

# Handling Image and Label Resolution Mismatch in Remote Sensing

Scott Workman  
DZYNE Technologies

Armin Hadzic  
DZYNE Technologies

M. Usman Rafique  
Kitware Inc.

## Abstract

*Though semantic segmentation has been heavily explored in vision literature, unique challenges remain in the remote sensing domain. One such challenge is how to handle resolution mismatch between overhead imagery and ground-truth label sources, due to differences in ground sample distance. To illustrate this problem, we introduce a new dataset and use it to showcase weaknesses inherent in existing strategies that naively upsample the target label to match the image resolution. Instead, we present a method that is supervised using low-resolution labels (without up-sampling), but takes advantage of an exemplar set of high-resolution labels to guide the learning process. Our method incorporates region aggregation, adversarial learning, and self-supervised pretraining to generate fine-grained predictions, without requiring high-resolution annotations. Extensive experiments demonstrate the real-world applicability of our approach.*

## 1. Introduction

Semantic segmentation is a fundamental computer vision problem where the goal is to assign each individual pixel of an image to a semantic class. This research area has been heavily explored for decades and is critical for many applications, such as autonomous driving [7, 32]. Recently, advances in machine learning have pushed performance levels to new heights. However, despite the success of these methods when applied to ground-level imagery, there remain numerous challenges to successfully applying semantic segmentation to imagery from the remote sensing domain [11].

Central to the issue is that overhead imagery comes in many diverse formats. Considering space-based remote sensing alone, there are hundreds of different imaging sensors orbiting the earth, each capturing information in different ways. These sensors can have different imaging modalities (e.g., multispectral, radar), measure light differently, varying ground resolutions (also known as ground sample distance), unique look angles in relation to the target, and much more. Beyond capture details, the appearance of a scene can vary drastically for many reasons, including sea-




Figure 1: Semantic segmentation in remote sensing has many unique challenges, such as differences in spatial resolution between overhead imagery and target labels. This is primarily due to the increased resource cost of collecting fine-grained annotations at high-resolutions. We propose a method for handling this resolution mismatch, without the need for high-resolution annotations.

sonal variations and artifacts such as clouds and cloud shadows. Many of these issues are unique to remote sensing and require novel solutions.

The diversity of overhead imagery creates downstream problems for semantic segmentation algorithms. For example, given off-nadir imagery (i.e., not captured from directly overhead), there is often misalignment between the imagery and ground-truth labels. There is a large body of work exploring how to address related issues in remote sensing. For example, Christie et al. [2] propose a method to regress the geocentric pose of an overhead image and show its utility for building segmentation and rectification. Deng et al. [4] propose a framework for handling differences in scale between overhead image datasets. In this work, we focus our efforts on how to handle resolution mismatch between overhead imagery and target annotations, a relatively unexplored problem (Figure 1).

Semantic segmentation depends heavily on having high quality, aligned labels. A recent study by Zlateski et al. [46] found that segmentation performance primarily depends on the time spent creating training labels. With overhead imagery, high-resolution labels do not exist at a large scale.

This is primarily due to practical limitations associated with the enormous resource costs needed for comprehensive annotation efforts. For example, the Chesapeake Conservancy spent 10 months and \$1.3 million to produce a high-resolution (1 meter) land cover map of the Chesapeake Bay watershed [28]. The result is a land cover dataset for a small portion of the globe, with a fixed spatial resolution, at only a single timestep.

In practice, it is often significantly easier to acquire low-resolution labels. For example, the National Land Cover Database [42] (NLCD) is freely available to the public, has complete coverage of the continental United States, is typically updated every five years, but has a spatial resolution of only 30 meters per pixel. In the case that an input overhead image is of higher resolution than the target label, the typical strategy is to simply upsample the label to match the native resolution of the image. Our experiments demonstrate that this approach is unsatisfactory and results in low quality output.

Instead, our goal is to develop a method capable of generating fine-grained predictions, but only using low-resolution ground-truth as a direct source of supervision. Our key insight is that even if high-resolution ground-truth is unavailable for the training imagery, examples of high-resolution annotations are often readily available for other locations. In other words, it is possible to observe what high-resolution output looks like, just not everywhere. A method should be able to take advantage of this auxiliary data, without the corresponding imagery, to aid in producing a fine-grained output.

We present a method that is supervised only using low-resolution labels, but takes advantage of an exemplar set of high-resolution labels to guide the learning process. Our method has several key components: 1) it incorporates the concept of region aggregation to allow the network to output native resolution predictions, without upsampling the low-resolution target label; 2) it uses adversarial learning combined with an exemplar set of high-resolution labels to encourage predictions to be fine-grained; and 3) it leverages self-supervised pretraining on a large set of unlabeled imagery to increase model generalization. The result is a method which bridges the performance gap between naively upsampling available low-resolution labels and assuming the existence of expensive high-resolution labels. Extensive experiments demonstrate the capability of our approach.

## 2. Related Work

Semantic segmentation is considered a foundational task in computer vision, a necessary stepping stone towards the larger goal of intelligent scene understanding [16]. Much of the recent progress in semantic segmentation can be attributed to: 1) the development of learning-based segmentation algorithms (see [22] for a comprehensive survey) and

2) the introduction of large-scale benchmark datasets such as Cityscapes [3], ADE20k [44], and Mapillary Vistas [24], which contain pixel-wise annotations that are important for enabling fully-supervised segmentation approaches.

Semantic segmentation has attracted a great deal of attention in the remote sensing community as well [45]. It is apt for traditional remote sensing tasks such as land cover and land use estimation [35], which seek to understand the physical cover of the Earth's surface and how it is utilized. Mátyus et al. [21] segment roads in overhead imagery as part of a method for directly extracting road topology. Girard et al. [8] leverage semantic segmentation to extract building polygons. Other applications include height estimation [19, 37], road safety assessment [33], traffic monitoring [9, 38], emissions estimation [23], near/remote sensing [40, 41], and more. However, similar to ground-level imagery, a major difficulty in applying semantic segmentation to overhead imagery is the cost of acquiring annotated training data [11].

Remote-sensed imagery presents many unique obstacles to contend with. Imagery can come from many different sensors, have different spatial resolutions, contain atmospheric artifacts such as clouds, etc. Deng et al. [4] propose a method for handling the scale variation commonly found between remote sensing image collections and demonstrate its application to domain adaptation. Workman et al. [39] show how multi-image fusion can be applied to detecting artifacts in imagery, such as clouds. These issues, and more, compound the difficulty of getting annotated training data. Ultimately, it is common for annotations to have a different spatial resolution, be captured at a different time, have spatial alignment errors, or a slew of other discrepancies.

In this work, we focus specifically on the resolution mismatch problem. As high-resolution imagery has become more widely available, for example WorldView-3 imagery at approximately 30 centimeters, it has become increasingly difficult to acquire high-resolution annotations to match the native image resolution. Instead, the typical strategy is to simply upsample the target label to match the image resolution. The 2020 IEEE GRSS Data Fusion Contest (DFC2020) [43] was of the first to consider the problem of resolution mismatch, but at an extreme scale (high-resolution labels of 10 meters and low-resolution labels of 500 meters), necessitating non-learning based methods. Robinson et al. [29] use an ensemble of individual solutions including: label super-resolution via iterative clustering, label super-resolution using epitomic representations [20], and post-processing inspired by deep image prior [34]. Instead, we explore an end-to-end learning based approach for addressing the mismatch problem by generating a high-resolution prediction in a single forward pass without an ensemble or any ad hoc post-processing steps.

Several weakly supervised methods have been proposed

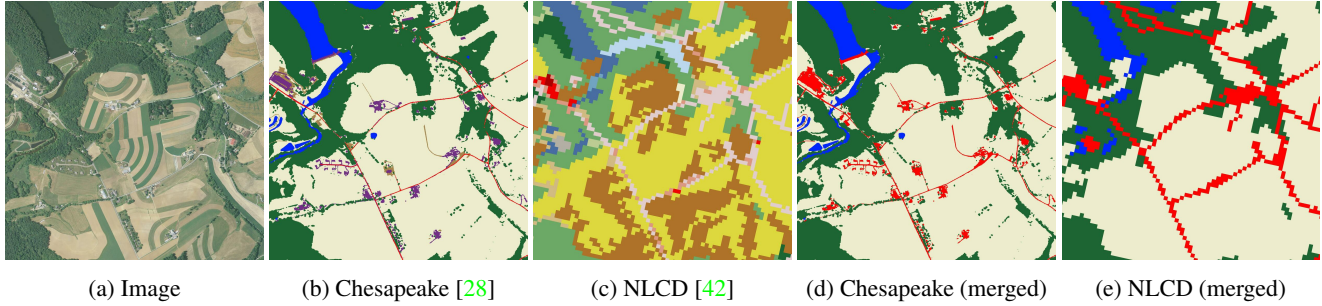


Figure 2: Example labels using our merged label taxonomy. Note that the low-resolution label (e) omits many fine-grained details, such as structures and secondary roads, that the high-resolution label (d) captures.

to take advantage of sparse ground-truth labels. Lu et al. [18] show how geotagged point labels can be used to identify bodies of water via the introduction of a feature aggregation network. Wang et al. [36] also consider geotagged points, as well as image-level labels, and show that standard segmentation architectures can outperform pixel-level algorithms. We consider an alternative, more common scenario, where annotations are dense but of a lower spatial resolution than the input imagery.

### 3. Low2High: A Dataset for Resolution Mismatch

To support our experiments, we introduce the Low2High dataset which extends the recent Chesapeake dataset [28]. The Chesapeake dataset contains high-resolution overhead imagery, at 1m per pixel, from the USDA National Agriculture Imagery Program (NAIP). In total, the dataset consists of over 700 non-overlapping tiles from six states, with each tile corresponding to an area of approximately  $6\text{km} \times 7.5\text{km}$ . This imagery is combined with land cover labels from two different sources. The first source is high-resolution labels (1m per pixel) obtained from the Chesapeake Conservancy land cover project covering approximately 100,000 square miles of land around the Chesapeake Bay watershed. The second source is low-resolution labels (30m per pixel) from the National Land Cover Database (NLCD). We extend this dataset to support our experiments related to the resolution mismatch problem.

#### 3.1. Dataset Generation

We start from the raw tiles in the Chesapeake dataset, update the land cover labels to use a merged label taxonomy, and generate an aligned dataset of non-overlapping images of size  $512 \times 512$ . We promote label diversity and omit any image with less than 3 unique land cover classes (either label source), or that has a single land cover class covering more than 75 percent of the image. This process results in 34791 images. We also generate a custom evaluation split, including a held-out set of images from each state. The final

split consists of training (75%), validation (10%), held-out (10%), and testing (5%).

**Merged Label Taxonomy** The high-resolution land cover labels from the Chesapeake Conservancy include six land cover classes: tree canopy / forest, low vegetation / field, barren land, impervious (other), and impervious (road). The low-resolution land cover labels from NLCD include twenty land cover classes with high-level categories spanning developed, barren, forest, and shrub land. For our purposes, we generate a merged label taxonomy consisting of four classes: water, forest, field, and impervious. Each label source is remapped to use the merged label taxonomy by assigning existing labels to the most relevant category (see the supplemental material for details). Figure 2 shows qualitative examples of this process. Though the labels often tend to agree there can be stark differences, with the low-resolution label missing certain details completely.

#### 3.2. Auxiliary Imagery & Labels

Self-supervised learning, where a supervisory signal is obtained directly from the data itself, is an alternative training paradigm capable of taking advantage of large amounts of unlabeled data. To support this scenario and improve geographic diversity, we supplement Low2High with an auxiliary set of images collected near rest stops across the contiguous United States [27]. Similar to the format of the Chesapeake imagery, we obtained overhead imagery ( $512 \times 512$ ) from NAIP between 01/01/2013 and 12/31/2014, resulting in 12,024 training and 603 validation images.

We also supplement the Chesapeake data with a new held-out test region to support domain adaptation experiments. Milwaukee, Wisconsin was selected as the new test site due to its diverse geography compared to the Northeastern USA. For land cover labels, we use annotations made available through the EPA EnviroAtlas [26] and, as before, the land cover classes were remapped to use our merged label taxonomy. The resulting test set contains 3262 samples at 1m ground sample distance.

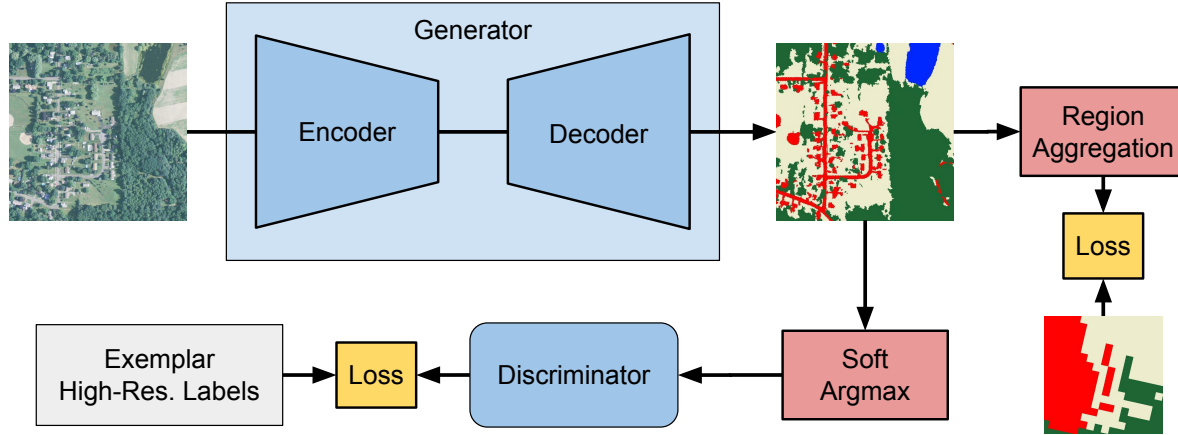


Figure 3: An overview of our architecture for handling resolution mismatch.

## 4. A Method for Semantic Segmentation with Resolution Mismatch

We propose an end-to-end architecture for semantic segmentation that is supervised using low-resolution annotations, but is capable of generating fine-grained predictions. Figure 3 provides a visual overview of our approach.

### 4.1. Approach Overview

We present a framework for semantic segmentation in the event of resolution mismatch. Our method is directly supervised with low-resolution labels, but takes advantage of an exemplar set of high-resolution labels (no correspondence to the input imagery) to guide the learning process. Our architecture has four primary components. First, we describe the general segmentation architecture (Section 4.2). Second, we incorporate the concept of region aggregation to allow the network to output native resolution predictions, without upsampling the low-resolution target label (Section 4.3). Third, we use adversarial learning combined with an exemplar set of high-resolution labels to encourage predictions to be fine-grained (Section 4.4). Finally, we leverage self-supervised pretraining on a large set of unlabeled imagery to increase model generalization (Section 4.5). The remainder of this section describes the high-level architecture. See the supplemental materials for additional details.

### 4.2. Segmentation Architecture

For our segmentation architecture, we use a variant of U-Net [30] with a ResNet-18 backbone. However, our approach is general enough to be conceivably combined with any modern segmentation architecture. U-Net is an encoder-decoder style architecture that propagates information from the encoder through the use of skip connections. For our feature encodings, we use the output activations from the first four stages of ResNet. The decoder (U-Net

style) expects four feature maps to be used for skip connections. Our variant has four upsampling blocks, each consisting of two convolutions ( $3 \times 3$ ) followed by a ReLU activation. The output is then upsampled and passed through a final convolutional layer with number of output channels that are task-specific. In our case, the segmentation architecture takes an overhead image as input and produces a native resolution output.

### 4.3. Region Aggregation

The typical strategy in the event of resolution mismatch between input image and target label is to simply upsample the label to the resolution of the input. Our experiments demonstrate that this approach is sub-optimal. Instead, we use a variant of region aggregation [12, 38] to allow our network to generate native resolution outputs, yet be supervised with low-resolution target labels (without any upsampling). This allows the underlying segmentation network to generate fine-grained predictions.

In the case of semantic segmentation of overhead imagery, the input imagery and target label are both georeferenced. This means that the ground sample distance of each is known, and enables computation of the geospatial correspondence between a single low-resolution pixel (from the target label) and many high-resolution prediction pixels (from the input image). See Figure 4 for a visual illustration of this geospatial relationship. The ratio of high-resolution pixels to low-resolution pixels ( $n$ ) can be directly computed from the respective ground sample distances.

We construct a region index matrix  $M \in \mathbb{R}^{H \times W}$ , which expresses a pixel mapping from high-resolution to low-resolution across a given image of height  $H$  and width  $W$ .  $M$  is composed of  $t$  number of  $n \times n$  pixel regions, where  $n$  is the ground sample distance scale-ratio from the high-resolution imagery to the low-resolution labels. For the case of mapping 1m NAIP imagery to 30m NLCD annotations,






Figure 4: Our region aggregation component takes advantage of the known geospatial context of an input image to aggregate high-resolution predictions such that they match the spatial resolution of a low-resolution target label.

$n = 30$ . We use  $M$  to select  $30 \times 30$  pixel regions from the high-resolution prediction, and sum them to produce a single value representative of the region. In practice, this operation is applied directly to the logits. This region-based summation operation results in a low-resolution prediction map. We then forward the low-resolution prediction map to a cross entropy loss:

$$L(Y, \hat{Y}) = -\frac{1}{N} \sum_{i=1}^N \log \left( \frac{e^{y_i}}{\sum_{c=1}^C e^{\hat{y}_{i,c}}} \right). \quad (1)$$

This allows newly down-scaled low-resolution predictions to be supervised by low-resolution annotations (bottom right of Figure 3). For the multi-class ( $c \in C$ ) cross entropy loss,  $y_i \in Y$  indicates that  $N$  low-resolution labels were used to supervise  $\hat{y} \in \hat{Y}$  aggregated predictions.

#### 4.4. Adversarial Learning

We use adversarial learning combined with an exemplar set of high-resolution labels to encourage predictions to be fine-grained. In other words, we treat the segmentation network (Section 4.2) as a generator ( $G$ ). The predictions from the generator are encouraged to match the appearance of an unrelated high-resolution annotation by having the discriminator ( $D$ ) predict whether an input is fake (a prediction) or real (a randomly selected high-resolution annotation). This is represented as a two-player minmax game:

$$\min_G \max_D L_{adv}(D, G). \quad (2)$$

To facilitate this training paradigm, predictions from the generator need to match the characteristics of randomly selected high-resolution annotations taken from a small held-out dataset. In our case, the discriminator operates on segmentation maps (indexed images) as input. The generator outputs logits ( $B \times C \times H \times W$ ), yet the discriminator expects inputs to be in the form of class indices ( $B \times 1 \times H \times W$ ). Transforming logits to class indices is

usually addressed by applying a *softmax* across the channels dimension followed by an *argmax* operation. However, *argmax* is non-differentiable, which would interrupt gradient flow from the discriminator to the generator during adversarial training.

Instead, we present a novel formulation using a differentiable soft *argmax* operation, denoted as *sargmax*. First, class probabilities are extracted from the input logits  $x$  by multiplying the logits by a temperature scaling parameter,  $\tau = 10^3$ , and applying a *softmax* operation across the channels dimension ( $c \in C$ ) of  $x$ , indicated as  $x_c$ . The *sargmax* can be expressed as:

$$sargmax := \sum_{c=1}^C c * p(x_c; \tau), \quad (3)$$

where class probabilities,  $p(x_c; \tau) = softmax(\tau x_c)$ , are subsequently passed to an expectation to produce an approximate class index per spatial position, corresponding to the highest predicted class probability. This formulation allows soft class assignments to each pixel of a segmentation output. Predictions and high-resolution exemplars are then passed to the discriminator,  $D$ , which will try to determine whether a given input is a real high-resolution annotation or a prediction synthesized by  $G$ . We follow this with the hinge loss [17],  $L_{adv}$ , to jointly optimize the generator and discriminator.

Kendall et al. [13] proposed a similar operation for differentiable aggregation denoted as *softargmin*. The key distinction between our *sargmax* and *softargmin* is that we use a temperature parameter  $\tau$  to make the predicted probability distributions have more pronounced peaks. Distinct peaks are desirable for allowing predictions to emulate class indices, as opposed to the more common trend of using *softargmin* for regression.

#### 4.5. Self-supervised Pretraining

Self-supervised learning has shown great promise for a variety of tasks [1, 15], including land cover segmentation [31]. The premise is to leverage various pretext tasks, which create supervisory signals from the data itself, for learning useful feature representations which can then be transferred to the downstream task. Masked Autoencoders (MAE) [10] introduced image reconstruction as a pretext task, showing that transfer performance on several downstream tasks outperformed supervised pretraining. Inspired by these promising results, we use the MAE reconstruction task to pretrain a Vision Transformer ViT [5] with a ResNet-18 patch embedding (masking ratio of 75%) followed by  $2 \times$  upsampling layer. The weights from the pretrained ResNet-18 embedding network are then used to initialize the feature backbone of our segmentation network (Section 4.2).

## 4.6. Implementation Details

We implement our methods using Pytorch [25] and Pytorch Lightning [6]. Our networks are optimized using Adam [14] with the initial learning rate set to  $1e^{-4}$ . The method is trained for 100 epochs and a validation set is used for model selection. For the adversarial learning component, the inputs passed to the discriminator are blurred using a Gaussian filter ( $3 \times 3$  kernel with  $\sigma = 0.6$ ).

## 5. Experiments

We evaluate our methods quantitatively and qualitatively through a variety of experiments. Results demonstrate that our approach which incorporates region aggregation, adversarial learning, and self-supervised pretraining, significantly reduces error compared to baselines.

**Baseline Methods** For evaluating our proposed architecture, we consider several baseline methods that share components of our full approach:

- *Oracle* uses the upsampled low-resolution ground-truth directly as the prediction. This represents having a perfect low-resolution estimator.
- *Low* is trained on the low-resolution ground-truth and naively upsamples to the native image resolution.
- *High* is trained on the high-resolution ground truth. This represents ideal performance in the event high-resolution ground-truth is always available.

Our full approach is outlined in Section 4 and is subsequently referred to as *Ours*. We also compare against the label super-resolution approach, *Self-Epitomic LSR* [20].

**Metrics** In our experiments, we evaluate all methods using the high-resolution labels. In other words, the goal is to examine how well each method performs relative to the high-resolution baseline, which represents optimal performance (direct supervision using high-resolution labels). We use the following standard evaluation metrics: F1 score (the harmonic mean of the precision and recall scores) and mean intersection-over-union (mIOU), also known as the Jaccard Index. For both metrics, we use the *macro* strategy which computes the metric for each class separately, and then averages across classes using equal weights for each class.

### 5.1. Case Study: Impact of Label Resolution

To begin, we study the impact of label resolution on prediction performance. For this experiment, we optimize the baselines outlined above, which share the same segmentation architecture as our approach (a modified U-Net with a ResNet-18 backbone). For each baseline, we train for 100 epochs using the Low2High dataset. Model selection is performed using a resolution-specific validation set, treating each label set in isolation.

Table 1: Performance of baseline methods.

	Acc	F1	mIOU
Oracle	80.17%	72.69%	60.28%
Low	84.40%	76.16%	64.25%
High	93.16%	90.87%	83.59%
<i>Self-Epitomic LSR</i> [20]	69.11%	50.90%	37.10%

Table 2: An ablation study highlighting the impact of different components. Our full approach significantly reduces the gap in performance to the high-resolution baseline.

			Acc	F1	mIOU
Agg.	App.	Pre.			
✓			82.63%	74.59%	62.52%
✓	✓		87.43%	81.12%	69.84%
✓	✓	✓	<b>88.61%</b>	<b>82.60%</b>	<b>71.69%</b>
Baseline: High			93.16%	90.87%	83.59%

Table 1 shows the result of this experiment. As expected, *High* performs best, due to being directly supervised using high-resolution ground truth. *Low* performs significantly worse, highlighting the performance cost associated with simply upsampling the low-resolution ground truth. Finally, *Oracle* performs worst, which represents a model capable of perfectly predicting the low-resolution ground truth. Together, these results show that high-resolution ground-truth is crucial for achieving high quality output in the form of fine-grained predictions.

For comparison, we also show the performance of the label super-resolution approach (*Self-Epitomic LSR*) proposed by Malkin et al. [20], adapting the implementation made available by the authors. This iterative method performs significantly worse than all other methods, despite requiring prior knowledge of statistics that describe how frequently high-resolution annotations occur within low-resolution annotations.

### 5.2. Ablation Study

Next, we evaluate the performance of our method, along with an ablation study showing the performance impact attributed to individual components. Table 2 shows the outcome of this experiment. The results indicate that using region aggregation (Agg.) alone with direct supervision from low-resolution annotations does not outperform the low-resolution baseline from Table 1. This is understandable, as the region aggregation step allows the segmentation network to output fine-grained predictions, but does not constrain the predictions to be realistic. In other words, the aggregation step simply encourages the aggregated prediction to agree with the target label.

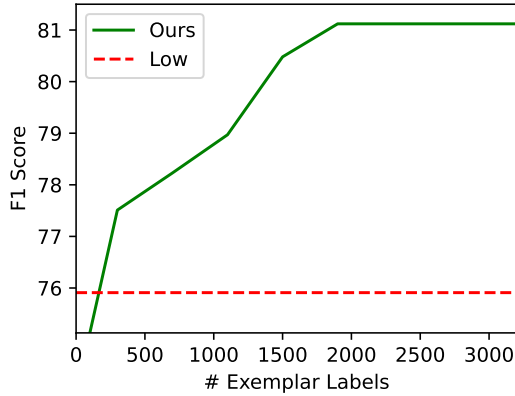


Figure 5: The performance of our proposed approach versus the quantity of exemplar labels used during model training.

However, including appearance-based adversarial training (App.) together with region aggregation (Agg.) results in performance that surpasses both the low-resolution and oracle baselines. This component encourages the output predictions to look similar in appearance to an exemplar set of high-resolution annotations. Finally, including self-supervised pretraining (Pre.) results in a 4.21% accuracy, 6.44% F1, and 7.44% mIOU gain over the low-resolution baseline. The full model’s performance approaches the high-resolution baseline, without requiring direct supervision from high-resolution labels. Figure 6 shows example outputs from our approach versus baselines.

### 5.3. Dependence on Exemplar Labels

Our method takes advantage of an exemplar set of high-resolution labels to guide the learning process. These labels have no correspondence with the input imagery, but encourage our method to produce realistic (i.e., fine-grained) predictions. Therefore, we evaluate how the performance of our proposed approach changes with respect to the quantity of exemplar labels used during model training. For this experiment, we use our full approach but omit the self-supervised pretraining component.

Figure 5 visualizes the results of this study. The x-axis represents the number of high-resolution labels used during model training, and the y-axis indicates the resulting model’s F1 score on the test set. As expected, model performance increases as the quantity of exemplar labels increases. Even with very few exemplar labels, our method is able to significantly outperform the low-resolution baseline.

### 5.4. Generalizing to Novel Locations

We use the Low2High dataset to investigate the impact of self-supervised pretraining on model generalization. For this experiment, we use the MAE image reconstruction strategy outlined in Section 4.5 to pretrain our method, but

Table 3: Domain adaption study with and without self-supervised pretraining (using our full model).

Pretrain	Test	Acc	F1	mIOU
None	Ches <sup>-</sup>	86.36%	80.07%	68.09%
Ches <sup>-</sup>	Ches <sup>-</sup>	88.38%	80.31%	68.68%
<b>USA*</b>	Ches <sup>-</sup>	<b>88.91%</b>	<b>82.19%</b>	<b>71.01%</b>
None	VA	85.50%	76.03%	63.35%
Ches <sup>-</sup>	VA	86.54%	77.03%	65.04%
<b>USA*</b>	VA	<b>88.40%</b>	<b>79.03%</b>	<b>67.81%</b>
None	MWI	55.93%	44.90%	30.81%
Ches <sup>-</sup>	MWI	55.41%	45.21%	30.99%
<b>USA*</b>	MWI	<b>63.32%</b>	<b>48.52%</b>	<b>35.51%</b>

vary the set of unlabeled images. The first set, Ches<sup>-</sup>, represents imagery from the training set minus Virginia (VA). VA was excluded to serve as a held-out test region from the Chesapeake Bay area. For the second set, we use the auxiliary imagery from Low2High (Section 3.2), denoted as USA\*. After pretraining for 100 epochs, the resulting models are fine-tuned for 100 epochs on the training portion of Ches<sup>-</sup> (states: DE, MD, NY, WV, and PA). Finally, we evaluate each of the two pretraining strategies and a baseline representing no pretraining (None) on three downstream held-out test sets: the test portion of Ches<sup>-</sup>, test samples from VA, and our new test set over Milwaukee, WI (MWI).

Table 3 summarizes the results of this experiment. The first three rows show performance for the Ches<sup>-</sup> test set. As expected, no pretraining (None) performs the worst. Notably, pretraining on USA\* outperforms pretraining on Ches<sup>-</sup>. As both pretraining sets have exposure to imagery similar to the testing region, this can likely be attributed to learning an improved feature representation from a more geographically diverse set of images.

The middle three rows show performance on the held-out VA test set. Unlike the prior three rows, no method had any previous exposure to VA imagery during model training. Despite this, the results follow the same trend as before. No pretraining (None) performs the worst and pretraining on USA\* outperforms pretraining on Ches<sup>-</sup>. This result shows that pretraining on a more diverse set of images (USA\*) is still beneficial for downstream learning, as compared to pretraining only on a set of images (Ches<sup>-</sup>) that is more geographically similar to the test set.

Finally, the bottom three rows show performance for the held-out MWI test set. In this scenario, the baseline (None) is competitive with Ches<sup>-</sup>. This is understandable, as MWI is geographically diverse compared to the Chesapeake Bay region. USA\* outperforms both alternatives, showcasing that self-supervised pretraining has a significant positive impact on generalizing a model to novel locations.



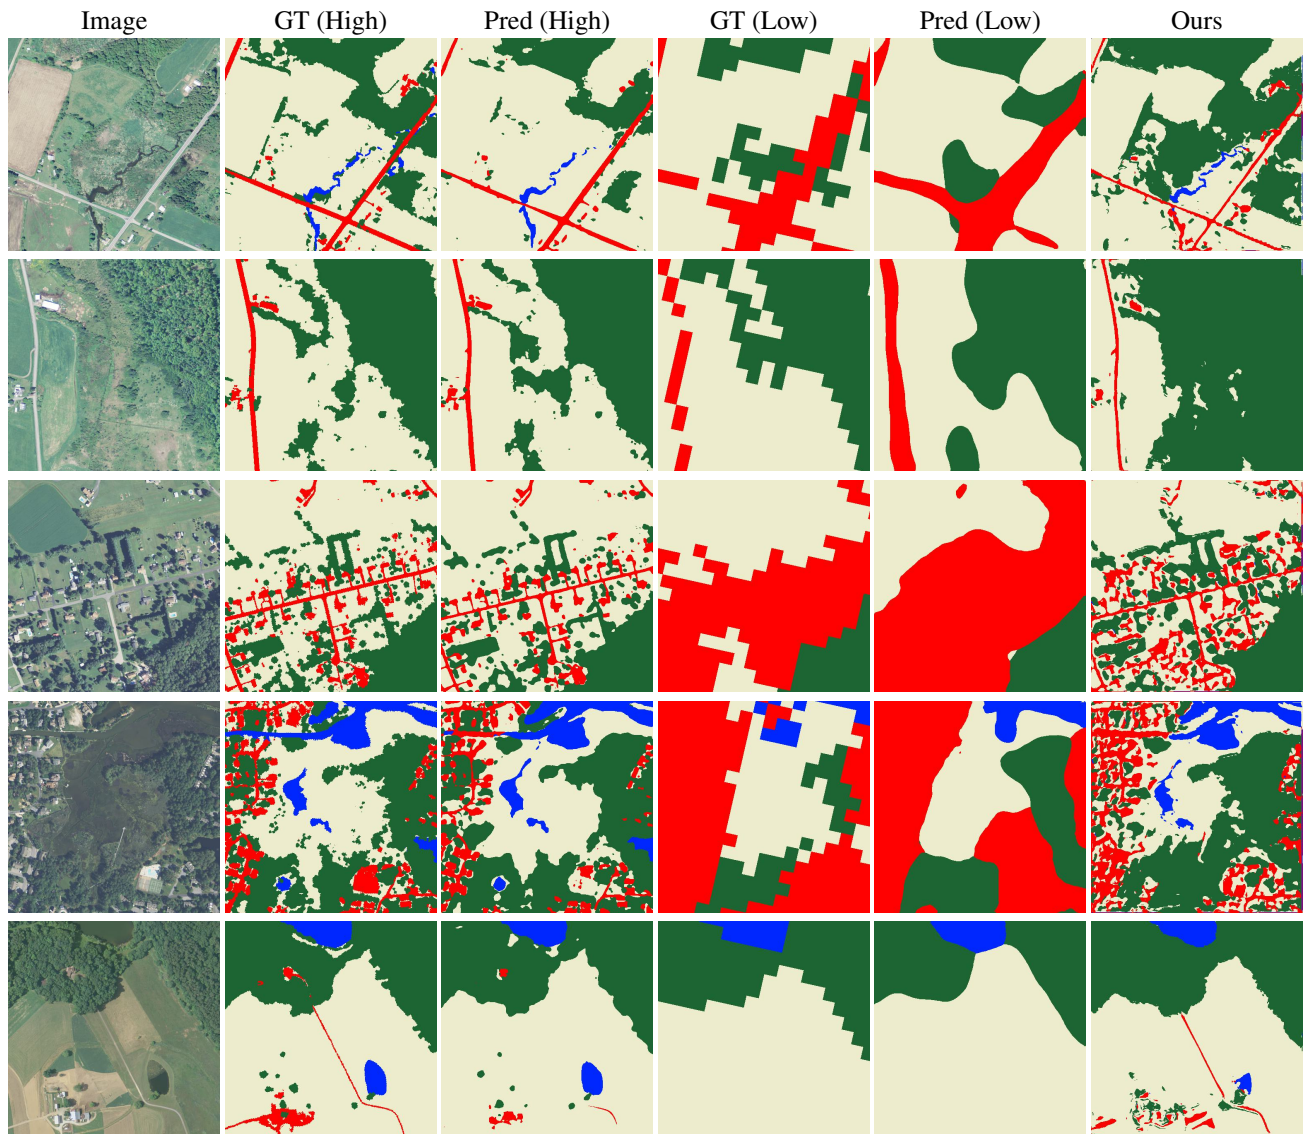


Figure 6: Example qualitative results. From left to right: input image, high-resolution ground-truth label, prediction from the high-resolution baseline, low-resolution ground-truth label, prediction from the low-resolution baseline, and our result.

## 6. Conclusion

Semantic segmentation has reached impressive performance levels, but in remote sensing, it is still extremely challenging to acquire high-resolution annotations. In practice, for many tasks it is much easier to get coarse, low-resolution labels. Our results demonstrate how naively up-sampling these labels results in low quality outputs that lack sufficient detail. Instead, we proposed a method that is supervised using low-resolution annotations, incorporates exemplar high-resolution labels, and generates fine-grained output. Experiments on a novel dataset demonstrate how our approach significantly reduces the performance gap to a

high-resolution baseline. This approach has the potential to have real-world applicability in the remote sensing domain.

**Acknowledgments** This research is based upon work supported in part by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via 2021-2011000005. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of ODNI, IARPA, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein.



## References

- [1] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International Conference on Machine Learning*, 2020.
- [2] Gordon Christie, Kevin Foster, Shea Hagstrom, Gregory D Hager, and Myron Z Brown. Single view geocentric pose in the wild. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2021.
- [3] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [4] Xueqing Deng, Yi Zhu, Yuxin Tian, and Shawn Newsam. Scale aware adaptation for land-cover classification in remote sensing imagery. In *IEEE Winter Conference on Applications of Computer Vision*, 2021.
- [5] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- [6] WA Falcon. Pytorch lightning. *GitHub. Note: <https://github.com/PyTorchLightning/pytorch-lightning>*, 3, 2019.
- [7] Di Feng, Christian Haase-Schütz, Lars Rosenbaum, Heinz Hertlein, Claudius Glaeser, Fabian Timm, Werner Wiesbeck, and Klaus Dietmayer. Deep multi-modal object detection and semantic segmentation for autonomous driving: Datasets, methods, and challenges. *IEEE Transactions on Intelligent Transportation Systems*, 22(3):1341–1360, 2020.
- [8] Nicolas Girard, Dmitriy Smirnov, Justin Solomon, and Yuliya Tarabalka. Polygonal building extraction by frame field learning. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2021.
- [9] Armin Hadzic, Hunter Blanton, Weilian Song, Mei Chen, Scott Workman, and Nathan Jacobs. Rasternet: Modeling free-flow speed using lidar and overhead imagery. In *IEEE Conference on Computer Vision and Pattern Recognition Workshop*, 2020.
- [10] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2022.
- [11] Yuansheng Hua, Diego Marcos, Lichao Mou, Xiao Xiang Zhu, and Devis Tuia. Semantic segmentation of remote sensing images with sparse annotations. *IEEE Geoscience and Remote Sensing Letters*, 19:1–5, 2021.
- [12] Nathan Jacobs, Adam Kraft, Muhammad Usman Rafique, and Ranti Dev Sharma. A weakly supervised approach for estimating spatial density functions from high-resolution satellite imagery. In *ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, 2018.
- [13] Alex Kendall, Hayk Martirosyan, Saumitro Dasgupta, Peter Henry, Ryan Kennedy, Abraham Bachrach, and Adam Bry. End-to-end learning of geometry and context for deep stereo regression. In *IEEE International Conference on Computer Vision*, 2017.
- [14] Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2014.
- [15] Junnan Li, Pan Zhou, Caiming Xiong, and Steven Hoi. Prototypical contrastive learning of unsupervised representations. In *International Conference on Learning Representations*, 2020.
- [16] Liulei Li, Tianfei Zhou, Wenguan Wang, Jianwu Li, and Yi Yang. Deep hierarchical semantic segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2022.
- [17] Jae Hyun Lim and Jong Chul Ye. Geometric gan. *arXiv preprint arXiv:1705.02894*, 2017.
- [18] Ming Lu, Leyuan Fang, and Yi Zhang. Point-based weakly supervised deep learning for water extraction from high-resolution remote sensing imagery. In *CIKM Workshops*, 2021.
- [19] Jisan Mahmud, True Price, Akash Bapat, and Jan-Michael Frahm. Boundary-aware 3d building reconstruction from a single overhead image. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2020.
- [20] Nikolay Malkin, Anthony Ortiz, and Nebojsa Jojic. Mining self-similarity: Label super-resolution with epitomic representations. In *European Conference on Computer Vision*, 2020.
- [21] Gellért Mátyus, Wenjie Luo, and Raquel Urtasun. Deep-roadmapper: Extracting road topology from aerial images. In *IEEE International Conference on Computer Vision*, 2017.
- [22] Shervin Minaee, Yuri Y Boykov, Fatih Porikli, Antonio J Plaza, Nasser Kehtarnavaz, and Demetri Terzopoulos. Image segmentation using deep learning: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021.
- [23] Ryan Mukherjee, Derek Rollend, Gordon Christie, Armin Hadzic, Sally Matson, Anshu Saksena, and Marisa Hughes. Towards indirect top-down road transport emissions estimation. In *IEEE Conference on Computer Vision and Pattern Recognition Workshop*, 2021.
- [24] Gerhard Neuhold, Tobias Ollmann, Samuel Rota Buló, and Peter Kotschieder. The mapillary vistas dataset for semantic understanding of street scenes. In *IEEE International Conference on Computer Vision*, 2017.
- [25] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*, 2019.
- [26] Andrew Pilant, Keith Endres, Daniel Rosenbaum, and Gillian Gundersen. Us epa enviroatlas meter-scale urban land cover (mulc): 1-m pixel land cover class definitions and guidance. *Remote sensing*, 12(12):1909, 2020.
- [27] Roundabout Publications. Interstate rest areas. <https://www.interstaterestareas.com/map-of-rest-areas>.

- [28] Caleb Robinson, Le Hou, Kolya Malkin, Rachel Soobitsky, Jacob Czawlytko, Bistra Dilkina, and Nebojsa Jojic. Large scale high-resolution land cover mapping with multi-resolution data. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2019.
- [29] Caleb Robinson, Kolya Malkin, Lucas Hu, Bistra Dilkina, and Nebojsa Jojic. Weakly supervised semantic segmentation in the 2020 ieee grss data fusion contest. In *IEEE International Geoscience and Remote Sensing Symposium*, 2020.
- [30] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 2015.
- [31] Linus Scheibenreif, Joëlle Hanna, Michael Mommert, and Damian Borth. Self-supervised vision transformers for land-cover segmentation and classification. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2022.
- [32] Mennatullah Siam, Mostafa Gamal, Moemen Abdel-Razek, Senthil Yogamani, Martin Jagersand, and Hong Zhang. A comparative study of real-time semantic segmentation for autonomous driving. In *IEEE Conference on Computer Vision and Pattern Recognition Workshop*, 2018.
- [33] Weilian Song, Scott Workman, Armin Hadzic, Xu Zhang, Eric Green, Mei Chen, Reginald Souleyrette, and Nathan Jacobs. Farsa: Fully automated roadway safety assessment. In *IEEE Winter Conference on Applications of Computer Vision*, 2018.
- [34] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Deep image prior. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- [35] Ava Vali, Sara Comai, and Matteo Matteucci. Deep learning for land use and land cover classification based on hyperspectral and multispectral earth observation data: A review. *Remote Sensing*, 12(15):2495, 2020.
- [36] Sherrie Wang, William Chen, Sang Michael Xie, George Azari, and David B Lobell. Weakly supervised deep learning for segmentation of remote sensing imagery. *Remote Sensing*, 12(2):207, 2020.
- [37] Scott Workman and Hunter Blanton. Augmenting depth estimation with geospatial context. In *IEEE International Conference on Computer Vision*, 2021.
- [38] Scott Workman and Nathan Jacobs. Dynamic traffic modeling from overhead imagery. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2020.
- [39] Scott Workman, M Usman Rafique, Hunter Blanton, Connor Greenwell, and Nathan Jacobs. Single image cloud detection via multi-image fusion. In *IEEE International Geoscience and Remote Sensing Symposium*, 2020.
- [40] Scott Workman, M. Usman Rafique, Hunter Blanton, and Nathan Jacobs. Revisiting near/remote sensing with geospatial attention. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2022.
- [41] Scott Workman, Menghua Zhai, David J. Crandall, and Nathan Jacobs. A unified model for near and remote sensing. In *IEEE International Conference on Computer Vision*, 2017.
- [42] Limin Yang, Suming Jin, Patrick Danielson, Collin Homer, Leila Gass, Stacie M Bender, Adam Case, Catherine Costello, Jon Dewitz, Joyce Fry, et al. A new generation of the united states national land cover database: Requirements, research priorities, design, and implementation strategies. *ISPRS Journal of Photogrammetry and Remote Sensing*, 146:108–123, 2018.
- [43] Naoto Yokoya, Pedram Ghamisi, Ronny Hänsch, and Michael Schmitt. 2020 ieee grss data fusion contest: Global land cover mapping with weak supervision. *IEEE Geoscience and Remote Sensing Magazine*, 8(1):154–157, 2020.
- [44] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ade20k dataset. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2017.
- [45] Xiao Xiang Zhu, Devis Tuia, Lichao Mou, Gui-Song Xia, Liangpei Zhang, Feng Xu, and Friedrich Fraundorfer. Deep learning in remote sensing: A comprehensive review and list of resources. *IEEE Geoscience and Remote Sensing Magazine*, 5(4):8–36, 2017.
- [46] Aleksandar Zlateski, Ronnachai Jaroensri, Prafull Sharma, and Frédo Durand. On the importance of label quality for semantic segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2018.