



HAL
open science

Get Your Embedding Space in Order: Domain-Adaptive Regression for Forest Monitoring

Sizhuo Li, Dimitri Gominski, Martin Brandt, Xiaoye Tong, Philippe Ciais

► **To cite this version:**

Sizhuo Li, Dimitri Gominski, Martin Brandt, Xiaoye Tong, Philippe Ciais. Get Your Embedding Space in Order: Domain-Adaptive Regression for Forest Monitoring. Computer Vision – ECCV 2024, 15135, Springer Nature Switzerland, pp.94-111, 2024, Lecture Notes in Computer Science, 10.1007/978-3-031-72980-5_6 . hal-04814425

HAL Id: hal-04814425

<https://hal.science/hal-04814425v1>

Submitted on 2 Dec 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.

Get Your Embedding Space in Order: Domain-Adaptive Regression for Forest Monitoring

Sizhuo Li^{1,2,*}, Dimitri Gominski^{1,*}, Martin Brandt¹, Xiaoye Tong¹, and
Philippe Ciais²

¹ University of Copenhagen, Denmark

² LSCE, Paris, France

{sizli,dg}@ign.ku.dk

Abstract. Image-level regression is an important task in Earth observation, where visual domain and label shifts are a core challenge hampering generalization. However, cross-domain regression within remote sensing data remains understudied due to the absence of suited datasets. We introduce a new dataset with aerial and satellite imagery in five countries with three forest-related regression tasks[†]. To match real-world applicative interests, we compare methods through a restrictive setup where no prior on the target domain is available during training, and models are adapted with limited information during testing. Building on the assumption that ordered relationships generalize better, we propose manifold diffusion for regression as a strong baseline for transduction in low-data regimes. Our comparison highlights the comparative advantages of inductive and transductive methods in cross-domain regression.

Keywords: Dataset · Regression · Remote sensing · Domain adaptation

1 Introduction

Distribution shifts between the training data and testing data are ubiquitous in computer vision. While it is generally accepted that a model trained on sufficient data has some ability to generalize to new target data, in practice visual variations induce significant performance drops [3].

A good example can be found in vegetation mapping and monitoring tasks. Biodiversity monitoring, carbon tracking, resource management, and food security analysis heavily depend on a regular assessment of forest resources through important parameters such as tree counts or height. A practical solution is to train image models to regress such values from satellite or aerial imagery [7, 28]. However, models trained on remote sensing imagery are usually limited to one area and perform poorly on others, due to different spectral ranges, ground resolutions, or visual patterns [34, 47, 49].

*Equal contribution

[†]Dataset and code available here: [dgominski.github.io/drift/](https://github.com/dgominski/drift/)

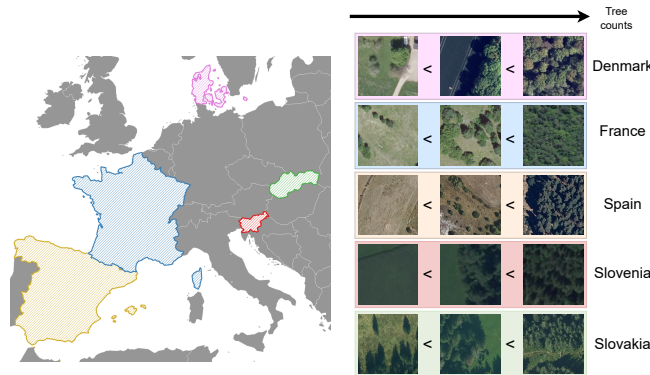


Fig. 1: Vegetation regression across countries. We hypothesize that order relations generalize better through visual and label domains than direct regression. Predicting tree counts directly requires domain-specific knowledge about local species, whereas predicting ordered chains remains intuitively easy.

Data availability and downstream applications define the scope of models (regional [29] – national [36] – continental [5]), but they don’t necessarily match the visual and semantic domains defined by biomes and environmental conditions. Intuitively, a model trained on a dataset should perform relatively well on datasets that are spatially close and visually similar, although it remains unclear to what extent.

To better assess generalization performance and go towards universal models that can be trained in countries with high data availability and applied to any country without prior knowledge or retraining, we introduce the DRIFT dataset for universal, low-data domain-adaptive regression.

Building on recent advances in order learning, we hypothesize that order relationships generalize well across domains (Fig. 1). We enforce a well-ordered embedding space in the source domain, then adapt to target domains in low-data setups. Notably, this strong regularization during training does not hamper source-domain performance **and** creates a clean embedding space to propagate similarities at test-time, which opens the way for advanced low-data inference in the form of manifold diffusion. Our contributions are as follows:

- We publish a large-scale dataset across five countries in Europe for domain-adaptive regression, using both aerial and satellite imagery, with a panel of tasks related to forest monitoring.
- We compare inductive and transductive methods across datasets and tasks, in a low-shot source-free domain adaptation setup. We show that transductive methods have an advantage when the domain gap becomes predominant.
- We introduce a simple but effective baseline for domain-adaptive regression. We propose a generalization of graph diffusion, a common technique for semi-supervised classification, to image-level regression. Our experiments demonstrate the applicative potential and superiority over common baselines.

2 Related Work

2.1 Order Learning

The seminal work of [32] introduced order learning as an alternative to direct regression, with the motivation that in certain cases, it is easier to compare images (this person is older than this other person) than to associate them independently with a scalar. Subsequent works built on this idea with iterative predictions [42], partial ordering [23] and further constraints on the geometry of the embedding space [24]. Our work builds on the latter. In a cross-domain setup, we hypothesize that the highly constrained embedding space with geometric constraints in the source domain generalizes better to other domains.

2.2 Cross-Domain Regression

Existing works have tackled domain adaptation through three main setups, depending on the available information about the target domain.

Unsupervised Domain Adaptation trains on unlabeled images of the target domain. Some works have focused on aligning features of the source and target domain with dedicated mechanisms to limit scale [8] or embedding structure [37] shifts, so that the regressor maintains its performance on the target domain. Adversarial training [35] has been proven effective as a tool to learn a domain shift metric.

Semi-Supervised Domain Adaptation, in contrast, assumes that some labeled images are available in the target domain. A common approach is to weight training points to minimize the performance gap [48], or as a form of boosting [39]. Few-shot domain adaptation can be considered a special case of semi-supervised adaptation, where the number of target domain labeled samples is very limited. A good example is the work of Teshima et al. [46], where authors generate synthetic target domain samples with the assumption that generative modeling offers more invariance.

These setups have limited applicability in many real-world problems due to the fact that utilizing labeled or unlabeled data from the target domain during training requires specific training for each target domain.

Source-Free Domain Adaptation removes the liberty of modifying the source domain training phase, which is not always possible in real settings. Models are only adapted at test-time, when applied on the target domain. In the context of image classification, the popular SHOT [31] method introduced information maximization and self-supervised learning as promising ideas for label-free target domain fine-tuning. If a few labeled samples are available, simple fine-tuning on the target domain with suited regularization remains a strong baseline [25]. We are not aware of works exploring this for regression.

Here, we aim at building a model that can be quickly adapted to any target domain, without any prior knowledge of that domain. As such, we prepare the model for generalization on the source domain, only with source domain data. At test-time, we explore low-data configurations, with very limited information

available on the target domain for quick adaptation. This setup has been labeled as **universal source-free domain adaptation** [21] in the context of classification.

2.3 Image-Level Regression in Earth Observation

The increasing availability of vast and diverse collections of remote sensing imagery and datasets presents both opportunities and challenges for computer vision tasks. To date, most works utilizing remote sensing imagery focus on image classification and semantic segmentation [2, 4, 40]. For example, to count trees, some works rely on the detection of individual trees [5, 26], but manual annotation relies on expert knowledge and is extremely time-consuming. Given the variations in ground resolution, image quality, and geographic areas, small objects and occlusion phenomena can make exhaustive annotation virtually impossible. Other labels on woody biomass or canopy height cannot be manually labeled. Therefore, labels often have to be sourced from field measures [9, 30], or derived from low-resolution dedicated sensors such as GEDI [13, 22], and remain noisy due to the varying viewing angles, technical inaccuracies in positioning systems [15] or temporal mismatch [6, 45]. As such, accurate labels are often found at a lower resolution than the ground resolution of the image. This makes image-level regression a core task in computer vision for Earth observation, as it offers a way to directly predict the information at the available resolution [7, 12]. To our knowledge, there is no existing dataset for cross-domain image-level regression using sub-meter resolution remote sensing imagery (Supplementary Tables 2 and 3).

3 DRIFT: Domain-Adaptive Regression for Forest Monitoring across Countries

We introduce the DRIFT dataset (Domain-adaptive Regression for Image-level Forest monIToring), including 25k images in five European countries from aerial and nanosatellite imagery, with three target variables to predict for each image:

1. Canopy height: average height value for pixels containing woody vegetation.
2. Tree count: number of overstory (visible from an overhead perspective) trees in the images.
3. Tree cover fraction: percentage of the image being covered by overstory tree crowns.

Dataset Collection. We built the dataset from national aerial orthophotography with various ground sampling distances (GSD) for Denmark (DK), Spain (SP), France (FR), and Slovakia (SK). All aerial images contain three spectral channels: Red (R), Green (G), and Blue (B). For Slovenia (SI), we used very high-resolution satellite imagery from commercial nanosatellites (SkySat). We acquired the orthorectified surface reflectance products at a spatial resolution of 50cm. Characteristics of each subset in the DRIFT are summarized in Table 1.

Refer to Supplementary section 1 for extended details regarding the dataset curation, data origins, and visual examples.

Table 1: DRIFT Dataset characteristics. GSD and Height in meters, Cover in percents.

| Country | # Images | Patch size | Sensor | GSD | Timeframe | Height | Range Count | Cover |
|----------|----------|------------|-----------|------|-----------|---------|----------------|----------|
| Denmark | 13094 | 200 | aerial | 0.2 | 2018 | [0, 38] | [0, 192] | [0, 100] |
| France | 10298 | 250 | aerial | 0.2 | 2018-20 | [0, 28] | [0, 204] | [0, 100] |
| Slovakia | 3129 | 250 | aerial | 0.2 | 2021 | [1, 38] | [0, 148] | [0, 100] |
| Spain | 1150 | 200 | aerial | 0.25 | 2020 | [0, 29] | [0, 135] | [0, 100] |
| Slovenia | 1903 | 100 | satellite | 0.5 | 2021 | [0, 36] | [0, 91] | [0, 100] |

Dataset Annotation. Aerial LiDAR is an accurate source of 3D structure information to characterize trees across various landscapes [19, 20, 50]. With the high-resolution height information, trees can be separated from other objects and measured in height and crown width. We first rasterized the LiDAR data to Canopy Height Models (CHMs), where each pixel indicates the top height of vegetation at this location, if there is any. Image-level canopy height was calculated by averaging all non-zero height values in the image, and tree cover fraction was extracted with 5-meter height thresholding to aggregate areas covered by vegetation taller than 5m [1]. Regarding tree counts, we used a public imagery-based tree counting model for Denmark, which was trained using point supervision with 22k tree annotations [27]. For the rest, we processed CHMs to extract tree crown centers with local peak identification over a sliding window with adaptive size [41]. CHMs were preprocessed with a conditional filter for higher separability of low foliage or artifacts [9]. We adjusted image-level counts with linear correction. All hyperparameters were fit on a small labeled set (\approx 1k trees) for each country, and we report the 3-fold R^2 score on random train/validation splits for correction in Supplementary Table 1.

Challenges. The DRIFT dataset includes significant shifts between label and visual distributions (see Fig. 2) due to sensor and area differences (see distribution of biomes in Supplementary Fig. 2). Moreover, within a tree species, individuals can exhibit different behaviors depending on their direct environment and living climate, which introduces concept drift: visually similar trees can vary in height [18]. Such variations can also be influenced by sensor properties such as the viewing angle.

Label Shift. The DRIFT dataset exhibits a substantial shift in label distribution across countries. We plot the label shift in Fig. 3 between all possible pairs of annotation subsets in DRIFT.

Regression Tasks. Canopy height, tree count, and tree cover heavily depend on tree species, local climate, and landscape characteristics. Height prediction from optical images is inherently challenging due to limited visual clues about depth from an aerial viewpoint, especially if the ground is not visible, but it

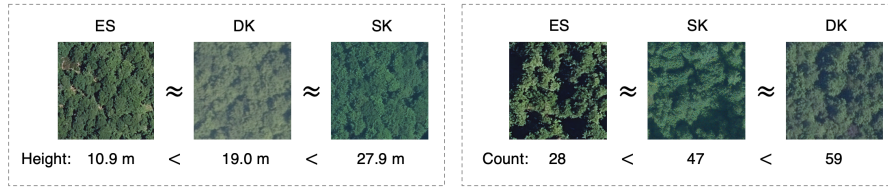


Fig. 2: Challenging examples in the DRIFT dataset: despite similar visual content, images can have different values in the label space. More visual examples can be found in Supplementary Fig. 1.

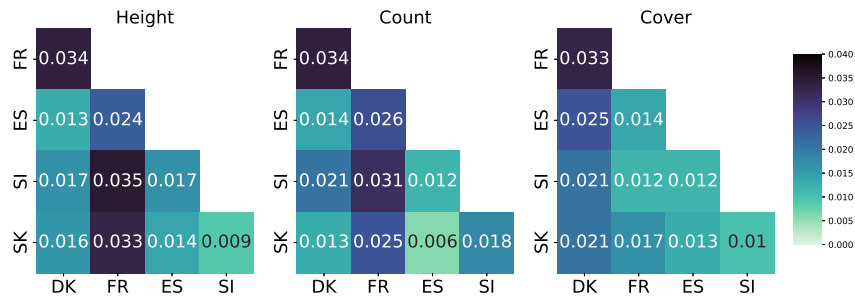


Fig. 3: Label shift varies across countries and tasks. We plot Wasserstein distances between label distributions in country subsets to indicate the level of label shift.

has been shown possible using satellite imagery [22, 33]. Counts depend on the model’s interpretation of individual crowns, which is subject to extreme variations even within a species [18]. Cover fraction is intuitively the easiest, as the model only needs to separate trees from background pixels. The difficulty depends on the level of similarity between short vegetation ($< 5\text{m}$) and trees.

Proposed Splits. We propose to use two countries (Denmark and France) with abundant high-quality CHM data as source domains for training. For testing, we apply models on three countries (Slovakia, Slovenia, Spain) representing different target domains. The Slovenia subset allows a comparison with another type of sensor (nanosatellites), and the Spain subset introduces a much drier ecosystem compared to the others. The France and Denmark subsets have different label distributions (Fig. 3), which introduces variation in the comparison of models trained on them.

Evaluation. We propose an evaluation framework for DRIFT that follows common constraints and represents real-world application contexts. First, we consider that no knowledge about the target domain is available during training. A model that can generalize to multiple areas is more valuable than models that are specific to a source/target dataset combination. As such, there is a high interest in going towards "universal" models that do not build on assumptions about their target domain. In a real-world scenario, a model that reached high

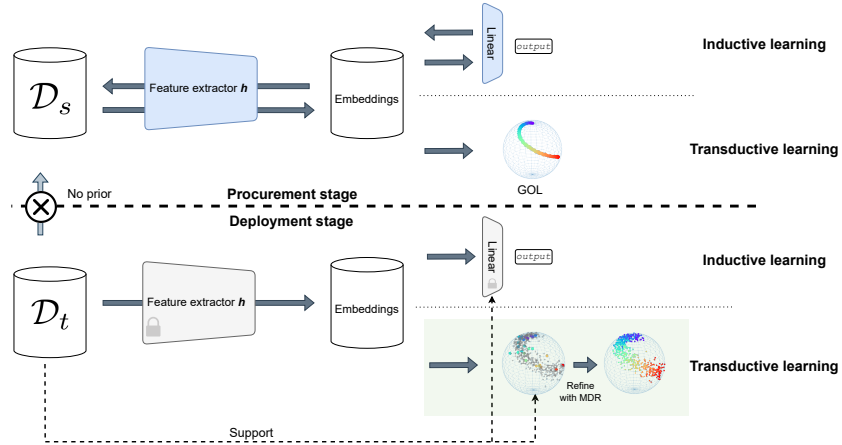


Fig. 4: Overview of our universal, source-free domain adaptation framework. The feature extractor (and linear layer for inductive regression) is trained on the source domain \mathcal{D}_s during the Procurement stage, then applied on the target domain \mathcal{D}_t during the Deployment stage. We do not use any prior regarding the target domain when training on the source domain, and do not assume sufficient data is available for training on the target domain. On the target domain, a few labeled examples are available to calibrate or fine-tune inductive methods, and support our proposed manifold diffusion for regression (MDR), a transductive approach.

performance on its source domain will be archived and considered for generalization, preferably with no or little supporting information about the target domain, and no retraining. Second, we consider that some knowledge about the target domain can be accessed at test-time. Following the terminology introduced in [21], we propose to follow a setup of universal source-free domain adaptation for regression, in two stages. In the Procurement stage, or source domain training, models are trained and prepared for generalization. During the Deployment stage, or target domain testing, models are adapted using a few labeled examples. In the following, we refer to this setup as N -shot adaptation.

4 Method

We explore the potential of ordered embeddings and simple transductive learning on the cross-domain challenges of the DRIFT dataset (see Fig. 4). We train the embedding extractor on the source domain, using geometric order learning to enforce structure. We extract embeddings on the target domain, and perform adaptation for regression using a few labels. Optionally, we refine predictions to take into account local and global manifold structure with diffusion.

4.1 Preliminaries

We first introduce notations for the problem of few-shot domain adaptation for regression.

Let $\mathcal{D}_s = \{(\mathbf{x}_i^s, y_i^s)\}$ and $\mathcal{D}_t = \{(\mathbf{x}_i^t, y_i^t)\}$ denote the source and target domains respectively. y^s and y^t take continuous values in \mathbb{R} , and we note that their distributions are different. Notably, it is common in vegetation mapping that their support $[y_{min}^s, y_{max}^s]$ and $[y_{min}^t, y_{max}^t]$ differs. For example, $y_{max}^s \neq y_{max}^t$ in height mapping due to different species.

At test-time, we assume that some information about the target domain is accessible, in the form of a few labeled samples, which we refer to as the target domain support set $\mathcal{D}_{ts} \subset \mathcal{D}_t$. Similarly to the common practice in classification [14, 44], we divide the target label set into k_r equally spaced value groups, and construct the target domain support set \mathcal{D}_{ts} by randomly drawing N samples from each group, forming a k_r -way N -shot adaptation setup.

A simple baseline method would be to train a regressor $f(\mathbf{x}) = y$ on \mathcal{D}_s and apply it to \mathcal{D}_t . Due to domain shifts, this performs generally poorly. Instead, we consider a feature extractor $h(\mathbf{x}) = \mathbf{v}$, with $\mathbf{v} \in \mathbb{R}^d$ and $\|\mathbf{v}\|_2 = 1$, and frame the problem as cross-domain transduction: y^t is inferred from the relationships between the embeddings in \mathcal{D}_t .

Geometric order learning (GOL) [24] enforces the notion of order in the embedding space, such that the direction and distance between embeddings reflect the relations between their labels. The framework is based on M reference points that are shared between embeddings of similar values. Consider $\mathbf{v}_a, \mathbf{v}_b$ and the corresponding reference points $r_{\theta(a)}, r_{\theta(b)}$, with θ the function assigning a sample to its reference point. GOL trains with the following composite loss function:

$$L = L_o + L_m + L_c \quad (1)$$

where

- L_o , the order loss, aligns the embeddings with the direction of their label relation. For the $r_{\theta(a)} < r_{\theta(b)}$ case, $L_o = -\log \frac{e^{\vec{u}_+ \cdot \vec{u}_{ab}}}{(e^{\vec{u}_+ \cdot \vec{u}_{ab}} + e^{\vec{u}_- \cdot \vec{u}_{ab}})}$ with $\vec{u}_+ = (r_{\theta(b)} - r_{\theta(a)}) / \|r_{\theta(b)} - r_{\theta(a)}\|$ the forward direction vector between the reference points of a and b , $\vec{u}_- = (r_{\theta(a)-1} - r_{\theta(a)}) / \|r_{\theta(a)-1} - r_{\theta(a)}\|$ the backward direction vector, and $\vec{u}_{ab} = (\mathbf{v}_b - \mathbf{v}_a) / \|\mathbf{v}_b - \mathbf{v}_a\|$. When $r_{\theta(a)} > r_{\theta(b)}$, the corresponding formula is obtained symmetrically.
- L_m , the metric loss, makes embedding space distances reflect label differences. When $r_a = r_b$, $L_m = \sum_{i=0}^{M-1} \max(|d(r_i, \mathbf{v}_a) - d(r_i, \mathbf{v}_b)| - \gamma, 0)$, with d the Euclidean distance. When $r_{\theta(a)} < r_{\theta(b)}$, $L_m = \sum_{i=0}^a \max(|d(r_i, \mathbf{v}_a) - d(r_i, \mathbf{v}_b)| + \gamma, 0) + \sum_{j=b}^{M-1} \max(|d(r_j, \mathbf{v}_b) - d(r_j, \mathbf{v}_a)| + \gamma, 0)$. That is, all reference points inferior to $r_{\theta(a)}$ pull \mathbf{v}_a closer, while all reference points superior to $r_{\theta(b)}$ push \mathbf{v}_a away, and the opposite for \mathbf{v}_b . The formula for $r_{\theta(a)} > r_{\theta(b)}$ is obtained symmetrically.
- L_c , the center loss, places reference points at the center of embeddings with similar values. $L_c = d(r_{\theta(a)}, \mathbf{v}_a) + d(r_{\theta(b)}, \mathbf{v}_b)$.

We invite interested readers to refer to the original paper for more details.

Lee et al. [24] report that GOL creates a well-ordered embedding space, which allows prediction with a simple distance-based k-nearest neighbors (kNN) assignment to estimate continuous values.

4.2 Manifold Diffusion for Regression

We introduce Manifold Diffusion for Regression (MDR) as an optional post-processing step to refine the predictions. Manifold diffusion [10, 17, 51] is a popular transductive method in semi-supervised learning and image retrieval to exploit the structure of the embedding space for better class predictions. Here, we modify it for regression in the context of cross-domain adaptation.

We start by introducing the diffusion framework originally proposed in [51]. Assuming a set $X = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n\}$ and $\mathbf{v}_i = h(\mathbf{x}_i)$, an affinity matrix $W \in \mathbb{R}^{n \times n}$ is constructed by calculating the pairwise cosine similarity between samples:

$$W_{ij} := \begin{cases} [\mathbf{v}_i^\top \mathbf{v}_j]_+^\gamma, & \text{if } i \neq j \wedge \mathbf{v}_i \in \text{NN}_k(\mathbf{v}_j) \\ 0, & \text{otherwise} \end{cases} \quad (2)$$

with NN_k the set of k nearest neighbors in X , and γ a hyperparameter. W is non-negative, sparse, and has zero diagonal.

The symmetric normalized form is computed:

$$W_n := D^{-1/2} W D^{1/2} \quad (3)$$

where D is the diagonal matrix $D := \text{diag}(W \mathbf{1}_n)$ with the row-wise sums of W . The normalized graph Laplacian of W is defined as $L := I_n - W_n$. Assuming that we have label information for some of the points, we can iteratively update the graph with $F(t+1) = \alpha W F(t) + (1-\alpha)Y$, where $Y \in \{0, 1\}^{n \times c}$ is the one-hot encoded label matrix for c classes, and α a hyperparameter. Zhou et al. [51] noted that this is equivalent to solving the linear system: $F^* := (I_n - \alpha W_n)^{-1} Y$. The final class prediction for a sample i is obtained as $\text{ARGMAX}(F_i)$.

We replace the label matrix Y with the support matrix $S \in \{0, 1\}^{n \times k}$, the one-hot encoded matrix for each of the k labeled samples. $S_{i,j} = 1$ for the first k rows and columns, and zero elsewhere. After solving for $S^* = (I_n - \alpha W_n)^{-1} S$, we obtain an updated similarity score between an image and each of the support elements. S^* captures the structure of the embedding space and creates a refined similarity measure taking into account local and global context.

We output the final regression estimate as the weighted sum of the k_v highest values in S^* and the known labels. Note that this step heavily relies on a well-ordered embedding space: predictions will be unstable if the embedding space does not maintain the order of the one-dimensional continuous label space.

5 Experiments

5.1 Implementation Details

During the Procurement stage, we train a ViT-B/16 [11] feature extractor with GOL with $M = k_r = 5$ and equal interval between label groups. For canopy height, the upper bound is set to 25m, and each group spans over 5m. For count, we use 0 as the lower bound and 90 (Denmark) or 70 (France) as the upper bound and split accordingly. For cover, we use 20% as the lower bound and 100% as the upper bound. We weight the order loss L_o , metric loss L_m , and center loss L_c with a factor of 1, 66, and 33, respectively, for similar convergence rates.

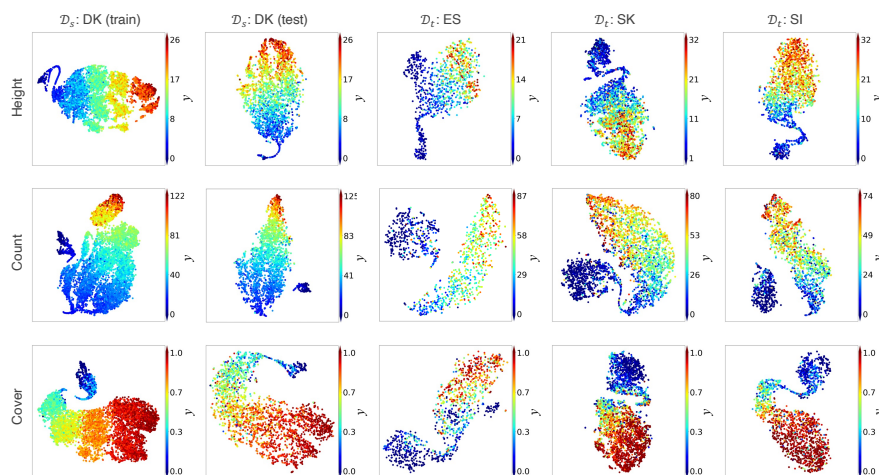


Fig. 5: Embeddings retain ordered relations across domains (t-SNE visualization). From left to right: source domain training data, source domain testing data, and three target domains, respectively.

During the Deployment stage, we pick the 1% and 99% percentile values of each target variable as the lower and upper bound and split the entire dataset into 5 groups.

We implement MDR in a few-shot setup. We use standard hyperparameters that are reported to work well across different setups ($\gamma = 3$, $\alpha = 0.99$) [16], and use the entire support set for diffusion, setting $k = N$ and assigning the final values with $k_v = 2N$, with the effects of k_v shown in Supplementary Fig. 6. We refer to this approach as GOL+MDR.

5.2 Qualitative Analysis

We plot projections of the embedding space in the source and target domains across tasks on the DRIFT dataset in Fig. 5. Label order in the embedding

space is roughly maintained across domains, but does not necessarily remain in a single well-connected cluster with a clear direction of label space increase. Consequently, kNN label assignment can pick up noisy neighbors from adjacent clusters and wrongly predict.

We plot the effect of MDR on DRIFT examples in Fig. 6. MDR activates more relevant support samples than kNN, *i.e.* the examples picked up for assigning the prediction are closer to the true value. This comes from taking into account both global and local manifold structure, rather than only the immediate neighbors. Assuming that embeddings remain overall ordered on the target domain, MDR captures higher-order similarities and discards potentially confusing support samples in the target domain. Example predictions can be found in Supplementary Fig. 3 and 4.

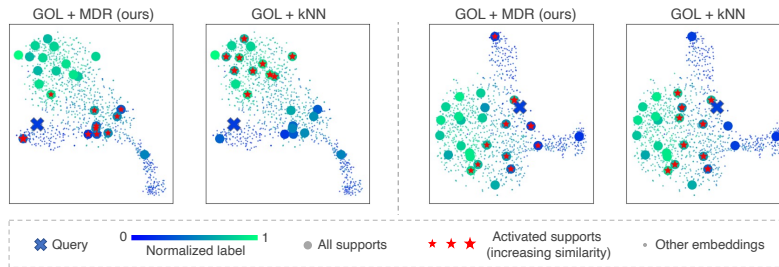


Fig. 6: Top 10 support samples activated by GOL+MDR and GOL+kNN in 5-shot setup (t-SNE visualization). Closer colors indicate higher adjacency in the label space.

5.3 Quantitative Analysis

Few-Shot Evaluation. We compare our approach with common baselines and state-of-the-art methods in Table 2. We compare against regression on the target domain without adaptation (Regression), regression on the target domain with linear calibration using the support set (Regression cal.), weighted k-nearest neighbor prediction from support examples (kNN), fine-tuning with the support set (FT) [25], geometric order learning with weighted kNN prediction (GOL) [24].

The off-the-shelf regressor has overall low and unstable performance, as expected. Calibration with the support set brings a significant but inconsistent boost in performance. This is particularly visible on the Denmark \rightarrow Spain setup, for which calibration helps only on the cover task. Similarly, fine-tuning sometimes improves performance, but it relies on how well the underlying model generalizes to the target dataset. All inductive methods experience failure cases, with R^2 scores below 0.3 in at least one setup.

Transductive approaches, on the other hand, decouple to some extent the performance on the target domain from the source domain. kNN performs overall

Table 2: Comparison of methods on the DRIFT dataset (5-shot). Average R^2 score across 3 random runs for FT and 10 random runs for the rest.

| Source | | Denmark | | | France | | | |
|--------|-----------------|------------------|------------------|------------------|------------------|------------------|------------------|-------------|
| Target | Spain | Slovakia | Slovenia | Spain | Slovakia | Slovenia | <i>avg</i> | |
| Height | Regression | 0.35 | -0.05 | 0.23 | 0.61 | 0.11 | 0.13 | 0.23 |
| | Regression cal. | 0.17±0.12 | 0.49±0.10 | 0.67±0.03 | 0.61±0.06 | 0.42±0.06 | 0.67±0.02 | 0.50 |
| | FT [25] | 0.34±0.00 | 0.45±0.02 | 0.60±0.02 | 0.69±0.01 | 0.49±0.00 | 0.50±0.09 | 0.51 |
| | kNN | 0.21±0.06 | 0.45±0.04 | 0.68±0.03 | 0.55±0.04 | 0.38±0.06 | 0.63±0.03 | 0.48 |
| | GOL [24] | 0.21±0.12 | 0.35±0.13 | 0.67±0.04 | 0.62±0.05 | 0.41±0.08 | 0.65±0.03 | 0.49 |
| | GOL+MDR | 0.60±0.04 | 0.49±0.06 | 0.65±0.05 | 0.65±0.04 | 0.46±0.05 | 0.66±0.04 | 0.59 |
| | GOL+FT+MDR | 0.56±0.07 | 0.57±0.03 | 0.69±0.01 | 0.70±0.01 | 0.47±0.01 | 0.71±0.01 | 0.62 |
| Count | Regression | 0.53 | 0.36 | 0.43 | 0.51 | 0.43 | 0.37 | 0.44 |
| | Regression cal. | 0.51±0.04 | 0.53±0.05 | 0.59±0.03 | 0.53±0.07 | 0.61±0.04 | 0.58±0.04 | 0.56 |
| | FT [25] | 0.53±0.00 | 0.48±0.00 | 0.60±0.01 | 0.54±0.00 | 0.53±0.01 | 0.40±0.00 | 0.51 |
| | kNN | 0.42±0.07 | 0.43±0.06 | 0.55±0.03 | 0.51±0.06 | 0.55±0.14 | 0.55±0.08 | 0.50 |
| | GOL [24] | 0.51±0.03 | 0.56±0.03 | 0.58±0.05 | 0.51±0.05 | 0.58±0.05 | 0.46±0.06 | 0.53 |
| | GOL+MDR | 0.51±0.04 | 0.55±0.06 | 0.59±0.03 | 0.54±0.04 | 0.61±0.03 | 0.54±0.04 | 0.56 |
| | GOL+FT+MDR | 0.53±0.01 | 0.49±0.10 | 0.59±0.05 | 0.51±0.09 | 0.63±0.01 | 0.53±0.03 | 0.55 |
| Cover | Regression | 0.11 | 0.70 | 0.81 | 0.80 | 0.77 | 0.76 | 0.66 |
| | Regression cal. | 0.41±0.04 | 0.75±0.01 | 0.77±0.04 | 0.83±0.02 | 0.79±0.01 | 0.77±0.04 | 0.72 |
| | FT [25] | 0.27±0.00 | 0.44±0.22 | 0.75±0.01 | 0.81±0.01 | 0.81±0.00 | 0.81±0.00 | 0.65 |
| | kNN | 0.47±0.04 | 0.69±0.02 | 0.71±0.06 | 0.75±0.04 | 0.73±0.05 | 0.68±0.07 | 0.67 |
| | GOL [24] | 0.51±0.06 | 0.76±0.02 | 0.69±0.10 | 0.81±0.03 | 0.82±0.02 | 0.71±0.09 | 0.72 |
| | GOL+MDR | 0.57±0.08 | 0.71±0.03 | 0.70±0.06 | 0.82±0.03 | 0.79±0.03 | 0.70±0.06 | 0.71 |
| | GOL+FT+MDR | 0.65±0.05 | 0.74±0.04 | 0.64±0.05 | 0.84±0.01 | 0.76±0.03 | 0.70±0.07 | 0.72 |

more consistently than inductive approaches. It however performs worse than calibrated regression on average. GOL has a similar behavior, with a minor increase in performance undoubtedly due to better regularization in the source domain. GOL+MDR performs the best among transductive approaches, and especially well on the difficult setups of Denmark \rightarrow Spain. Combining with fine-tuning on the support set further increases performance on some setups.

We note that the performance of direct cross-domain regression can serve as an indicator of a multifactorial domain gap. For example, on cover regression, where the label shift is limited, direct regression performs well, except for Denmark \rightarrow Spain which is arguably the largest visual and semantic domain shift (Supplementary Fig. 2). We plot in Fig. 7 a pairwise comparison of the two best-performing inductive and transductive methods against off-the-shelf regression performance drop. When the implied domain gap becomes large enough for inductive methods to fail, transductive methods are favored.

We repeat experiments with VGG16 [43], another popular feature extractor. GOL+MDR and GOL outperformed regression baselines for height and count with larger margins than using the ViT-B/16 backbone (Supplementary Table 4).

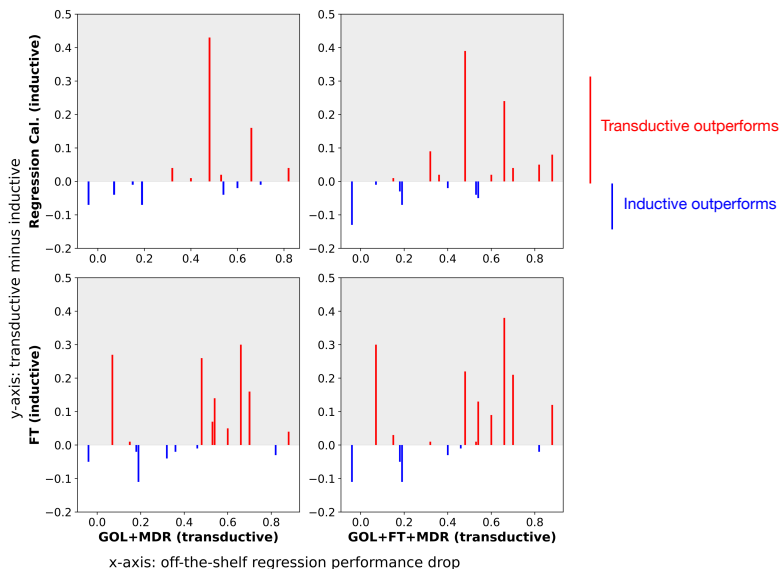


Fig. 7: Transductive methods outperform inductive methods when the domain gap is larger. We plot pairwise performance differences between inductive and transductive methods on the y-axis, with off-the-shelf regression performance drop (source domain performance minus target domain performance) on the x-axis, an ad-hoc indicator of domain gap.

In-Domain Performance. We report in Table 3 the performance of GOL in the source domain. As reported in the original publication [24], GOL performs on par with ordinary regression.

Table 3: Source domain test performance. Average R^2 score across two source domains.

| Backbone | Model | Height | Count | Cover |
|----------|------------|-----------------|-----------------|-----------------|
| ViT-B/16 | GOL [24] | 0.86 ± 0.05 | 0.92 ± 0.01 | 0.87 ± 0.08 |
| | Regression | 0.88 ± 0.05 | 0.90 ± 0.01 | 0.86 ± 0.09 |
| VGG16 | GOL [24] | 0.80 ± 0.07 | 0.91 ± 0.01 | 0.80 ± 0.13 |
| | Regression | 0.84 ± 0.06 | 0.93 ± 0.01 | 0.86 ± 0.09 |

5.4 Ablation Studies

We report in Table 4 the effect of removing components of GOL+MDR, averaging across countries in the DRIFT dataset. MDR does not improve performance compared to the baseline when applied without ordered embeddings: the

manifold in the target domain is too unstructured to give meaningful similarities. Ordered embeddings (GOL) generalize better in all tasks, and combining GOL+MDR gives the best results in 5 out of 6 experiments.

Table 4: Ablation study for combining GOL with the proposed MDR. Average R^2 score across two source domains and three target domains. We report the performance of a model with only the center loss and weighted kNN prediction as the baseline.

| Backbone | GOL | MDR | Height | Count | Cover |
|----------|-----|-----|------------------|------------------|------------------|
| ViT-B/16 | ✓ | ✓ | 0.59±0.08 | 0.56±0.03 | 0.71±0.08 |
| | ✓ | | 0.49±0.17 | 0.53±0.04 | 0.72±0.10 |
| | | ✓ | 0.40±0.17 | 0.43±0.11 | 0.65±0.10 |
| | | | 0.48±0.16 | 0.50±0.06 | 0.67±0.09 |
| VGG16 | ✓ | ✓ | 0.55±0.09 | 0.53±0.05 | 0.67±0.07 |
| | ✓ | | 0.47±0.16 | 0.51±0.05 | 0.64±0.09 |
| | | ✓ | 0.25±0.16 | 0.39±0.11 | 0.59±0.08 |
| | | | 0.31±0.13 | 0.34±0.11 | 0.61±0.08 |

We test the sensitivity of different methods to the number of shots (Supplementary Fig. 5), *i.e.* the number of support samples per group.

As often observed in few-shot setups, performance saturates after a certain number of shots [38, 44], here 20. Our proposed GOL+MDR outperforms the other methods on most of the setups, except for the easiest task of tree cover, where finetuning is needed to reach comparable performance as other baselines.

An important parameter in our approach is k_v , the number of support samples used for weighted sum prediction after MDR. We tested a set of values and observed a safe range around $k_v = 2N$ (Supplementary Fig. 6).

6 Conclusion

We introduce the DRIFT dataset with forest-related image regression tasks across five European countries. We reveal different behaviors depending on tasks and target countries, which underlines the importance of taking into account the domain gap when choosing between inductive and transductive approaches. DRIFT covers a panel of situations with varying levels of difficulty and domain gaps. As such, it is close to real-world conditions, where those factors are difficult to estimate beforehand. In this context, we argue that transductive methods are a good starting point, as they allow qualitative analysis of the embedding space, even before having to label samples for quantitative evaluation. In contrast, inductive approaches remain "black boxes" with poor explainability.

We verify qualitatively and quantitatively that the embedding order can be maintained across domains, which unlocks advanced adaptation with diffusion approaches. This works particularly well in low-data regimes, and when the domain gap is high.

Acknowledgements

PC acknowledges support from the European Space Agency Climate Change Initiative (ESA-CCI) Biomass project (ESA ESRIN/ 4000123662) and RECCAP2 project 1190 (ESA ESRIN/ 4000123002/ 18/I-NB) and the ANR BMBF French German AI4FOREST (ANR-22-FAI1-0002).

MB and SL were supported by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement no. 947757 TOFDRI) and a DFF Sapere Aude grant (no. 9064-00049B).

DG acknowledge support by the European Union’s Eurostars programme through the C-Trees project, grant number E114613.

XT acknowledges funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement no. 947757 TOFDRI).

Image Credits. Figures 1, 2 and Supplementary Figures 1, 3, and 4 contain data from the Danish Agency for Data Supply and Infrastructure (2018), the BD ORTHO© product from the French Institute for Geography (2018-2020), the Orthophoto and ALS products from the Slovakian authorities GKÚ Bratislava and ÚGKK SR (2021), the PNOA product from the Spanish Institute of Geography (scne.es, 2020), the SkySat sensor from Planet labs PBC, and the Slovenian Environment Agency (ARSO, 2021).

References

1. Trees outside forests - towards a better awareness, <https://www.fao.org/3/y2328e/y2328e25.htm>
2. Akiva, P., Purri, M., Leotta, M.: Self-supervised material and texture representation learning for remote sensing tasks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 8203–8215 (2022)
3. Beery, S., Van Horn, G., Perona, P.: Recognition in terra incognita. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 472–489 (2018)
4. Beery, S., Wu, G., Edwards, T., Pavetic, F., Majewski, B., Mukherjee, S., Chan, S., Morgan, J., Rathod, V., Huang, J.: The auto arborist dataset: A large-scale benchmark for multiview urban forest monitoring under domain shift. In: 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 21262–21275 (2022)
5. Brandt, M., Tucker, C.J., Kariryaa, A., Rasmussen, K., Abel, C., Small, J., Chave, J., Rasmussen, L.V., Hiernaux, P., Diouf, A.A., Kergoat, L., Mertz, O., Igel, C., Gieseke, F., Schöning, J., Li, S., Melocik, K., Meyer, J., Sinno, S., Romero, E., Glennie, E., Montagu, A., Dendoncker, M., Fensholt, R.: An unexpectedly large count of trees in the west african sahara and sahel. *Nature* **587**(7832), 78–82 (2020)
6. Burgert, T., Ravanbakhsh, M., Demir, B.: On the effects of different types of label noise in multi-label remote sensing image classification. *IEEE Transactions on Geoscience and Remote Sensing* (2022)
7. Chen, G., Shang, Y.: Transformer for tree counting in aerial images. *Remote Sensing* **14**(3), 476 (2022)

8. Chen, X., Wang, S., Wang, J., Long, M.: Representation subspace distance for domain adaptation regression. In: International Conference on Machine Learning. pp. 1749–1759 (2021)
9. Deng, S., Katoh, M., Yu, X., Hyyppä, J., Gao, T.: Comparison of tree species classifications at the individual tree level by combining ALS data and RGB images using different algorithms. *Remote Sensing* **8**(12), 1034 (2016)
10. Donoser, M., Bischof, H.: Diffusion processes for retrieval revisited. In: 2013 IEEE Conference on Computer Vision and Pattern Recognition. pp. 1320–1327 (2013)
11. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houlsby, N.: An image is worth 16x16 words: Transformers for image recognition at scale. In: International Conference on Learning Representations (2021)
12. DrivenData: The BioMassters, <https://www.drivendata.org/competitions/99/biomass-estimation/page/534/>
13. Fayad, I., Ciais, P., Schwartz, M., Wigneron, J.P., Baghdadi, N., de Truchis, A., d’Aspremont, A., Frappart, F., Saatchi, S., Sean, E., Pellissier-Tanon, A., Bazzi, H.: Hy-tec: a hybrid vision transformer model for high-resolution and large-scale mapping of canopy height. *Remote Sensing of Environment* **302**, 113945 (2024)
14. Finn, C., Abbeel, P., Levine, S.: Model-agnostic meta-learning for fast adaptation of deep networks. In: Proceedings of the 34th International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 70, pp. 1126–1135. PMLR (2017)
15. Girard, N., Charpiat, G., Tarabalka, Y.: Noisy supervision for correcting misaligned cadaster maps without perfect ground truth data. In: 2019 IEEE International Geoscience and Remote Sensing Symposium. pp. 10103–10106 (2019)
16. Iscen, A., Tolia, G., Avrithis, Y., Chum, O.: Label propagation for deep semi-supervised learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
17. Iscen, A., Tolia, G., Avrithis, Y., Furon, T., Chum, O.: Efficient diffusion on region manifolds: Recovering small objects with compact CNN representations. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 926–935 (2017)
18. Jucker, T., Caspersen, J., Chave, J., Antin, C., Barbier, N., Bongers, F., Dalponte, M., van Ewijk, K.Y., Forrester, D.I., Haeni, M., Higgins, S.I., Holdaway, R.J., Iida, Y., Lorimer, C., Marshall, P.L., Momo, S., Moncrieff, G.R., Ploton, P., Poorter, L., Rahman, K.A., Schlund, M., Sonké, B., Sterck, F.J., Trugman, A.T., Usoltsev, V.A., Vanderwel, M.C., Waldner, P., Wedeux, B.M.M., Wirth, C., Wöll, H., Woods, M., Xiang, W., Zimmermann, N.E., Coomes, D.A.: Allometric equations for integrating remote sensing imagery into forest monitoring programmes. *Global Change Biology* **23**(1), 177–190 (2017)
19. Kalinicheva, E., Landrieu, L., Mallet, C., Chehata, N.: Multi-layer modeling of dense vegetation from aerial LiDAR scans. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops. pp. 1342–1351 (2022)
20. Knapp, N., Fischer, R., Huth, A.: Linking lidar and forest modeling to assess biomass estimation across scales and disturbance states. *Remote Sensing of Environment* **205**, 199–209 (2018)
21. Kundu, J.N., Venkat, N., M V, R., Babu, R.V.: Universal source-free domain adaptation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. (2020)

22. Lang, N., Jetz, W., Schindler, K., Wegner, J.D.: A high-resolution canopy height model of the Earth. *Nature Ecology & Evolution* **7**(11), 1778–1789 (2023)
23. Lee, S.H., Kim, C.S.: Order learning using partially ordered data via chainization. In: *ECCV 2022*. pp. 196–211 (2022)
24. Lee, S.H., Shin, N.H., Kim, C.S.: Geometric order learning for rank estimation. In: *Advances in Neural Information Processing Systems*. vol. 35, pp. 27–39 (2022)
25. Lee, S., Seo, S., Kim, J., Lee, Y., Hwang, S.: Few-shot fine-tuning is all you need for source-free domain adaptation (2023), arXiv preprint arXiv:2304.00792
26. Li, S., Brandt, M., Fensholt, R., Kariryaa, A., Igel, C., Gieseke, F., Nord-Larsen, T., Oehmcke, S., Carlsen, A.H., Junntila, S., Tong, X., d’Aspremont, A., Ciais, P.: Deep learning enables image-based tree counting, crown segmentation and height prediction at national scale. *PNAS Nexus* **2**(4) (2023)
27. Li, S., Brandt, M., Fensholt, R., Kariryaa, A., Igel, C., Gieseke, F., Nord-Larsen, T., Oehmcke, S., Carlsen, A.H., Junntila, S., et al.: Deep learning enables image-based tree counting, crown segmentation, and height prediction at national scale. *PNAS nexus* **2**(4) (2023)
28. Li, S., Brandt, M., Tong, X., Oehmcke, S., Igel, C., Gieseke, F., Nord-Larsen, T., Fensholt, R., Chave, J., Ciais, P.: Deep learning tree and forest biomass from sub-meter resolution images (2023), researchSquare preprint 10.21203/rs.3.rs-3335298
29. Li, W., Buitenwerf, R., Munk, M., Bøcher, P.K., Svenning, J.C.: Deep-learning based high-resolution mapping shows woody vegetation densification in greater maasai mara ecosystem. *Remote Sensing of Environment* **247**, 111953 (2020)
30. Li, Y., Li, M., Li, C., Liu, Z.: Forest aboveground biomass estimation using landsat 8 and sentinel-1a data with machine learning algorithms. *Scientific Reports* **10**(1), 9952 (2020)
31. Liang, J., Hu, D., Feng, J.: Do we really need to access the source data? source hypothesis transfer for unsupervised domain adaptation. In: *International Conference on Machine Learning (ICML)*. pp. 6028–6039 (2020)
32. Lim, K., Shin, N.H., Lee, Y.Y., Kim, C.S.: Order learning and its application to age estimation. In: *International Conference on Learning Representations* (2020)
33. Liu, S., Brandt, M., Nord-Larsen, T., Chave, J., Reiner, F., Lang, N., Tong, X., Ciais, P., Igel, C., Pascual, A., Guerra-Hernandez, J., Li, S., Mugabowindekwe, M., Saatchi, S., Yue, Y., Chen, Z., Fensholt, R.: The overlooked contribution of trees outside forests to tree cover and woody biomass across europe. *Science Advances* **9**(37) (2023)
34. Marsocci, V., Gonthier, N., Garioud, A., Scardapane, S., Mallet, C.: Geomulti-tasknet: Remote sensing unsupervised domain adaptation using geographical coordinates. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*. pp. 2075–2085 (2023)
35. Mathelin, A.d., Richard, G., Deheeger, F., Mougeot, M., Vayatis, N.: Adversarial weighting for domain adaptation in regression. In: *2021 IEEE 33rd International Conference on Tools with Artificial Intelligence (ICTAI)*. pp. 49–56. IEEE Computer Society (2021)
36. Mugabowindekwe, M., Brandt, M., Chave, J., Reiner, F., Skole, D.L., Kariryaa, A., Igel, C., Hiernaux, P., Ciais, P., Mertz, O., Tong, X., Li, S., Rwanyiziri, G., Dushimiyimana, T., Ndoli, A., Uwizeyimana, V., Lillesø, J.P.B., Gieseke, F., Tucker, C.J., Saatchi, S., Fensholt, R.: Nation-wide mapping of tree-level above-ground carbon stocks in rwanda. *Nature Climate Change* **13**(1), 91–97 (2023)
37. Nejjar, I., Wang, Q., Fink, O.: DARE-GRAM : Unsupervised domain adaptation regression by aligning inversed gram matrices. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. (2023)

38. Nguyen, K.D., Tran, Q.H., Nguyen, K., Hua, B.S., Nguyen, R.: Inductive and transductive few-shot video classification via appearance and temporal alignments. In: ECCV 2022. pp. 471–487 (2022)
39. Pardoe, D., Stone, P.: Boosting for regression transfer. In: International Conference on Machine Learning (2010)
40. Robinson, C., Hou, L., Malkin, K., Soobitsky, R., Czawlytko, J., Dilkina, B., Jojic, N.: Large scale high-resolution land cover mapping with multi-resolution data. In: 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 12718–12727 (2019)
41. Roussel, J.R., Auty, D., Coops, N.C., Tompalski, P., Goodbody, T.R.H., Meador, A.S., Bourdon, J.F., de Boissieu, F., Achim, A.: lidR: An r package for analysis of airborne laser scanning (ALS) data. *Remote Sensing of Environment* **251**, 112061 (2020)
42. Shin, N.H., Lee, S.H., Kim, C.S.: Moving window regression: a novel approach to ordinal regression. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (2022)
43. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. 3rd International Conference on Learning Representations (ICLR 2015) pp. 1–14 (2015)
44. Snell, J., Swersky, K., Zemel, R.: Prototypical networks for few-shot learning. In: Advances in Neural Information Processing Systems (2017)
45. Sumbul, G., Demir, B.: Label noise robust image representation learning based on supervised variational autoencoders in remote sensing. In: IEEE International Geoscience and Remote Sensing Symposium (IGARSS) (2023)
46. Teshima, T., Sato, I., Sugiyama, M.: Few-shot domain adaptation by causal mechanism transfer. In: Proceedings of the 37th International Conference on Machine Learning (2020)
47. Voulgaris, G., Philippides, A., Dolley, J., Reffin, J., Marshall, F., Quadrianto, N.: Seasonal domain shift in the global south: Dataset and deep features analysis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops. pp. 2116–2124 (2023)
48. Wang, B., Mendez, J., Cai, M., Eaton, E.: Transfer learning via minimizing the performance gap between domains. In: Advances in Neural Information Processing Systems. vol. 32. Curran Associates, Inc. (2019)
49. Wang, S., Han, W., Huang, X., Zhang, X., Wang, L., Li, J.: Trustworthy remote sensing interpretation: Concepts, technologies, and applications. *ISPRS Journal of Photogrammetry and Remote Sensing* **209**, 150–172 (2024)
50. Weinstein, B.G., Graves, S.J., Marconi, S., Singh, A., Zare, A., Stewart, D., Bohlman, S.A., White, E.P.: A benchmark dataset for canopy crown detection and delineation in co-registered airborne RGB, LiDAR and hyperspectral imagery from the national ecological observation network. *PLoS computational biology* **17**(7), e1009180 (2021)
51. Zhou, D., Bousquet, O., Lal, T., Weston, J., Schölkopf, B.: Learning with local and global consistency. In: Advances in Neural Information Processing Systems. vol. 16. MIT Press (2003)