

# Domain Adaptation for the Classification of Remote Sensing Data

*An overview of recent advances*



IMAGE LICENSED BY INGRAM PUBLISHING

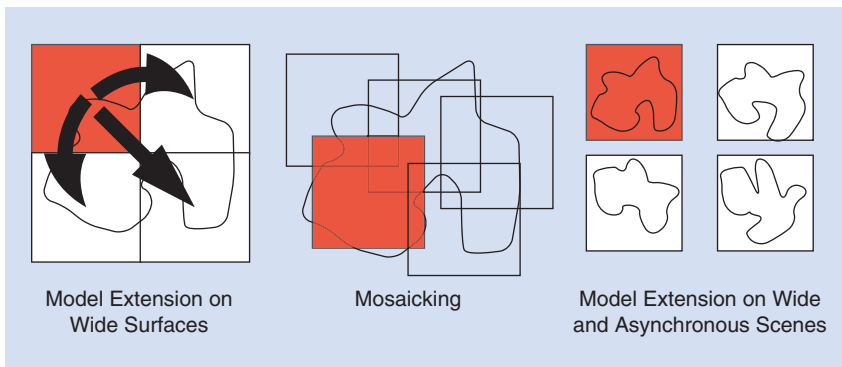
## Advances in Machine Learning for Remote Sensing and Geosciences

**DEVIS TUIA, CLAUDIO PERSELLO,  
AND LORENZO BRUZZONE**

The success of the supervised classification of remotely sensed images acquired over large geographical areas or at short time intervals strongly depends on the representativity of the samples used to train the classification algorithm and to define the model. When training samples are collected from an image or a spatial region that is different from the one used for mapping, spectral shifts between the two distributions are likely to make the model fail. Such shifts are generally due to differences in acquisition and

atmospheric conditions or to changes in the nature of the object observed. To design classification methods that are robust to data set shifts, recent remote sensing literature has considered solutions based on domain adaptation (DA) approaches. Inspired by machine-learning literature, several DA methods have been proposed to solve specific problems in remote sensing data classification. This article provides a critical review of the recent advances in DA approaches for remote sensing and presents an overview of DA methods divided into four categories: 1) invariant feature selection, 2) representation matching, 3) adaptation of classifiers, and 4) selective sampling. We provide an overview of recent

.....  
Digital Object Identifier 10.1109/MGRS.2016.2548504  
Date of publication: 13 June 2016



**FIGURE 1.** Examples of cases where DA is necessary to extend a model to new image acquisitions. In all three cases, the images can be from different sensors, but only the images in red have extensive reference labels (i.e., can be used for training an accurate supervised model).

methodologies, examples of applications of the considered techniques to real remote sensing images characterized by very high spatial and spectral resolution as well as possible guidelines for the selection of the method to use in real application scenarios.

#### REMOTE SENSING FACING NEW OPPORTUNITIES

With the advent of the new generation of satellite missions, which are often made up of constellations of satellites with short revisit time and very high-resolution sensors, the amount of remote sensing images available has increased significantly. Nowadays, the monitoring of dynamic processes has become possible [1], [2], and biophysical parameter estimation and classification problems can be addressed with the use of several data sources [3]–[6]. As a consequence, analysts have the opportunity to use multitemporal and multisource images for tasks such as repetitive monitoring of the territory, change detection, image mosaicking,

and large-scale processing (i.e., processing involving many image tiles) [7].

Remote sensing is therefore facing new opportunities. However, such opportunities cannot be seized unless they come with the capability to provide accurate products in a timely manner. A bottleneck of supervised image-processing-based pipelines is the need of training the model on reference points that are specific to every acquisition. To be accurate, most models need to be trained on the known samples coming from the image under study. Since obtaining new ground samples of high quality for each image acquisition is not realistic, to retrain or adapt an

existing model without such ground samples becomes mandatory. Figure 1 illustrates situations where adaptive models might be of great use. In these situations (which correspond to those considered in this article), only one image, i.e., shown in red in the figure, has sufficient reference labels (e.g., obtained in an extensive ground campaign), whereas the others have no labeled samples or have them only in an insufficient number. This setting is more realistic, since, on the one hand, extensive labeling cannot follow the pace of image acquisitions, and, on the other hand,

repetitive ground campaigns are simply not often an option, mainly for economic and manpower reasons. Indeed, gathering ground information is costly and cannot always be performed by photointerpretation. This is particularly true when the task concerns very large areas or considers quantities that cannot be photointerpreted by an analyst, such as chlorophyll concentrations [8], plant water stress [9], or tree species [10].

To address the cases described in Figure 1, a possible solution would be to bypass the problem and assume that the model that is already available is robust enough to process the new images accurately. Despite the fact that this is possible only in cases where the new image is acquired by the same sensor as the previous one, it is well known that the direct application of a pretrained model on a new data set often provides poor results because the spectra observed in the new scene, even though representing the same types of objects, are different from those of the scene used for training. The differences can be related to a series of deformations (or shifts) related to a variety of effects, such as a biased sampling in the spatial domain (typically if the ground sampling has been focused on a region nonrepresentative of the new scene), changes in the acquisition conditions (including the illumination or acquisition angle), or seasonal changes. When the new data are acquired by a different sensor, the strategy explained previously is simply not applicable, as most models require that all images (or domains) provide samples of the same dimensionality (and where each dimension has the same meaning) at test time. In this case, fusion strategies exist but generally only apply to certain combinations of sensors, and they only use the bands that these sensors have in common or that are reasonably similar [3], which prevents reusing the models that are already trained on the first image of the region that becomes available (which can be crucial, for example, in postcatastrophe interventions) and exploiting sensors' synergies in multisensorial schemes.

To process remote sensing images efficiently and accurately, modern processing systems must be designed to be

**TO PROCESS REMOTE SENSING IMAGES EFFICIENTLY AND ACCURATELY, MODERN PROCESSING SYSTEMS MUST BE DESIGNED TO BE ROBUST IN THE FACE OF CHANGES IN ACQUISITION CONDITIONS AND TEMPORAL SHIFTS AND, IDEALLY, TO BE ADAPTIVE TO SENSOR DIFFERENCES.**


robust in the face of changes in acquisition conditions and temporal shifts and, ideally, to be adaptive to sensor differences. The need for adapting existing models has been acknowledged for many years, as shown by the signature extension community [11], [12], but the change in the amount and nature of data (as well as their resolution) created the need for a new research direction. In this article, we advocate that the solution can be found in DA strategies, a field that is deeply rooted in statistical and machine learning [13], [14].

In general, DA aims to adapt models trained to solve a specific task to a new, yet related, task, for which the knowledge of the initial model is sufficient, although not perfect. As a traditional example in computer vision, DA methods have been deployed to take classifiers that recognize objects in pictures from commercial websites and adapt them to the recognition of objects photographed by simple web cameras [15]. In this example, the classifier is presented a problem with the same objective (i.e., classifying pictures into a limited set of object classes) and the same features but where the data relations are slightly different. For example, at Amazon.com the pictures have no background and the object is mostly in the center of the image, while this is not the case in webcam images. DA is therefore used to adapt the classifier that is accurate on Amazon.com to the new data distribution. Of course, this is just one example. In the DA literature, models are modified to adapt to new data spaces (multimodal), related tasks (multitask), or subtle changes in probability distributions (see a recent review in [16]). The connections to the multitemporal, multisensor, and multiresolution image classification tasks discussed previously are strong [7].

The aim of this review article is to provide an introduction to the DA field and to provide examples of applications of DA techniques in remote sensing. With this in mind, we draw a taxonomy of the DA strategies that have been proposed in recent remote sensing literature and discuss their strengths and weaknesses. We also provide a series of practical examples about the use of DA in high- to very high-resolution image-processing tasks. However, we will not enter into the technical details of specific DA literature; for interested readers, we refer to the recent surveys published in [13], [14], [16], and [17].

## TRANSFER LEARNING AND DOMAIN ADAPTATION

Transfer-learning problems arise when inferences have to be made on processes that are not stationary over time or space. As previously discussed, this is the case in the analysis of remote sensing images where different acquisitions are typically subject to different conditions (e.g., illumination, viewing angle, soil moisture, and topography). Such differences can affect the observed spectral signatures of the land-cover types and, therefore, the distribution of the information classes in the feature space [18]. Different transfer-learning problems (and the techniques to tackle them) have been considered in the literature,



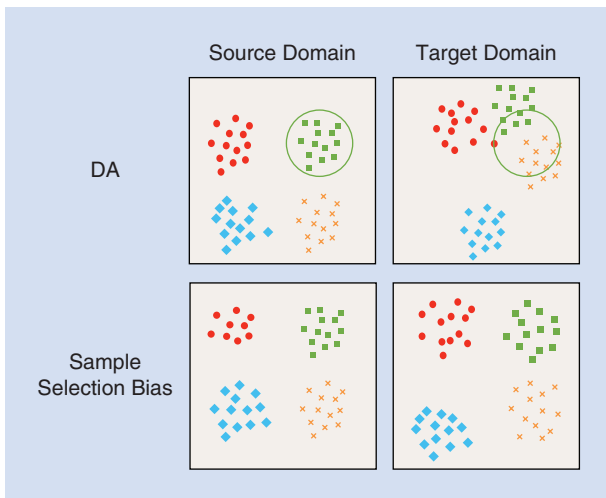
**FIGURE 2.** A graphical representation of the DA problem in the context of remote sensing image classification. Source and target images can be acquired on different geographical areas (but with similar land-cover characteristics) or on the same area at different times. The two images are associated with two different joint distributions, which characterize the two classification problems. The two distributions can differ due to different acquisition conditions (e.g., illumination, viewing angle, soil moisture, and topography).

including DA, multitask learning, domain generalization, sample selection bias, and covariate shift [14]. In this article, we will focus on DA, which is a particular form of transfer learning.

Let us consider two domains, called *source domain* and *target domain*, that are associated with two images acquired on different geographical areas (but with similar land-cover characteristics) or on the same area at different time instants. Figure 2 shows the DA problem in the context of remote sensing image classification. The source and target domains are associated with the joint probability distributions  $P^s(X, Y)$  and  $P^t(X, Y)$ , respectively. The two joint probabilities define the classification problems on the two domains, where  $X$  is the input (vector) variable (i.e., spectral bands of the source image with possible additional features used to characterize the contextual information of the single pixel) and  $Y$  is the output variable associated with a set of classes (i.e., land-cover or land-use information). The aim of DA methods is to adapt a classifier trained on the source domain to make predictions on the target domain.

**TRANSFER-LEARNING PROBLEMS ARISE WHEN INFERENCES HAVE TO BE MADE ON PROCESSES THAT ARE NOT STATIONARY OVER TIME OR SPACE.**

Supervised DA assumes that labeled samples are available for both domains. The labeled sets  $T^s = \{(\mathbf{x}_1, \gamma_1), \dots, (\mathbf{x}_m, \gamma_m)\}$  and  $T^t = \{(\mathbf{x}_1, \gamma_1), \dots, (\mathbf{x}_n, \gamma_n)\}$  are the source- and target-domain training sets, respectively. Supervised DA methods focus on challenging situations where labeled target-domain samples are less numerous than those available in the source domain (i.e.,  $m \ll n$ ). In such conditions, the proper use of source-domain information is very important in solving the target problem. Most of the work in DA assumes that source and target domains share the same set



**FIGURE 3.** Explanatory examples of DA and sample selection bias problems. The plots represent labeled samples in a bidimensional feature space on the source and target domain. A four-class classification problem is considered. In the case of DA, the class-conditional densities may vary from the source domain to the target domain, resulting in a significantly different classification problem (the green circles correspond to the location of the green classes in the source domain). In the case of sample selection bias, the aforementioned issue will not occur, resulting in a milder type of shift from the source domain.

of classes. There are only a few papers that address DA considering differences in the set of classes between source and target [19]–[22].

Semisupervised DA methods assume that a training set is available only for the source domain, whereas target-domain information is limited to a set of unlabeled samples  $U^t = \{\mathbf{x}_1, \dots, \mathbf{x}_m\}$ . This DA setting is more challenging than the supervised case and requires important assumptions on the relationship between source and target domains to make the algorithm converge to a consistent solution on the target domain. All DA methods are based on the assumption that  $P^s(X, Y)$  and  $P^t(X, Y)$  are different but close enough to ensure that the source-domain information can be of help for solving the target-domain learning problem. On the one hand, if the source and target domain are arbitrarily different, there is no hope that the source-domain information will provide an advantage in solving the task in the target domain. On the other hand, if  $P^s(X, Y) = P^t(X, Y)$ , no adaptation is necessary, and the model trained on the source can be readily applied to the target. Semisupervised DA methods are effective in situations that lie in between these two extreme cases.

Unsupervised DA methods are the last family, and they assume that two unlabeled domains have to be matched. This is the most difficult case because label information is not available for any domain. In this situation, DA methods aim to match the marginal distributions of the two domains  $P^s(X)$  and  $P^t(X)$  without knowledge on the learning task (classification or regression). Unsupervised methods can be used as preprocessing of any analysis task (e.g., clustering or density estimation), but they imperatively need to have data sets with similar structural properties before adaptation. Unsupervised DA models are generally feature extractors or matching algorithms that exploit the geometrical structure of the data.

Problems related to DA are the sample selection bias and covariate shift [23]–[25]. The sample selection bias originates when the available training samples are not independently and randomly selected from the underlying distribution (i.e., the training set is not a random sample of the population). This situation is very likely to occur in remote sensing problems where training points are usually selected and labeled through photointerpretation or field surveys. The covariate shift is a particular case of sample selection bias where the bias depends only on the input variable  $X$  (and not on  $Y$ ). Different operational conditions that result in biased training samples are discussed in [26]. Clearly, a training set  $T = \{(\mathbf{x}_1, \gamma_1), \dots, (\mathbf{x}_n, \gamma_n)\}$  obtained under a sample selection bias leads to skewed estimation of the true underlying distribution of the classes, resulting in an estimated  $\hat{P}(X, Y) \neq P(X, Y)$ . For this reason, the effect of a sample selection bias is similar to the DA problem described previously. The covariate shift problem is generally not as severe as the general sample selection bias and the DA problem. Let us consider that both training sets on the source and target are samples with bias (a bias depending on  $X$  only, i.e., covariate shift) from the same joint distribution. The joint probabilities on the two domains can be factorized as  $P^s(X, Y) = P^s(X)P^s(Y|X)$  and  $P^t(X, Y) = P^t(X)P^t(Y|X)$ . In this particular case of covariate shift, the estimated conditional probabilities will be approximately equal, while the marginal distributions will generally be different, i.e.,  $\hat{P}^s(Y|X) \approx \hat{P}^t(Y|X)$  and  $\hat{P}^s(X) \neq \hat{P}^t(X)$ .

Figure 3 illustrates the two different issues of DA and the sample selection bias with two explanatory examples. The plots report the distributions of the labeled samples from two domains in a bidimensional feature space. A four-class classification problem is considered. In the case of DA, the class-conditional densities may change from the source to the target domain, possibly resulting in a significantly different classification problem. Note that the distribution of the classes in the target domain may overlap with different classes on the source domain (see the green circle in the target domain that represents the location of the green class in the source). In this situation, training samples of the source domain can be misleading for the classification of the target domain [27], [28]. In the case of sample selection bias,

**THE SAMPLE SELECTION BIAS ORIGINATES WHEN THE AVAILABLE TRAINING SAMPLES ARE NOT INDEPENDENTLY AND RANDOMLY SELECTED FROM THE UNDERLYING DISTRIBUTION.**



source- and target-domain samples are drawn (with bias) from the same underlying distribution, generally resulting in a milder type of shift with respect to DA. (The problem of cross-domain class overlap is not likely to occur in sample selection bias problems.)

In real remote sensing problems, it is expected that the two aforementioned issues may occur at the same time and can be profoundly entangled, meaning that 1) the two theoretical underlying distributions  $P^s(X, Y)$  and  $P^t(X, Y)$  associated with the source and target domains, respectively, may differ because of changes in the image acquisition conditions (e.g., radiometric conditions), and 2) the available training samples could be not fully representative of the statistical populations in their respective domains and will, therefore, lead to biased estimations of the class probabilities and to inaccurate classification models. In the remote sensing literature, several techniques have been presented to solve the transfer-learning problem irrespective of the cause of the data set shift between source and target domains. The next section will present the different families of techniques that have been proposed in remote sensing image processing.

## A TAXONOMY OF ADAPTATION METHODS

Adapting a model trained on one image to another (or a series of new images) can be performed in different ways. In this section, we detail recent approaches proposed in the remote sensing literature by grouping them into the following four categories:

- ▮ **The selection of invariant features:** A set of the input features (i.e., the original bands or additional features extracted from the remote sensing image) that are not affected by shifting factors are identified and selected before training the classification algorithm. Accordingly, the features affected by the most severe data set shift are removed, and the classifier considers a feature space showing higher stability across domains. An alternative approach is to encode invariance by including additional synthetic labeled samples in the training set to extract features that better model the intraclass variability across the domains.
- ▮ **The adaptation of the data distributions:** The data distributions of the target and source domains are made as similar as possible to keep the classifier unchanged. With respect to the previous family, these methods work on the original input spaces and try to extract a common space where all domains can be treated equally. This is generally achieved by means of joint feature extraction.
- ▮ **The adaptation of the classifier:** Here, the classification model defined by training on the source domain is adapted to the target domain by considering unlabeled samples of the target domain. In this case, the data distributions remain unchanged, and the classifier is adapted to the target data distribution using strategies based on semisupervised learning.

- ▮ **The adaptation of the classifier by active learning (AL):** This adaptation is performed by providing a limited amount of well-chosen labeled samples from the target domain. This is a special case of the previous family, where we allow some new labeled examples to be sampled in the target domain to retrain the model iteratively. Due to their acquisition cost, these samples need to be selected well, according to their potential to lead the model toward the desired target classifier.

The rest of this section details recent advances in these four families. We will limit the discussion to approaches specific to the remote sensing literature and invite the interested reader to consult the specialized machine-learning and computer vision literature in [13], [14], and [16].

## SELECTING INVARIANT FEATURES

The first family of DA methods is based on the selection of invariant features that are usually a subset of the original set of features that are the most robust in the face of changes from the source to the target domain. The main idea of the approach is to select features to reduce the difference between  $P^s(X, Y)$  and  $P^t(X, Y)$ . An alternative strategy for encoding the invariance is based on the inclusion of additional synthetic labeled samples in the training set, a procedure known in machine learning as *data augmentation*. A method adopting this strategy was studied in [29], where sample-selection bias problems are addressed by enriching the training set with artificial examples that correspond to physically consistent variations of the training samples (e.g., illumination, size, and rotation). To limit the number of additional examples to be used by the support vector machine (SVM), variations are generated only for the training samples considered as support vectors by the classifier trained on the source domain only.


Let us consider the first strategy and focus on the analysis of hyperspectral images as an application of particular interest. Hyperspectral sensors are capable of capturing hundreds of narrow spectral bands from a wide range of the electromagnetic spectrum, which is why they are particularly sensitive to subtle changes in image-acquisition conditions, leading to a nonstationary behavior of the spectral signature of the classes and, therefore, to problems that should be solved by transfer learning and DA approaches. Figure 4 shows an example of a shift in the signature of a hyperspectral image acquired by the Hyperion sensor over two areas of the Okavango Delta in Botswana.

In [30], the authors propose an approach for selecting subsets of features that are both discriminative of the land-cover classes and invariant between the source and the target

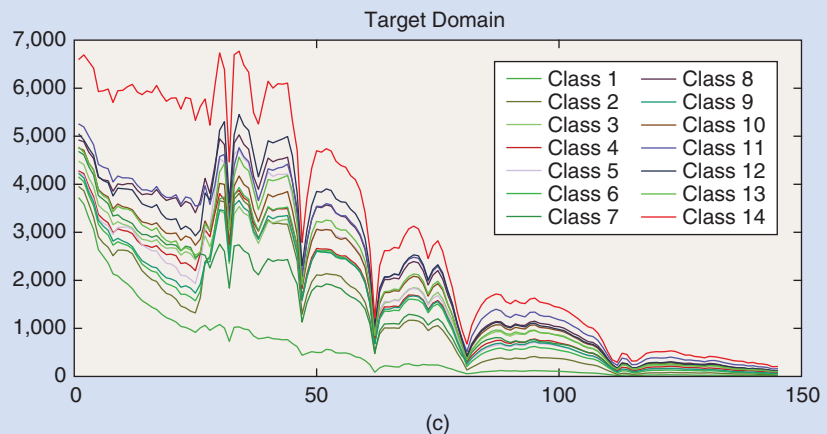
**THE COVARIATE SHIFT IS A PARTICULAR CASE OF SAMPLE SELECTION BIAS WHERE THE BIAS DEPENDS ONLY ON THE INPUT VARIABLE X (AND NOT ON Y).**



(a)



(b)



(c)

**FIGURE 4.** (a) A false color composition of a portion of the hyperspectral data set. (b) The mean spectral signature of the classes on the source domain. (c) The mean spectral signature of the classes on the target domain.

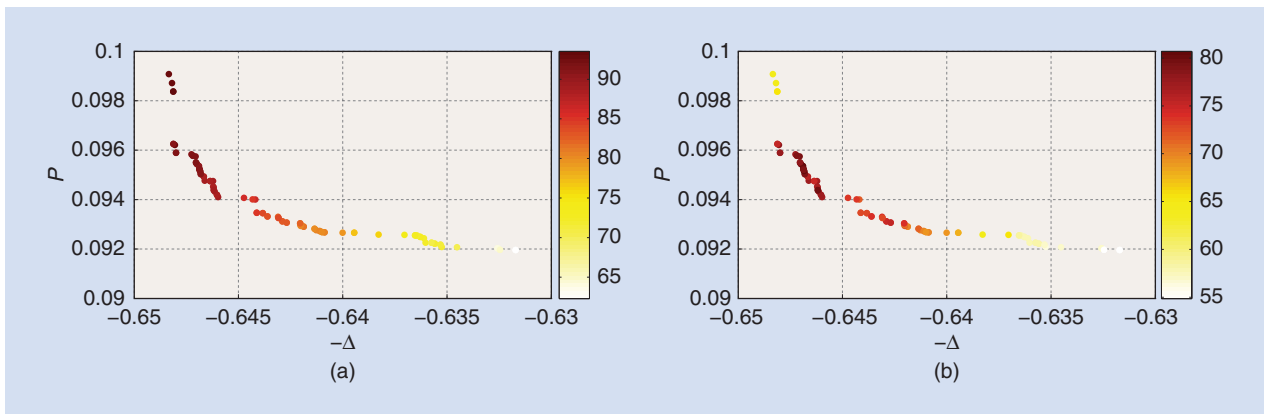
**IN THE REMOTE SENSING LITERATURE, SEVERAL TECHNIQUES HAVE BEEN PRESENTED TO SOLVE THE TRANSFER-LEARNING PROBLEM IRRESPECTIVE OF THE CAUSE OF THE DATA SET SHIFT BETWEEN SOURCE AND TARGET DOMAINS.**

domain. The main idea of this approach is to explicitly consider two distinct terms in the criterion function for evaluating both the discrimination capability  $\Delta$  of the feature subset and the data set shift  $P$  of the features between the source and target domain. The first term is standard in filter methods for feature selection and provides high scores when the features selected show some kind of dependency with the desired output (e.g., the classes to be predicted). The

second term has been introduced to evaluate the invariance of the feature subset between the two domains. The subset of features  $F$  is selected by jointly optimizing the two terms  $\Delta$  and  $P$ , i.e., by solving the following multiobjective optimization problem:

$$\operatorname{argmin}_{|F|=l} (-\Delta(F), P(F)), \quad (1)$$

where  $l$  is the size of the feature subset. Both  $\Delta$  and  $P$  are treated as functions of the subset of considered features  $F$ . The specific definitions of the terms  $\Delta$  and  $P$  are reported in [30], considering their parametric estimation (assuming Gaussian distribution of the classes) in both the supervised and semisupervised DA settings. In [31], the two terms are defined considering kernel-based dependence estimators and kernel embedding of conditional distributions, resulting in a nonparametric approach that does not require the estimation of the class distributions as an intermediate step. Equation (1) is solved by adopting a genetic multiobjective optimization algorithm. The solution results in features with high capability to discriminate classes (with a small value of  $-\Delta$ ) and high stability on the two domains (with a small data set shift  $P$ ). Adopting a multiobjective optimization approach instead of considering a linear combination of the two terms frees the user from specifying in advance the relative importance of the two terms  $\Delta$  and  $P$ . The solution of the multiobjective problem allows one to find the solutions that represent the best tradeoffs of discriminative and stable feature subsets for the specific transfer-learning problem at hand.



**FIGURE 5.** The Pareto front estimated using a multiobjective genetic algorithm for the selection of six features. Each dot corresponds to a feature set minimizing (1). The color indicates the OA on (a) the source test set  $\mathcal{T}^s$  and (b) the target test set  $\mathcal{T}^t$ , according to the reported color scale bar (adapted from [30]).

In the following, we report the experimental results obtained on a hyperspectral image acquired by the Hyperion sensor of the Earth Observation 1 satellite in an area of the Okavango Delta, Botswana [32] (see Figure 4). For more information about the experimental setting and the obtained results, see [30]. The labeled reference samples were collected on two spatially disjoint areas with slightly different characteristics, thus representing two different domains. The samples taken on the first area, which was considered as the source domain, were partitioned into a training set  $T^s$  and a test set  $\mathcal{T}^s$  by random sampling. Samples taken on the second area (i.e., the target domain) were used to derive a training set  $T^t$  and test set  $\mathcal{T}^t$  according to the same procedure. The estimated Pareto front for the selection of six features is reported in Figure 5.

Each point on the two graphs corresponds to a different selected feature subset  $F$ , i.e., a feature set minimizing (1). In Figure 5(a), the color of the points indicates the overall accuracy (OA) obtained on the source-domain test set  $\mathcal{T}^s$  using an SVM classifier trained using  $T^s$  (according to the reported color scale bar). In Figure 5(b), the color indicates the OA obtained by the same SVM classifier on the target-domain test set  $\mathcal{T}^t$ .

The results show that the solutions with higher relevance  $\Delta$  result in better classification accuracies on the source domain. However, relevance alone is not sufficient for selecting features that are stable for the classification on the target domain. We observe that the most accurate solutions on the target domain  $\mathcal{T}^t$  are those that exhibit a good tradeoff between the relevance and invariance terms, which confirms the importance of the invariance term and shows that the  $P$  measure is able to capture the information of feature stability. To select the subset of features that leads to good generalization capabilities on different domains, tradeoff solutions between the two competing objectives should be identified. The selected subset of features results in an OA of 91% on the source domain and 80.7% on the target. The set of features selected according to the optimization of  $\Delta$  results in an OA of 92.7% on the source but only 64.4% on

the target, showing that the features selected by accounting for the data set shift between the domains can significantly improve the generalization capability on the target domain.

#### ADAPTING DATA DISTRIBUTIONS

The second family reviewed considers DA methods that aim to adapt the representation of the original data, regardless of the model that will process them afterward. A review of the methods proposed in computer vision and machine learning can be found in [16]. Here, we will focus on the approaches proposed in the remote sensing literature. This type of adaptation is often done by relative normalization methods, i.e., methods that do not provide physical units as an output but that instead provide similarly distributed digital numbers. Their aim is to make the data distributions more similar across the domains to train a single model that can simultaneously classify the source and target domains.

In general, a data representation transformation with the aim of making data sources more similar should have the following desirable properties:

- ▶ The method should be able to align unpaired data (see the “Unpaired” column in Table 1), which allows the alignment of noncoregistered data (not even imaging the same location) or data with different spatial resolutions.
- ▶ The method should be able to align data of different dimensionality (see the “ $\Delta$  Dimensionality” column in Table 1) to allow multisource classification.
- ▶ The method should be able to align several domains at the same time (see the “Multisource” column in Table 1) to enhance multitemporal adaptation instead of pairwise adaptation.
- ▶ The method should be able to align in a nonlinear way (see the “Nonlinear” column in Table 1), since the

**THE AIM OF DA METHODS IS TO ADAPT A CLASSIFIER TRAINED ON THE SOURCE DOMAIN TO MAKE PREDICTIONS ON THE TARGET DOMAIN.**

**TABLE 1. THE REPRESENTATION ALIGNMENT METHODS USED IN REMOTE SENSING.**

METHOD	LABELS IN $s$	NO LABELS IN $t$	MULTISOURCE	UNPAIRED	$\Delta$ DIMENSIONALITY	NONLINEAR
PCA	×	✓	×	✓	×	×
KPCA [34]	×	✓	×	✓	×	✓
(SS)TCA [35]	× ✓	✓	×	✓	×	✓
CCA [36]	× ✓	× ✓	✓	×	✓	×
KCCA [37]	× ✓	× ✓	✓	×	✓	✓
MA [43]	✓	✓	✓	✓	✓	×
SSMA [38]	✓	×	✓	✓	✓	×
KERNEL METHOD FOR MANIFOLD ALIGNMENT [39]	✓	×	✓	✓	✓	✓
GM [50]	×	✓	×	✓	×	×

transformation between domains can be nonlinear because of atmospheric or illumination effects.

- ▶ The method should be able to use labeled information from the source domain when available (see the “Labels in  $s$ ” column in Table 1). A discriminative transform tends to align the data sets better because it aligns the data according to the semantic classes required by the user.
- ▶ The method should avoid being forced to use labeled information in all domains (see the “No labels in  $t$ ” column in Table 1), as labels might not be available in all domains or their acquisition might have a high cost, typically through terrestrial campaigns (see the “Adaptation of the Classifier by Active Learning” section).

Several methods have been proposed in recent remote sensing literature. We provide a brief review in this section and a summary in Table 1. Depending on the specific situation, the analyst can use this table to select the most suitable approach.

Most of the recent literature focuses on feature-extraction strategies, where the extracted features align the data spaces with each other. In that space, the same classifier (or regressor) can be applied to all the domains. Beyond the works dealing with traditional or multidimensional histogram matching [33] or data alignment with principal component analysis (PCA) or kernel PCA (KPCA) [34], the authors of [35] propose to minimize the statistical distance between domains, which is assessed through a kernel-based dependence estimator, the maximum mean discrepancy (MMD) [18]. Other studies still focus on feature extraction, but based on multiview models. In [36], Nielsen aligns domains with canonical correlation analysis (CCA) and performs change detection therein. The approach is extended to a kernel and semisupervised version in [37], where the authors perform change detection with different sensors. In [38], the domains are matched in a latent space defined through an eigenproblem aiming at preserving label (dis)similarities and the geometric structure of the single manifolds. A nonlinear (kernelized) version

of the algorithm is proposed in [39], where the approach is particularly appealing because it can align an arbitrary number of domains of different dimensionality, as do CCA and kernel CCA (KCCA), but without requiring paired examples. However, it has the disadvantage of requiring labeled samples in all domains.

In [40], the authors relax this requirement by working on semantic ties, i.e., samples issued from the same object but whose class is unknown. This last method therefore requires at least a partial overlap between the images to find the ties, either manually or by stereo matching, as in [41]. The authors in [42] regularize the manifold alignment (MA) solution with spatial information, leading to a more stable feature representation transfer. In [43], they propose a multiscale approach, considering the preservation of both local and global geometric characteristics and relying on clustering pairs rather than labeled correspondences.

Other recent methods rely on eigendecompositions, such as those proposed in [44] and [45]. In [44], two PCA eigensystems (i.e., one for the source domain and another for the target domain) are aligned by minimizing their divergence. In [45], the authors consider a sparse representation approach where they reduce the difference between domains again by minimizing the MMD. In both [44] and [45], the authors aim to transfer category models that are learned on landscape views to aerial views from very high-resolution remote sensing images. In [46], the authors propose a set of techniques based on sample reweighing and transformation to address different DA situations. The study also offers a causal interpretation of the different forms of domain shift. The adaptation strategies are developed on the basis of the embedding of sample distributions in the reproducing kernel Hilbert space.

Beyond classical feature extraction, the authors in [47] align multitemporal sequences based on a measure of similarity between sequence barycenters, which corresponds to a global alignment of the spectra in a time series of images. In [48], the authors consider spatial shifts in large image



acquisitions, in which the spectra are spatially detrended using Gaussian processes to avoid shifts related to localized class variability. In [49], the authors perform anomaly detection by a sparse discriminative transform that maximizes the distance between the anomaly class and the background classes (defined as a set of endmembers) and minimizes the distance between the source and target distributions after reduction by PCA. In [50], the authors consider the domains as multidimensional graphs and propose to align the domains by solving a graph-matching problem. Finally, the authors in [51] find a multispectral mapping between source and target spectra to project the labeled pixels of the source domain into the target domain. Tie points are found between the labeled source pixels and the pixels in the target by registration, and then the mapping between the source and target is learned by regression between the corresponding pairs. Then the labeled pixels are projected into the target domain and are used to train a classifier therein. As for [40], partial overlap between the images is required.

As one can see in Table 1, some methods will be more suitable than others, depending on the problem. For example, canonical correlation-based methods can be used only for coregistered data, while nonmultiview methods such as KPCA and (semisupervised) transfer component analysis [(SS)TCA] cannot align more than two domains at a time.

In the following, we compare a series of methods on the challenging problem of transferring a classifier over a multiangular sequence of images over Rio de Janeiro [52], illustrated in Figure 6. More details on this example can be found in [38]. The images are not coregistered but are all acquired from a single pass of the WorldView2 sensor. For this reason, the only shifts observed are due to angular effects. The problem is an 11-classes problem, and a separate ground truth is provided per each image (Table 2).

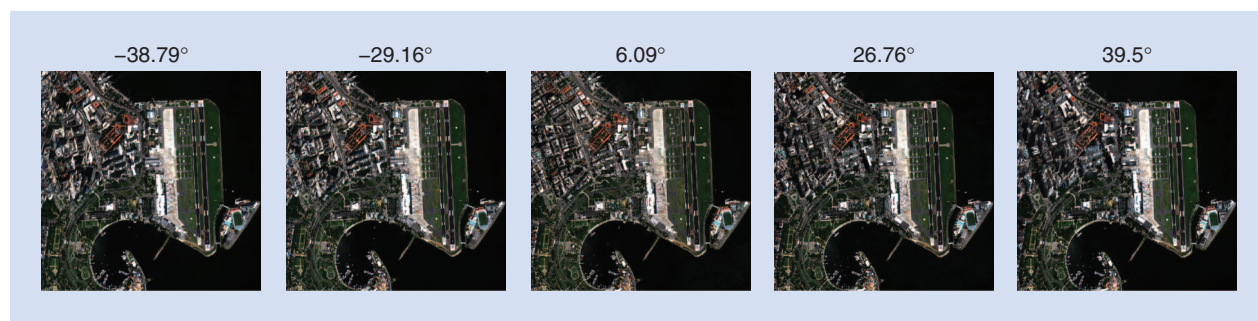
The adaptation experiment is designed by taking the nadir image (off-nadir angle  $\theta = 6.09^\circ$ ) as the source image and using all the others as target images. We apply the PCA, KPCA, graph matching (GM), and semisupervised manifold alignment (SSMA) transforms and then train a classifier using 100 labeled pixels from the source domain and predict all of the target domains using that classifier, without further modifications. For PCA, KPCA,

**TABLE 2. THE NUMBER OF LABELED PIXELS AVAILABLE FOR EACH DATA SET IN THE MULTIANGULAR EXPERIMENTS ( $\theta =$  OFF-NADIR ANGLE).**

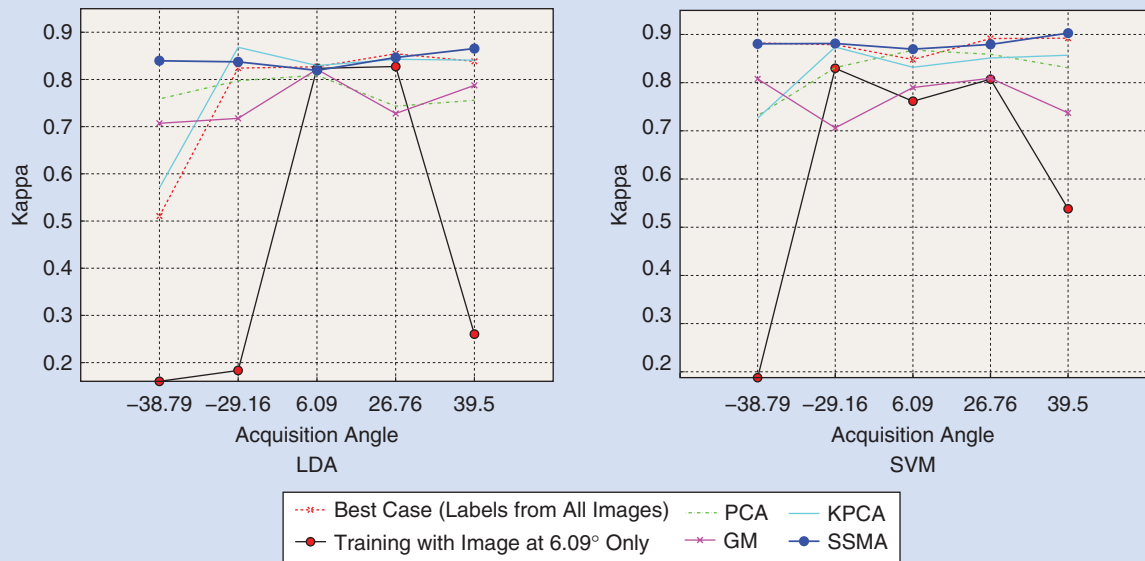
CLASS \ $\theta$	$-38.79^\circ$	$-29.16^\circ$	$6.09^\circ$	$26.76^\circ$	$39.5^\circ$
Water	83,260	79,937	66,084	63,492	54,769
Grass	8,127	8,127	8,127	8,127	8,127
Pools	244	244	223	195	195
Trees	4,231	4,074	3,066	3,046	3,046
Concrete	707	719	719	719	696
Bare soil	790	790	790	790	811
Asphalt	2,949	2,949	2,949	2,827	2,827
Gray buildings	6,291	6,061	5,936	4,375	4,527
Red buildings	1,147	1,080	1,070	1,046	1,042
White buildings	1,683	1,683	1,571	1,571	1,571
Shadows	1,829	1,056	705	512	525
Tarmac	5,179	5,179	5,179	2,166	2,758

and GM, the adaptation is done for each target domain separately, while, for SSMA, a single adaptation projection is obtained for all domains at once. For SSMA, we also used 50 labeled pixels per class from each target domain. To be fair in the evaluation, the projections for the PCA, KPCA, and GM methods are obtained in an unsupervised way, but then the classifier is trained using the original training points from the nadir acquisition, stacked to the transformed labeled pixels of the domain to be tested. We also add a best-case scenario, where we directly use labeled samples from the target domains for the classification. The results are illustrated in Figure 7.

The prediction in the off-nadir images using the original training samples from the nadir image leads to poor results, especially for strong off-nadir angles. All of the methods considered leverage the decrease in performance and lead to a quasi-flat prediction surface (meaning that the model can predict correctly, regardless of the angular configuration) with particularly good performance by the SSMA method, which seems to best align the data



**FIGURE 6.** The five images of the Rio de Janeiro angular sequence [52].



**FIGURE 7.** The classification results over the five Rio de Janeiro acquisitions (adapted from [38]). There are 100 labeled samples per class from the nadir image used to train a classifier and then to test on the others. In the SSMA experiment, 50 labeled pixels per class are used from the other acquisitions.

distributions. This is not surprising, since, among the tested methods, SSMA is the only one with a clear discriminative component that uses labels in all domains to define the projections.

### ADAPTING CLASSIFIERS WITH SEMISUPERVISED APPROACHES

A widely used approach to DA is based on adaptation of the model of the classifier derived on the source domain to the target domain. In the literature, the approach is defined as semisupervised if this adaptation is based only on unlabeled samples of the target domain (no target training samples are used). The rationale of semisupervised adaptation is to use the relations between the distributions of the source and target domains to infer a reliable solution to the problem described in the target domain. The common assumption of most of the methods is that the source and target domains share the same set of classes and features.

The first attempts to address semisupervised DA in remote sensing image classification were presented in [53], where a DA technique is proposed that updates the parameters of an already trained parametric maximum-likelihood (ML) classifier on the basis of the distribution of a new image for which no labeled samples are available. In [54], the ML-based DA technique is extended to the framework of the Bayesian rule for cascade classification (i.e., the classification process is performed by jointly considering information contained in the source and the target domains). The basic idea in both methods is modeling the observed spaces by a mixture of distributions, whose components can be estimated by the use of unlabeled target data. This is achieved

by using the expectation maximization (EM) algorithm with the finite Gaussian mixture model. In [55] and [56], DA approaches based on multiple-classifier and multiple-cascade-classifier architectures were defined. Gaussian ML classifiers, radial basis function neural networks, and hybrid cascade classifiers are used as base classifiers. The decision tree update and randomization are used for DA in [57]. A set of decision trees is made robust in the face of data set shift either by training it with the EM using density functions from the target domain (similar to [54]) or by randomizing the decision trees (but, in this case, without control of the adaptation objective). The authors also propose a semisupervised extension where the classifiers performing poorly on the few labeled samples in the target domain are downweighted in the final decision. Finally, the DA technique proposed in [20] for multitemporal images addresses the challenging situation where the source and target domains have a different set of classes. The sets of classes of the target and the source domains are automatically identified in the DA step by the joint use of unsupervised change detection and the Jeffries–Matusita statistical distance measure. This process results in the detection of classes that appeared or disappeared between the domains.

The semisupervised problem has also been extensively studied in the framework of kernel methods with SVM classifiers, which has been done especially for addressing sample selection bias problems (see the “Transfer Learning and Domain Adaption” section). Most of the semisupervised techniques proposed with SVM exploit the cluster assumption, i.e., adapt the position of the hyperplane estimated on the source domain to the target domain, assuming that it should be located in low-density regions

of the feature space. In [58], the authors employ the transductive SVM, a method that iteratively moves the decision boundary of the SVM classifier toward low-density areas of the unlabeled target domain. Other semisupervised approaches are imported in remote sensing in [59], where the SVM semisupervised learning is addressed in the primal formulation of the cost function. In [60], the authors regularize the SVM solution by adding a new term in the optimization accounting for the divergence between source and target domains (i.e., the MMD discussed in the “Adapting Data Distributions” section). By doing so, the decision function selected depends on a kernel that both projects in a discriminative space and minimizes the shift between training and test data. In [61], the authors cast the DA problem as a multitask learning problem, where each source–domain pair (i.e., each task) is solved by deforming the kernel by sharing information among the tasks.

The Laplacian SVM technique applied to the classification of multispectral remote sensing images is presented in [62], which exploits an additional regularization term on the geometry of both the labeled and the unlabeled samples by using the Laplacian graph. In [63], the authors also use a manifold-regularized classifier in a semisupervised setting, where the adaptation is performed by adding semi-labeled examples from the target domain.

A specific semisupervised SVM that is defined for addressing DA problems is presented in [64]. The domain adaptation SVM (DASVM) starts from a standard supervised learning on the training samples of the source domain, which is followed by an iterative procedure. At each iteration, it includes in the learning cost function a subset of unlabeled samples of the target domain adequately selected while gradually removing the training samples of the source domain. At convergence, the DASVM can accurately classify the samples of the target domain.

## ADAPTATION OF THE CLASSIFIER BY ACTIVE LEARNING

In most of the previously mentioned approaches, it is assumed that no label information can be obtained in the newly acquired target domains (i.e., semisupervised DA). This assumption may hinder the success of classification in the case of very strong deformations or when new classes that are unseen during training appear in the test data. A small amount of labeled data issued from the target domain may solve this problem efficiently. However, since the acquisition is timely and can be costly, it becomes mandatory to choose the samples well. AL strategies have been proposed to tackle this challenging task and guide the DA process with the selection of the most informative target samples [19], [27], [28], [65]–[67].

AL is the name of a set of methodologies aiming at the interaction between a user and a prediction model, where the user provides labels based on knowledge of the task to be solved and the model performs the prediction and highlights samples for which it has the highest uncertainty [68]. By focusing on these samples, the user provides the labels where they help the most and, thus, allows the classifier to migrate in a fast way toward the optimal model. Surveys on AL methods applied to remote sensing can be found in [69]–[71]. In the case of DA, the user provides examples coming from the target domain alone, and the optimal classifier is the one that would have been obtained with several examples in the target domain. The process starts with a classifier that is optimal for the source domain and gradually evolves to model the data distribution in the target domain. Figure 8 summarizes the AL process for DA.

One could apply classical AL strategies in transfer-learning problems under the sample selection bias assumption with successful results, because classical AL will point out samples close to the current decision boundaries, and the

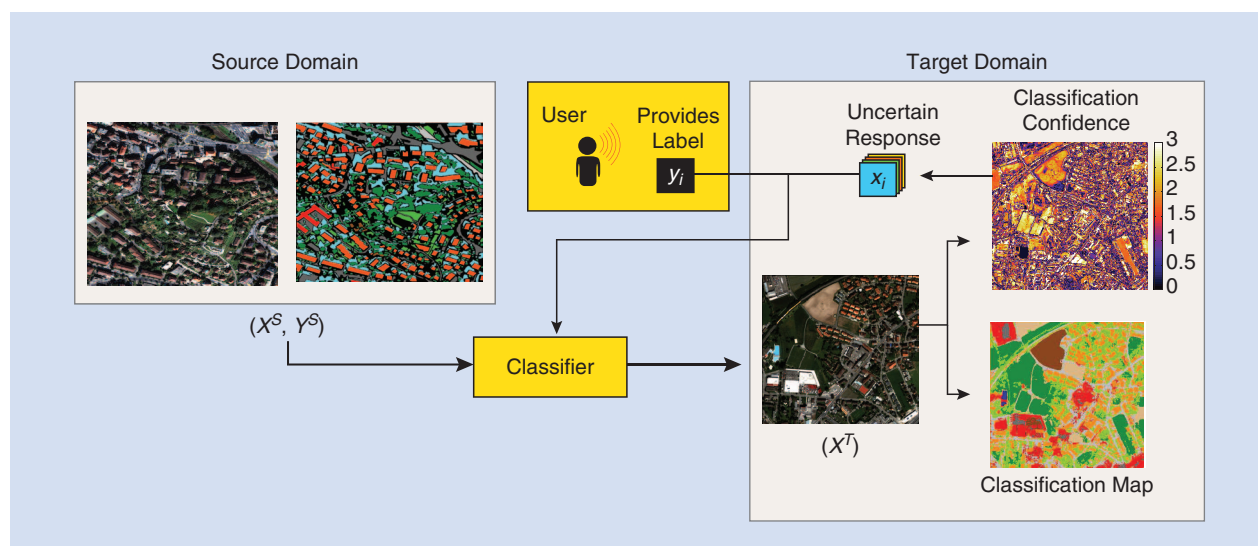
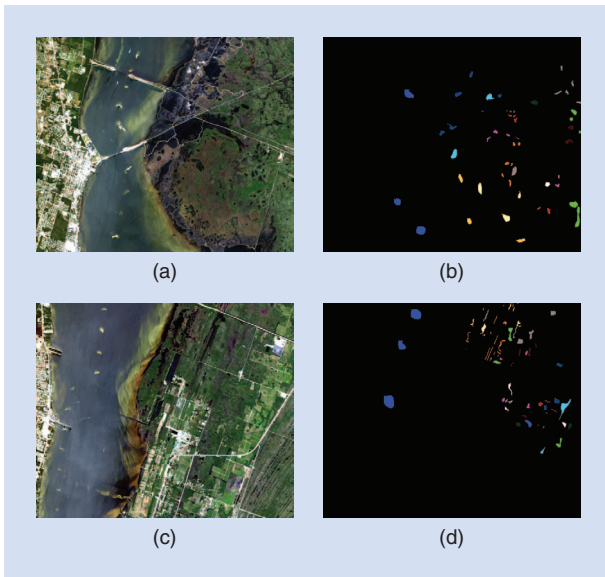


FIGURE 8. A flow chart of the AL paradigm for DA.



**FIGURE 9.** KSC data used in the AL DA experiment: (a) the source image, (b) the source GT, (c) the target image, and (d) the target GT (not available).

user, by the labels provided, will disclose the shifted areas where the next iterations will focus. However, depending on the degree of transformation between the domains, one can use more sophisticated strategies that take into account measures of the deformation between the domains. In this respect, the following problems have driven DA-related research in AL:

- ▶ When it is expected that new classes will appear in the target domain, AL can be used to highlight the areas of the feature space where these classes could be. In [19], by using the reasoning of sample selection bias, the feature space in the target domain is screened using clustering, and dense clusters with no labeled samples are presented to the user, who can then provide labels if new classes are present. In [72], the detection of new classes is set as a change-detection problem where an uncertainty of changes is assessed with an information theoretic criterion. Image time series are analyzed in [21].
- ▶ When significant differences between source and target domains are expected (i.e., when the sample selection bias assumption does not hold), the presence of labeled source samples, although beneficial at the beginning of the process, can be harmful for the classification of the target domain [27], [28] (see the examples in the “Transfer Learning and Domain Adaptation” section and Figure 3). If the class distributions in the target domain overlap with those in the source domain, relying on the labels from the source will lead to classifier errors. Accordingly, the approaches in [27], [28], and [67] consider reweighting of the samples in the training set enriched by AL. When samples from the source domain become less relevant or misleading for the correct classification

of the target domain, they are downweighted in the adapted classifier or completely removed. Accordingly, the classifier specializes to the target domain through the inclusion of target samples and gradually forgets the initial source domain.

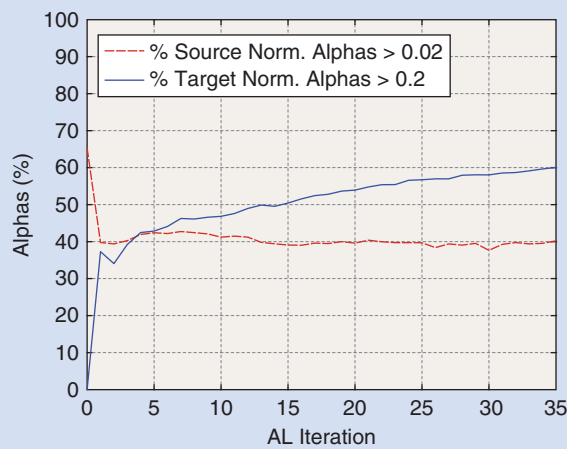
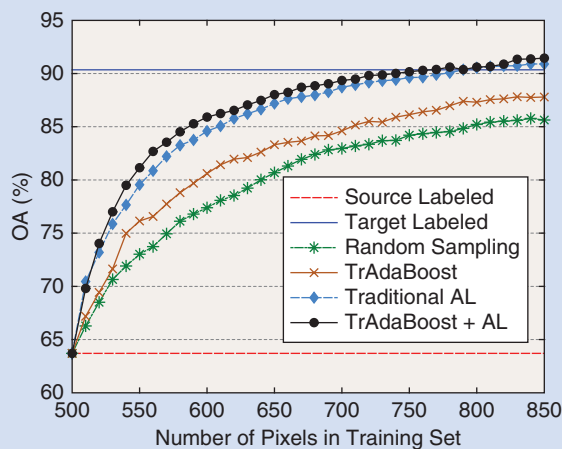
- ▶ When the areas to be processed become very large, specific solutions must be designed to avoid too many iterations of the AL process. In this respect, solutions based on the selection of clusters [73], compressed sensing [74], or geographically distributed search strategies [75]–[78] have been considered.

In the following, we focus on one example related to the second point above, i.e., the reweighting of source samples. This example is adapted from [67]. We study the feasibility of the migration of a model optimized for land-cover mapping in a geographical area to another spatially disjoint region. To do so, we consider the well-known Kennedy Space Center (KSC) hyperspectral image acquired by the airborne visible/infrared imaging spectrometer [Figure 9(a) and (b)] and try to adapt the model learned therein to be accurate in a spatially disjoint section of the same flightline [Figure 9(c) and (d)]. We consider only the ten classes present in all images. The starting model is learned using a training set composed of 50 labeled pixels per class and is then enriched by new samples that are either added randomly or using the breaking ties AL strategy [79]. The classifier is an SVM, either standard (when no other mention is done) or adaptive using the TrAdaBoost model, which is a DA method based on the reweighting of the SVM sample weights after the inclusion of the new labeled points from the target domain [66].

When using the source SVM without adaptation, we reach an OA lower than 65%, while the results obtained by an SVM that is trained directly on the target labeled samples (which are available for testing) would provide an accuracy of 90% (Figure 10). Here, the shift is clearly visible and relates to a loss in accuracy of 25%. Using a random sampling in the target domain, we get a constant increase in performance [shown in Figure 10(a) by the green line with \* markers]; however, after 300 queries, we are still 5% away from the classifier learned using only 500 samples from the target domain. Moreover, the learning rate is slow and the gain is almost linear with the number of queries. We then assess different DA strategies.

First, TrAdaBoost is applied to the set enriched by the random samples [shown in Figure 10(a) by the brown line with × markers]. By forgetting the source domain, i.e., by downweighting the source samples that are contradictory with respect to the new samples from the target domain, we already see a significant improvement that fills half of the gap between the best case and the random sampling. However, when using AL (shown by the blue line with diamond markers) and even more when using it in conjunction with the TrAdaBoost model (shown by the black line with circle markers), the learning rate is much higher in the first iterations, which means the first queries are much more





**FIGURE 10.** The AL results over the KSC data of Figure 9 (adapted from [67]). (a) The learning rates (OA) for different methods. (b) The percentage of source (dashed red line) and target (solid blue line)  $\alpha$  weights larger than 0.02 (source) and 0.2 (target) along the iterative AL process in the TrAdaBoost + AL experiment, which is shown by the black line in (a).

effective than in the random sampling experiments, and the model converges to the result obtained with 500 random target queries (shown by the solid blue line) with only 250 active queries, corresponding to a total of 750 samples in the model since we still have the 500 initial samples from the source. Figure 10(b) shows the percentage of the support vectors from each domain with nonzero weights. In the target domain, this share increases constantly [shown in Figure 10(b) by the solid blue line], while it stabilizes for the samples from the source to 40% of the original training samples (shown by the dashed red line). This means that the importance of the source in the model is strongly reduced in the first iterations and then remains constant, while each new sample from the target becomes immediately important and receives a strong weight from the boosted SVM classifier.

### GUIDELINES FOR CHOOSING THE ADAPTATION STRATEGY

In this section, we will first provide guidelines for the selection of the most appropriate adaptation strategy and then discuss the issue of the validation of the adapted models.

#### HOW TO CHOOSE

In the previous sections, we presented different approaches to DA that were grouped into four families. Depending on the problem considered, an analyst can favor one or another. However, there are some guidelines that should be taken into account that will depend mainly on the data available (e.g., the sensors to be used) and on the effort already provided by the analyst (e.g., whether a classifier is already available or if labels in the target domain are available or can be acquired easily). The guidelines are as follows:

- ▶ If the data to be used are acquired by different sensors, they are associated with different feature spaces. In this case, only heterogeneous DA (i.e., methods that allow for aligning spaces of different dimensionality) should be considered. Accordingly, multiview feature-representation transfer methods such as CCA and KCCA or MA (see the “Adapting Data Distributions” section) are the recommended choices.
- ▶ If a classifier trained on the source is already available and the effort of training is considerable, methods of the third and fourth families (i.e., the adaptation of classifiers and the adaptation by selective sampling, respectively) should be preferred. These methods build on the model that is already defined on the source domain, while those of the two other families imply the definition of a new classifier that is successful in all domains.
- ▶ Whenever it is possible to acquire new labeled samples in the target, it should be done. There is no better way to correct for a data set shift than by having examples of the class-conditional distribution in the target. The AL and MA methods are to be preferred in that case.
- ▶ The level of data set shift the methods can cope with goes along with their level of flexibility. Representation methods relying on labeled samples from the target can cope with strong nonlinear deformations because they allow for a kind of feature registration between the domains, while those that do not use target samples (e.g., PCA, TCA, and GM) are successful only if the data distributions are already prealigned and have not undergone drastic shifts, such as the cases where the signature of a target class becomes identical with one of the others in the source. Among the unlabeled methods, the differences in their flexibility should be considered, going from

linear global methods (e.g., PCA representation transfer) to local methods (e.g., based on clustering, as with GM and MA). If the first can address only rotations, translations, and to some extent scalings of the data clouds, the others can model the per-sample transformation and allow more flexibility of the transform. The same type of reasoning holds for semisupervised methods, which will be able to correct for smaller shifts than methods based on AL. When deployed in a DA setting, AL methods can collect target labeled samples that provide evidence of the real target class distributions, while the semisupervised method uses only unlabeled data in the target and, therefore, cannot easily discover drastic changes in the class distributions.

- ▶ A combination of methods from different families is also possible. For example, selecting invariant features can be a preprocessing step to kernel MA, where the labels in the target domain have been acquired by AL using the labels from the source domain.

With these simple guidelines in mind, the analyst can select the most appropriate strategy (or combine a series of them) according to the considered data and application.

#### HOW TO VALIDATE

A typical bottleneck for the employment of an adaptation strategy is the validation of the adaptation process itself, since it is assumed that no (or only few) labeled data are available for the target domain. Nonetheless, one should assess whether the adaptation was successful in the processing of the target image, even though no labeled samples are available for such validation. To address this crucial issue, a circular validation strategy is presented and applied to remote sensing images in [64]. The strategy is based on the idea that an intrinsic structure relates the solutions that are consistent with the source and the target domains. A solution for the target domain, for which no prior information is available, is assumed to be consistent if the solution to the source-domain data is associated with an acceptable accuracy. The solution to the

source-domain data should be obtained by applying the same DA algorithm in the reverse sense, i.e., by using the classification labels in place of missing prior knowledge for target-domain instances. The source-domain data is considered as unlabeled in the reverse DA learning, and the accuracy of the source-domain data can be evaluated due to the available true labels for source-domain samples. This strategy can be effective for both understanding

if the adaptation is feasible in the considered data set and selecting the most effective strategy.

#### CONCLUSIONS

In this article, we reviewed the recent DA advances for remote sensing image analysis. DA is a rising field of investigation in remote sensing, as it answers the need for reusing available ground reference samples to classify or further process new image acquisitions that may be covering different areas, at different time instants, and possibly with different sensors. The increasing satellite-data availability trend observed in the last few years (in particular, thanks to satellite constellations such as the Sentinels or the NASA A-Train) and the commercialization of drone-mounted cameras have pushed these problems to the forefront of researchers' and analysts' priorities.

We have reviewed the recent models proposed in the literature, which were grouped in four main families: 1) the approaches based on the selection of invariant features, 2) those based on the matching of the data representation, 3) those based on the adaptation of the classifier trained on the source domain, and 4) those based on limited but effective sampling of labeled samples in the target domain. With practical examples, we have provided the reader with a thorough introduction to the field and some guidelines for the selection of the approaches to use in real application scenarios.

We believe that DA is of the highest importance to future Earth observation since multimodality and repeated imaging have become unavoidable [7]. The data are already there, and new, challenging problems can now be tackled with remote sensing. The discipline has succeeded in entering many new sectors of society, and it is now time to provide the tools to the users to perform a trustable monitoring that can be obtained in different sensor configurations or modalities. We think that DA and machine learning in general can contribute to providing an answer to this call.

#### AUTHOR INFORMATION

**Devis Tuia** (devis.tuia@geo.uzh.ch) received a diploma in geography at the University of Lausanne (UNIL) in 2004, a master of advanced studies degree in environmental engineering at the Federal Institute of Technology of Lausanne (EPFL) in 2005, and a Ph.D. degree in environmental sciences at UNIL in 2009. Subsequently, he worked as a visiting postdoctoral researcher at the University of Valencia, Spain, and at the University of Colorado, Boulder. He then worked as a senior research associate at EPFL under a Swiss National Foundation (SNF) program. Since 2014, he has been an SNF assistant professor in the Department of Geography at the University of Zurich. His research interests include the development of algorithms for information extraction and data fusion of remote sensing images using machine-learning algorithms. He serves as chair of the Image Analysis and Data Fusion Technical Committee of the IEEE Geoscience and Remote Sensing Society. He is

**DA IS A RISING FIELD OF INVESTIGATION IN REMOTE SENSING, AS IT ANSWERS THE NEED FOR REUSING AVAILABLE GROUND REFERENCE SAMPLES TO CLASSIFY OR FURTHER PROCESS NEW IMAGE ACQUISITIONS THAT MAY BE COVERING DIFFERENT AREAS.**

an associate editor of *IEEE Journal of Selected Topics in Applied Earth Observation and Remote Sensing*. He is a Senior Member of the IEEE.

**Claudio Persello** (c.persello@utwente.nl) received the Laurea (B.S.) and Laurea Specialistica (M.S.) degrees in telecommunications engineering and the Ph.D. degree in communication and information technologies from the University of Trento, Italy, in 2003, 2005, and 2010, respectively. He is currently an assistant professor in the Faculty of Geo-Information Science and Earth Observation, University of Twente, The Netherlands. In 2011–2014, he was a Marie Curie research fellow, conducting research activity at the Max Planck Institute for Intelligent Systems and the Remote Sensing Laboratory of the University of Trento. His main research interests are the analysis of remote sensing data, machine learning, image classification, pattern recognition, and unmanned aerial vehicles. He is a referee for multiple journals, including *IEEE Transactions on Geoscience and Remote Sensing*, *Pattern Recognition Letters*, and *Image and Vision Computing Journal*. He served on the Scientific Committee of the Sixth International Workshop on the Analysis of Multitemporal Remote Sensing Images. He is a Member of the IEEE.

**Lorenzo Bruzzone** (lorenzo.bruzzone@disi.unitn.it) received the Laurea (M.S.) degree in electronic engineering (summa cum laude) and the Ph.D. degree in telecommunications from the University of Genoa, Italy, in 1993 and 1998, respectively. He is currently a full professor of telecommunications at the University of Trento, Italy. His current research interests are in the areas of remote sensing, radar and synthetic aperture radar, signal processing, and pattern recognition. He is the principal investigator of the Radar for Icy Moon Exploration instrument in the framework of the Jupiter Icy moons Explorer mission of the European Space Agency. He is the author (or coauthor) of 180 scientific publications in refereed international journals (129 in IEEE journals), more than 250 papers in conference proceedings, and 20 book chapters. He is the editor-in-chief of *IEEE Geoscience and Remote Sensing Magazine* and is an associate editor of other journals. He is a Fellow of the IEEE.

## REFERENCES

- [1] R. de Jong, S. de Bruin, A. de Wit, M. E. Schaepman, and D. L. Dent, "Analysis of monotonic greening and browning trends from global NDVI time-series," *Remote Sens. Environ.*, vol. 115, no. 2, pp. 692–702, 2011.
- [2] F. Pacifici, N. Longbotham, and W. Emery, "The importance of physical quantities for the analysis of multitemporal and multangular optical very high spatial resolution images," *IEEE Trans. Geosci. Remote Sens.*, vol. 52, no. 10, pp. 6241–6256, 2014.
- [3] J. J. Walker, K. M. de Beurs, R. H. Wynne, and F. Gao, "Evaluation of Landsat and MODIS data fusion products for analysis of dryland forest phenology," *Remote Sens. Environ.*, vol. 117, pp. 381–393, Feb. 2012.
- [4] W. Liao, X. Huang, F. Van Collie, A. Gautama, W. Philips, H. Liu, T. Zhu, M. Shimoni, G. Moser, and D. Tuia, "Processing of thermal hyperspectral and digital color cameras: Outcome of the 2014 data fusion contest," *IEEE J. Sel. Topics Appl. Earth Observ.*, vol. 8, no. 6, pp. 2984–2996, 2015.
- [5] J. Amorós-López, L. Gómez-Chova, L. Alonso, L. Guanter, R. Zurita-Milla, J. Moreno, and G. Camps-Valls, "Multitemporal fusion of Landsat/TM and ENVISAT/MERIS for crop monitoring," *Int. J. Appl. Earth Obs. Geoinfo.*, vol. 23, pp. 132–141, Aug. 2013.
- [6] I. Walde, S. Hese, C. Berger, and C. Schmullius, "From land cover-graphs to urban structure types," *Int. J. Geogr. Info. Sci.*, vol. 28, no. 3, pp. 584–609, 2014.
- [7] L. Gómez-Chova, D. Tuia, G. Moser, and G. Camps-Valls, "Multimodal classification of remote sensing images: A review and future directions," *Proc. IEEE*, vol. 103, no. 9, pp. 1560–1584, 2015.
- [8] J. Verrelst, L. Alonso, J. P. R. Caicedo, J. Moreno, and G. Camps-Valls, "Gaussian Process retrieval of chlorophyll content from imaging spectroscopy data," *IEEE J. Sel. Topics Appl. Earth Observ.*, vol. 6, no. 2, pp. 867–874, 2013.
- [9] J. Behmann, J. Steinrücken, and L. Plümer, "Detection of early plant stress responses in hyperspectral images," *ISPRS J. Int. Soc. Photo. Remote Sens.*, vol. 93, pp. 98–111, July 2014.
- [10] M. Dalponte, L. Bruzzone, and D. Gianelle, "Tree species classification in the Southern Alps based on the fusion of very high geometrical resolution multispectral/hyperspectral images and LiDAR data," *Remote Sens. Environ.*, vol. 123, pp. 258–270, Aug. 2012.
- [11] M. D. Fleming, J. S. Berkebile, and R. M. Hoffer, "Computer-aided analysis of LANDSAT-I MSS data: A comparison of three approaches, including a 'modified clustering' approach," Purdue University, West Lafayette, IN, Lab. Applicat. Remote Sensing Inform. Note 072475, 1975.
- [12] I. Olthof, C. Butson, and R. Fraser, "Signature extension through space for northern landcover classification: A comparison of radiometric correction methods," *Remote Sens. Environ.*, vol. 95, no. 3, pp. 290–302, 2005.
- [13] J. Quiñero-Candela, M. Sugiyama, A. Schwaighofer, and N. D. Lawrence, Eds., *Dataset Shift in Machine Learning*, Cambridge, MA: MIT Press, 2008.
- [14] S. J. Pan and Q. Yang, "A survey on transfer learning," *IEEE Trans. Know. Data Eng.*, vol. 22, no. 10, pp. 1345–1359, 2010.
- [15] K. Saenko, B. Kulis, M. Fritz, and T. Darrell, "Adapting visual category models to new domains," in *Proc. 11th Europ. Conf. Comput. Vision*, Heraklion, Greece, 2010, pp. 213–226.
- [16] V. M. Patel, R. Gopalan, R. Li, and R. Chellappa, "Visual domain adaptation: A survey of recent advances," *IEEE Signal Processing Mag.*, vol. 32, no. 3, pp. 53–69, 2015.
- [17] M. Sugiyama and M. Kawanabe, *Machine Learning in Non-stationary Environments*, Cambridge, MA: MIT Press, 2012.
- [18] G. Matasci, N. Longbotham, F. Pacifici, M. Kanevski, and D. Tuia, "Understanding angular effects in VHR imagery and their significance for urban land-cover model portability: A study of two multi-angle in-track image sequences," *ISPRS J. Int. Soc. Photo. Remote Sens.*, vol. 107, pp. 99–111, Sept. 2015.

- [19] D. Tuia, E. Pasolli, and W. J. Emery, "Using active learning to adapt remote sensing image classifiers," *Remote Sens. Environ.*, vol. 115, no. 9, pp. 2232–2242, 2011.
- [20] K. Bahirat, F. Bovolo, L. Bruzzone, and S. Chaudhuri, "A novel domain adaptation Bayesian classifier for updating land-cover maps with class differences in source and target domains," *IEEE Trans. Geosci. Remote Sens.*, vol. 50, no. 7, pp. 2810–2826, 2012.
- [21] B. Demir, F. Bovolo, and L. Bruzzone, "Updating land-cover maps by classification of image time series: A novel change-detection-driven transfer learning approach," *IEEE Trans. Geosci. Remote Sens.*, vol. 51, no. 1, pp. 300–312, 2013.
- [22] G. Jun and J. Ghosh, "Semisupervised learning of hyperspectral data with unknown land-cover classes," *IEEE Trans. Geosci. Remote Sens.*, vol. 51, no. 1, pp. 273–282, 2013.
- [23] J. Heckman, "Sample selection bias as a specification error," *Econometrica*, vol. 47, no. 1, pp. 153–161, 1979.
- [24] B. Zadrozny, "Learning and evaluating classifiers under sample selection bias," in *Proc. 21st Int. Conf. on Machine Learn.*, Banff, Canada, 2004, p. 114.
- [25] J. Huang, A. Gretton, B. Schölkopf, A. J. Smola, and K. M. Borgwardt, "Correcting sample selection bias by unlabeled data," in *Proc. 21st Annu. Conf. Neural Inform. Proc. Syst.*, Vancouver, Canada, 2007, pp. 601–608.
- [26] C. Persello and L. Bruzzone, "Active and semisupervised learning for the classification of remote sensing images," *IEEE Trans. Geosci. Remote Sens.*, vol. 52, no. 11, pp. 6937–6956, 2014.
- [27] C. Persello and L. Bruzzone, "Active learning for domain adaptation in the supervised classification of remote sensing images," *IEEE Trans. Geosci. Remote Sens.*, vol. 50, no. 11, pp. 4468–4483, 2012.
- [28] C. Persello, "Interactive domain adaptation for the classification of remote sensing images using active learning," *IEEE Geosci. Remote Sens. Lett.*, vol. 10, no. 4, pp. 736–740, 2013.
- [29] E. Izquierdo-Verdiguier, V. Laparra, L. Gómez-Chova, and G. Camps-Valls, "Encoding invariances in remote sensing image classification with SVM," *IEEE Geosci. Remote Sens. Lett.*, vol. 10, no. 5, pp. 981–985, 2013.
- [30] L. Bruzzone and C. Persello, "A novel approach to the selection of spatially invariant features for the classification of hyperspectral images with improved generalization capability," *IEEE Trans. Geosci. Remote Sens.*, vol. 47, no. 9, pp. 3180–3191, 2009.
- [31] C. Persello and L. Bruzzone, "Kernel-based domain invariant feature selection in hyperspectral images for transfer learning," *IEEE Trans. Geosci. Remote Sens.*, vol. 54, no. 5, pp. 2615–2626, 2016.
- [32] J. Ham, Y. Chen, M. M. Crawford, and J. Ghosh, "Investigation of the random forest framework for classification of hyperspectral data," *IEEE Trans. Geosci. Remote Sens.*, vol. 43, no. 3, pp. 492–501, 2005.
- [33] S. Inamdar, F. Bovolo, L. Bruzzone, and S. Chaudhuri, "Multidimensional probability density function matching for preprocessing of multitemporal remote sensing images," *IEEE Trans. Geosci. Remote Sens.*, vol. 46, no. 4, pp. 1243–1252, 2008.
- [34] A. Nielsen and M. J. Canty, "Kernel principal component and maximum autocorrelation factor analyses for change detection," *Proc. SPIE*, vol. 7477, 74770T, 2009.
- [35] G. Matasci, M. Volpi, M. Kanevski, L. Bruzzone, and D. Tuia, "Semisupervised transfer component analysis for domain adaptation in remote sensing image classification," *IEEE Trans. Geosci. Remote Sens.*, vol. 53, no. 7, pp. 3550–3564, 2015.
- [36] A. A. Nielsen, "The regularized iteratively reweighted MAD method for change detection in multi- and hyperspectral data," *IEEE Trans. Image Proc.*, vol. 16, no. 2, pp. 463–478, 2007.
- [37] M. Volpi, G. Camps-Valls, and D. Tuia, "Spectral alignment of cross-sensor images with automated kernel canonical correlation analysis," *ISPRS J. Int. Soc. Photo. Remote Sens.*, vol. 107, pp. 50–63, Sept. 2015.
- [38] D. Tuia, M. Volpi, M. Trollet, and G. Camps-Valls, "Semisupervised manifold alignment of multimodal remote sensing images," *IEEE Trans. Geosci. Remote Sens.*, vol. 52, no. 12, pp. 7708–7720, 2014.
- [39] D. Tuia and G. Camps-Valls, "Kernel manifold alignment for domain adaptation," *PLoS One*, vol. 11, p. e0148655, Feb. 2016.
- [40] D. Marcos Gonzalez, G. Camps-Valls, and D. Tuia, "Weakly supervised alignment of multisensor images," in *Proc. IEEE Int. Geosci. and Remote Sens. Symp (IGARSS)*, Milan, Italy, 2015, pp. 2588–2591.
- [41] J. Montoya-Zegarra, C. Leistner, and K. Schindler, "Semantic tie points," in *Proc. IEEE Workshop Appl. Comput. Vision (WACV)*, Clearwater Beach, FL, 2013, pp. 377–384.
- [42] H. L. Yang and M. M. Crawford, "Spectral and spatial proximity-based manifold alignment for multitemporal hyperspectral image classification," *IEEE Trans. Geosci. Remote Sens.*, vol. 54, no. 1, pp. 51–64, 2016.
- [43] H. L. Yang and M. M. Crawford, "Domain adaptation with preservation of manifold geometry for hyperspectral image classification," *IEEE J. Sel. Topics Appl. Earth Observ.*, vol. 9, no. 2, pp. 543–555, 2016.
- [44] H. Sun, S. Liu, S. Zhou, and H. Zou, "Unsupervised cross-view semantic transfer for remote sensing image classification," *IEEE Geosci. Remote Sens. Lett.*, vol. 13, no. 1, pp. 13–17, 2016.
- [45] H. Sun, S. Liu, S. Zhou, and H. Zou, "Transfer sparse subspace analysis for unsupervised cross-view scene model adaptation," *IEEE J. Sel. Topics Appl. Earth Observ.*, 2016. DOI: 10.1109/JSTARS.2015.2500961
- [46] K. Zhang, B. Schölkopf, K. Muandet, Z. Wang, Z. Zhou, and C. Persello, "Single-source domain adaptation with target and conditional shift," in *Regularization, Optimization, Kernels, and Support Vector Machines*, J. A. K. Suykens, M. Signoretto, and A. Argyriou, Eds. Boca Raton, FL: CRC, 2014.
- [47] F. Petitjean, A. Ketterlin, and P. Gancarski, "A global averaging method for dynamic time warping, with applications to clustering," *Pattern Recogn.*, vol. 44, no. 3, pp. 678–693, 2011.
- [48] G. Jun and J. Ghosh, "Spatially adaptive classification of land cover with remote sensing data," *IEEE Trans. Geosci. Remote Sens.*, vol. 49, no. 7, pp. 2662–2673, 2011.



- [49] L. Zhang, L. Zhang, D. Tao, and X. Huang, "Sparse transfer manifold embedding for hyperspectral target detection," *IEEE Trans. Geosci. Remote Sens.*, vol. 52, no. 2, pp. 1030–1043, 2014.
- [50] D. Tuia, J. Muñoz-Marí, L. Gómez-Chova, and J. Malo, "Graph matching for adaptation in remote sensing," *IEEE Trans. Geosci. Remote Sens.*, vol. 51, no. 1, pp. 329–341, 2013.
- [51] E. Othman, Y. Bazi, N. Alajlan, H. AlHichri, and F. Melgani, "Three-layer convex network for domain adaptation in multitemporal VHR images," *IEEE Geosci. Remote Sens. Lett.*, vol. 13, no. 3, 2016.
- [52] F. Pacifici, J. Chanussot, and Q. Du, "2011 GRSS Data Fusion Contest: Exploiting WorldView-2 multi-angular acquisitions," in *Proc. IEEE Int. Geosci. and Remote Sens. Symp (IGARSS)*, Vancouver, Canada, 2011, pp. 1163–1166.
- [53] L. Bruzzone and D. F. Prieto, "Unsupervised retraining of a maximum-likelihood classifier for the analysis of multitemporal remote-sensing images," *IEEE Trans. Geosci. Remote Sens.*, vol. 39, no. 2, pp. 456–460, 2001.
- [54] L. Bruzzone and D. F. Prieto, "A partially unsupervised cascade classifier for the analysis of multitemporal remote-sensing images," *Pattern Recogn. Lett.*, vol. 23, no. 9, pp. 1063–1071, 2002.
- [55] L. Bruzzone and R. Cossu, "A multiple cascade-classifier system for a robust a partially unsupervised updating of land-cover maps," *IEEE Trans. Geosci. Remote Sens.*, vol. 40, no. 9, pp. 1984–1996, 2002.
- [56] L. Bruzzone, R. Cossu, and G. Vernazza, "Combining parametric and non-parametric algorithms for a partially unsupervised classification of multitemporal remote-sensing images," *Info. Fusion*, vol. 3, no. 4, pp. 289–297, 2002.
- [57] S. Rajan, J. Ghosh, and M. M. Crawford, "Exploiting class hierarchy for knowledge transfer in hyperspectral data," *IEEE Trans. Geosci. Remote Sens.*, vol. 44, no. 11, pp. 3408–3417, 2006.
- [58] L. Bruzzone, M. Chi, and M. Marconcini, "A novel transductive SVM for semisupervised classification of remote-sensing images," *IEEE Trans. Geosci. Remote Sens.*, vol. 44, no. 11, pp. 3363–3373, 2006.
- [59] M. Chi and L. Bruzzone, "Semi-supervised classification of hyperspectral images by SVMs optimized in the primal," *IEEE Trans. Geosci. Remote Sens.*, vol. 45, no. 6, pp. 1870–1880, 2007.
- [60] Z. Sun, C. Wang, H. Wang, and J. Li, "Learn multiple-kernel SVMs for domain adaptation in hyperspectral data," *IEEE Geosci. Remote Sens. Lett.*, vol. 10, no. 5, pp. 1224–1228, 2013.
- [61] J. M. Leiva-Murillo, L. Gómez-Chova, and G. Camps-Valls, "Multitask remote sensing data classification," *IEEE Trans. Geosci. Remote Sens.*, vol. 51, no. 1, pp. 151–161, 2013.
- [62] L. Gómez-Chova, G. Camps-Valls, J. Muñoz-Marí, and J. Calpe-Maravilla, "Semi-supervised image classification with Laplacian support vector machines," *IEEE Geosci. Remote Sens. Lett.*, vol. 5, no. 3, pp. 336–340, 2008.
- [63] W. Kim and M. M. Crawford, "Adaptive classification for hyperspectral image data using manifold regularization kernel machines," *IEEE Trans. Geosci. Remote Sens.*, vol. 48, no. 11, pp. 4110–4121, 2010.
- [64] L. Bruzzone and M. Marconcini, "Domain adaptation problems: A DASVM classification technique and a circular validation strategy," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 32, no. 5, pp. 770–787, 2010.
- [65] G. Jun and J. Ghosh, "An efficient active learning algorithm with knowledge transfer for hyperspectral data analysis," in *Proc. IEEE Int. Geosci. and Remote Sens. Symp (IGARSS)*, Boston, MA, 2008, pp. I-52–I-55.
- [66] S. Rajan, J. Ghosh, and M. M. Crawford, "An active learning approach to hyperspectral data classification," *IEEE Trans. Geosci. Remote Sens.*, vol. 46, no. 4, pp. 1231–1242, 2008.
- [67] G. Matasci, D. Tuia, and M. Kanevski, "SVM-based boosting of active learning strategies for efficient domain adaptation," *IEEE J. Sel. Topics Appl. Earth Observ.*, vol. 5, no. 5, pp. 1335–1343, 2012.
- [68] B. Settles, *Active Learning*. San Rafael, CA: Morgan & Claypool, 2012.
- [69] D. Tuia, M. Volpi, L. Copa, M. Kanevski, and J. Muñoz-Marí, "A survey of active learning algorithms for supervised remote sensing image classification," *IEEE J. Select Topics Signal Processing*, vol. 5, no. 3, pp. 606–617, 2011.
- [70] M. M. Crawford, D. Tuia, and H. L. Yang, "Active learning: Any value for classification of remotely sensed data?" *Proc. IEEE*, vol. 101, no. 3, pp. 593–608, 2013.
- [71] L. Bruzzone, C. Persello, and B. Demir, "Active learning methods in classification of remote sensing images," in *Signal and Image Processing for Remote Sensing*, C. Chen, Ed. Boca Raton, FL: CRC, 2012.
- [72] B. Demir, F. Bovolo, and L. Bruzzone, "Detection of land-cover transitions in multitemporal remote sensing images with active learning based compound classification," *IEEE Trans. Geosci. Remote Sens.*, vol. 50, no. 5, pp. 1930–1941, 2012.
- [73] A. Stumpf, N. Lachiche, J. P. Malet, N. Kerle, and A. Puissant, "Active learning in the spatial-domain for landslide mapping in remote sensing images," in *Proc. Euro. Conf. Machine Learn. (ECML)*, Active Learning in Real-World Applications Workshop, Bristol, U.K., 2012.
- [74] M. Roy, F. Melgani, A. Ghosh, E. Blanzieri, and S. Ghosh, "Land-cover classification of remotely sensed images using compressive sensing having severe scarcity of labeled patterns," *IEEE Geosci. Remote Sens. Lett.*, vol. 12, no. 6, pp. 1257–1261, 2015.
- [75] N. Alajlan, E. Pasolli, F. Melgani, and A. Franzoso, "Large-scale image classification using active learning," *IEEE Geosci. Remote Sens. Lett.*, vol. 11, no. 1, 259–263, 2014.
- [76] P. Blanchart, M. Ferecatu, S. Cui, and M. Datcu, "Pattern retrieval in large image databases using multiscale coarse-to-fine cascaded active learning," *IEEE J. Sel. Topics Appl. Earth Observ.*, vol. 7, no. 4, 1127–1141, 2014.
- [77] B. Demir, L. Minello, and L. Bruzzone, "Definition of effective training sets for supervised classification of remote sensing images by a novel cost-sensitive active learning method," *IEEE Trans. Geosci. Remote Sens.*, vol. 52, pp. 1272–1284, Feb. 2014.
- [78] C. Persello, A. Boularis, M. Dalponte, T. Gobakken, E. Naesset, and B. Schölkopf, "Cost-sensitive active learning with lookahead: Optimizing field surveys for remote sensing data classification," *IEEE Trans. Geosci. Remote Sens.*, vol. 52, no. 10, pp. 6652–6664, 2014.
- [79] T. Luo, K. Kramer, D. B. Goldgof, L. O. Hall, S. Samson, A. Remsen, and T. Hopkins, "Active learning to recognize multiple types of plankton," *J. Mach. Learn. Res.*, vol. 6, 589–613, Apr. 2005.