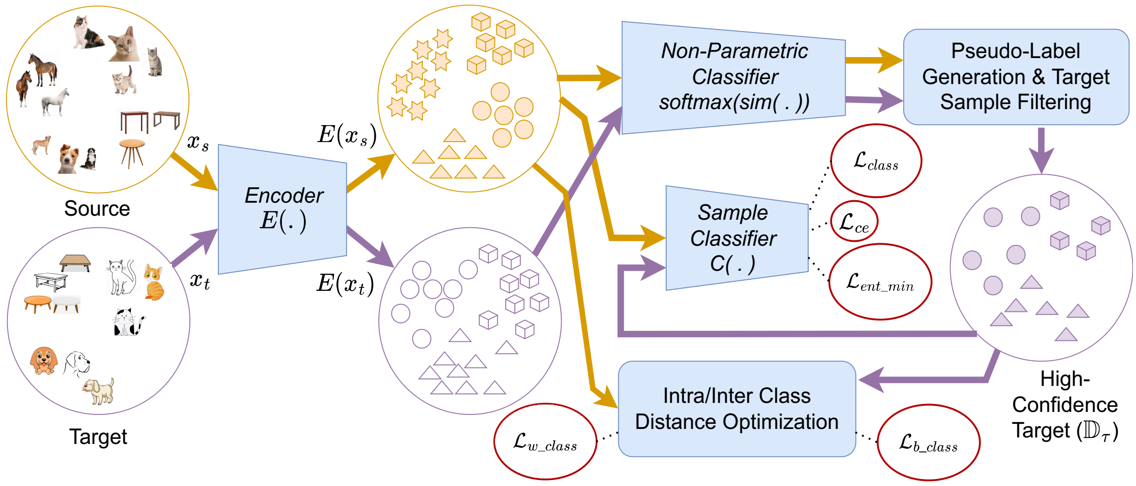

Section 3.2.3.3). The overall classification objective $\mathcal{L}_{class}$ is represented as ($|\cdot|$ represents set-cardinality):

$$\mathcal{L}_{class} = \begin{cases} \frac{1}{|\mathbb{D}_s|} \sum\limits_{\langle x_s^i, y_s^i \rangle \in \mathbb{D}_s} l_{cr\_ent}(C(E(x_s^i)), y_s^i) + \\ \frac{1}{|\mathbb{D}_\tau|} \sum\limits_{\langle x_\tau^j, \bar{y}_\tau^j \rangle \in \mathbb{D}_\tau} l_{cr\_ent}(C(E(x_\tau^j)), \bar{y}_\tau^j), & \text{if } \mathbb{D}_\tau \neq \emptyset \\ \frac{1}{|\mathbb{D}_s|} \sum\limits_{\langle x_s^i, y_s^i \rangle \in \mathbb{D}_s} l_{cr\_ent}(C(E(x_s^i)), y_s^i), & \text{otherwise} \end{cases} \tag{1}$$

3.2.2. Reducing uncertainty with complement entropy objective

Prior domain adaptation approaches (Cao et al., 2018a; Ganin et al., 2016; Zhang et al., 2018) mostly disregard the feature discriminability and merely apply the standard cross-entropy loss in the labeled source domain to learn transferable features across domains, besides ensuring domain alignment. This does not pose a sufficient condition for improving the target classification performance as there is no explicit addressing of the inter-class separation and suppression of incorrect class probabilities. For instance, in a three-class classification problem, an output probability of [0.6, 0.25, 0.15] is more uncertain than [0.6, 0.2, 0.2] despite having the same cross-entropy loss. Some source and target samples, from $\mathbb{D}_s$ and $\mathbb{D}_\tau$ respectively, might reside at the decision boundaries of classes in close proximity. These hard-to-classify samples eventually lead to a propagation of confusion (uncertainty) when classifying target data, thereby thwarting the classification process. In this approach, we have addressed this issue by enforcing the classifier $C(\cdot)$ to yield more certain predictions. Since the cross-entropy loss supervises the ground-truth class solely and ignores the scores of the incorrect classes, we follow an approach on similar lines as Chen et al. (2019) and Liang et al. (2020) and utilize a complement entropy objective to ensure uniformity and low-prediction probabilities for incorrect classes. The proposed complement entropy objective is presented as follows:

$$\mathcal{L}_{ce} = \begin{cases} \frac{1}{|\mathbb{D}_s|} \sum\limits_{\langle x_s^i, y_s^i \rangle \in \mathbb{D}_s} l_{ce}(C(E(x_s^i)), y_s^i) + \\ \frac{1}{|\mathbb{D}_\tau|} \sum\limits_{\langle x_\tau^j, \bar{y}_\tau^j \rangle \in \mathbb{D}_\tau} l_{ce}(C(E(x_\tau^j)), \bar{y}_\tau^j), & \text{if } \mathbb{D}_\tau \neq \emptyset \\ \frac{1}{|\mathbb{D}_s|} \sum\limits_{\langle x_s^i, y_s^i \rangle \in \mathbb{D}_s} l_{ce}(C(E(x_s^i)), y_s^i), & \text{otherwise} \end{cases} \tag{2}$$

$$\text{where} \quad l_{ce}(\hat{y}, y) = (1 - \hat{y}_g)^\gamma \sum_{k \neq g} \frac{\hat{y}_k}{1-\hat{y}_g} \log \frac{\hat{y}_k}{1-\hat{y}_g}$$

Here, γ is a hyperparameter, k represents the indices of all classes except the ground-truth class, and g is the index of the ground-truth class in y. The process of constructing the subset $\mathbb{D}_\tau$ of high-confidence target samples is illustrated in the subsequent section.

3.2.3. Pseudo-label generation and consensus-based target supervision

Domain alignment and *elimination of the adverse effects of* x_s with classes private to s ($\langle x_s, y_s \rangle \in \mathbb{D}_s, \quad y_s \in Y_s - Y_t$) is achieved through the utilization of a nonparametric classifier that maps the target samples to the nearest source cluster centers and assigns confidence scores accordingly. Inspired by Jing et al. (2020), we build the classifier by coupling the *cosine similarity* metric $cos(\cdot||\cdot)$ with *softmax*. The dataset $\mathbb{D}_\tau$, thus created, consists of those target samples with confidently predicted probability labels and is allowed to participate in the cross-domain alignment updating process. The adaptive pseudo-labeling strategy is as follows:

- **Step 1:** The latent representation $E(x) \in \mathcal{Z}$ is obtained $\forall x \in \mathbb{D}_s \cup \mathbb{D}_t$ by encoding them using $E(\cdot)$.

- **Step 2:** The representative cluster center (mean embedding) μ_s^c of a source class $c \in Y_s$ is computed on samples $x_s \in \mathbb{D}_s$, using the term given below:

$$\mu_s^c = \frac{1}{|\mathbb{D}_s^c|} \sum_{x_s^i \in \mathbb{D}_s^c} E(x_s^i)$$
$$\text{where } \mathbb{D}_s^c = \{\langle x_s^i, y_s^i \rangle \mid \forall i, \langle x_s^i, y_s^i \rangle \in \mathbb{D}_s, y_s^i = c\} \tag{3}$$

- **Step 3:** The encoded representation $E(x_t)$ of a target sample $x_t \in \mathbb{D}_t$ is processed by a nonparametric similarity function $sim(\cdot)$ to generate a vector of $|Y_s|$ similarity values that quantify x_t s closeness to the representative cluster centers of the source classes. An entry in vector $sim(E(x_t^j))$ corresponding to the source class c is represented as:

$$sim(E(x_t^j))^c = cos(E(x_t^j)||\mu_s^c) \tag{4}$$

In the equation above, $sim(E(x_t^j))^c$ ranges between [0, 1] with a higher value signifying greater similarity.

- **Step 4:** Pseudo-label $\bar{y}_t^j$ for a target sample x_t^j is generated by computing the probability scores using $softmax(\cdot)$ function over $sim(E(x_t^j))$, followed by the selection of class that yields the highest probability:

$$\mathbf{p}_t^j = softmax(sim(E(x_t^j))) \tag{5}$$

The vector $\mathbf{p}_t^j$ of *softmax* probabilities might be susceptible to noise due to the absence of target labels. To mitigate this effect, a subsequent label refinement procedure is carried out. In particular, to obtain robust and cleaner pseudo-labels, we employ the moving average on the *current*, and $N-1$ *prior output predictions* aggregated over the latest N epochs. This is represented below ($[\cdot]_k$ indicates value at k^{th} epoch):

$$\begin{aligned} &[\bar{y}_t^j]_T \leftarrow \underset{c}{\text{argmax}} \quad [\mathbf{p}_t^j]_T \\ &\text{where} \quad [\mathbf{P}_t^j]_T = \tfrac{1}{N} \sum_{k=T-N-1}^{T} [\mathbf{p}_t^j]_k \end{aligned} \tag{6}$$

A target sample's likelihood of being mapped with the closest cluster center can be tracked using the confidence probability value $max([\mathbf{P}_t^j]_T)$, in an epoch T. A low confidence score implies significant confusion exists in the model when assigning the sample to a category. The training process would impede by deflecting focus away from the intended objective of accurately categorizing target data when using such potentially inaccurate pseudo-labels. To circumvent this issue, we regulate target supervision by allowing a portion of the target samples in $\mathbb{D}_t$ (ones with above-average confidence scores) to partake in the training process. The dataset $\mathbb{D}_\tau$, thus constructed using a adaptive threshold τ, is mathematically expressed as follows:

$$[\mathbb{D}_\tau]_T = \{\langle x_t^j, \bar{y}_t^j \rangle \mid \forall j, x_t^j \in \mathbb{D}_t, \bar{y}_t^j \leftarrow [\bar{y}_t^j]_T, max([\mathbf{P}_t^j]_T)[\tau]_T\} \tag{7}$$

The threshold parameter $[\tau]_T$ (in equation (8)) is computed over the values obtained in the latest N epochs and indicates the average probability of source samples belonging to the ground-truth class and is computed across the predicted outputs of the nonparametric classifier $softmax(sim(\cdot))$, at epoch T:

$$[\tau]_T = \frac{1}{N \cdot |\mathbb{D}_s|} \sum_{k=T-N-1}^{T} \sum_{x_s^i \in \mathbb{D}_s} max(softmax(sim([E(x_s^i)]_k))) \tag{8}$$

3.2.4. Maximizing inter-class separation

An essential requirement for improving classification accuracy is *strategically creating a sufficiently "well-organized" arrangement of samples in the encoded latent space with regard to the class distributions in a domain-agnostic fashion*. In other words, data with various class labels must be assigned to different class distributions, whereas samples with the same class label must be clustered to their respective distribution, alongside eliminating discrepancy between the two domains. Citing this, we emphasize separating two distinct classes by maximizing their mean embeddings. It is worth mentioning that the approach is deployed from a domain-invariant standpoint, as distances between distinct classes of the same domain are simultaneously maximized (captured in the first and second terms of equation (9)). The *between-class loss* $\mathcal{L}_{b\_class}$ that aids in maximizing the inter-class distance is represented as:

$$\begin{aligned} \mathcal{L}_{b\_class} = -\Bigg[&\alpha\Big(\frac{1}{|Y_s|(|Y_s|-1)} \sum_{c \in Y_s} \sum_{\substack{c' \in Y_s \\ c' \neq c}} ||\mu_s^c - \mu_s^{c'}||_2 \\ &+ \frac{1}{|Y_\tau|(|Y_\tau|-1)} \sum_{c \in Y_\tau} \sum_{\substack{c' \in Y_\tau \\ c' \neq c}} ||\mu_\tau^c - \mu_\tau^{c'}||_2\Big) \\ &+ \frac{\beta}{|Y_\tau|(|Y_\tau|-1)} \sum_{c \in Y_\tau} \sum_{\substack{c' \in Y_\tau \\ c' \neq c}} ||\mu_s^c - \mu_\tau^{c'}||_2 \Bigg] \end{aligned} \tag{9}$$

We utilize L_2-norm (represented by $||\cdot||_2$) for computing the distance between two mean embeddings in the equation above. The label set of $\mathbb{D}_\tau$ (refer to equation (7)) is represented by Y_τ, where $Y_\tau \subseteq Y_s$. For computing the mean embeddings $\mu_{s/\tau}^c$ over samples from a class c, we follow a similar strategy as highlighted in equation (3). The contribution of the cross-domain and within-domain terms is regulated through the hyperparameters α and β.

3.2.5. Minimizing within-class separation

As mentioned above, a "well-organized" arrangement of samples in the latent space is essential for refining the classification process. The previous section addresses maximizing the separation of distinct classes to avoid classifier confusion. In this section, we propose an objective that groups samples of the same class together to generate more compact category distributions. This is realized by minimizing the distance between the encoded representations of any two samples belonging to the same category in a domain-neutral arrangement. The *within-class loss* $\mathcal{L}_{w\_class}$, hence formulated, is represented as follows:

$$\mathcal{L}_{w\_class} = \frac{1}{|Y_s|} \sum_{c \in Y_s} \left[\frac{1}{|\mathbb{D}^c|(|\mathbb{D}^c|-1)} \sum_{x^i \in \mathbb{D}^c} \sum_{\substack{x^j \in \mathbb{D}^c \\ x^j \neq x^i}} ||E(x^i) - E(x^j)||^2 \right] \tag{10}$$

Where dataset $\mathbb{D}^c = \{\langle x^k, y^k \rangle | \forall k, \langle x^k, y^k \rangle \in \mathbb{D}_s \cup \mathbb{D}_\tau, y^k = c\}$ contains all samples from $\mathbb{D}_s$ and $\mathbb{D}_\tau$ with class label c, for $c \in Y_s$. Indices i, j are used in equation (10) to represent distinct samples in $\mathbb{D}^c$.

3.2.6. Entropy minimization of target samples

Two significant negative effects are observed in the early phases of a classification process in a domain adaptation setup: (a) difficulty in transferring sample information due to significant domain shifts and (b) unfavorable reduction of certainty in the classifier. As an effort to counter such impacts, we incorporate the entropy minimization principle on the target samples in $\mathbb{D}_t$, which is represented as:

$$L_{ent\_min} = -\frac{1}{|\mathbb{D}_t|} \sum_{x_t^j \in \mathbb{D}_t} \sum_{c \in Y_s} [C(E(x_t^j))]^c \quad log([C(E(x_t^j))]^c) \tag{11}$$

$[C(E(x_t^j))]^c$, in equation (11), represents the classifier prediction probability of the target sample x_t belonging to class c.

3.2.7. Overall objective

To summarize, the overall objective function is formulated as follows:

$$\mathcal{L} = \mathcal{L}_{class} + \eta\mathcal{L}_{ce} + \mathcal{L}_{b\_class} + \delta\mathcal{L}_{w\_class} + \mathcal{L}_{ent\_min} \tag{12}$$

η and δ are user-defined hyperparameters regulating the contribution of each objective in the learning process.

4. Experiments

To conduct a thorough evaluation, we compare the proposed model with the state-of-the-art techniques on two benchmark datasets (Office-Home (Venkateswara et al., 2017) and Office-31 (Saenko et al., 2010)) under a variety of *pda* settings with multiple adaptation tasks. We adhere to the accepted evaluation standard as highlighted in Cao et al. (2018a, 2018b) and Tzeng et al. (2017) and utilize all labeled source data and unlabeled target data for partial domain adaptation. A target classification accuracy percentage is computed as follows:

$$\frac{\sum_{<x_t, y_t> \in \mathbb{D}_t} \mathbb{1}_{\hat{y}_t = y_t}}{|\mathbb{D}_t| \times 100\%} \tag{13}$$

Here, $\hat{y}_t = C(E(x_t))$ represents the trainable classifier output for the target sample x_t. The indicator function $\mathbb{1}$ is set to 1 when the predicted output is equal to the ground-truth label ($\hat{y}_t = y_t$), and to 0 otherwise. It is to be noted that the ground-truth labels y_t for the target samples $x_t \in \mathbb{D}_t$ are only utilized during evaluation and not during model training. We further provide a comprehensive analysis of the model performance with the addition of *complement entropy objective*, *consensus-based target supervision*, and the *intra-/inter-class distribution optimization module*. Ablation analysis of the mentioned modules, in addition to the experimental results, are presented in the following sections.

4.1. Datasets

For the assessment of domain information transferability and target classification accuracy, we utilize two standard image datasets for domain adaptation, namely *Office-Home* and *Office-31*.

Office-31: The Office-31 (Saenko et al., 2010) dataset consists of RGB images sampled from three different domains: *Amazon (A), DSLR (D)*, and *Webcam (W)*. The images are categorized into 31 distinct classes. For establishing a *pda* setup, we follow the standard protocol adopted by Cao et al. (2018) in which the target dataset contains samples from 10 different categories. For a thorough assessment, we conducted an evaluation of the proposed model for multiple adaptation tasks on the following source–target combinations: A→D, A→W, D→A, D→W, W→A, and W→D.

Office-Home: The larger Office-Home dataset (Venkateswara et al., 2017) consists of RGB images from four distinct domains, namely Artistic (Ar), Clip Art (Cl), Product (Pr), and Real world (Rw). Along similar lines to the aforementioned *pda* setup, we follow the arrangement proposed by Cao et al. (2018) and build the source and target datasets with 65 and 25 distinct categories, respectively. For an in-depth evaluation, we design 12 different source → target adaptation tasks, namely: Ar→Cl, Ar→Pr, Ar→Rw, Cl→Ar, Cl→Pr, Cl→Rw, Pr→Ar, Pr→Cl, Pr→Rw, Rw→Ar, Rw→Cl, and Rw→Pr.

4.2. Implementation details

All models in the experiment are implemented with *PyTorch* using a *Nvidia 3090-Ti* GPU (24 GB memory). For adaptation tasks involving each source–target pair, we utilize Resnet-50 (He et al., 2016), pretrained on the Imagenet dataset (Deng et al., 2009) and fine-tuned on the source samples, as the backbone. The feature encoder $E(\cdot)$ is built on the backbone network by dropping the last dense layer and augmenting it with two fully connected layers, with a layer output size of 1,024, followed by a ReLU activation and a 0.1 dropout probability. The feature encoder generates latent representations of size 512, which are processed by the sample classifier $C(\cdot)$ and the nonparametric classifier $softmax(sim(\cdot))$. $C(\cdot)$ is a multilayer perceptron model and is built using two fully connected layers with hidden layer output dimensions of 512. The output dimension is set according to the number of source categories, i.e., 31 and 65 when training on Office-31 and Office-Home, respectively. The model is trained for 4500 epochs using ADAM (Kingma & Ba, 2014) optimizer, with a learning rate set to 1e-4. For a robust estimation of the output predictions of the nonparametric classifier and the threshold parameter τ, we compute their moving average over $N = 10$ epochs. During parameter sensitivity analysis, it is observed that N is nonsensitive and yields satisfactory results for values > 7. The empty set $\mathbb{D}_\tau$ receives its first update after the completion of the initial N epochs. Parameters λ and η, which regulate the learning from *adaptive complement entropy objective*, are set to 1, 8, and 0.3, 1 for Office-31 and Office-Home, respectively. Similarly, parameters α, β, and δ control the *intra-class compactness* and *inter-class separation objectives* and are set to 0.2, 0.9, 0.6, and 0.1, 0.9, 1.5 for Office-31, and Office-Home, respectively (check Table 1). For target classification during model evaluation, we report the outputs from the trainable classifier network $C(\cdot)$.

Table 1
Model parameter settings for evaluation

Dataset	λ	η	α	β	δ
Office-31	1	8	0.2	0.9	0.6
Office-Home	0.3	1	0.1	0.9	1.5

4.3. Comparison methods

We utilize all the samples present in $\mathbb{D}_s$ and $\mathbb{D}_t$ and the *target classification accuracy* metric to evaluate the proposed approach against the state-of-the-art models for closed-set and partial domain adaptation tasks: Domain Adversarial Neural Network (DANN) (Ganin et al., 2016), Adversarial Discriminative Domain Adaptation (ADDA) network (Tzeng et al., 2017), Partial Adversarial Domain Adaptation (PADA) (Cao et al., 2018a) and ETN (Cao et al., 2019), IWAN (Zhang et al., 2018), Deep Residual Correction Network (DRCN) (Li et al., 2020), and SAN (Cao et al., 2018b). Besides this, we report the classification accuracy on Resnet-50 (He et al., 2016), trained directly on the target data in a supervised fashion, to highlight the existence of negative transfer in the DANN model that is limited to solving a closed-set adaptation task.

4.4. Classification results

The target classification accuracies for the two benchmark datasets are summarized in Tables 2 and 3 (highest values highlighted in bold). The existence of the negative transfer problem is evident from the accuracy values of Resnet-50 (He et al., 2016) and DANN (Ganin et al., 2016), reported for tasks $A \to W, A \to D, D \to A$ in Table 2 and for $Ar \to Cl, Cl \to Pr, Pr \to Ar, Pr \to Cl$, and $Rw \to Cl$ in Table 3. Since the standard

Table 2
Classification accuracy (%) for partial domain adaptation tasks on Office-31 dataset (backbone: Resnet-50)

Method	A → D	A → W	D → A	D → W	W → A	W → D	Avg.
Resnet-50 (He et al., 2016)	83.44	75.59	83.92	96.27	84.97	98.09	87.05
DANN (Ganin et al., 2016)	81.53	73.56	82.78	96.27	86.12	98.73	86.50
ADDA (Tzeng et al., 2017)	83.41	75.67	83.62	95.38	84.25	99.85	87.03
PADA (Cao et al., 2018a)	82.17	86.54	92.69	99.32	95.41	**100.00**	92.69
IWAN (Zhang et al., 2018)	90.45	89.15	95.62	99.32	94.26	99.36	94.69
SAN (Cao et al., 2018b)	94.27	93.90	94.15	99.32	88.73	99.36	94.96
ETN (Cao et al., 2019)	95.03	94.52	**96.21**	**100.00**	94.64	**100.00**	96.73
Proposed model	**96.47**	**94.91**	96.02	**100.00**	**96.09**	**100.00**	**97.25**
W/o $\mathcal{L}_{ce}$	93.26	91.33	91.96	95.32	90.13	96.83	93.14
W/o $\mathcal{L}_{b\_class}$ & $\mathcal{L}_{w\_class}$	86.19	77.03	86.81	97.02	87.06	98.36	88.74
W/o adaptive consensus	88.97	86.72	88.91	97.61	90.53	98.87	91.94

Table 3
Classification accuracy (%) for partial domain adaptation tasks on Office-Home dataset (backbone: Resnet-50)

Method	Ar → Cl	Ar → Pr	Ar → Rw	Cl → Ar	Cl → Pr	Cl → Rw	Pr → Ar
Resnet-50 (He et al., 2016)	46.33	67.51	75.87	59.14	59.94	62.73	58.22
DANN (Ganin et al., 2016)	43.76	67.90	77.47	63.73	58.99	67.59	56.84
ADDA (Tzeng et al., 2017)	45.23	68.79	79.21	64.56	60.01	68.29	57.56
PADA (Cao et al., 2018a)	51.95	67.00	78.74	52.16	53.78	59.03	52.61
DRCN (Li et al., 2020)	54.00	76.40	83.00	62.10	64.50	71.00	70.80
IWAN (Zhang et al., 2018)	53.94	54.45	78.12	61.31	47.95	63.32	54.17
SAN (Cao et al., 2018b)	44.42	68.68	74.60	**67.49**	64.99	**77.80**	59.78
ETN (Cao et al., 2019)	59.24	77.03	79.54	62.92	65.73	75.01	**68.29**
Proposed model	**61.97**	**82.91**	**86.84**	63.97	**73.89**	74.86	67.04
W/o $\mathcal{L}_{ce}$	57.83	78.36	83.10	62.09	70.63	70.98	65.02
W/o $\mathcal{L}_{b\_class}$ & $\mathcal{L}_{w\_class}$	46.72	66.91	74.21	54.92	59.18	62.31	54.71
W/o adaptive consensus	53.96	74.73	77.89	58.17	64.92	67.27	61.17
Method	**Pr → Cl**	**Pr → Rw**	**Rw → Ar**	**Rw → Cl**	**Rw → Pr**	**Avg.**	
Resnet-50 (He et al., 2016)	41.79	74.88	67.40	48.18	74.17	61.35	
DANN (Ganin et al., 2016)	37.07	76.37	69.15	44.30	77.48	61.72	
ADDA (Tzeng et al., 2017)	38.89	77.45	70.28	45.23	78.32	62.82	
PADA (Cao et al., 2018a)	43.22	78.79	73.73	56.60	77.09	62.06	
DRCN (Li et al., 2020)	49.80	80.50	**77.50**	59.10	79.90	69.00	
IWAN (Zhang et al., 2018)	52.02	81.28	76.46	56.75	82.90	63.56	
SAN (Cao et al., 2018b)	44.72	80.07	72.18	50.21	78.66	65.30	
ETN (Cao et al., 2019)	55.37	**84.37**	75.72	57.66	**84.54**	70.45	
Proposed model	**55.73**	83.59	74.02	**61.07**	82.96	**72.40**	
W/o $\mathcal{L}_{ce}$	51.86	80.47	70.95	58.72	79.41	68.45	
W/o $\mathcal{L}_{b\_class}$ & $\mathcal{L}_{w\_class}$	41.18	73.02	65.92	45.81	71.69	59.71	
W/o adaptive consensus	49.72	77.29	70.22	56.59	75.36	65.61	

DANN model is limited to addressing a closed-set domain adaptation problem, it does not exhibit a filtration mechanism to diminish the effect of samples from classes private to the source domain ($Y_s - Y_t$). In fact, this unwanted transfer of superfluous information poisons the model to the extent that it performs worse than the regular Resnet-50 model trained in a supervised fashion on the relatively small target data. This bolsters the necessity for a domain adaptation model that is specifically tailored for the *pda* task.

Unlike other methods (Cao et al., 2018a, 2018b, 2019; Tzeng et al., 2017), which operate through class importance weight/sample weight estimation from the training initiation, we addressed the negative transfer problem by aiming to construct a well-organized "latent space" that separates the private class information from that of the overlapping classes. Our hypothesis is empirically justified in Tables 2 and 3, where we have bagged the highest classification accuracies in 5 out of 6 tasks and 6 out of 12 tasks, respectively. The proposed model has also yielded the highest average accuracies on the two datasets.

4.5. Parameter analysis

In this section, we analyze the trade-off parameters λ and η that control the complement entropy objective (presented in equations (2) and (12)). η regulates the contribution of $\mathcal{L}_{ce}$ that evenly distributes the complement classes' expected probability and reduces the likelihood of an incorrect category to having an adequately high-prediction probability. In contrast, λ regulates the level of emphasis on samples based on classification confidence; it pays more attention to uncertain samples that yield smaller cross-entropy loss. Setting λ to 0 places equal emphasis on both the highly confident and the

Table 4
Average accuracy (%) for η values

$\eta \rightarrow$	0.0	0.5	1.0	2	4	6	8	10
Office-31	93.14	96.51	96.83	97.04	97.13	97.19	97.25	97.17
Office-home	68.45	71.82	72.40	72.16	71.98	71.69	71.53	71.48

Table 5
Average accuracy (%) for λ values

$\lambda \rightarrow$	0.0	0.1	0.3	0.5	0.7	0.9	1
Office-31	96.31	96.63	96.71	96.98	97.07	97.19	97.25
Office-home	71.48	72.29	72.40	72.33	72.27	72.16	71.98

less confident samples. Table 4 and 5 report the mean accuracy for different values of parameters λ and η on Office-31 and Office-Home datasets. As observed, the accuracy values vary within an acceptable range for different nonzero values of these parameters, indicating that the suggested approach is less sensitive to them.

4.6. Ablation analysis

To test the significance of the proposed components in the network, we conducted an ablation study by suppressing them one at a time. The analysis thus conducted is described below:

• **W/o** $\mathcal{L}_{ce}$: The complement entropy objective captured by $\mathcal{L}_{ce}$ removes uncertainty from the classification process by enforcing the incorrect classes to have low-prediction probabilities that are distributed uniformly. To evaluate its contribution, we prohibit $\mathcal{L}_{ce}$ from contributing to the overall loss objective (by setting η to 0).

• **W/o** $\mathcal{L}_{b\_class}$ **and** $\mathcal{L}_{w\_class}$: We aim for a category-level arrangement in the latent space that is "well-organized". Instead of regulating the training process by assigning importance weights to various classes, our objective is to optimize the intra-/inter-class distributions such that samples of the same class are bundled close to each other while that from distinct classes are well separated in the latent space, regardless of the sample domains. To realize this, we adopted $\mathcal{L}_{b\_class}$ and $\mathcal{L}_{w\_class}$ into the overall objective. To validate their contribution, we suppressed them by setting α, β, and δ to 0.

• **W/o adaptive consensus**: In the proposed method, we leverage the pseudo-labeling technique for target supervision. Since the majority of the pseudo-labels generated at the initial stages of model training are unreliable and can divert the learning process from its intended course, we mitigate this by utilizing a subset of the target samples that yield prediction probabilities higher than an estimated adaptive threshold parameter, aggregated over a steady number of latest iterations. To assess its effectiveness, we suppress this adaptive consensus-based pseudo-label generation technique, i.e., instead of an elitist approach, we apply supervision on all the target samples using pseudo-labels generated by the nonparametric classifier. The moving average estimate is, however, continued to obtain a robust estimate of the classification probabilities. As a result, the dataset $\mathbb{D}_\tau$ is replaced with $\mathbb{D}_t$ in equations (1) and (2) in $\mathcal{L}_{class}$ and $\mathcal{L}_{ce}$, respectively.

The classification accuracies (as presented in Tables 2 and 3) obtained from different training strategies, as mentioned above, demonstrate the effectiveness of these modules and our motivation to incorporate them in the proposed framework. From the reported values, it is observed that the objectives centering around class distribution optimization (maximizing intra-class compactness and inter-class separation) exhibit the highest influence, followed by the component for adaptive consensus-based pseudo-label generation and refinement. The complement entropy objective contributes significantly, as its suppression has adversely affected model performance on all the tasks designed with the two datasets.

5. Conclusion

In this work, we propose a novel classification framework for a partial domain adaptation setup that operates by aligning domain and class distributions through (a) domain-agnostic intra- and inter-class distance optimization, (b) suppression of uncertainty in classifier predictions, and (c) adaptive consensus-based target supervision. Unlike recent approaches that utilize class/sample reweighting schemes to emphasize the learning on sample categories common to both domains, we aim to restructure the sample arrangement in the latent space where data points from identical classes are enforced to reside in close proximity while that from distinct classes are well separated, in a domain-neutral fashion. The proposed model addresses a challenging problem of uncertainty propagation into the classifier by utilizing a complement entropy objective that enforces the incorrect classes to have uniformly distributed low-prediction probabilities. For target supervision, we engineered a robust technique for pseudo-label generation by utilizing a nonparametric classifier over the target samples and adaptively selecting a portion of the target data that yield prediction probabilities greater than an estimated adaptive threshold parameter, aggregated over a steady number of iterations. For a thorough assessment, we conducted experiments on two benchmark datasets to evaluate the proposed model under a variety of *pda* tasks against the state-of-the-art models addressing closed-set and partial domain adaptation problems. To confirm the significance of the highlighted modules and validate their contribution to the proposed framework, we conducted an ablation analysis. The experimental results demonstrate the proposed model's efficacy, where it has outperformed the state-of-the-art average performances in all the challenging tasks designed with the two benchmarks.

Conflicts of Interest

The authors declare that they have no conflicts of interest to this work.

References

Cao, Z., *et al.* (2019). Learning to transfer examples for partial domain adaptation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 2985–2994).

Cao, Z., *et al.* (2018a). Partial adversarial domain adaptation In *Proceedings of the European conference on computer vision (ECCV)* (pp. 135–150).

Cao, Z., *et al.* (2018b). Partial transfer learning with selective adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 2724–2732).

Chen, H.-Y., *et al.* (2018). Complement objective training. *In International conference on learning representations*.

Chen, H.-Y., *et al.* (2019). Complement objective training. ArXiv preprint arXiv:1903.01182.

Chen, M., Weinberger, K. Q., & Blitzer, J. (2011). Cotraining for domain adaptation. In *Advances in neural information processing systems*, vol. 24.

Choudhuri, S., Venkateswara, H., & Sen, A. (2022). Coupling adversarial learning with selective voting strategy for distribution alignment in partial domain adaptation. *Journal of Computational and Cognitive Engineering*, *1*, 181–186.

Choudhuri, S., *et al.* (2018). Object localization on natural scenes: A survey. *International Journal of Pattern Recognition and Artificial Intelligence*, *32*, 1855001. https://dx.doi.org/10.1142/S0218001418550017

Choudhuri, S., *et al.* (2020). Partial domain adaptation using selective representation learning for class–weight computation. In 2020 *54th Asilomar Conference on Signals, Systems, and Computers*. IEEE (pp. 289–293).

Dang, Q., *et al.* (2019). Deep learning based 2D human pose estimation: A survey. *Tsinghua Science and Technology*, *24*, 663–676. https://dx.doi.org/10.26599/TST.2018.9010100

Deng, J., *et al.* (2009). Imagenet: A large–scale hierarchical image database. In 2009 IEEE *conference on computer vision and pattern recognition*. IEEE (pp. 248–255).

Ganin, Y. & Lempitsky, V. (2015). Unsupervised domain adaptation by backpropagation. In *International conference on machine learning*. PMLR (pp. 1180–1189).

Ganin, Y., *et al.* (2016). Domain–adversarial training of neural networks. *The Journal of Machine Learning Research*, *17*, 2096–2030.

Ghifary, M., *et al.* (2016). Scatter component analysis: A unified framework for domain adaptation and domain generalization. *IEEE Transactions on Pattern Analysis and Machine Intelligence 39*, 1414–1430. https://dx.doi.org/10.1109/TPAMI.2016.2599532

Guo, Z., *et al.* (2021). A survey on deep learning based approaches for scene understanding in autonomous driving. *Electronics*, *10*, 471.

He, K., *et al.* (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 770–778).

Hoffman, J., *et al.* (2014). LSDA: Large scale detection through adaptation. In *Advances in neural information processing systems*, vol. 27.

Huang, J., *et al.* (2006). Correcting sample selection bias by unlabeled data. In *Advances in neural information processing systems*, vol. 19.

Jing, T., Xia, H., & Ding, Z. (2020). Adaptively–accumulated knowledge transfer for partial domain adaptation. In *Proceedings of the 28th ACM International Conference on Multimedia* (pp. 1606–1614).

Kingma, D. P. & Ba, J. (2014). Adam: A method for stochastic optimization. *ArXiv preprint arXiv:1412.6980*.

Kumar, A., *et al.* (2018). Co–regularized alignment for unsupervised domain adaptation". In: *Advances in neural information processing systems*, vol. 31.

Li, S., *et al.* (2020). Deep residual correction network for partial domain adaptation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *43*, 2329–2344. https://dx.doi.org/10.1109/TPAMI.2020.2964173

Li, S., *et al.* (2019). Joint adversarial domain adaptation. In *Proceedings of the 27th ACM International Conference on Multimedia* (pp. 729–737).

Liang, J., *et al.* (2020). A balanced and uncertainty–aware approach for partial domain adaptation. In *European Conference on Computer Vision*. Springer (pp. 123–140).

Liu, X., *et al.* (2021). A review of deep–learning–based medical image segmentation methods. *Sustainability*, *13*, 1224. https://dx.doi.org/10.3390/su13031224

Long, M., *et al.* (2013). Transfer feature learning with joint distribution adaptation. In *Proceedings of the IEEE international conference on computer vision* (pp. 2200–2207).

Miyato, T., *et al.* (2018). Virtual adversarial training: A regularization method for supervised and semi–supervised learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *41*, 1979–1993. https://dx.doi.org/10.1109/TPAMI.2018.2858821

Oquab, M., *et al.* (2014). Learning and transferring mid–level image representations using convolutional neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 1717–1724).

Pan, S. J., & Yang, Q. (2010). A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering*, *22*, 1345–1359. https://dx.doi.org/10.1109/TKDE.2009.191

Pan, S. J., *et al.* (2010). Domain adaptation via transfer component analysis. *IEEE Transactions on Neural Networks*, *22*, 199–210. https://dx.doi.org/10.1109/TNN.2010.2091281

Saenko, K., *et al.* (2010). Adapting visual category models to new domains. In *European conference on computer vision*. Springer (pp. 213–226).

Shu, R., *et al.* (2018). A dirt–t approach to unsupervised domain adaptation. *ArXiv preprint arXiv:1802.08735*.

Tzeng, E., *et al.* (2017). Adversarial discriminative domain adaptation. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 7167–7176).

Tzeng, E., *et al.* (2014). Deep domain confusion: Maximizing for domain invariance. *ArXiv preprint arXiv:1412.3474*.

Venkateswara, H., *et al.* (2017). Deep hashing network for unsupervised domain adaptation. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 5018–5027).

Wang, W., *et al.* (2019). Development of convolutional neural network and its application in image classification: a survey. *Optical Engineering*, *58*, 040901.

Yosinski, J., *et al.* (2014). How transferable are features in deep neural networks? In *Advances in neural information processing systems*, vol. 27.

Zhang, J., *et al.* (2018). Importance weighted adversarial nets for partial domain adaptation". In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 8156–8164).

Zhang, L., *et al.* (2018). Unsupervised domain adaptation using robust class–wise matching. *IEEE Transactions on Circuits and Systems for Video Technology*, *29*, 1339–1349. https://dx.doi.org/10.1109/TCSVT.2018.2842206

How to Cite: Choudhuri, S., Adeniye, S., & Sen, A. (2023). Distribution Alignment Using Complement Entropy Objective and Adaptive Consensus-Based Label Refinement for Partial Domain Adaptation. *Artificial Intelligence and Applications* 1(1), 43–51, https://doi.org/10.47852/bonviewAIA2202524