Randomized Histogram Matching: A Simple Augmentation for Unsupervised Domain Adaptation in Overhead Imagery

Can Yaras, Kaleb Kassaw, *Student Member, IEEE*, Bohao Huang, *Student Member, IEEE*, Kyle Bradbury, *Member, IEEE*, and Jordan M. Malof, *Member, IEEE*

***Abstract*—Modern deep neural networks (DNNs) are highly accurate on many recognition tasks for overhead (e.g., satellite) imagery. However, visual domain shifts (e.g., statistical changes due to geography, sensor, or atmospheric conditions) remain a challenge, causing the accuracy of DNNs to degrade substantially and unpredictably when testing on new sets of imagery. In this work, we model domain shifts caused by variations in imaging hardware, lighting, and other conditions as nonlinear pixel-wise transformations, and we perform a systematic study indicating that modern DNNs can become largely robust to these types of transformations, if provided with appropriate training data augmentation. In general, however, we do not know the transformation between two sets of imagery. To overcome this, we propose a fast real-time unsupervised training augmentation technique, termed randomized histogram matching (RHM). We conduct experiments with two large benchmark datasets for building segmentation and find that despite its simplicity, RHM consistently yields similar or superior performance compared to state-of-the-art unsupervised domain adaptation approaches, while being significantly simpler and more computationally efficient. RHM also offers substantially better performance than other comparably simple approaches that are widely used for overhead imagery.**

***Index Terms*—Augmentation, domain adaptation, segmentation.**

Fig. 1. Illustration of the domain shifts between different collections of overhead imagery. These are representative images from two cities from the Inria and DG datasets. Both Inria and DG serve as our experimental datasets in this work.

I. Introduction

MODERN deep neural networks (DNNs) can now achieve accurate recognition on a variety of tasks involving overhead imagery (e.g., satellite imagery, aerial photography), such as classification, object detection, and semantic segmentation [1], [2], [3]. One emergent limitation of DNNs in remote sensing, however, is their sensitivity to the statistics of their training imagery. Recent research has shown that DNNs often perform unpredictably, and often much more poorly, when they are applied to novel collections with respect to their training data [4], [5], [6], [7]. Furthermore, this performance degradation seems to occur even if DNNs are trained on relatively large and diverse datasets, encompassing large and diverse geographic regions [5], [7].

One cause of the performance degradation of DNNs on new sets of imagery involves visual domain shift (i.e., distribution shift); these are statistical differences between the training imagery and new collections of imagery [4], [5]. Fig. 1 presents images from different collections of imagery where the domain shift is readily visible. These domain shifts are caused by variations in a diverse set of factors that influence the appearance (i.e., statistics) of the overhead imagery including scene geography, the built environment (e.g., building and road styles), imaging hardware, weather, time-of-day, among others. Each of these factors influences the imagery in a manner that is generally complex and unknown in advance and, therefore, challenging to address.

Manuscript received 4 August 2023; revised 9 November 2023; accepted 23 November 2023. Date of publication 7 December 2023; date of current version 28 December 2023. *(Corresponding author: Can Yaras.)*

Can Yaras is with the Department of Electrical and Computer Engineering, University of Michigan, Ann Arbor, MI 48109 USA (e-mail: cjyaras@gmail.com).

Kaleb Kassaw and Bohao Huang are with the Department of Electrical and Computer Engineering, Duke University, Durham, NC 27708 USA (e-mail: kaleb.kassaw@duke.edu; bohao.huang@duke.edu).

Kyle Bradbury is with the Nicholas Institute for Energy, Environment, and Sustainability, Duke University, Durham, NC 27708 USA (e-mail: kyle.bradbury@duke.edu).

Jordan M. Malof is with the Department of Computer Science, University of Montana, Missoula, MT 59812 USA (e-mail: jordan.malof@umt.edu).

Digital Object Identifier 10.1109/JSTARS.2023.3340412

One straightforward solution to these domain shifts is to label a subset of each new collection of imagery and then retrain the DNN; however, this solution is costly and time consuming [4], [5]. Instead, we would ideally have a model that performs well across many different collections of imagery and does so without the need for labels from each one. This setting is a special case of a broader problem in machine learning known as *unsupervised domain adaptation*, wherein it is assumed that we are given a "source domain" dataset with ground truth labels and that we aim to maximize recognition performance on one (or more) sets of unlabeled "target domain" data [8].

A. Spectral Domain Shift and Adaptation

The unsupervised domain adaptation problem has been studied extensively in recent years [9], [10], and has also recently received growing attention in the remote sensing community due to the aforementioned challenges of domain shifts [4], [8], [11], [12], [13]. Most recent domain adaptation approaches for overhead imagery attempt to address all sources of domain shift simultaneously. In this work, however, we attempt to simplify the problem by focusing on a subset of domain shifts that can be modeled as purely spectral (single-pixel) transformations. Mathematically, this is given by

$$p^t = T(p^s) \tag{1}$$

where p^s and p^t are the source and target domain pixel intensities, respectively, of an otherwise identical scene. We hypothesize that domain shifts of this type arise from variations across imagery collections in several specific factors: e.g., camera specifications and calibration, time of day, and lighting conditions. Variations in these factors are likely to occur, to varying degrees, between almost any two collections of imagery so that domain shifts of the kind in (1) are common. Our experiments here suggest that this is not only the case, but that spectral domain shifts appear to be responsible for a significant proportion of the performance degradation of DNNs when applied to novel collections of imagery.

B. Contributions of This Work

In this work, we begin by investigating whether spectral domain shifts of the kind in (1) can be addressed simply through data augmentation during training. We perform a systematic study where we train DNNs of varying capacity (i.e., number of parameters) with image augmentation comprising different classes of spectral augmentations (e.g., gamma, affine, etc). We then test the performance of these networks on collections of imagery that have been augmented with one of these same classes of spectral transformations. We find that modern DNNs with large encoders (e.g., ResNet-18, 50, 100 [14]) can become largely robust to several different classes of spectral transformations if provided with a *matching* training augmentation strategy. In general, however, we do not know the transformation between any two collections of imagery, or even the class of transformations from which it may be drawn (e.g., affine, gamma), so it is unclear, which augmentation should be adopted.

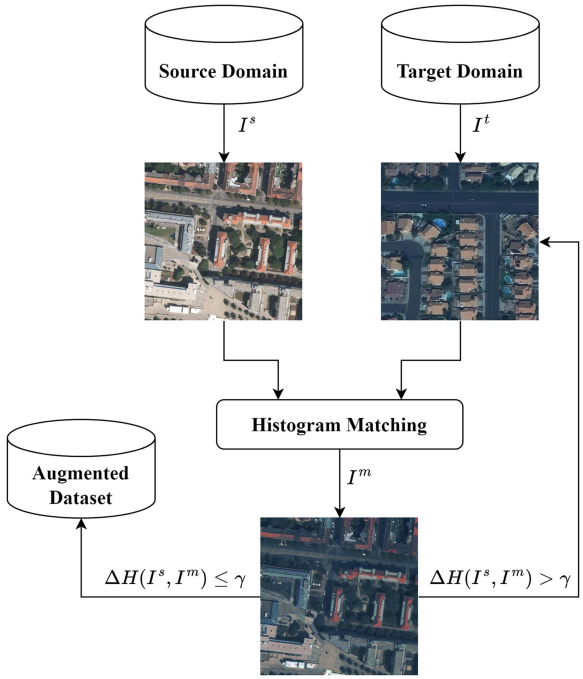

Fig. 2. Illustration of the RHM concept. To produce an augmented training dataset, we repeat the following process. For each image I^s in the training dataset (source domain), an image I^t is drawn randomly from the testing dataset (unlabeled target domain). Then, the histogram of I^s is matched to the histogram of I^t, which yields the modified image I^m. If the information loss ΔH (defined in Section IV-A) between I^m and I^s is below a set threshold γ, I^m is added to the augmented training dataset. Otherwise, another random image is drawn from the target domain and the process is repeated. This resampling is only performed at most once.

To overcome this problem, we propose a simple augmentation technique, termed *randomized histogram matching* (RHM) that matches the histogram of each training image to a randomly-chosen (unlabeled) target domain image, as illustrated in Fig. 2. This approach results in a random spectral shift being applied to *each* training image, and we hypothesize that this occasionally (by chance) approximates the true spectral shift between the source and target domains (see Section IV). Consequently, the training data are occasionally augmented (again, by chance) with the true spectral shift and can thereby become more robust to it. Since RHM only requires the unlabeled testing data, it can be viewed as a simple unsupervised domain adaptation approach.

To demonstrate the efficacy of RHM, we conduct benchmark testing with two large publicly-available datasets for building segmentation in two settings: 1) training on one collection and testing on one collection (one-to-one adaptation, following [13]), and 2) a more real-world scenario where we train on multiple domains and test on multiple domains (many-to-many, also following [13]). We focus on building segmentation because it is a challenging task that has received substantial attention in recent years, with large and diverse benchmark datasets to support our multidomain experiments (e.g., Inria [6] and DG [1]). We now summarize our contributions as follows.

1) *Can augmentation confer spectral robustness in DNNs?* We provide, to the best of authors' knowledge, the first

systematic empirical evidence that modern DNNs with large encoders are capable of becoming robust to complex spectral transformations via data augmentation during training. We show that the ability of DNNs to become robust depends upon their capacity, especially for more complex classes of transformations. We also find that DNNs only become robust to the precise class of spectral transforms that were used for augmentation, rather than becoming robust to generic spectral transforms. This suggests that augmentation is an effective mechanism to address spectral domain shifts, if a class of spectral transforms is used that includes in it the particular transforms that are often encountered in real-world imagery (e.g., between independent collections of imagery).

2) *RHM augmentation:* RHM is a simple, yet highly effective unsupervised domain adaptation approach via spectral augmentation. We show that RHM almost always offers substantial performance benefits in unsupervised cross-domain settings (e.g., we wish to apply a model in a new geo-location with no labeled data). Our results also indicate that RHM usually offers substantially greater performance benefits than other common types of spectral augmentation [e.g., Affine, Gamma, or hue-saturation-value (HSV)], and additionally performs competitively with (or even better than) two state-of-the-art unsupervised domain adaptation approaches (i.e., CycleGAN or ColorMapGAN [4]), despite being substantially simpler and faster in many cases (e.g., RHM only has one hyperparameter, and does not require training additional models).

Next, in Section II, we provide further details about related work and how our contributions differ from them.

II. Related Work

In this section, we review related work for unsupervised domain adaptation in overhead imagery and how our work relates and builds upon it. Unsupervised adaptation methods can broadly be divided into two main groups: model adaptation and data adaptation.

A. Unsupervised Model Adaptation

In many of these approaches, the goal is to obtain features (e.g., through selection or learning) that are invariant across the source and target domains, but still useful to discriminate between the target classes. Some approaches have focused on a feature selection strategy, often using a curriculum learning approach, where "easier" (i.e., far-from-margin or high-confidence) samples are progressively given as training data until convergence [15], [16], [17], [18]. Other approaches attempt to learn the desired feature representation in target domains using feature alignment techniques [19] or discriminators to distinguish between domains [20], [21]. These approaches have been widely used for domain adaptation of remote sensing imagery, e.g., curriculum learning in [22], [23], discriminators used in [24], and feature alignment in [25], [26], and [27].

B. Unsupervised Data Adaptation

Our work here builds directly upon recently proposed methods of this kind. These methods are designed to modify the source and/or target domain data so that they are statistically more similar to each other. If successful, a recognition model that is trained and evaluated on the modified source and target data should be more accurate. These methods can be subdivided into the following two main categories: (i) domain standardization and (ii) domain matching. In (i), the goal is to map the source and target domains into some common domain. Some well-known examples of such approaches are normalization (or z-scoring) [28]; histogram equalization [28], color invariance approaches (e.g., [29], [30], [31], [32]), and recent approaches using DNNs [11].

In (ii), the goal is to match the source domain to the target domain. Graph matching [33], [34] and especially histogram matching (HM) [11] are common approaches for this. Based on the CycleGAN model [35], a large number of approaches have been proposed to train a DNN to map source domain data to be more similar to the target domain [4], [36], [37], [38], [39], [40]. One challenge with many of these approaches is that they can alter the semantic content of the source domain imagery [4] (e.g., changing object shapes or even their semantic class).

More recently, ColorMapGAN [4] was proposed to address this challenge by restricting the DNN to perform pixel-wise intensity transformations, preventing the model from making more complex semantic changes to the imagery. The authors show that ColorMapGAN (along with CycleGAN [35]) outperformed a variety of types of unsupervised domain adaptation approaches for segmentation on overhead imagery. One limitation of ColorMapGAN, however, is that a separate model has to be learned between each *pair* of source and target domains, which is impractical for a large number of source and target domains (many-to-many testing). In this work, we propose RHM as a simple and fast alternative to recent DNN-based unsupervised domain adaptation approaches.

C. Data Augmentation

In this approach, the original training dataset is supplemented with transformed yet semantically consistent views of the training data that create variations in the training imagery. Some commonly utilized classes of transformations used for augmentation in remote sensing are Gamma corrections [41] and contrast changes (e.g., via HSV shifts [41]). These approaches are designed to build invariance to different classes of spectral shift. For this reason, we will compare RHM to Gamma and HSV augmentation approaches, and investigate whether DNNs can indeed become robust to these transformations, as is implicitly assumed in their application.

III. Experimental Materials and Methods

A. Experimental Datasets

In our experiments, we employ two large publicly-available datasets for building segmentation: Inria [6] and DeepGlobe (DG) [1]. Both datasets are composed of high-resolution (0.3 m)

TABLE I
CITIES INCLUDED IN INRIA AND DEEPGLOBE WITH CORRESPONDING SURFACE AREA OF IMAGERY

Dataset	City	Surface area (km^2)
Inria [6]	Austin, TX, USA	81
	Chicago, IL, USA	81
	Kitsap, WA, USA	81
	West Tyrol, Austria	81
	Vienna, Austria	81
DeepGlobe [1]	Las Vegas, NV, USA	113
	Paris, France	33
	Shanghai, China	133
	Khartoum, Sudan	29

TABLE II
PARAMETERIZED FORMS OF TRANSFORMATIONS FOR EACH BASELINE AUGMENTATION

Augmentation	Transformation	Parameters
Affine	$X^{(c)} \mapsto \alpha^{(c)} X^{(c)} + \mu^{(c)}$ $c \in \{R, G, B\}$	$\alpha^{(c)} \overset{\text{iid}}{\sim} \mathcal{U}(0.82, 1.18)$ $\mu^{(c)} \overset{\text{iid}}{\sim} \mathcal{U}(-0.38, 0.38)$
Gamma	$X^{(c)} \mapsto \left(X^{(c)}\right)^{\gamma^{(c)}}$ $c \in \{R, G, B\}$	$\gamma^{(c)} \overset{\text{iid}}{\sim} \mathcal{U}(0.32, 1.68)$
HSV	$Y^c \mapsto \alpha^{(c)} Y^{(c)} + \mu^{(c)}$ $Y = \text{HSV}(X)$ $c \in \{H, S, V\}$	$\alpha^{(S)}, \alpha^{(V)} \overset{\text{iid}}{\sim} \mathcal{U}(0.63, 1.37)$ $\alpha^{(H)} = 1$ $\mu^{(c)} \overset{\text{iid}}{\sim} \mathcal{U}(-0.27, 0.27)$

color overhead imagery and have accompanying pixel-wise building labels. Both datasets include large quantities of imagery from several distant geographic locations, summarized in Table I. Importantly, each collection varies greatly in both their scene content and their spectral characteristics—Fig. 1 presents examples of imagery from DG and Inria illustrating these differences.

B. Segmentation Model and Training

In recent years, U-Net [42] and its variants (e.g., [43]) have achieved state-of-the-art performance for building segmentation in overhead imagery (e.g., [1], [6]). Following [43], we modify U-Net by using ResNet encoders of varying size that have been pretrained on the ImageNet dataset. In our experiments, we will use a ResNet-50 encoder unless otherwise noted, to balance training speed and performance. We also make the following specific design choices for our models:

1) cross-entropy loss between the pixel-wise ground truth and predictions;
2) the SGD optimizer;
3) 90 epochs of training;
4) a batch size of 8.

We also use a learning rate of 0.001 and 0.01, respectively, for the encoder and decoder of the U-Net models. A smaller learning rate is applied to the encoder since it is already pretrained on ImageNet. For both the encoder and decoder, we drop the learning rate by one order of magnitude after 50 and 80 epochs. These settings are chosen to be nearly identical to that those in [44]—the only variation is that we additionally drop the learning rate after 80 epochs to ensure that the validation loss converges by the end of training.

C. Baseline Adaptation Methods

For baselines, we focused upon spectral adaptation methods, allowing us to compare RHM to methods with similar complexity (e.g., only altering spectral content of the imagery). We include computationally simple spectral augmentations that are widely-used in remote sensing (e.g., Gamma, HSV, Affine intensity augmentations), as well as ColorMapGAN, a recent *data-driven* spectral adaptation method. We also consider one state-of-the-art method that is not restricted to spectral transforms, CycleGAN, to determine how spectral methods compare with more general adaptation methods.

1) Augmentation: We consider three parameterized transformations as baseline augmentations for comparison to RHM: *Affine*, *Gamma*, and *HSV*. These are common transformations for modeling spectral transformations and as such are most relevant for comparison to our proposed method. Table II contains the parameterized functional forms and their respective distributions. The distribution of each parameter is chosen via a standard Bayesian hyperparameter optimization procedure using Gaussian processes, as described in [45]. For the HSV augmentation, we fix the scaling factor of the hue channel $\alpha^{(H)}$ to be 1, since the hue value corresponds to the angular dimension in the cylindrical geometry of HSV space. During training, we apply these augmentations in real time throughout training, with uniquely sampled augmentation parameters for each mini-batch iteration.

2) Standardization. Histogram equalization [28]: One approach to standardizing each domain is to ensure that the contrast of all images are the same. Histogram equalization achieves this by adjusting the histogram of pixel intensities of each image to be uniform. We transform images in both the source and target domain in this manner and use the transformed images for training and testing the U-Net model, respectively.

Gray world [29]: This approach attempts to find a standardized domain in which each image's average color is gray, and therefore, invariant to illumination conditions that may affect each color channel independently. By modeling the deviation in color illumination of each channel from gray as a linear scaling, we may remove the scaling factor by normalizing each channel by its average pixel intensity. We transform images in both the source and target domain in this manner and use the transformed images for training and testing of the U-Net model, respectively.

3) Image-to-Image Translation. HM [28]: A naive method for matching the distribution of the source domain to the target domain is to match the histogram of each source image to the aggregate histogram of the target domain. This matching is done independently for each channel. We transform only the images in the source domain in this manner and use the transformed images for training the U-Net model.

ColorMapGAN [4]: This state-of-the-art approach aims to learn an unconstrained pixel-wise mapping from the source to target domain, modeled as a generator in an unsupervised adversarial setting. As with most GAN set-ups, there is a generator G and a discriminator D, where D attempts to differentiate images generated by G from the images in the target domain. On the

other hand, G learns a unique pixel-to-pixel mapping for every possible RGB triple. G and D are trained simultaneously with the LSGAN [46] loss. After training, G is used to generate fake images from the source domain that look like the target domain—these fake source images are then used to train the U-Net model. We use the same hyperparameters as given in [4] in our own experiments.

CycleGAN. [35]: Similar to ColorMapGAN, this method learns a transformation between domains in an unsupervised adversarial setting. However, we now have two generators G and F where G attempts to transform the source domain S to the target domain T and F attempts to transform the target domain T to the source domain S. Unlike ColorMapGAN, both G and F are multilayer networks that can realize more complicated functions than pixel-wise transforms. We also have two domain-specific discriminators D_S and D_T that attempt to differentiate the real and fake images in their respective domains. G and D are trained simultaneously via an objective that combines the conditional GAN [47] loss in both directions with a cycle-consistency loss term, which forces the compositions $F \circ G$ and $G \circ F$ to be the identity mapping. After training, G is used to generate fake images from the source domain that look like the target domain—these fake source images are then used to train the U-Net model. We use the same hyperparameters as given in [35] in our own experiments.

IV. RHM Augmentation

RHM is a modification of conventional HM. In conventional HM, one matches the histogram of pixel intensities of one set of imagery to the corresponding histogram of another set of imagery. For example, for cross-domain adaptation, we can transform a single source image to match the histogram of created from the full collection of target domain imagery (see Section III-C). This approach works well if the scene content of the source imagery and the target imagery are similar, in which case the histograms of the two image collections *should* be similar as well; any differences must be due to other factors (e.g., variations in lighting, imaging hardware, etc), which are often modeled as spectral domain shifts. Consequently, matching the histograms removes any existing spectral domain shift between the two image sets, if their scene content is similar. However, we hypothesize that it is unlikely that two random sets of imagery will contain similar scene content. In such cases, the differences in histograms will be due to content differences, in which case the two histograms will not generally be the same, even if other imaging conditions are similar (e.g., lighting, hardware). When such content differences are present, therefore, conventional HM produces undesirable results. For example, our experimental results in Section V-B indicate that conventional HM often works well, but that it is also inconsistent, and can sometimes fail badly.

RHM is intended to mitigate this limitation of HM by relying upon matching many random pairs of imagery, making it likely that the content between the pairs will sometimes (by chance) be similar, causing HM to approximate the true underlying spectral shift for those pairs. Specifically, as outlined in Fig. 2, RHM matches the histogram of each training image with the histogram of a randomly-sampled target domain image. We hypothesize that this approach often results in image pairs with dissimilar scene content (like conventional HM), but that it also sometimes creates pairs with similar scene content. Consequently, some training images are augmented in a way that approximates the true underlying spectral shift between the source and target domain imagery. Furthermore, we hypothesize that largely inaccurate augmentations will be effective for training robust DNNs as long as it *periodically* augments with the correct transformations. This is motivated by the experiments in Section V-A suggesting that a given class of spectral augmentations will work well as long as the true augmentation is a special case of the class (e.g., augmentation with *random* spectral Affine transforms works well for a given target-domain if that target domain is shifted by *any specific* Affine transform).

A. RHM Algorithm

The detailed RHM training procedure is summarized in Algorithm 1, and we describe two of the major components in more detail: HM and entropy-based resampling.

HM: Mathematically, the HM is performed as follows: for a source image $X \in \mathbb{R}^{C\times H\times W}$, let $F_c : \mathbb{R} \to [0, 1]$ be the normalized cumulative histogram of each channel $c \in [C]$, i.e., $F_c(x)$ is the proportion of pixels in channel c with magnitude no more than x. Similarly, define G_c to be the normalized cumulative histogram of channel $c \in [C]$ for a randomly selected target image $\widetilde{X} \in \mathbb{R}^{C\times H'\times W'}$. Then, the RHM augmented version $T(X; \widetilde{X})$ of source X with target $\widetilde{X}$ is defined as

$$T(X; \widetilde{X})_{c,i,j} = G_c^{-1}(F_c(X_{c,i,j})) \tag{2}$$

where $G_c^{-1}(y) \triangleq \min\{x : G_c(x) \geq y\}$.

Entropy-based resampling: For some pairings of source and target images, the transform in (2) can result in a large loss of image information needed to perform the building segmentation task, which can be detrimental to model training. To limit the amount of image compression in RHM augmentations, we discard augmentations that lead to large compressions of the image intensity values. We measure the compression via the change in Shannon entropy of their histograms, here denoted $H(X)$ for an image X and $H(T(X))$ for a transformed image, where $H(X)$ is defined as

$$H(X) = -\frac{1}{c}\sum_{c\in[C]}\int f_c(x)\log(f_c(x))\,dx \tag{3}$$

where f_c is the normalized histogram of each channel $c \in [C]$ of X. We then define the quantity ΔH, the change in image information, as

$$\Delta H \triangleq H(X) - H(T(X)). \tag{4}$$

The distribution of values of ΔH for RHM transforms is shown in Fig. 3 along with the associated augmented images. Here, the loss of image information is apparent as ΔH increases.

By evaluating the value of ΔH for each source-target image pair during training, and excluding those pairs with high ΔH, we limit the resulting image compression from RHM transforms. We set a threshold γ, and in cases where RHM transforms

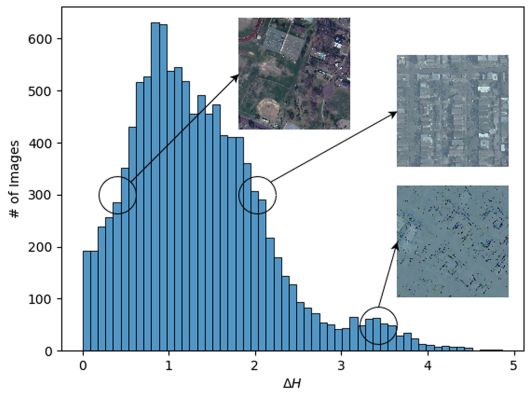

Fig. 3. Distribution of changes in image information ΔH resulting from RHM transforms. Higher values of ΔH correspond to greater degrees of loss in image information.

Algorithm 1: Randomized Histogram Matching.

Input: Labeled source domain $\mathcal{D}_S = \{(X_i, Y_i)\}$, unlabeled target domain $\mathcal{D}_T = \{\widetilde{X}_j\}$, training epochs T, batch size B, learning rate schedule $(\eta_t)_{t\geq 1}$, segmentation loss $\ell(\cdot,\cdot)$
Output: Segmentation model $f_{\theta^*}(\cdot)$ adapted to $\mathcal{D}_T$
Initialize (pretrained) segmentation model $f_\theta(\cdot)$
for epoch $t = 1, \ldots, T$ **do**
  $\mathcal{B}_t \leftarrow$ shuffle and group $\mathcal{D}_S$ into batches of size B
  for batch $\{(X_i, Y_i)\}_{i=1}^{B}$ in $\mathcal{B}_t$ **do**
    for $i = 1, \ldots, B$ **do**
      $\widetilde{X}_i \leftarrow$ uniformly sampled from $\mathcal{D}_T$
      $Z_i \leftarrow T(X_i; \widetilde{X}_i)$
      if $H(X_i) - H(Z_i) > \gamma$ **then**
        $\widetilde{X}_i \leftarrow$ uniformly sampled from $\mathcal{D}_T$
        $Z_i \leftarrow T(X_i; \widetilde{X}_i)$
      end if
    end for
    $\theta \leftarrow \theta - \eta_t \nabla_\theta \sum_i \ell(f_\theta(Z_i), Y_i)$
  end for
end for
$\theta^* \leftarrow \theta$

produce $\Delta H > \gamma$, the first pairing is discarded, a new randomly selected target image is chosen, and the transformation in (2) is computed once more. To limit overall computation time, we only repeat this resampling step once for each source image. We show in our experiments in Section V-B and V-C that this filtering step is consistently beneficial.

As a result, we are able to utilize variations in the target domain to apply random spectral shifts to the training imagery that we hypothesize periodically coincide with the true spectral shift between the source and target domains. Like the baseline augmentations described in Section III-C, we apply RHM as an online augmentation to each training image, where a new target image is sampled every iteration for matching. See Algorithm 1 for a full description of training using RHM with entropy resampling. RHM is not applied to the target domain during testing.

V. Experiments and Results

A. Does Spectral Augmentation Confer Robustness to Spectral Transformations?

In this section, we investigate the extent to which modern DNNs with high-capacity feature encoders can become robust to spectral transformations of the form in (1). Many popular augmentation approaches for remote sensing (and elsewhere) apply random intensity transformations (e.g., HSV, Gamma) with the implicit assumption that DNNs will become robust to domain shifts with a similar functional form. To the best of authors' knowledge, there have been no controlled experiments investigating these assumptions or evaluating how they depend either on the complexity of the spectral transformations, or the capacity of the DNNs involved. We also investigate whether augmentation with one class of spectral functions (e.g., HSV) confers general invariance to spectral transformations (e.g., the DNN learns to ignore spectral shifts of any kind), or that the DNN only becomes robust to the particular class of functions it was trained upon.

To address these questions, we emulate several different kinds of spectral domain shifts on the Inria dataset. We use the Inria test partition to create five different testing datasets. Each testing dataset is the result of applying just one of five possible types of spectral augmentation to the original Inria testing dataset: Original (no augmentation), Affine, Gamma, HSV, and RHM. Each of these test datasets then represents one *class* of spectral domain shifts. We then train five different models on the Inria training partition, however, each models is trained with just one of the five aforementioned augmentation strategies (including "Original"). We consider a DNN to be robust to a particular class of spectral domain shift (e.g., HSV) if its performance does not degrade significantly—compared to the unaltered Inria "Original" test dataset—when it is evaluated on that transformed test dataset.

In these experiments, we can examine the robustness of models when their augmentation is perfectly matched to the test datasets spectral domain shift (e.g., evaluate the HSV-augmented model on the HSV-augmented testing dataset)—an ideal scenario. In this setting, performance degradation (relative to testing on the "Original" unaugmented Inria test data) should arise only due to: i) inability of the model to become robust to the spectral transformation, or ii) loss of image information due to the augmentations (e.g., some spectral augmentations compress the imagery, by mapping several pixel intensities into a single intensity).

The results of this experiment are presented on the diagonal (bolded) entries in Table III. As we see, the performance degradation for all domain shifts is relatively low, suggesting that the DNNs do achieve relative robustness when trained with a matching augmentation. We also applied each trained model to all of the other test datasets (i.e., that have different augmentations), and the results of this are presented in the unbolded entries in Table III. As expected, substantial performance degradation

TABLE III
PERFORMANCE IN TERMS OF IOU OF A U-NET MODEL WITH A RESNET-50 ENCODER–DECODER STRUCTURE WHEN TRAINED AND TESTED ON DATA WITH DIFFERENT CLASSES OF AUGMENTATION

Training	Testing				
	Original	Affine	Gamma	HSV	RHM
Original	**0.746**	0.367	0.496	0.524	0.547
Affine	0.739	**0.725**	0.728	0.669	0.646
Gamma	0.739	0.708	**0.734**	0.636	0.650
HSV	0.725	0.681	0.703	**0.708**	0.631
RHM	0.720	0.507	0.652	0.635	**0.713**

Bolded text indicates that the model was trained and tested on imagery with the same augmentation.

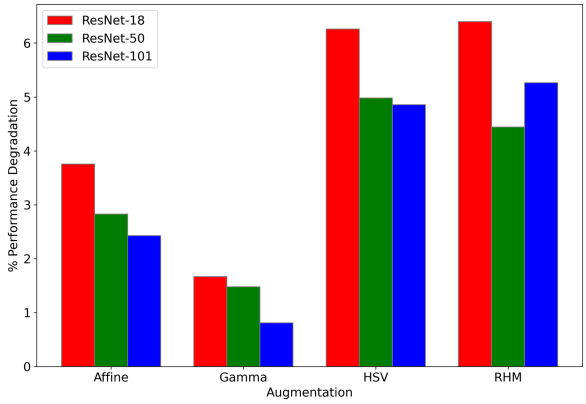

Fig. 4. Percent performance degradation as a function of model size and augmentation type. For each combination of model size and training augmentation type, we measure the percentage performance degradation when testing on the augmented test set compared to the un-augmented test set.

is observed when applying the "Original" model to any of the augmented testing datasets. This confirms the importance of spectral augmentation of some kind when training DNNs with overhead imagery. Furthermore, when an augmentation is applied to the test set, we see that the best model for any testing dataset is the model that was trained with a matching augmentation.

Interestingly, we find that training with an augmentation that does not match the test set augmentation usually results in further performance degradation (i.e., compared to a matching augmentation strategy), and sometimes a substantial degradation. This has several implications. First, these results suggest that spectral augmentations do not result in *general* robustness to spectral domain shifts, so that the DNN learns to ignore many or most kinds of spectral shifts. Instead, it appears that they confer robustness only for the class of domain shifts that were presented in training. A corollary of this is that it is important to choose an augmentation strategy that does indeed emulate the domain shifts that can be expected in real-world imagery, and that failing to do so can result in substantial loss of otherwise recoverable performance.

We also investigate the extent to which DNN robustness depends upon the capacity of the model (e.g,. the number of free parameters it has). Therefore, we repeated our experiments using segmentation models with three different encoder sizes: ResNet-18, ResNet-50, and ResNet-101. The results of this experiment are presented in Fig. 4, where we only report results when we train and test with the same augmentations. The results indicate that a larger model does seem to enable a greater level of invariance; except for RHM where ResNet-50 has slightly better performance than ResNet-101, the ResNet-18, and ResNet-101 models consistently perform the worst and best, respectively.

B. One-to-One Domain Adaptation

In this section, we compare RHM to other unsupervised domain adaptation methods when evaluated in a one-to-one scenario, i.e., we are given a single source domain, and we must maximize performance on a single (unlabeled) target domain [4], [12]. Following the practice of recent work [4], we treat the imagery over a single city as a single domain, and we randomly chose two pairs of cities (i.e., four total cities) for our one-to-one experiments. For each pair, we alternately trained on one of the two cities, and tested on the other. The only constraint on the selection of the city pairs is that they must contain one city from the Inria dataset, and one city from DG dataset. These two datasets were produced by different groups and at different times, and therefore, we reason they are more likely to exhibit domain shifts. All experiments were conducted with a U-Net model with a ResNet-50 encoder, as described in Section III-B. All results are reported in terms of intersection-over-union (IoU).

Descriptions of our baseline methods can be found in Section III-C. As baselines, we included a variety of methods that are comparable in their simplicity and speed to RHM (e.g., HSV and Gamma augmentation, gray-world standardization, etc.). We also included ColorMapGAN [4], which is a more sophisticated approach, which recently reported superior results to a large number of other state-of-the-art unsupervised domain adaptation methods when evaluated in the one-to-one scenario. As an ablation study, we also test RHM *without* the entropy-based resampling step described in Section IV-A, which is termed "RHM w/o EBR" in Table IV.

The results of the benchmark experiments are presented in Table IV. RHM gives a slightly better IoU than RHM w/o EBR, while both RHM models perform substantially better than all other baselines (on average). RHM achieves the highest IoU on two of the four individual test cities (Vienna $\rightarrow$ Vegas and Tyrol-w $\rightarrow$ Shanghai). Although it does not achieve the highest performance on Vegas $\rightarrow$ Vienna or Shanghai $\rightarrow$ Tyrol-w, in both cases, it is the second best performing model, and achieves very similar performance to the top-performing approach.

C. Many-to-Many Domain Adaptation

In this section, we compare RHM to other unsupervised domain adaptation methods when evaluated in a many-to-many scenario, i.e., we are given multiple source domains, and we must maximize performance on multiple (unlabeled) target domains [4], [12]. In contrast to the one-to-one scenario, the many-to-many setting is more likely to reflect real-world testing conditions in which a model is trained on a large and diverse

TABLE IV
BENCHMARK OF DOMAIN ADAPTATION METHODS FOR ONE-TO-ONE USING A RESNET-50 ENCODER

Method	Type	Vienna → Vegas	Vegas → Vienna	Tyrol-w → Shanghai	Shanghai → Tyrol-w	Average
Original	Naive	0.342	0.503	0.416	0.398	0.415
Affine	Augmentation	0.637	**0.642**	0.337	0.472	0.522
Gamma		0.583	0.459	0.336	0.351	0.432
HSV		0.600	0.565	0.250	0.516	0.483
RHM w/o EBR		0.627	0.584	0.398	0.557	0.542
RHM		**0.646**	0.597	**0.433**	0.534	**0.553**
Hist-Eq [28]	Standardization	0.500	0.531	0.260	0.505	0.449
Gray-World [29]		0.436	0.517	0.336	0.469	0.440
Hist-Match [28]	I2I Translation	0.558	0.527	0.317	0.501	0.476
ColorMapGAN[4]		0.365	0.429	0.291	**0.562**	0.412

The method "RHM w/o EBR" refers to RHM without entropy-based resampling. All results are reported in terms of IoU.
Bold values indicate highest value in the column.

TABLE V
BENCHMARK OF MANY-TO-MANY DOMAIN ADAPTATION METHODS FOR INRIA → DG USING A RESNET-50 ENCODER

Method	Type	Khartoum	Paris	Shanghai	Vegas	Overall	City Average
Original	Naive	0.270	0.365	0.467	0.684	0.533	0.447
Affine	Augmentation	0.215	0.484	0.443	0.736	0.542	0.470
Gamma		0.206	0.422	0.435	0.735	0.539	0.450
HSV		0.159	0.427	0.391	0.740	0.519	0.429
RHM w/o EBR		0.244	0.525	0.494	0.746	0.583	0.502
RHM		0.260	**0.556**	0.518	**0.746**	0.590	**0.520**
Hist-Eq [28]	Standardization	0.166	0.316	0.367	0.663	0.477	0.378
Gray-World [29]		0.209	0.379	0.399	0.716	0.519	0.426
Hist-Match [28]	I2I Translation	0.144	0.282	0.279	0.734	0.432	0.360
ColorMapGAN [4]		0.087	0.260	0.285	0.706	0.450	0.335
CycleGAN [35]		**0.340**	0.470	**0.534**	0.730	**0.595**	0.519

The method "RHM w/o EBR" refers to RHM without entropy-based resampling. All results are reported in terms of IoU.
Bold values indicate highest value in the column.

TABLE VI
BENCHMARK OF MANY-TO-MANY DOMAIN ADAPTATION METHODS FOR DG → INRIA USING A RESNET-50 ENCODER

Method	Type	Austin	Chicago	Kitsap	West Tyrol	Vienna	Overall	City Average
Original	Naive	0.396	0.310	0.491	0.472	0.592	0.424	0.452
Affine	Augmentation	0.431	0.394	0.588	0.553	0.647	0.492	0.523
Gamma		0.314	0.302	0.502	0.491	0.564	0.392	0.435
HSV		0.493	0.377	**0.605**	0.559	0.635	0.494	0.534
RHM w/o EBR		0.562	0.576	0.579	0.649	0.664	0.613	0.606
RHM		0.574	0.573	0.601	**0.658**	**0.671**	0.618	**0.615**
Hist-Eq [28]	Standardization	0.560	0.566	0.568	0.629	0.634	0.597	0.591
Gray-World [29]		0.343	0.287	0.560	0.493	0.614	0.401	0.459
Hist-Match [28]	I2I Translation	0.544	0.568	0.577	0.566	0.663	0.599	0.584
ColorMapGAN [4]		0.540	0.561	0.448	0.496	0.637	0.576	0.536
CycleGAN [35]		**0.625**	**0.585**	0.590	0.595	0.656	**0.619**	0.610

All results are reported in terms of IoU. The method "RHM w/o EBR" refers to RHM without entropy-based resampling.
Bold values indicate highest value in the column.

training set and then tested on multiple new collections of imagery (i.e., multiple target domains). For these experiments, we train each model on one of our two multicity benchmark datasets (Inria and DG), and test on the other.

Because ColorMapGAN was designed specifically for the one-to-one task and, therefore, may be at a disadvantage [4], for these experiments, we utilized an additional benchmark method, CycleGAN [35]. CycleGAN recently achieved comparable performance to ColorMapGAN in [4], while being better-suited for the many-to-many testing scenario. As in the previous section, we also test RHM with, and without, the entropy-based resampling step (see Section IV-A for details). The RHM model without resampling is denoted "RHM w/o EBR" in Tables V and VI.

Our many-to-many experimental results are reported in Tables V and VI, respectively. In each case, the IoU for each testing city is provided along with an "Overall" IoU (computed after aggregating all test city predictions) and a "City Average" (computed by averaging the IoUs of each test city). As with the one-to-one setting, we find that entropy-based resampling improves both the overall and per-city performance of RHM (in all nine cities, except for Chicago). Moreover, RHM outperforms CycleGAN on 5 of the 9 cities, and achieves better city average performance than all other baselines, while having a

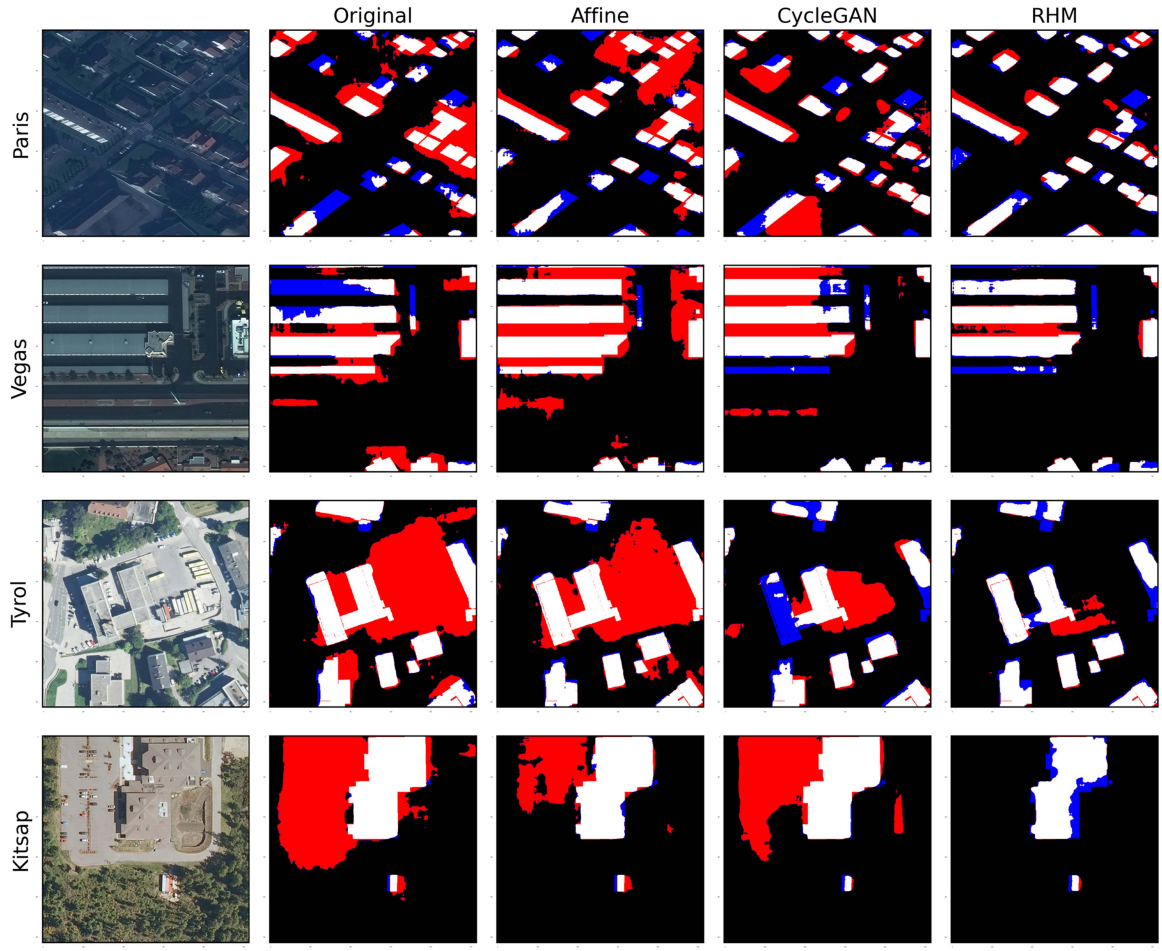

Fig. 5. Visualization of segmentation masks for several domain adaptation methods across various cities in Inria and DG. The masks are colored as follows: black for true negative, blue for false negative, red for false positive, white for true positive.

very similar overall IoU to CycleGAN in both training/testing directions. In Fig. 5, we illustrate example predictions for various methods, including RHM—across diverse conditions, we see that RHM yields considerably fewer false positives (encoded in red), which can explain the improvement in IoU. Notably, RHM gives comparable or better performance than CycleGAN without the need to train an auxiliary DNN (e.g., our CycleGAN typically required over three days to train on an NVIDIA Titan RTX), or tune many hyperparameters. Compared to other simple unsupervised domain adaptation approaches, RHM provides substantially better average performances on both benchmark test sets. Finally, we note that applying other augmentations (e.g., HSV or Gamma) along with RHM does not improve performance, and in fact substantially degrades the performance of RHM—we postpone the results and discussion to the Appendix.

D. Run-Time Analysis

In this section, we demonstrate that employing RHM during training only incurs modest computational costs compared to similar online augmentation approaches. Following the model and optimizer configurations outlined in Section III-B, we train a U-Net with a ResNet-18 encoder over a single epoch of the Inria dataset and benchmark the wall times of an average single training iteration for each baseline augmentation given in Table II, as well as RHM with entropy-based resampling. In Table VII, we report the increase in wall time compared to no augmentation used during training.

TABLE VII
INCREASE IN TRAINING TIME DUE TO USE OF VARIOUS ONLINE AUGMENTATION APPROACHES WITH A RESNET-18 ENCODER

Augmentation	% Increase
Affine	0.62
Gamma	1.25
HSV	4.08
RHM	1.88

We see that the use of RHM only results in a marginal increase in training time that is comparable to common baselines, and is in fact considerably faster than the widely-used HSV augmentation. We note that our implementation of RHM computes histograms on-the-fly during training, and does not reuse previously computed histograms in future iterations. At the cost of additional memory usage, it is possible to make RHM even faster by precomputing all histograms before training so

TABLE VIII
COMBINING AUGMENTATIONS IN MANY-TO-MANY DOMAIN ADAPTATION FOR INRIA → DG USING A RESNET-50 ENCODER

Method	K.	P.	S.	V.	O.	C.A.
Gamma+RHM	0.246	0.351	0.496	0.708	0.547	0.450
HSV+RHM	0.248	0.356	0.496	0.706	0.548	0.451
RHM	**0.260**	**0.556**	**0.518**	**0.746**	**0.590**	**0.520**

All results are reported in terms of IoU. Table headings are the same as in Table V. Bold values indicate highest value in the column.

that they are readily available for the matching process, but this is beyond the scope of our work.

VI. CONCLUSION

In this work, we address the problem of unsupervised domain adaptation in overhead imagery. To do so, we model domain shifts caused by variations in imaging hardware, lighting conditions (e.g., due to time-of-day), or atmospheric conditions as nonlinear pixel-wise transformations, and we show that DNNs can become largely robust to these types of transformations if they are provided with the appropriate training augmentation. In general, however, we do not know the transformation between any two sets of imagery. To overcome this problem, we propose RHM, a simple real-time training data augmentation approach. We then conduct experiments with two large benchmark datasets for building segmentation and we find that RHM consistently yields comparable performance to recent state-of-the-art unsupervised domain adaptation approaches for overhead imagery, despite being substantially easier and faster to use. RHM also offers substantially better performance than other comparably simple and widely-used unsupervised approaches for overhead imagery. This new approach to training augmentation has the ability to expand the efficacy of automated analysis of remote sensing data to more applications while reducing the burden of expensive labeled imagery from target domains.

APPENDIX

As briefly mentioned in Section V-C, applying spectral augmentations, such as Gamma and HSV in conjunction with RHM does not improve performance over standalone RHM—results for Inria to DG many-to-many domain augmentation are shown in Table VIII. We note that for all target cities, RHM alone achieves a better IoU than either of the combined methods, suggesting that this addition is not beneficial but detrimental for performance.

ACKNOWLEDGMENT

The authors would like to thank the Energy Initiative at Duke University for supporting this work.

REFERENCES

[1] I. Demir et al., "DeepGlobe 2018: A challenge to parse the Earth through satellite images," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. Workshops*, 2018, pp. 172–181.
[2] A. V. Etten, D. Lindenbaum, and T. M. Bacastow, "SpaceNet: A remote sensing dataset and challenge series," 2018, *arXiv:1807.01232*.
[3] N. Sergievskiy and A. Ponamarev, "Reduced focal loss: 1st place solution to xView object detection in satellite imagery," 2019, *arXiv:1903.01347*.
[4] O. Tasar, S. Happy, Y. Tarabalka, and P. Alliez, "ColorMapGAN: Unsupervised domain adaptation for semantic segmentation using color mapping generative adversarial networks," *IEEE Trans. Geosci. Remote Sens.*, vol. 58, no. 10, pp. 7178–7193, Oct. 2020.
[5] F. Kong, B. Huang, K. Bradbury, and J. Malof, "The Synthinel-1 dataset: A collection of high resolution synthetic overhead imagery for building segmentation," in *Proc. IEEE/CVF Winter Conf. Appl. Comput. Vis.*, 2020, pp. 1814–1823.
[6] E. Maggiori, Y. Tarabalka, G. Charpiat, and P. Alliez, "Can semantic labeling methods generalize to any city? The inria aerial image labeling benchmark," in *Proc. IEEE Int. Geosci. Remote Sens. Symp.*, 2017, pp. 3226–3229.
[7] B. Huang, K. Bradbury, L. M. Collins, and J. M. Malof, "Do deep learning models generalize to overhead imagery from novel geographic domains? The xGD benchmark problem," in *Proc. IEEE Int. Geosci. Remote Sens. Symp.*, 2020, pp. 1476–1479.
[8] D. Tuia, C. Persello, and L. Bruzzone, "Domain adaptation for the classification of remote sensing data: An overview of recent advances," *IEEE Geosci. Remote Sens. Mag.*, vol. 4, no. 2, pp. 41–57, Jun. 2016.
[9] M. Wang and W. Deng, "Deep visual domain adaptation: A survey," *Neurocomputing*, vol. 312, pp. 135–153, 2018.
[10] V. M. Patel, R. Gopalan, R. Li, and R. Chellappa, "Visual domain adaptation: A survey of recent advances," *IEEE Signal Process. Mag.*, vol. 32, no. 3, pp. 53–69, May 2015.
[11] O. Tasar, Y. Tarabalka, A. Giros, P. Alliez, and S. Clerc, "StandardGAN: Multi-source domain adaptation for semantic segmentation of very high resolution satellite images by data standardization," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Workshops*, 2020, pp. 747–756.
[12] O. Tasar, S. Happy, Y. Tarabalka, and P. Alliez, "SEMI2I: Semantically consistent image-to-image translation for domain adaptation of remote sensing data," in *Proc. IEEE Int. Geosci. Remote Sens. Symp.*, 2020, pp. 1837–1840.
[13] O. Tasar, A. Giros, Y. Tarabalka, P. Alliez, and S. Clerc, "DAugNet: Unsupervised, multisource, multitarget, and life-long domain adaptation for semantic segmentation of satellite images," *IEEE Trans. Geosci. Remote Sens.*, vol. 59, no. 2, pp. 1067–1081, Feb. 2021.
[14] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 770–778.
[15] 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, May 2010.
[16] 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, May 2016.
[17] C. Chen et al., "Progressive feature alignment for unsupervised domain adaptation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 627–636. [Online]. Available: https://ieeexplore.ieee.org/document/8953748/
[18] J. Choi, M. Jeong, T. Kim, and C. Kim, "Pseudo-labeling curriculum for unsupervised domain adaptation," in *Proc. Brit. Mach. Vis. Conf.*, 2019, pp. 69.1–69.13, doi: 10.5244/C.33.69.
[19] H. Huang, Q. Huang, and P. Krahenbuhl, "Domain transfer through deep activation matching," in *Proc. Eur. Conf. Comput. Vis.*, 2018, pp. 590–605.
[20] Y.-H. Tsai, W.-C. Hung, S. Schulter, K. Sohn, M.-H. Yang, and M. Chandraker, "Learning to adapt structured output space for semantic segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 7472–7481.
[21] J. Hoffman, D. Wang, F. Yu, and T. Darrell, "FCNS in the wild: Pixel-level adversarial and constraint-based adaptation," 2016, *arXiv:1612.02649*.
[22] J. Zhang, C. Liang, and C.-C. J. Kuo, "A fully convolutional tri-branch network (FCTN) for domain adaptation," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2018, pp. 3001–3005.
[23] Y. Zou, Z. Yu, B. Kumar, and J. Wang, "Unsupervised domain adaptation for semantic segmentation via class-balanced self-training," in *Proc. Eur. Conf. Comput. Vis.*, 2018, pp. 289–305.
[24] X. Deng, H. L. Yang, N. Makkar, and D. Lunga, "Large scale unsupervised domain adaptation of segmentation networks with adversarial learning," in *Proc. IEEE Int. Geosci. Remote Sens. Symp.*, 2019, pp. 4955–4958.
[25] 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, Feb. 2001.
[26] L. Bruzzone and D. F. Prieto, "A partially unsupervised cascade classifier for the analysis of multitemporal remote-sensing images," *Pattern Recognit. Lett.*, vol. 23, pp. 1063–1071, 2002.

[27] Y. Qin, L. Bruzzone, and B. Li, "Tensor alignment based domain adaptation for hyperspectral image classification," *IEEE Trans. Geosci. Remote Sens.*, vol. 57, no. 11, pp. 9290–9307, Nov. 2019.
[28] R. C. Gonzalez and R. E. Woods, *Digital Image Processing*. Hoboken, NJ, USA: Prentice-Hall, 2002.
[29] G. Buchsbaum, "A spatial processor model for object colour perception," *J. Franklin Inst.*, vol. 310, no. 1, pp. 1–26, 1980.
[30] F. Pacifici, N. Longbotham, and W. J. Emery, "The importance of physical quantities for the analysis of multitemporal and multiangular optical very high spatial resolution images," *IEEE Trans. Geosci. Remote Sens.*, vol. 52, no. 10, pp. 6241–6256, Oct. 2014.
[31] D. A. Forsyth, "A novel algorithm for color constancy," *Int. J. Comput. Vis.*, vol. 5, no. 1, pp. 5–35, 1990.
[32] K. I. Itten and P. Meyer, "Geometric and radiometric correction of TM data of mountainous forested areas," *IEEE Trans. Geosci. Remote Sens.*, vol. 31, no. 4, pp. 764–770, Jul. 1993.
[33] D. Tuia, J. Munoz-Mari, L. Gomez-Chova, and J. Malo, "Graph matching for adaptation in remote sensing," *IEEE Trans. Geosci. Remote Sens.*, vol. 51, no. 1, pp. 329–341, Jan. 2013.
[34] D. Das and C. G. Lee, "Unsupervised domain adaptation using regularized hyper-graph matching," in *Proc. IEEE 25th Int. Conf. Image Process.*, 2018, pp. 3758–3762.
[35] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 2242–2251.
[36] J. Hoffman et al., "CyCADA: Cycle-consistent adversarial domain adaptation," in *Proc. Int. Conf. Mach. Learn.*, 2018, pp. 1989–1998.
[37] M.-Y. Liu, T. Breuel, and J. Kautz, "Unsupervised image-to-image translation networks," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2017, pp. 700–708.
[38] X. Huang, M.-Y. Liu, S. Belongie, and J. Kautz, "Multimodal unsupervised image-to-image translation," in *Proc. Eur. Conf. Comput. Vis.*, 2018, pp. 172–189.
[39] H.-Y. Lee, H.-Y. Tseng, J.-B. Huang, M. Singh, and M.-H. Yang, "Diverse image-to-image translation via disentangled representations," in *Proc. Eur. Conf. Comput. Vis.*, 2018, pp. 35–51.
[40] B. Benjdira, Y. Bazi, A. Koubaa, and K. Ouni, "Unsupervised domain adaptation using generative adversarial networks for semantic segmentation of aerial images," *Remote Sens.*, vol. 11, no. 11, 2019, Art. no. 1369.
[41] A. Buslaev, V. I. Iglovikov, E. Khvedchenya, A. Parinov, M. Druzhinin, and A. A. Kalinin, "Albumentations: Fast and flexible image augmentations," *Information*, vol. 11, no. 2, 2020, Art. no. 125.
[42] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in *Proc. Int. Conf. Med. Image Comput. Comput.- Assist. Interv.*, 2015, pp. 234–241.
[43] V. Iglovikov and A. Shvets, "TernausNet: U-Net with VGG11 encoder pre-trained on ImageNet for image segmentation," 2018, *arXiv:1801.05746*.
[44] V. Nair, P. Rhee, J. Yang, B. Huang, K. Bradbury, and J. M. Malof, "Designing synthetic overhead imagery to match a target geographic region: Preliminary results training deep learning models," in *Proc. IEEE Int. Geosci. Remote Sens. Symp.*, 2020, pp. 948–951.
[45] J. Snoek, H. Larochelle, and R. P. Adams, "Practical Bayesian optimization of machine learning algorithms," in *Proc. Adv. Neural Inf. Process. Syst.*, 2012, pp. 2951–2959.
[46] X. Mao, Q. Li, H. Xie, R. Y. Lau, Z. Wang, and S. P. Smolley, "Least squares generative adversarial networks," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 2813–2821.
[47] M. Mirza and S. Osindero, "Conditional generative adversarial nets," 2014, *arXiv:1411.1784*.

Can Yaras received the B.S.E. degree in electrical and computer engineering and the B.S. degree in mathematics from Duke University, Durham, NC, USA, in 2021, and the M.S. degree in electrical and computer engineering in 2023 from the University of Michigan, Ann Arbor, Ann Arbor, MI, USA, where he is currently working toward the Ph.D. degree in ECE with the Electrical Engineering and Computer Science Department.

His research interests include representation learning, efficient deep learning, and nonconvex optimization.

Kaleb Kassaw (Student Member, IEEE) received the B.S.E.E. degree in electrical engineering from the University of Arkansas, Fayetteville, AR, USA, in 2020 and the M.S. degree in electrical and computer engineering in 2023 from Duke University, Durham, NC, USA, where he is currently working toward the Ph.D. degree in electrical and computer engineering, advised by Dr. Leslie Collins and Dr. Jordan Malof.

His research focuses on applied machine learning and computer vision.

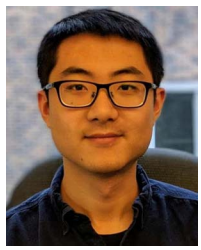

Bohao Huang (Student Member, IEEE) received the B.S. degree in electrical, electronics and communications engineering from the University of Electronic Science and Technology of China, Chengdu, China, in 2011 and the M.S. and Ph.D. degrees in electrical and computer engineering from Duke University, Durham, NC, USA, in 2017 and 2021, respectively. His work in Ph.D. focuses on the application of computer vision and machine learning in remote sensing.

Kyle Bradbury (Member, IEEE) received the B.S.E.E. degree in electrical engineering from Tufts University, Medford, MA, USA, in 2007, the M.S. degree in electrical and computer engineering and the Ph.D. degree in energy systems modeling from Duke University, Durham, NC, USA, in 2008 and 2013, respectively.

He is currently an Assistant Research Professor in Electrical and Computer Engineering with Duke University and the Director of the Energy Data Analytics Lab, Nicholas Institute for Energy, Environment and Sustainability, Durham, NC, USA. His research focuses on developing and applying machine learning techniques to better understand, plan, and manage energy infrastructure, scarce energy resources, and climate impacts. In particular, his work focuses on the use of remotely sensed data including satellite and aerial imagery.

Jordan M. Malof (Member, IEEE) received the B.S. degree in electrical and computer engineering (ECE) from the University of Louisville, Louisville, KY, USA, in 2008 and the Ph.D. degree in ECE from Duke University, Durham, NC, USA, in 2015.

He is currently an Assistant Professor with the Department of Computer Science, University of Montana, Missoula, MT, USA, where he develops advanced computer vision, machine learning, especially deep learning approaches to solve challenging real-world problems in fields such as materials science, remote sensing, and defense.