



HAL
open science

Semi Supervised Heterogeneous Domain Adaptation via Disentanglement and Pseudo-Labeling

Cassio F. Dantas, Raffaele Gaetano, Dino Ienco

► **To cite this version:**

Cassio F. Dantas, Raffaele Gaetano, Dino Ienco. Semi Supervised Heterogeneous Domain Adaptation via Disentanglement and Pseudo-Labeling. ECML PKDD 2024 - European Conference on Machine Learning and Knowledge Discovery in Databases, Sep 2024, Vilnius, Lithuania. pp.440-456, 10.1007/978-3-031-70352-2_26 . hal-04617917

HAL Id: hal-04617917

<https://hal.science/hal-04617917v1>

Submitted on 19 Jun 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Semi Supervised Heterogeneous Domain Adaptation via Disentanglement and Pseudo-Labeling

Cassio F. Dantas^{1,3}, Raffaele Gaetano^{2,3}, and Dino Ienco^{1,3,4}

¹ INRAE, UMR TETIS, Univ. Montpellier, Montpellier, France
{cassio.fraga-dantas,dino.ienco}@inrae.fr

² CIRAD, UMR TETIS, Univ. Montpellier, Montpellier, France
raffaele.gaetano@cirad.fr

³ INRIA, Univ. Montpellier, Montpellier, France

⁴ LIRMM, Univ. Montpellier, CNRS, Montpellier, France

Abstract. Semi-supervised domain adaptation methods leverage information from a source labelled domain with the goal of generalizing over a scarcely labelled target domain. While this setting already poses challenges due to potential distribution shifts between domains, an even more complex scenario arises when source and target data differs in modality representation (e.g. they are acquired by sensors with different characteristics). For instance, in remote sensing, images may be collected via various acquisition modes (e.g. optical or radar), different spectral characteristics (e.g. RGB or multi-spectral) and spatial resolutions. Such a setting is denoted as Semi-Supervised Heterogeneous Domain Adaptation (SSHDA) and it exhibits an even more severe distribution shift due to modality heterogeneity across domains.

To cope with the challenging SSHDA setting, here we introduce *SHeDD* (Semi-supervised Heterogeneous Domain Adaptation via Disentanglement) an end-to-end neural framework tailored to learning a target domain classifier by leveraging both labelled and unlabelled data from heterogeneous data sources. *SHeDD* is designed to effectively disentangle domain-invariant representations, relevant for the downstream task, from domain-specific information, that can hinder the cross-modality transfer. Additionally, *SHeDD* adopts an augmentation-based consistency regularization mechanism that takes advantages of reliable pseudo-labels on the unlabelled target samples to further boost its generalization ability on the target domain. Empirical evaluations on two remote sensing benchmarks, encompassing heterogeneous data in terms of acquisition modes and spectral/spatial resolutions, demonstrate the quality of *SHeDD* compared to both baseline and state-of-the-art competing approaches. Our code is publicly available [here](#).

Keywords: Domain Adaptation · Heterogeneous data · Feature disentanglement · Pseudo-labeling · Consistency regularization.

1 Introduction

When it comes to real-world applications of machine learning, disposing of a vast amount of labelled samples remains a major issue in many domains, especially those featured by costly and time-consuming labelling processes. Consequently, make value of already available data, covering similar downstream tasks, is of paramount importance to enhance the classification performances on target domains where labelled data are scarce. Nonetheless, this process is not straightforward due to potential differences or shifts in their underlying data distributions between a rich source labelled domain and the target one [26]. To cope with data distribution shifts between source and target domains, Domain Adaptation (DA) techniques have been proposed [21]. The main objective of this family of machine learning methods is to learn a classification model across different domains, generally sharing the same set of classes, with the aim to transfer information from a source to a target one.

Many research efforts have focused on addressing situations wherein the target domain lacks completely of associated labels, while only the source domain disposes of labelled information [21]. This scenario is commonly referred as Un-supervised Domain Adaptation (UDA). However, a more practical assumption for real-world applications is to have access to a small amount of labelled information from the target domain, enabling the simultaneous exploitation of abundant labelled samples from the source domain and limited labelled samples from the target domain. Such a setting is generally termed as Semi-Supervised Domain Adaptation [15], and existing literature has highlighted that directly using UDA approaches fails to exploit the label information associated with the target domain, thus requiring tailored solutions for this setting [16].

Nevertheless, most of the aforementioned research strategies rely on the strong assumption that data coming from source and target domains share a similar (homogeneous) data modality representation. However, in real-world applications data can be collected by means of heterogeneous sensors, as it is the case for remote sensing imagery exhibiting differences in acquisition modes (e.g. optical and radar), spectral characteristics (RGB, multi-/hyper-spectral), and spatial resolution. Consequently, it is increasingly common to encounter label-abundant source domains and label-scarce target domains that are heterogeneous in terms of data modality representation, further exacerbating potential data distribution shifts between domains. To address this challenging scenario, Semi-Supervised Heterogeneous Domain Adaptation (SSHDA) methods are gaining increasing attention within the research community [2]. However, the majority of the proposed approaches rely on pre-trained deep learning models that are only available for standard modalities (e.g. RGB imagery, text data), limiting their applicability in scenarios involving non-standard sensor data, such as those used in the medical [3] and remote sensing [13] fields.

With the aim to address the challenging SSHDA setting, in this research work we introduce a new end-to-end deep learning framework especially tailored to learn a target domain classifier by leveraging both labelled and unlabelled data from heterogeneous data sources. Our framework, *SHeDD* (Semi-supervised

Heterogeneous Domain Adaptation via Disentanglement), tackles data modality heterogeneity by extracting, via a feature disentanglement approach, domain-invariant representations, relevant for the downstream task, and domain-specific information, that can prevent cross-modality transfer. To this end, invariant and domain specific features are enforced to be orthogonal to each other with the latter carrying domain-discriminant information. Furthermore, *SHeDD* harnesses unlabelled target data to enhance its generalization ability, aiming to transfer discriminative information from the labelled source domain to the scarcely-labelled target domain. This last point is achieved via consistency learning where confident pseudo-labels are derived on the target domain by the classification model and subsequently exploited in the training process. Empirical evaluations on two remote sensing benchmarks, encompassing heterogeneous data domains in terms of acquisition modes and spectral/spatial characteristics/resolutions, clearly demonstrate the quality of *SHeDD* compared to both baseline and state-of-the-art competing methods.

This paper is organized as follows: related works are discussed in Section 2; the proposed method is described Section 3; experimental results are presented and discussed in Section 4, followed by concluding remarks in Section 5.

2 Related Works

Domain adaption [21] (DA) methods belong to the family of transfer learning approaches [26] which have the main objective to transfer a model trained on a labelled source domain to a target domain. When the target domain is completely unlabelled, Unsupervised Domain Adaptation (UDA) strategies are designed in order to align domains through data transformation and/or extract domain-invariant features to reduce the distribution gap between the labelled source and the unlabelled target domain [10].

When, for the target domain, a limited amount of labelled samples are available, Semi Supervised Domain Adaptation (SSDA) strategies have been proposed to combine both labelled (source and target) with unlabelled (target) information [16,14,8,9,5,22,23,24]. In [16], a framework based on Minimax Entropy is introduced to exploit the available target supervision. The same research work clearly illustrates that directly use UDA methods performs poorly when small amount of labeled samples are accessible from the target domain, thus emphasizing the need to design specialized approaches for the SSDA setting. In [14], an adversarial learning paradigm is leveraged in order to obtain two contradictory classifiers (source and target), enforcing well-scattered source features and compact target features respectively. A slightly different approach is proposed in [8], where a cross-domain adaptive clustering algorithm is presented to achieve cluster-wise feature alignment across domains, still employing an adversarial learning strategy. In [5], additional adversarial examples are introduced to fill the gap between domains and model robustness. The work in [23] decomposes SSDA into an SSL problem within the target domain coupled with an inter-domain UDA problem, then optimizes both tasks simultaneously using co-

training. Moving away from the adversarial paradigm, [17] proposes a contrastive learning framework operating both at a class level to reduce inter-domain gap and at the instance level with strong augmentations to minimize intra-domain discrepancy. Mitigating discrepancy within the target domain is also the main goal in [6], which proposes a feature alignment approach for achieving it.

Despite the effectiveness demonstrated by these methods, they are especially designed for managing homogeneous domains since they capitalize on the fact that source and the target domains share a similar modality representation (e.g., both involving RGB images). Therefore, the direct extension of these approaches to handle a heterogeneous setting, where source and target data differ in modality representation, is challenging.

In recent years, research efforts have been devoted towards addressing DA in an heterogeneous setting [1]. These efforts primarily focus on aligning the source and target domains through heterogeneous feature transformations. For instance, [25] learns feature transformations to map source and target data into a common latent space, where both marginal and class-conditional distributions are matched. In addition, self-training via pseudo labelling is used to update the target label set. Another approach proposed in [2] introduces a joint mean embedding alignment method where a neural network based approach aligns source and target data distribution via domain discrepancy minimization. However, these methods rely on features derived either through hand-crafted processes or from modality/domain specific pre-trained models (e.g. RGB pre-trained model) thus lacking end-to-end behaviour. Such reliance can prevent their applicability in scenarios involving data beyond the standard RGB (three channels) format. Recently, [12] has introduced an end-to-end SSHDA method, addressing the aforementioned limitations. This method adopts per-domain encoders sequentially connected to a shared backbone, with a classification head used for the final decision. During training, the neural network is optimized for simultaneously classifying and align the embedding representations coming from the different heterogeneous domains with standard cross entropy and domain critic discrimination based on wasserstein distance, respectively.

In this work we propose a different framework for SSHDA based on feature *disentanglement*, intended as the capacity of a network to identify domain-invariant representations, relevant for the downstream task, by explicitly seeking in parallel to isolate domain-specific information which may hinder the cross-modality transfer.

3 Proposed Method

The proposed architecture, summarized in Figure 1, is given by two independent encoder branches with specialized backbones (one dedicated to the source data modality and another to the target data), followed by two parallel classifiers (a task classifier and a domain classifier).

A given input data \mathbf{x} is firstly encoded by its matching backbone and the obtained embedding vector $\mathbf{z} = g(\mathbf{x}) \in \mathbb{R}^{2D}$ is then split into two equal parts:

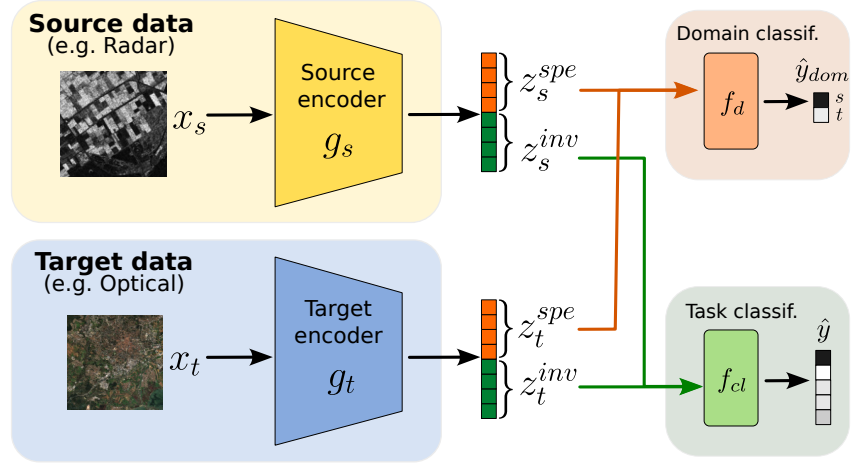


Fig. 1. Schematic view of the proposed method architecture with a separate encoder for each of the data modalities (source and target). Feature disentanglement enables domain-specific and domain-invariant information to be encoded separately into each half of the generated embedding vectors (depicted in orange and green respectively). The domain-invariant information (z^{inv}) is used by the task classifier, while the domain-specific portion of the embedding vector (z^{spe}) is used by the domain classifier. At inference time, only the bottom part of the architecture is used, the top part being instrumental in the training stage to enable the feature disentanglement procedure.

$\mathbf{z}^{spe} \in \mathbb{R}^D$ and $\mathbf{z}^{inv} \in \mathbb{R}^D$. While the former vector is fed into the domain classifier f_d , a binary classifier that tries to predict from which branch (source or target) the sample originates, the latter one is sent to the task classifier f_{cl} that outputs class probabilities $\hat{y} = f_{cl}(\mathbf{z}^{inv}) \in \mathbb{R}^C$ for the C existing classes.

At training time, guided by the losses described in Section 3.2, the weights of these four modules—source and target encoders, task and domain classifiers—are optimized on the available supporting data composed of the following sets:

$$\begin{aligned} S &:= \{(\mathbf{x}_s, y_s)^{(i)}\}_{i=1}^{N_s} && \text{labelled source data.} \\ T &:= \{(\mathbf{x}_t, y_t)^{(i)}\}_{i=1}^{N_t} && \text{labelled target data.} \\ U &:= \{\mathbf{x}_u^{(i)}\}_{i=1}^{N_u} && \text{unlabelled target data.} \end{aligned}$$

From each unlabeled target sample \mathbf{x}_u , we generate a corresponding augmented counterpart (see details in Section 3.2) denoted $\mathbf{x}_{\hat{u}}$ that form the set below:

$$\hat{U} := \{\mathbf{x}_{\hat{u}}^{(i)}\}_{i=1}^{N_u} \quad \text{augmented unlabelled target data.}$$

where we denote N_s , N_t and N_u the corresponding dataset sizes.

At inference time, only the target encoder is retained. Similarly, only the task classifier is required. The two dropped modules, however, are crucial as supporting elements during training in order to fully guide the network’s ability to effectively disentangle domain-invariant from domain-specific information.

This ability, acquired during training and carried over to inference time in the two retained modules, helps enhancing the generalization capabilities of the final network.

3.1 Training losses

In case of a labeled training sample, from either source or target domain, the output of the task classifier $f_{cl}(\mathbf{z}^{inv})$ is compared to its ground-truth annotation y in the following cross-entropy classification loss:

$$\mathcal{L}_{cl} = \text{CE}(f_{cl}(\mathbf{z}^{inv}), y). \quad (1)$$

Because the provenance domain $y_{dom} \in \{s, t\}$ of any given data sample is always known (even for unlabeled target samples), the domain classifier prediction $f_d(\mathbf{z}^{spe})$ can be systematically taken into account in the following cross-entropy loss:

$$\mathcal{L}_{dom} = \text{CE}(f_d(\mathbf{z}^{spe}), y_{dom}). \quad (2)$$

To further enforce disentanglement between domain-invariant and domain-specific information, we enforce orthogonality between the two embedding types for any given input sample (source and target, labelled and unlabelled):

$$\mathcal{L}_{\perp} = \frac{\langle \mathbf{z}^{inv}, \mathbf{z}^{spe} \rangle}{\|\mathbf{z}^{inv}\|_2 \|\mathbf{z}^{spe}\|_2}. \quad (3)$$

To fully exploit the available target unlabelled data (set U), for each sample $\mathbf{x}_u \in U$ we first generate an associated augmented sample $\mathbf{x}_{\hat{u}} = \text{Augment}(\mathbf{x}_u)$ (see details in section 3.2) and then employ an unsupervised loss *à la* FixMatch [18] that enforces consistency between predictions obtained from the unlabeled sample \mathbf{x}_u and its augmentation $\mathbf{x}_{\hat{u}}$ via pseudo-labelling procedure:

$$\mathcal{L}_{pl} = m_u^\tau \text{CE}(f(\mathbf{z}_{\hat{u}}^{inv}), y_{\hat{u}}) \quad (4)$$

where pseudo-labels $y_{\hat{u}} := \text{argmax}(f_{cl}(\mathbf{z}_{\hat{u}}^{inv}))$, given by the classifier predictions on the unlabelled target data $x_u \in U$, are used as ground-truth for the corresponding augmentation $x_{\hat{u}}$. Only a subset of the pseudo-labels (those with higher confidence) are retained. This is expressed through the multiplying binary factor

$$m_u^\tau := \mathbb{1}(\max(f_{cl}(\mathbf{z}_u^{inv})) > \tau) \quad (5)$$

where $\tau \in [0, 1]$ is a scalar hyper-parameter denoting the confidence threshold. and $\mathbb{1}(\text{condition})$ denotes the indicator function, which is equal to 1 if *condition* holds and 0 otherwise.

3.2 Training procedure

The proposed training scheme is summarized in Figure 2, where we show the different input data paths through the network during training as well as the

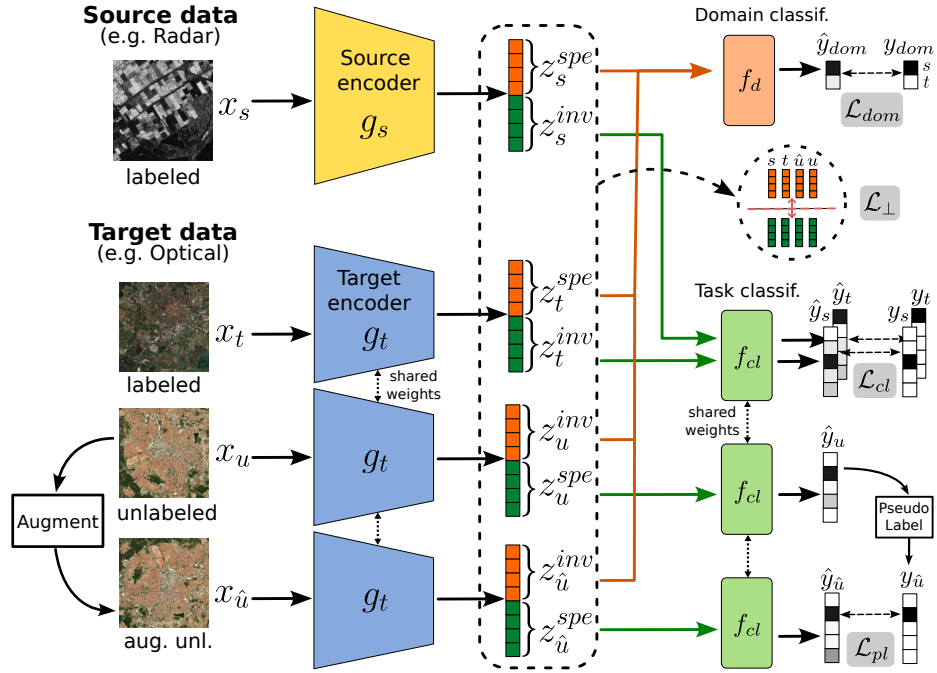


Fig. 2. Schematic view of the data flow during the training phase. The four proposed loss terms (framed in grey) are illustrated with their corresponding inputs.

inputs used by each of the four proposed losses. A more detailed and formalized description of the training procedure is given in Algorithm 1.

For each epoch, we go through the source dataset sequentially (as it is usually the dataset with the highest number of samples $N_s > N_u > N_t$). This is done by batches in practice, even if in Algorithm 1 we illustrate the sample-wise case (unitary batch) for simplicity⁵. At each iteration, we sample uniformly at random the same number of samples (batch size) from the set labeled and unlabeled target data —lines 4 and 5. Each sample is then passed through their matching encoder at lines 6–9 (note that the target encoder g_t is used not only for the labeled target samples \mathbf{x}_t with matching subscript, but also for the unlabeled samples \mathbf{x}_u and $\mathbf{x}_{\hat{u}}$). Finally, in lines 10–16, each of the loss terms defined in the previous section are computed with respect to the all relevant input data and, subsequently (line 17), backpropagated through the entire architecture to update its composing modules (g_s, g_t, f_{cl}, f_d) weights.

For convenience, we introduce superscripts on a loss term, say \mathcal{L}^V , to specify its application on input data coming from a certain dataset $V \in \{S, T, U, \hat{U}\}$ (or several datasets in case of multiple superscripts). For instance, we denote by $\mathcal{L}_{cl}^{S,T}$

⁵ The generic minibatch version of Algorithm 1 is obtained simply by additionally averaging each of the loss terms over the batch dimension.

Algorithm 1 *SHeDD* Train procedure**Require:** Datasets S, T, U ; Pseudo-labeling threshold τ .

```

1: for epoch  $\in \{1, \dots, N_{ep}\}$  do
2:   for all  $(\mathbf{x}_s, y_s) \in S$  do
3:      $(\mathbf{x}_t, y_t) \sim \mathcal{U}(T)$ 
4:      $\mathbf{x}_u \sim \mathcal{U}(U)$ 
5:      $x_{\hat{u}} = \text{Augment}(x_u)$ 
6:      $\mathbf{z}_s^{inv}, \mathbf{z}_s^{spe} = g_s(x_s)$ 
7:      $\mathbf{z}_t^{inv}, \mathbf{z}_t^{spe} = g_t(x_t)$ 
8:      $\mathbf{z}_u^{inv}, \mathbf{z}_u^{spe} = g_t(x_u)$ 
9:      $\mathbf{z}_{\hat{u}}^{inv}, \mathbf{z}_{\hat{u}}^{spe} = g_t(x_{\hat{u}})$ 
10:     $\mathcal{L}_{cl}^{S,T} = \frac{1}{2} \sum_{v \in \{s,t\}} CE(f_{cl}(\mathbf{z}_v^{inv}), y_v)$ 
11:     $\mathcal{L}_{dom}^{S,T} = \frac{1}{2} \sum_{v \in \{s,t\}} CE(f_d(\mathbf{z}_v^{spe}), v)$ 
12:     $\mathcal{L}_{dom}^{U,\hat{U}} = \frac{1}{2} \sum_{v \in \{u,\hat{u}\}} CE(f_d(\mathbf{z}_v^{spe}), t)$ 
13:     $\mathcal{L}_{\perp}^{S,T} = \frac{1}{2} \sum_{v \in \{s,t\}} \frac{\langle \mathbf{z}_v^{inv}, \mathbf{z}_v^{spe} \rangle}{\|\mathbf{z}_v^{inv}\|_2 \|\mathbf{z}_v^{spe}\|_2}$ 
14:     $\mathcal{L}_{\perp}^{U,\hat{U}} = \frac{1}{2} \sum_{v \in \{u,\hat{u}\}} \frac{\langle \mathbf{z}_v^{inv}, \mathbf{z}_v^{spe} \rangle}{\|\mathbf{z}_v^{inv}\|_2 \|\mathbf{z}_v^{spe}\|_2}$ 
15:     $y_{\hat{u}}, m_u^{\tau} = \text{PseudoLabel}(f_{cl}(\mathbf{z}_u^{inv}), \tau)$  // cf. equations (4) and (5)
16:     $\mathcal{L}_{pl}^{\hat{U}} = m_u^{\tau} \cdot CE(f_{cl}(\mathbf{z}_{\hat{u}}^{inv}), y_{\hat{u}})$ 
17:    Update weights of  $(g_s, g_t, f_{cl}, f_d)$  by back-propagating the loss:
       $\mathcal{L}_{cl}^{S,T} + \mathcal{L}_{dom}^{S,T} + \mathcal{L}_{dom}^{U,\hat{U}} + \mathcal{L}_{\perp}^{S,T} + \mathcal{L}_{\perp}^{U,\hat{U}} + \mathcal{L}_{pl}^{\hat{U}}$ 
18:   end for
19: end for
20: return  $g_T, f_{cl}$ 

```

the classification loss defined in eq. (1) applied on (and averaged over) samples from labeled source and target datasets. This notation has the merit of making more explicit to which dataset each loss applies and will prove particularly useful for our ablation analysis in Table 6.

Hence, while the classification loss \mathcal{L}_{cl} naturally applies only to labelled data (S, T), the domain classification \mathcal{L}_{dom} and orthogonality \mathcal{L}_{\perp} losses can be evaluated for both labeled (S, T) and unlabelled data (U, \hat{U}). Finally, the Fix-Match loss \mathcal{L}_{pl} applies to the augmented unsupervised data (\hat{U}) while leveraging pseudo-labels obtained for the corresponding non-augmented samples (in U). These multiple data paths involved in the proposed training scheme are summarized in Figure 2. In the figure, we replicate the encoder and classification modules to properly outline each separate data flow, but we emphasize that these modules are characterized by an unique set of shared weights.

Data augmentation: The employed augmentation operation $\text{Augment}(\cdot)$ (line 5 in Algorithm 1) consists of a series of possible transformations with 50% of occurrence probability each, among the following: horizontal flip; vertical flip;

rotation with random angle on the set $\{0^\circ, 90^\circ, 180^\circ, 270^\circ\}$; color jitter (random changes in the image brightness, contrast, saturation and hue)⁶.

4 Experiments

In order to assess the performance of *SHeDD*, we consider two different benchmarks covering heterogeneous data coming from the remote sensing field.

Benchmark	Volume	Modality	Spatial Res.	# Classes
<i>RESISC45-Euro</i>	$5\,600 \times 3 \times 256 \times 256$	RGB	0.2m–30m	8
	$24\,000 \times 13 \times 64 \times 64$	MS	10m	
<i>EuroSat-MS-SAR</i>	$27\,000 \times 13 \times 64 \times 64$	MS	10m	10
	$27\,000 \times 2 \times 64 \times 64$	SAR	10m	

Table 1. Benchmark statistics and description. Each benchmark covers two heterogeneous domains. *EuroSat-MS-SAR* involves MS and SAR images, both with a spatial resolution of 10m, for a classification task with 10 classes. *RESISC45-Euro* includes RGB and MS images, with varying spatial resolutions, for a classification task with 8 classes. The **Volume** column reports per-domain statistics in the format (# images) \times (# channels) \times (image height) \times (image width).

As our first dataset, we adopt the *RESISC45-Euro* benchmark previously introduced in [12]. This dataset contains 5 600 RGB images at different spatial resolutions and 24 000 multispectral (MS) images, with 13 channels, spanning eight different land cover classes. Here, the heterogeneity is related to domains covering imagery with different spatial and spectral resolutions. As our second dataset, we use the *EuroSat-MS-SAR* benchmark [20]. This dataset contains 54 000 pairs of MS and synthetic aperture radar (SAR) images, with 13 and 2 channels respectively. With the aim to avoid possible data biases and spurious correlations, for each sample of the dataset we only retain one of the two modalities. This leads to a benchmark including 27 000 MS and 27 000 SAR unaligned images over the set of ten land cover classes. Here, the heterogeneity corresponds to imagery collected via different acquisition modes (optical and radar). Details about benchmarks are reported in Table 1. For each benchmark we set up two transfer tasks where each transfer task is denoted as $(\mathcal{D}_s \rightarrow \mathcal{D}_t)$ where the right arrow indicates the transfer direction from the fully labelled source domain (\mathcal{D}_s) to the scarce labelled target domain (\mathcal{D}_t).

Considering the competing approaches, we include in our experimental evaluation a fully supervised baseline that only exploits available target labelled data, referred as *Target Only*. As a state-of-the-art semi-supervised framework that exploits both labelled and unlabelled target samples in order to leverage

⁶ For this transformation, we used the PyTorch implementation `torchvision.transforms.ColorJitter()` with default parameters.

the full amount of available target data, here we adopt the well-known *FixMatch* framework [18]. Finally, according to SSHDA literature, we include the *SS-HIDA* approach recently introduced in [12].

For all the competing approaches, as well as our proposed *SHeDD*, to set up a fair comparison, we adopt the same backbone architecture, ResNet-18 [4]. In the particular case of our proposed *SHeDD*, the final fully-connected layer (with softmax activation) of ResNet-18 is employed as our task classifier module and the same structure is used for the domain classifier. For *FixMatch* and *SHeDD* we fix the pseudo-labeling threshold τ to 0.95 and we use as weak augmentation the identity function and as strong augmentation a random combination of geometrical (flipping and rotation) and radiometric (color jitter) transformations. For the *SS-HIDA*, according to the original work, we used half of the backbone network as specific per-domain encoder and the rest of the backbone as shared encoder. For all the competing approaches we adopt a number of training epochs equal to 300, a batch size of 128, AdamW [11] as parameter optimizer with a learning rate of 10^{-4} . Additionally, based on recent literature [7], for all the methods we adopted an exponential moving average (EMA) of the weight parameters, with momentum equals to 0.95, since we experimentally observed that all the approaches took advantage from it.

For the experimental assessment, we set up two different transfer tasks for each benchmark, considering each of the available domains firstly as source and then as target. While for the source domain all the available data are labelled, for the target domain we varied the amount of available supervision, ranging in the set $\{25, 50, 100, 200\}$ samples per class. This means that, for instance, if the supervision value is equal to 25, then 25 labelled samples are accessible per class for the target domain. The rest of the target samples constitute the test set, which is also assumed to be available at training time as additional unlabelled target data. The assessment of the models performance, on the test set, is done considering the weighted F1-Score, subsequently referred simply as F1-Score. We repeat each experiment five times and report average and standard deviation results.

All the methods are implemented in Pytorch and available [here](#). Experiments are carried out on a workstation equipped with an Intel(R) Xeon(R) Gold 6226R CPU @ 2.90GHz, with 377Gb of RAM and four RTX3090 GPU. All the methods require only one GPU for training.

4.1 Results

Tables 2, 3, 4 and 5 report the results of all the competing methods, in terms of F1-Score, varying the amount of labelled target sample in the set $\{25, 50, 100, 200\}$ for the *RESISC45-Euro* and *EuroSat-MS-SAR* benchmarks, respectively.

Concerning the *RESISC45-Euro* benchmark, we evaluate two transfer tasks: (RGB \rightarrow MS) and (MS \rightarrow RGB). Here, the two domains differ in terms of radiometric content (imagery with 3 or 13 channels) and spatial resolution as outlined in Table 1. For the first transfer task (RGB \rightarrow MS) the results are

presented in Table 2. Notably, *SHeDD* systematically outperforms all the competing approaches. Although *SS-HIDA* also exhibits improvements over baseline approaches, it achieves lower performances compared to our method.

In the second transfer task (MS \rightarrow RGB), as illustrated in Table 3, our method continues to outperform competing approaches in the majority of cases, exception made for the case with 200 labelled target samples per class where our proposed approach, nonetheless, still achieves comparable performance to *SS-HIDA*. Generally, the use of the source data clearly enables our framework to achieve a gain of over 2 points in terms of F1-Score, regardless of the amount of target labelled samples, compared to strategies relying solely on target information (*Target Only* and *FixMatch*).

Method	25	50	100	200
<i>Target Only</i>	79.48 \pm 1.34	85.05 \pm 1.01	88.99 \pm 0.77	92.34 \pm 0.52
<i>FixMatch</i>	81.74 \pm 1.38	85.60 \pm 0.37	89.37 \pm 0.55	92.57 \pm 0.79
<i>SS-HIDA</i>	82.29 \pm 0.68	88.81 \pm 0.95	91.64 \pm 1.67	93.59 \pm 1.39
<i>SHeDD</i>	84.06 \pm 0.73	89.12 \pm 0.84	92.84 \pm 0.18	95.29 \pm 0.38

Table 2. Average and standard deviation F1-Score results, over 5 runs, on *RESISC45-Euro* with RGB as source and MS as target domain (RGB \rightarrow MS) varying the amount of per-class target supervision in the range {25, 50, 100, 200}.

Method	25	50	100	200
<i>Target Only</i>	75.19 \pm 1.67	82.52 \pm 0.91	87.45 \pm 0.82	91.74 \pm 0.65
<i>FixMatch</i>	77.17 \pm 1.29	82.80 \pm 0.81	87.86 \pm 0.56	91.93 \pm 0.53
<i>SS-HIDA</i>	79.78 \pm 1.06	85.00 \pm 1.07	89.56 \pm 2.34	93.83 \pm 0.18
<i>SHeDD</i>	81.72 \pm 1.93	86.65 \pm 0.82	91.00 \pm 0.55	93.79 \pm 0.32

Table 3. Average and standard deviation F1-Score results, over 5 runs, on *RESISC45-Euro* with MS as source and RGB as target domain (MS \rightarrow RGB) varying the amount of per-class target supervision in the range {25, 50, 100, 200}.

Regarding the *EuroSat-MS-SAR* benchmark, we consider the transfer tasks: (MS \rightarrow SAR) and (SAR \rightarrow MS). Here, the two domains differ in terms of acquisition modes (Optical vs. Radar), thus providing a more challenging transfer scenario in term of source/target domain heterogeneity. The results for the first transfer task (MS \rightarrow SAR) are reported in Table 4 while the results for the second transfer task (SAR \rightarrow MS) are outlined in Table 5. Irrespective of the amount of labeled samples in the target domain, *SHeDD* consistently outperforms all the competing approaches by a notable margin. Differences are generally more

Method	25	50	100	200
<i>Target Only</i>	60.08 ± 1.25	62.52 ± 0.38	64.93 ± 0.31	67.80 ± 0.65
<i>FixMatch</i>	59.07 ± 1.29	64.45 ± 0.35	67.26 ± 0.95	70.38 ± 0.59
<i>SS-HIDA</i>	60.24 ± 1.65	62.96 ± 0.86	66.63 ± 1.00	70.40 ± 0.87
<i>SHeDD</i>	63.66 ± 1.53	67.91 ± 1.83	70.64 ± 1.50	73.97 ± 0.67

Table 4. Average and standard deviation F1-Score results, over 5 runs, on *EuroSat-MS-SAR* with MS as source and SAR as target domain (MS → SAR) varying the amount of per-class target supervision in the range {25, 50, 100, 200}.

Method	25	50	100	200
<i>Target Only</i>	75.85 ± 0.28	82.94 ± 0.45	87.08 ± 0.83	90.92 ± 0.23
<i>FixMatch</i>	76.87 ± 1.32	83.25 ± 0.65	87.67 ± 0.57	91.74 ± 0.39
<i>SS-HIDA</i>	76.49 ± 0.81	80.52 ± 1.49	85.33 ± 0.73	89.38 ± 0.52
<i>SHeDD</i>	82.30 ± 1.12	88.16 ± 0.85	91.67 ± 0.23	94.52 ± 0.14

Table 5. Average and standard deviation F1-Score results, over 5 runs, on *EuroSat-MS-SAR* with SAR as source and MS as target domain (SAR → MS) varying the amount of per-class target supervision in the range {25, 50, 100, 200}.

pronounced for low amount of target labelled samples. For instance, when only 25 target labeled samples per-class are considered for the transfer task (SAR → MS), *SHeDD* demonstrates nearly a 6-point increase in F1-Score over the second-best competitor.

It is worth noting that, differently from the case of *RESISC45-Euro* benchmark, here *SS-HIDA* only performs on-par with the baseline methods (*Target Only* and *FixMatch*). This point can be partly related to the architectural structure of *SS-HIDA*. While *SHeDD* employs distinct per-domain encoders, *SS-HIDA* shared a portion of its encoder between the two domains.

If on the one hand this architectural choice can prove advantageous in scenarios where domains exhibit a limited degree of heterogeneity (e.g. transferring between RGB and MS data, where one modality can be considered as a subset or superset of the other), on the other hand it may hinder transfer performance in more challenging scenarios characterized by a high degree of heterogeneity, as for the *EuroSat-MS-SAR* benchmark. Consequently, it may fail to establish an effective strategy for general heterogeneous domain adaptation. This result further supports the flexibility of our method in modeling a wide range of heterogeneous data transfer scenarios owing to its inherent structural design.

Ablation Analysis: Table 6 reports the ablation analysis of *SHeDD* on the *EuroSat-MS-SAR* benchmark where MS images serve as source domain and SAR images as target domain. Here we consider the case in which 50 labelled samples per class are available from the target domain. Six different ablations were

devised from the complete model to comprehensively assess the various components upon which *SHeDD* relies. Firstly, we observe a clear positive impact of enforcing disentanglement between domain-invariant and domain-specific features ($L_{\perp}^{S,T}$ and $L_{dom}^{S,T}$) over the scenario where only the supervised classification loss is optimized (*Abla*₁ vs *Abla*₂). Secondly, we can underline that the use of unlabelled target data, through the $L_{\perp}^{U,\hat{U}}$, $L_{dom}^{U,\hat{U}}$ and $L_{pl}^{\hat{U}}$ losses, systematically enhances the performances compared to using the labelled information alone (*Abla*₁, *Abla*₂ vs *Abla*₃, *Abla*₄, *Abla*₅ and *Abla*₆). Thirdly, the highest performances are generally attained when consistency regularization, through pseudo-labelling, is considered (*Abla*₄, *Abla*₅ and *Abla*₆). Fourthly, when either $L_{\perp}^{U,\hat{U}}$ and $L_{dom}^{U,\hat{U}}$ or $L_{pl}^{\hat{U}}$ are employed separately (*Abla*₃ and *Abla*₄), performances are still far from the ones achieved by the whole framework. This indicates that the combined use of these three losses, to leverage unlabelled target data, synergistically enhances the final outcome. Finally, the performed ablations indicate that *SHeDD* clearly benefits from all the components it is built on, thus exhibiting the best performance overall in terms of F1-Score.

Ablation	$L_{clf}^{S,T}$	$L_{\perp}^{S,T}$	$L_{dom}^{S,T}$	$L_{\perp}^{U,\hat{U}}$	$L_{dom}^{U,\hat{U}}$	$L_{pl}^{\hat{U}}$	F1-score
<i>Abla</i> ₁	✓						63.84 ± 0.34
<i>Abla</i> ₂	✓	✓	✓				64.58 ± 0.85
<i>Abla</i> ₃	✓	✓	✓	✓	✓		65.04 ± 0.74
<i>Abla</i> ₄	✓	✓	✓			✓	66.00 ± 0.87
<i>Abla</i> ₅	✓		✓		✓	✓	66.54 ± 0.77
<i>Abla</i> ₆	✓	✓		✓		✓	67.47 ± 1.63
<i>SHeDD</i>	✓	✓	✓	✓	✓	✓	67.91 ± 1.83

Table 6. Ablation study of *SHeDD* on the *EuroSat-MS-SAR* benchmark with MS as source and SAR as target domain when 50 samples per class are considered as labelled target data. F1-Score results, in terms of mean and standard deviation over 5 runs, are reported.

Visual Inspection of learnt representations: Figure 3 visually depicts the internal representation learnt by the different competing methods on the *RESISC45-Euro* benchmark for the transfer task (RGB → MS) when only 25 labelled samples per class for the target domain are considered. To this end, we randomly chose 50 samples per class on the target domain and we extracted the corresponding feature representation per method, that is, the embedding vector used as input to the classifier module—in the case our proposed approach, no-




Fig. 3. Visualization of the embeddings extracted from the different competing approaches: (a) *Target Only* (b) *FixMatch* (c) *SS-HIDA* and (d) *SHeDD* when trained on the *RESISC45-Euro* benchmark with RGB as source and MS as target domain (RGB \rightarrow MS) and only 25 labelled samples per class are considered for the target domain. For this visual inspection, 50 random samples per class from the test set (coming from the target domain) are sampled. The two dimensional representation is obtained via the T-SNE algorithm [19].

tably, the domain-invariant embeddings \mathbf{z}^{inv} are used. Subsequently, we applied t-SNE [19] to reduce the feature dimensionality for visualisation purposes.

When only a limited amount of labelled target data is employed to learn the underlying classification models, as for *Target Only* and *FixMatch* methods, the 2D spatial arrangement of the generated embeddings demonstrates evident visual cluttering, with samples coming from different classes overlapping. Although this phenomenon is partially alleviated on *SS-HIDA* embeddings, the resulting manifold still struggles to accurately recover the underlying eight-cluster structure. In contrast, *SHeDD* produces embeddings that depict a more distinct class-aware manifold, visually aligning better with the underlying data distribution compared to competing approaches.

Overall, the visualisation of internal features representation confirms the quantitative findings we previously discussed.

5 Conclusions and Perspectives

In this paper we have presented *SHeDD*, a deep learning based framework to cope with the challenging scenario of semi-supervised domain adaptation when source and target data are heterogeneous in terms of modality representation. Our end-to-end framework has the objective to learn a target domain classifier by leveraging labelled and unlabelled data from both source and target domain via consistency regularized pseudo-labelling and disentanglement learning. While the former mechanism allows to fully leverage the available unlabelled data, the latter allows to simultaneously extract domain-invariant representations, relevant for the downstream task, while retrieving domain-specific information, that can hinder the cross-modality transfer.

The evaluation on two real-world benchmarks, spanning different degrees of source/target domain heterogeneity, has demonstrated the effectiveness of *SHeDD* compared to baselines and recent competing approaches.

While the proposed experimental evaluation clearly demonstrates the effectiveness of *SHeDD* on challenging remote sensing benchmarks, further assessment on general computer vision tasks involving heterogeneous data sources, such as RGB/Depth, RGB/Thermal, or RGB/LIDAR data, still represents a concrete opportunity. Additional evaluations on these benchmarks could further emphasize the value of *SHeDD* in the broader field of computer vision.

In the short term, we aim to enhance the quality of *SHeDD* by drawing inspiration from recent semi-supervised learning strategies, such as *FlexMatch*, and by exploring the impact of various augmentation techniques on consistency regularization and pseudo-labeling to improve the model’s performance in data-scarce environments. In the medium term, we plan to extend our framework towards a multi-source domain adaptation setting, enabling the use of multiple heterogeneous domains as source data. This could lead to a more robust classifier and potentially improved performance on the target domain. Additionally, further exploration could involve adapting the proposed framework to more structured

classification tasks, such as semantic segmentation or object recognition, where data spanning heterogeneous modalities are abundant.

6 Acknowledgment

This work was supported by the French National Research Agency under the grant ANR-23-IAS1-0002 (ANR GEO ReSeT).

References

1. Oscar Day and Taghi M. Khoshgoftaar. A survey on heterogeneous transfer learning. *J. Big Data*, 4:29, 2017. [4](#)
2. Zhen Fang, Jie Lu, Feng Liu, and Guangquan Zhang. Semi-supervised heterogeneous domain adaptation: Theory and algorithms. *IEEE Trans. Pattern Anal. Mach. Intell.*, 45(1):1087–1105, 2023. [2](#), [4](#)
3. Hao Guan and Mingxia Liu. Domain adaptation for medical image analysis: A survey. *IEEE Trans. Biomed. Eng.*, 69(3):1173–1185, 2022. [2](#)
4. Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 770–778. IEEE Computer Society, 2016. [10](#)
5. Pin Jiang, Aming Wu, Yahong Han, Yunfeng Shao, Meiyu Qi, and Bingshuai Li. Bidirectional adversarial training for semi-supervised domain adaptation. In *IJ-CAI*, pages 934–940, 2020. [3](#)
6. Taekyung Kim and Changick Kim. Attract, perturb, and explore: Learning a feature alignment network for semi-supervised domain adaptation. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIV 16*, pages 591–607. Springer, 2020. [4](#)
7. Anish Lakkapragada, Essam Sleiman, Saimourya Surabhi, and Dennis P. Wall. Mitigating negative transfer in multi-task learning with exponential moving average loss weighting strategies (student abstract). In *AAAI*, pages 16246–16247, 2023. [10](#)
8. Jichang Li, Guanbin Li, Yemin Shi, and Yizhou Yu. Cross-domain adaptive clustering for semi-supervised domain adaptation. In *CVPR*, pages 2505–2514, 2021. [3](#)
9. Jichang Li, Guanbin Li, and Yizhou Yu. Inter-domain mixup for semi-supervised domain adaptation. *Pattern Recognit.*, 146:110023, 2024. [3](#)
10. Yang Liu, Zhipeng Zhou, and Baigui Sun. COT: unsupervised domain adaptation with clustering and optimal transport. In *CVPR*, pages 19998–20007. IEEE, 2023. [3](#)
11. Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *ICLR*. OpenReview.net, 2019. [10](#)
12. Mihailo Obrenovic, Thomas Andrew Lampert, Milos R. Ivanovic, and Pierre Gançarski. Learning domain invariant representations of heterogeneous image data. *Mach. Learn.*, 112(10):3659–3684, 2023. [4](#), [9](#), [10](#)
13. Jiangtao Peng, Yi Huang, Weiwei Sun, Na Chen, Yujie Ning, and Qian Du. Domain adaptation in remote sensing image classification: A survey. *IEEE J. Sel. Top. Appl. Earth Obs. Remote. Sens.*, 15:9842–9859, 2022. [2](#)

14. Can Qin, Lichen Wang, Qianqian Ma, Yu Yin, Huan Wang, and Yun Fu. Contradictory structure learning for semi-supervised domain adaptation. In *SDM*, pages 576–584, 2021. [3](#)
15. Can Qin, Lichen Wang, Qianqian Ma, Yu Yin, Huan Wang, and Yun Fu. Semi-supervised domain adaptive structure learning. *IEEE Trans. Image Process.*, 31:7179–7190, 2022. [2](#)
16. Kuniaki Saito, Donghyun Kim, Stan Sclaroff, Trevor Darrell, and Kate Saenko. Semi-supervised domain adaptation via minimax entropy. In *ICCV*, pages 8049–8057. IEEE, 2019. [2](#), [3](#)
17. Ankit Singh. CLDA: Contrastive learning for semi-supervised domain adaptation. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, 2021. [4](#)
18. Kihyuk Sohn, David Berthelot, Nicholas Carlini, Zizhao Zhang, Han Zhang, Colin Raffel, Ekin Dogus Cubuk, Alexey Kurakin, and Chun-Liang Li. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. In *NeurIPS*, 2020. [6](#), [10](#)
19. L. van der Maaten and G. Hinton. Visualizing Data Using t-SNE. *Journal of Machine Learning Research*, 9:2579–2605, 2008. [14](#), [15](#)
20. Yi Wang, Hugo Hernández Hernández, Conrad M. Albrecht, and Xiao Xiang Zhu. Feature guided masked autoencoder for self-supervised learning in remote sensing. *CoRR*, abs/2310.18653, 2023. [9](#)
21. G. Wilson and D. J. Cook. A survey of unsupervised deep domain adaptation. *ACM Trans. Intell. Syst. Technol.*, 11(5):51:1–51:46, 2020. [2](#), [3](#)
22. Zizheng Yan, Yushuang Wu, Guanbin Li, Yipeng Qin, Xiaoguang Han, and Shuguang Cui. Multi-level consistency learning for semi-supervised domain adaptation. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, (IJCAI-22)*, pages 1530–1536, July 2022. [3](#)
23. Luyu Yang, Yan Wang, Mingfei Gao, Abhinav Shrivastava, Kilian Q Weinberger, Wei-Lun Chao, and Ser-Nam Lim. Deep co-training with task decomposition for semi-supervised domain adaptation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 8906–8916, 2021. [3](#)
24. Ting Yao, Yingwei Pan, Chong-Wah Ngo, Houqiang Li, and Tao Mei. Semi-supervised domain adaptation with subspace learning for visual recognition. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 2142–2150, 2015. [3](#)
25. Yuan Yao, Yu Zhang, Xutao Li, and Yunming Ye. Heterogeneous domain adaptation via soft transfer network. In *Multimedia*, pages 1578–1586. ACM, 2019. [4](#)
26. F. Zhuang, Z. Qi, K. Duan, D. Xi, Y. Zhu, H. Zhu, H. Xiong, and Q. He. A comprehensive survey on transfer learning. *Proc. IEEE*, 109(1):43–76, 2021. [2](#), [3](#)