
Leveraging Domain Adaptation for Low-Resource Geospatial Machine Learning

Jack Lynch^{1,2} Sam Wookey²

Abstract

Machine learning in remote sensing has matured alongside a proliferation in availability and resolution of geospatial imagery, but its utility is bottlenecked by the need for labeled data. What's more, many labeled geospatial datasets are specific to certain regions, instruments, or extreme weather events. We investigate the application of modern domain-adaptation to multiple proposed geospatial benchmarks, uncovering unique challenges and proposing solutions to them.

1. Introduction

The use of machine learning for remote sensing has matured alongside an increase in the availability and resolution of satellite imagery, enabling advances in such tasks as land use classification (Campos-Taberner et al., 2020), natural risk estimation (Cerrón et al., 2020), disaster damage assessment (Xu et al., 2019), and agricultural forecasting (Lees et al., 2020).

However, many labeled geospatial datasets—including all of those used in the previously cited works—only cover *specific regions* of the world, limiting their utility and adoption elsewhere. Xu et al. explicitly cite this limitation as a focus of future work. Furthermore, these datasets are often limited to a single instrument and method of acquisition, limiting their generalization to acquisitions made by other instruments and at other angles.

Domain adaptation attempts to leverage labeled “source” data toward learning on some separate “target” dataset, for which only unlabeled data exists. In the context of climate-focused machine learning, it provides a potential framework for scaffolding innovations beyond their limiting datasets to other areas, instruments, and contexts.

¹Department of Electrical Engineering, NC State University, Raleigh, North Carolina, USA ²Masterful.AI, San Francisco, California, USA. Correspondence to: Jack Lynch <jmlynch3@ncsu.edu, jack@masterful.ai>.

In this proposal, we claim domain adaptation will be necessary to widen the equitable adoption of machine learning for geospatial applications, enabling its use for environmental monitoring and forecasting in areas of the world without labeled data. We propose multiple geospatial domain-adaptation benchmarks, explore the challenges they pose to traditional domain-adaptation algorithms, and suggest methods of improvement.

2. Geospatial Domain Adaptation

The application of domain adaptation to geospatial machine learning requires both datasets and methods tailored to the context of remote sensing.

2.1. Proposed Benchmarks

For our initial experiments we propose two domain-adaptation benchmarks derived from the well-known SpaceNet datasets of labeled satellite imagery (Etten et al., 2019). As noted by Dao et al. (2019), land use determines a significant portion of human carbon emission; as such, we focus on building segmentation as a useful tool in the monitoring and prediction of emissions. We measure the Jaccard index (Jaccard, 1912), or intersection over union (IoU), of target-dataset segmentations as our primary performance metric.

SpaceNet-2 contains imagery and corresponding building annotations for four cities: Las Vegas, Paris, Shanghai, and Khartoum. We may treat one subset of cities as the “source” and another subset as the “target” and attempt adaptation from one to the other: for example, from Las Vegas to Khartoum. The ability to generalize building segmentation from well-labeled cities to growing ones in different areas of the world would greatly increase the utility of such methods. Given compute constraints, we restrict our experiments here to those treating Khartoum as the target dataset.

SpaceNet-4 contains imagery of specific regions for varying angles off-nadir, loosely grouped into “on-nadir” for low angles, “off nadir” for moderate angles, and “very off-nadir” for large angles. Off-nadir imagery introduces a lower effective resolution and is marked in urban settings by building “tilt” relative to on-nadir imagery. We can attempt adaptation



Figure 1. Example acquisitions from the SpaceNet-2 (left) and SpaceNet-4 (right) datasets. The shift from Las Vegas to Khartoum is characterized by different colors, textures, building styles and arrangements, and (occasionally) the angle off-nadir. When moving from On- to Off-Nadir data, buildings move from “top-down” footprints to three-dimensional objects, and the effective resolution degrades.

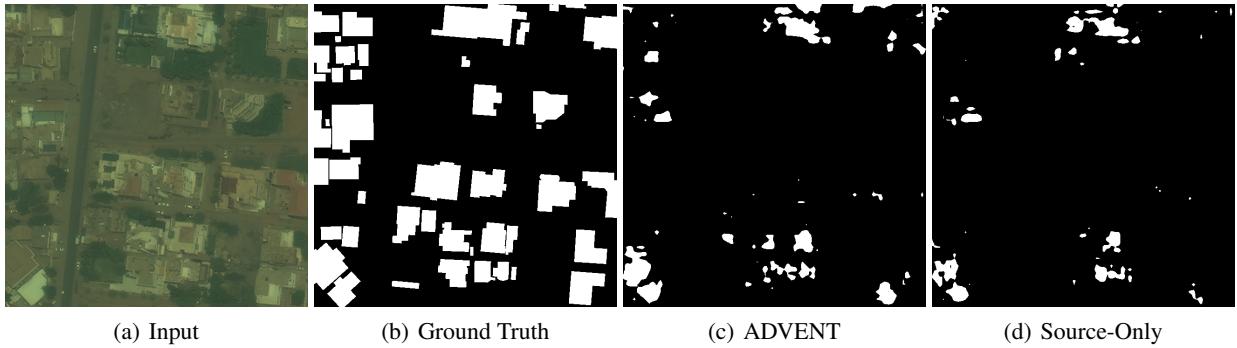


Figure 2. Example predictions from both ADVENT and source-only (non-adaptive) Vegas-to-Khartoum training for (a). Both ADVENT and source-only approaches fail to identify most “building” pixels, and ADVENT does not improve over source-only.

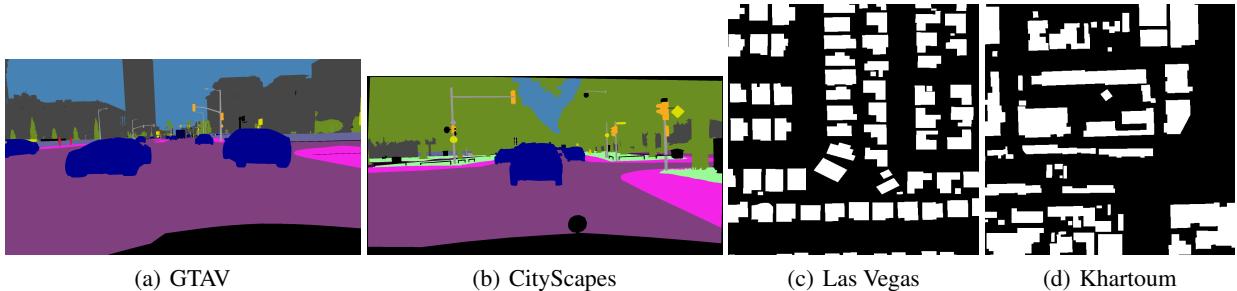


Figure 3. Example labels for GTA, CityScapes, and the Las Vegas and Khartoum subsets of SpaceNet-2. Though GTA features some finer-grain labeling detail, it might otherwise be difficult to tell its labels apart from those of CityScapes. In contrast, the binary labels of Las Vegas and Khartoum are easily distinguished, most notably by the difference in building shapes and arrangement.

from on-nadir to off-nadir or very off-nadir imagery, from off-nadir to very off-nadir, etc. As off-nadir imagery is often the quickest path to imaging during disasters (M. Bacastow et al., 2019), the ability to extend more-common “on-nadir” performance to off-nadir scenarios could aid in accelerating disaster response.

Samples from these datasets are shown in Figure 1. Notably, off-nadir acquisition angle varies significantly between cities in SpaceNet-2, suggesting the two benchmarks feature some overlap in sources of difficulty.

2.2. Methods

For our experiments we adopt the ADVENT approach of Vu et al. (2019), a relatively simple adversarial domain adaptation method whose precursor (Tsai et al., 2020) serves as a first stage for the current state of the art (Zhang et al., 2021). We perform all experiments using the same hyperparameters provided in Vu et al. (2019).

Notably, the ADVENT method, alongside precursors like Tsai et al. (2020); Hoffman et al. (2017), assumes that source and target label distributions are significantly more similar than their corresponding input distributions; otherwise, the discriminator would be able to distinguish between even exceptionally domain-general source and target predictions, diminishing its utility as a regularizer.

2.3. Challenges

2.3.1. LIMITED SOURCE DATA

Most common domain-adaptation benchmarks—for example, VisDA (Peng et al., 2017) in classification or GTA-to-CityScapes (Richter et al., 2016; Cordts et al., 2016) in segmentation—provide a large amount of labeled source data; as such, many recent domain-adaptation algorithms have not been examined in the context of limited source data (what could be considered “few-shot” domain adaptation).

The few-shot limitations specific to adversarial domain adaptation are not well understood, but they are especially relevant to geospatial machine learning, where new unlabeled data appears daily but labeled datasets are often small. For example, in our city-to-city benchmarks, no single city features more than 5,000 labeled samples, while none of on-nadir, off-nadir, or very-off-nadir collections feature more than 10,000.

We observe that this limited source data hampers the adaptive ability of ADVENT, often performing worse than source-only training as shown in Table 1 and Figure 2¹.

¹For GTA-to-CityScapes, we report the ADVENT metric from Vu et al. (2019) and the source-only metric from Tsai et al. (2020). Though we didn’t replicate the results of Vu et al. due to compute constraints, we extended their open-sourced code for our own

Though ADVENT trains stably for hundreds of thousands of iterations on traditional domain-adaptation benchmarks, on most city-to-city and nadir-angle paired datasets it can only train for thousands before beginning to diverge. We attribute this poor performance to discriminator overfitting (Karras et al., 2020), leading to the injection of unhelpful gradients during training.

2.3.2. DISPARATE LABEL DISTRIBUTIONS

As mentioned previously, adversarial methods like ADVENT and its predecessors assume that source and target label distributions are significantly more similar than their corresponding image/input distributions. When label distributions are dissimilar, source and target predictions can be more easily distinguished at the same level of domain-generality.

Visual inspection suggests that the cities in SpaceNet-2 feature larger gaps between their label distributions than do GTA and CityScapes (Figure 3, in the appendix), but this remains to be rigorously examined. We plan to consult structure-preserving dimensionality reduction through methods like UMAP (McInnes et al., 2020; Sainburg et al., 2021) to assess the relative “gaps” between pairs of source and target label distributions.

More specifically, a relative, quantitative assessment could be made by applying persistent homology to UMAP complexes of mixed source and target labels, observing average cluster “label” purity (here, referring to the binary distinction of source versus target) as a function of increasing similarity. Integrals of these purity curves—one for every pair of source/target label distributions—could be used to compare cluster separability between pairs of label distributions, in turn suggesting their inherent ease of discrimination during adversarial training.

3. Proposed Solutions

To address the problems described above, we propose the use of an adaptive discriminator as introduced in Karras et al. (2020). Though our needs differ from theirs—we do not fear augmentation leakage—an adaptive discriminator should address both issues encountered:

- It has been shown effective at avoiding discriminator overfitting in low-data regimes; and
- The geometric class of adaptive augmentations (rotations, affine transforms, etc.) can improve overlap in label distributions, potentially further improving adversarial performance.

Early results from the introduction of aggressive augmentation experiments.

Table 1. Improvement in target Intersection over Union (IoU) using ADVENT, relative to source-only training. V, S, P, and K represent SpaceNet-2 data for cities Las Vegas, Shanghai, Paris, and Khartoum, respectively. On and V. Off represent SpaceNet-4 “on-nadir” and “very off-nadir” subsets, respectively. GTA and CS represent the common domain-adaptation benchmark datasets of Grand Theft Auto and CityScapes, respectively. Certain combinations were omitted due to compute constraints.

	GTA → CS	V → K	V, P → K	P, S → K	V, S, P → K	ON → V. OFF
IoU (ADVENT)	47.6	13.59	9.95	26.36	25.05	11.03
IoU (SRC-ONLY)	36.6	15.09	17.56	23.62	30.09	14.77
Δ IOU	+11.0	-1.50	-7.61	+2.74	-5.04	-3.74

tions are promising (see Appendix).

With respect specifically to adaptation across instruments or collection methods, as with on-to-off-nadir, we will investigate modeling instrument parameters (e.g. off-nadir angle) through matrix capsules (Hinton et al., 2018), adapted to segmentation per LaLonde & Bagci (2018).

There is also work to be done in the construction of more targeted geospatial domain-adaptation benchmark datasets. For example, as off-nadir imagery is commonly used in disaster response, a dataset for on-to-off-nadir damage estimation would be more directly related to near-term applications.

Finally, geospatial domain adaptation has the unique facet of near-unlimited unlabeled data; as such, in situations where an appropriately delimited target dataset can be automatically aggregated, domain-adaptive methods with limited source data might benefit greatly from self-supervised pre-training, e.g. with Zbontar et al. (2021).

Acknowledgements

The authors would like to thank Leland McInnes for his direction in determining how best to apply structure-preserving dimensionality reduction to comparisons of distribution similarity, as well as Vu et al. (2019) and Karras et al. (2020) for their open-source code, which we extended. Finally, the authors would like to thank our anonymous reviewers for their feedback and Kartik Chandra for his suggestions while revising.

References

Campos-Taberner, M., García-Haro, F. J., Martínez, B., Izquierdo-Verdiguier, E., Atzberger, C., Camps-Valls, G., and Gilabert, M. A. Understanding deep learning in land use classification based on sentinel-2 time series. *Scientific Reports*, 10(1):17188, Oct 2020. ISSN 2045-2322. doi: 10.1038/s41598-020-74215-5. URL <https://doi.org/10.1038/s41598-020-74215-5>.

Cerrón, B., Bazan, C., and Coronado, A. Detection of housing and agriculture areas on dry-riverbeds for the evaluating and risk by landslides using low-resolution satellite imagery based on deep learning. Study zone: Lima, Peru. In *ICML 2020 Workshop: Tackling Climate Change with Machine Learning*, 2020.

Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., and Schiele, B. The cityscapes dataset for semantic urban scene understanding. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.

Dao, D., Rausch, J., and Zhang, C. GeoLabels: Towards Efficient Ecosystem Monitoring using Data Programming on Geospatial Information. In *NeurIPS 2019 Workshop: Tackling Climate Change with Machine Learning*, 2019.

Etten, A. V., Lindenbaum, D., and Bacastow, T. M. Spacenet: A remote sensing dataset and challenge series, 2019.

Hinton, G. E., Sabour, S., and Frosst, N. Matrix capsules with EM routing. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=HJWLfGWRb>.

Hoffman, J., Tzeng, E., Park, T., Zhu, J.-Y., Isola, P., Saenko, K., Efros, A. A., and Darrell, T. Cycada: Cycle-consistent adversarial domain adaptation, 2017.

Jaccard, P. The distribution of the flora in the alpine zone.1. *New Phytologist*, 11(2):37–50, 1912. doi: <https://doi.org/10.1111/j.1469-8137.1912.tb05611.x>. URL <https://nph.onlinelibrary.wiley.com/doi/abs/10.1111/j.1469-8137.1912.tb05611.x>.

Karras, T., Aittala, M., Hellsten, J., Laine, S., Lehtinen, J., and Aila, T. Training generative adversarial networks with limited data, 2020.

LaLonde, R. and Bagci, U. Capsules for object segmentation, 2018.

Lees, T., Tseng, G., Dadson, S., Hernández, A., G. Atzberger, C., and Reece, S. A Machine Learning Pipeline to Predict Vegetation Health. In *ICML 2020 Workshop: Tackling Climate Change with Machine Learning*, 2020.

M. Bacastow, T. Van Etten, A., and Weir, N. Automated Feature Extraction Using High Off-Nadir Satellite Imagery for Humanitarian Assistance and Disaster Response (HADR). In *GTC-DC 2019*, 2019. URL <https://developer.nvidia.com/gtc-dc/2019/video/dc91263-vid>.

McInnes, L., Healy, J., and Melville, J. Umap: Uniform manifold approximation and projection for dimension reduction, 2020.

Peng, X., Usman, B., Kaushik, N., Hoffman, J., Wang, D., and Saenko, K. Visda: The visual domain adaptation challenge, 2017.

Richter, S. R., Vineet, V., Roth, S., and Koltun, V. Playing for data: Ground truth from computer games, 2016.

Sainburg, T., McInnes, L., and Gentner, T. Q. Parametric umap embeddings for representation and semi-supervised learning, 2021.

Tsai, Y.-H., Hung, W.-C., Schulter, S., Sohn, K., Yang, M.-H., and Chandraker, M. Learning to adapt structured output space for semantic segmentation, 2020.

Vu, T.-H., Jain, H., Bucher, M., Cord, M., and Pérez, P. Advent: Adversarial entropy minimization for domain adaptation in semantic segmentation, 2019.

Xu, J. Z., Lu, W., Li, Z., Khaitan, P., and Zaytseva, V. Building damage detection in satellite imagery using convolutional neural networks, 2019.

Zbontar, J., Jing, L., Misra, I., LeCun, Y., and Deny, S. Barlow twins: Self-supervised learning via redundancy reduction, 2021.

Zhang, P., Zhang, B., Zhang, T., Chen, D., Wang, Y., and Wen, F. Prototypical pseudo label denoising and target structure learning for domain adaptive semantic segmentation, 2021.

A. Effect of Augmentations on Domain Adaptation

Early results using augmentations have been promising. For example, applying the augmentations of Karras et al. (2020) with a fixed per-augmentation probability significantly improves both source-only and ADVENT performance, and improves ADVENT’s performance *relative* to source-only, as shown in Figure 4 and Table 2 for Vegas-to-Khartoum adaptation. It remains to be seen if this effect persists across different benchmarks for multiple runs and, if so, how it varies across per-augmentation probabilities.

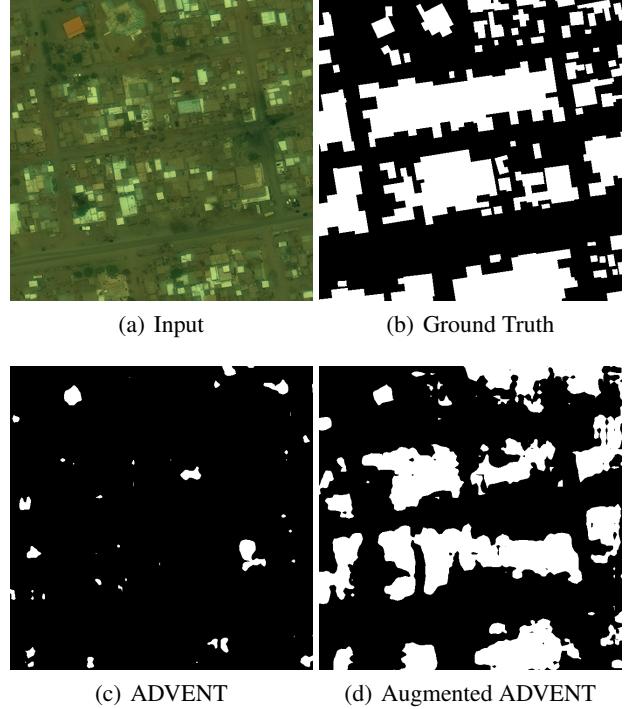


Figure 4. Example Vegas-to-Khartoum predictions for (a), from ADVENT, with and without augmentations (fixed per-augmentation probability of 0.6).

Table 2. Intersection over Union (IoU) for source-only and ADVENT training on Vegas-to-Khartoum adaptation, with and without augmentations (fixed per-augmentation probability of 0.6).

	SOURCE-ONLY	ADVENT	Δ IOU
NO AUG.	15.09	13.59	-1.50
AUG.	33.81	36.00	+2.19

A.1. Effect of Augmentations on Adversarial and Source-Only Adaptation

The observed improvement augmentation brings to both source-only and ADVENT-based adaptation is significantly greater than the *relative* improvement it confers to ADVENT over source-only. As these augmentations are comparatively simple to apply, they represent a reasonable, unobtrusive starting-point for improving the domain adaptation of relevant supervised tasks in geospatial machine learning.

A.2. Role of Adaptive Augmentation Probability

Though we do not require adaptive augmentation probability for the reasons described by Karras et al. (2020), it remains of interest whether adaptive probability schedules can prove more useful during training than fixed probabilities.