

# Monitoring Vegetation From Space at Extremely Fine Resolutions via Coarsely-Supervised Smooth U-Net

Joshua Fan<sup>1\*</sup>, Di Chen<sup>1</sup>, Jiaming Wen<sup>2</sup>, Ying Sun<sup>2</sup> and Carla P. Gomes<sup>1</sup>

<sup>1</sup>Department of Computer Science, Cornell University

<sup>2</sup>School of Integrative Plant Science – Soil and Crop Sciences Section, Cornell University

jyf6@cornell.edu, di@cs.cornell.edu, jw2495@cornell.edu, ys776@cornell.edu, gomes@cs.cornell.edu

## Abstract

Monitoring vegetation productivity at extremely fine resolutions is valuable for real-world agricultural applications, such as detecting crop stress and providing early warning of food insecurity. Solar-Induced Chlorophyll Fluorescence (SIF) provides a promising way to directly measure plant productivity from space. However, satellite SIF observations are only available at a coarse spatial resolution, making it impossible to monitor how individual crop types or farms are doing. This poses a challenging *coarsely-supervised regression* (or *downscaling*) task; at training time, we only have SIF labels at a coarse resolution (3km), but we want to predict SIF at much finer spatial resolutions (e.g. 30m, a 100x increase). We also have additional fine-resolution input features, but the relationship between these features and SIF is unknown. To address this, we propose *Coarsely-Supervised Smooth U-Net (CS-SUNet)*, a novel method for this coarse supervision setting. CS-SUNet combines the expressive power of deep convolutional networks with novel regularization methods based on prior knowledge (such as a smoothness loss) that are crucial for preventing overfitting. Experiments show that CS-SUNet resolves fine-grained variations in SIF more accurately than existing methods.

## 1 Introduction

Monitoring crop growth is vital for preventing food insecurity and supporting humanitarian efforts to address the UN Sustainable Development Goal 2 (“Zero Hunger”) [Peng *et al.*, 2020].<sup>1</sup> This is particularly crucial due to the impacts of climate change on agriculture [Ortiz-Bobea *et al.*, 2021]. Thus, there have been many attempts to observe vegetation progress from space via remote sensing. A common approach is to use vegetation indices that can be computed from satellite imagery, such as NDVI. However, these indices

are only indicative of vegetation “greenness”, and cannot directly measure photosynthetic activity [Li and Xiao, 2019]. On the other hand, Solar-Induced Chlorophyll Fluorescence (SIF) holds the potential to directly monitor plant productivity, as it has mechanistic linkages to plant photosynthesis, and provides more accurate and timely information about plant growth [Kira and Sun, 2020]. Satellite-based Solar-Induced Chlorophyll Fluorescence (SIF) can outperform traditional vegetation indices for crop yield prediction [Peng *et al.*, 2020], as well as for monitoring the effects of drought and heat stress on vegetation productivity [Song *et al.*, 2020].

However, despite recent progress, SIF continues to be a very difficult signal to measure from space [Duveiller *et al.*, 2020]. Most satellite SIF measurements are noisy, and are thus only available at a coarse spatial resolution, such as 3-5km [Yu *et al.*, 2019]. Within a large region of many square kilometers, SIF can vary dramatically, depending on crop type, management practices, and genetic varieties. For example, the soybean fields in a particular region may be growing very well, while the corn fields may be experiencing stress. Yet existing satellite measurements cannot resolve this important variation [Kira and Sun, 2020].

There have been several attempts to predict SIF at finer spatial resolutions, given coarse-resolution SIF measurements and additional fine-resolution variables [Yu *et al.*, 2019; Duveiller *et al.*, 2020; Wen *et al.*, 2020]. This task is sometimes known as “statistical downscaling.” These works train a supervised machine learning model at a coarse scale, to map from auxiliary input variables (averaged across an entire coarse tile) to the SIF of that tile. Then this trained model is applied at a fine scale, using the fine-resolution input variables to predict fine-resolution SIF. Since the input variables are averaged across a large region during training, information about fine-grained spatial variation is lost. Thus, these works only attempt to resolve SIF down to a 0.05-degree ( $\approx 5$  km) resolution. They do not attempt to resolve SIF at super-fine resolutions such as 30 meters; our results show that these methods are inaccurate at such fine resolutions.

This problem poses a challenging *coarsely-supervised regression* task. We want to learn a model that maps from fine-resolution input features to fine-resolution SIF, yet the only SIF labels available at training time are at a much coarser resolution. We note that this task (inferring a fine-resolution map of a variable, given only coarse-resolution observations and

\*Contact Author

<sup>1</sup>Appendices can be found at [http://joshuafan.github.io/files/SIF\\_supp.pdf](http://joshuafan.github.io/files/SIF_supp.pdf)

other fine-resolution auxiliary data) is not exclusive to SIF, but is applicable to many problems in diverse fields, such as improving the resolution of climate model outputs [Wood *et al.*, 2004], soil moisture [Peng *et al.*, 2017], and malaria risk [Sturrock *et al.*, 2014].

To address this issue, we propose a novel technique, *Coarsely-Supervised Smooth U-Net (CS-SUNet)*, which can produce SIF predictions at a 30-meter resolution, even though it is only trained with SIF labels at a 3km resolution (100x coarser). CS-SUNet takes in a fine-resolution input image, and outputs a SIF prediction for each pixel; the average of the pixel predictions is trained to match the true coarse SIF for the entire 3km image. The U-Net architecture is good at producing localized predictions, as the prediction for a pixel primarily depends on nearby pixels' features [Wang *et al.*, 2020]. We further enhance the architecture to make the localization more precise and reduce instability. Still, the coarse-resolution SIF labels are an extremely weak form of supervision, and do not provide enough constraints for the pixel predictions. However, we show that CS-SUNet can still learn to predict accurately through regularization techniques inspired by prior knowledge.

**Our contributions.** (1) We propose CS-SUNet, a novel approach to coarsely-supervised regression problems, which builds on the expressive power and localization abilities of U-Net. To our knowledge, this type of approach has never been applied to resolving fine-resolution SIF or similar “coarsely-supervised regression” problems with spatial structure. (2) We find that naively training a U-Net can result in a unique form of overfitting (by outputting extreme pixel predictions that happen to average to the right value). Thus, to regularize the model, we design a smoothness loss that encourages pixels with similar features to map to similar SIF. We also use early stopping based on a small fine-resolution validation set. (3) We evaluate our method on real-world remote sensing and SIF data. CS-SUNet predicts SIF at fine resolutions more accurately than all baselines, reducing RMSE by almost 10%.

## 2 Related Work

**SIF Downscaling: Mechanistic Approaches.** There have been several attempts to downscale coarse-resolution SIF measurements. One line of work uses a mechanistic Light Use Efficiency modeling approach to improve the spatial resolution of the GOME-2 satellite's SIF measurements from 0.5-degree to 0.05-degree [Duveiller *et al.*, 2020]. Hand-crafted theoretical models are fitted to establish semi-empirical relationships between SIF and explanatory variables (vegetation indices, evapotranspiration, and temperature). These methods impose strong assumptions on the nature of these relationships, and assume that they hold across all land cover types.

**SIF Downscaling: Data-Driven Approaches.** There have also been attempts to use completely data-driven *statistical downscaling* methods. These studies use standard machine learning algorithms, such as Artificial Neural Network [Yu *et al.*, 2019] or regression trees [Li and Xiao, 2019], to learn an empirical relationship between averaged coarse-resolution input features (e.g. reflectance) and coarse-resolution SIF. This

relationship is then applied on fine-resolution reflectance to estimate fine-resolution SIF. These approaches work well at a relatively coarse resolution (such as downscaling from 0.5 to 0.05 degrees), but become less accurate at finer resolutions, since coarse-resolution SIF labels used in training typically represent the SIF of a large region with multiple land cover types. However, we are interested in predicting SIF for individual fields of a single crop type. Because reflectance data for the entire large tile is averaged together, the model cannot observe variation between crop types within a tile during training, or learn crop-specific patterns. Recently, [Kira and Sun, 2020] addressed this “SIF unmixing” problem by estimating crop-specific SIF for corn and soybean portions of a larger region. They produce a training dataset of artificial “mixed pixels”, which were generated by mixing pure-corn and pure-soy tiles together in random proportions. Then they train a neural network to estimate how much of the SIF was from the corn areas vs. the soybean areas. However, this approach requires that pure tiles for each crop type exist, which is not the case for most crops other than corn and soybean.

**Statistical Downscaling in Other Fields.** The “statistical downscaling” approach has also been applied to numerous problems in geoscience and other fields, including downscaling soil moisture [Peng *et al.*, 2017] and other remotely sensed variables. As previously mentioned, these methods discard information about spatial variability within the large regions while training. Bias-corrected spatial disaggregation [Wood *et al.*, 2004] is another method to downscale projections from climate models, but it requires extensive ground-truth fine-resolution labels. Hierarchical Bayesian models have also been used for downscaling, such as for species distributions and malaria risk [Sturrock *et al.*, 2014]. However, they model fine-resolution probabilities as a simple linear function of the inputs, which limits the expressiveness of the learned functions. Some works have proposed methods based on Gaussian Processes [Tanaka *et al.*, 2019], but they are computationally intractable on large datasets (such as our setting where we want to predict for millions of pixels).

**Multiple-Instance Learning.** In the machine learning literature, there are some prior works on multiple-instance learning, where each training label is over a group of examples. While most of these works focus on classification, a few papers consider regression tasks (*aggregate output regression*). [Musicant *et al.*, 2007] presented adaptations of supervised learning algorithms to this setting. [Kotzias *et al.*, 2015] studied this in the context of fine-grained sentiment analysis of online reviews, and proposed adding a smoothness loss term, which is similar to the idea in our paper. [Zhang *et al.*, 2020] proposed probabilistic models for this setting, but the method they derive for regression tasks is just a simple L2 loss between the predicted and true aggregated outputs, similar to the coarse loss term in our approach. Unlike these works, our approach utilizes the spatial structure in our data via U-Net.

**Supervised Single-Image Super-Resolution.** Many papers have studied *single-image super-resolution* using convolutional neural networks (CNNs), where the goal is to learn a function mapping directly from a coarse-resolution image to the fine-resolution version. For example, [Liu *et al.*, 2020]

use a novel fusion CNN architecture to downscale climate model variables to finer resolution. However, these methods require abundant fine-resolution data for training, which is not available for SIF. By contrast, CS-SUNet does not use fine-resolution labels at training time, so it is usable in settings where fine-resolution labels are scarce.

**Image Segmentation.** CS-SUNet is inspired by techniques originally developed for image segmentation, specifically the U-Net architecture [Ronneberger *et al.*, 2015]. The U-Net produces a prediction for each pixel in the original image, based on the contents of that pixel and neighboring pixels. The U-Net is typically used for classification, but CS-SUNet is adapted for regression.

**Weakly-Supervised Pixel Predictions.** While image segmentation tasks are usually trained using pixel-level labels, some works use weaker forms of supervision (e.g. image-level labels). Recently, [Wang *et al.*, 2020] showed that the U-Net can accurately segment remote sensing imagery (e.g. classifying each pixel as cropland or not), even if it is only trained with image-level labels. This is largely due to the U-Net’s impressive localization abilities, which we harness in our approach. [Wang and Yang, 2018] show that CNNs can provide localization abilities for regression tasks also. They train a CNN to predict diabetic retinopathy from a retina image, and propose a method, Regression Activation Mapping, to determine which regions of the image have the greatest influence on the final prediction. By contrast, our paper explicitly predicts the label (SIF) for every pixel in the image.

### 3 Datasets

We summarize the datasets used in this paper. More details (including citations) can be found in Appendices E and F.

#### 3.1 Input Features

**Landsat Surface Reflectance.** Satellite imagery is available from Landsat at a fine 30m resolution, every 16 days. For each pixel, there are 7 bands (channels). We zero out pixels with missing data and create a mask to mark this.

**FLDAS land data.** Temperature, rainfall, and surface radiation data are available from FLDAS (Famine Early Warning Systems Network (FEWS NET) Land Data Assimilation System), at a 10km resolution. We linearly interpolate these variables to a 30m resolution to match the other variables.

**Land cover data.** Land cover data is available from the Cropland Data Layer at 30m resolution. Each pixel is labelled with a land cover type. For each land cover type that makes up more than 1% of the dataset, we create a binary mask of which pixels are covered by that type.

#### 3.2 SIF Labels

To evaluate model performance, we need some ground-truth fine-resolution SIF data. The Chlorophyll Fluorescence Imaging Spectrometer (CFIS) provides SIF observations at fine resolutions for extremely limited areas. We aggregate the CFIS observations to a 30m resolution, matching the Landsat grid. Then, we produce datasets at different resolutions: {30, 90, 150, 300, 600} meters for evaluation, and 3km for

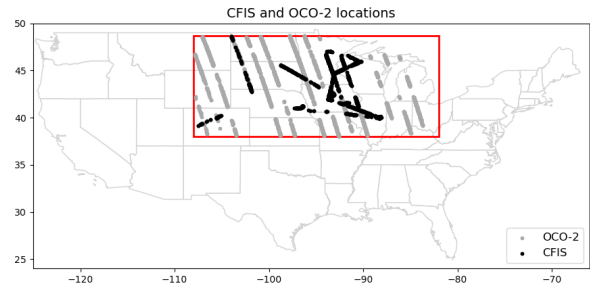


Figure 1: Locations of CFIS (black) and OCO-2 (dark gray) tiles in our study region (red box).

training (to simulate the realistic setting in which only coarse-resolution SIF observations are available).

OCO-2 [Sun *et al.*, 2017] provides additional SIF labels at a coarse spatial resolution, which we can also use for coarse-resolution training. The original footprints cover areas of around  $1.3 \times 2.25$  km at nadir. We grid these to a 3km resolution (on the same grid as our CFIS 3km dataset). To correct between the differing wavelengths of the different instruments, we multiply OCO-2 SIF values by 1.11 to match CFIS. (This scaling factor was empirically determined by fitting a model on OCO-2, evaluating it on CFIS, and determining what scaling factor produced the optimal fit.)

Since SIF observations are noisy, we filter our dataset to only include pixels and tiles with enough observations (so that noise can be reduced through averaging). We also remove tiles with many missing features. See Appendix G for details.

#### 3.3 Study Region and Time Range

All tiles in our dataset correspond to a coarse-resolution CFIS or OCO-2 SIF label that lies within the Midwest US, from 38 to 48.7 degrees N and 108 to 82 degrees W. The region is plotted in Figure 1. As shown in the figure, even coarse-resolution SIF datasets are sparse geographically, so machine learning is needed to fill in the gaps and predict at a higher resolution. We extract input feature data (reflectance, crop cover maps, etc.) from the same time periods as the CFIS SIF observations: June 15-29, 2016 and August 1-16, 2016. We also use OCO-2 observations from those time periods.

#### 3.4 Data Preprocessing

We use a total of 22 input features: 7 Landsat reflectance bands (plus 1 binary mask indicating missing Landsat data), 3 FLDAS variables (temperature, rainfall, surface radiation), and 11 land cover binary masks. All features were resampled to a fine resolution of 30 meters. For each coarse (3km) SIF label, we extracted the corresponding tile of feature data. These tiles are 3-dimensional tensors, of size 22 (features)  $\times$  100  $\times$  100 (pixels). For the continuous variables, we standardized them to have mean 0 and variance 1, and clipped them to  $[-3, 3]$  standard deviations from the mean.

We randomly assign each tile to one of 3 sets: train (60%), validation (20%), and test (20%); we ensure that CFIS and OCO-2 tiles that overlap end up in the same set. The model is trained on coarse-resolution (tile-level) SIF labels from the train tiles; there are a total of 712 CFIS and 1390 OCO-2 tiles

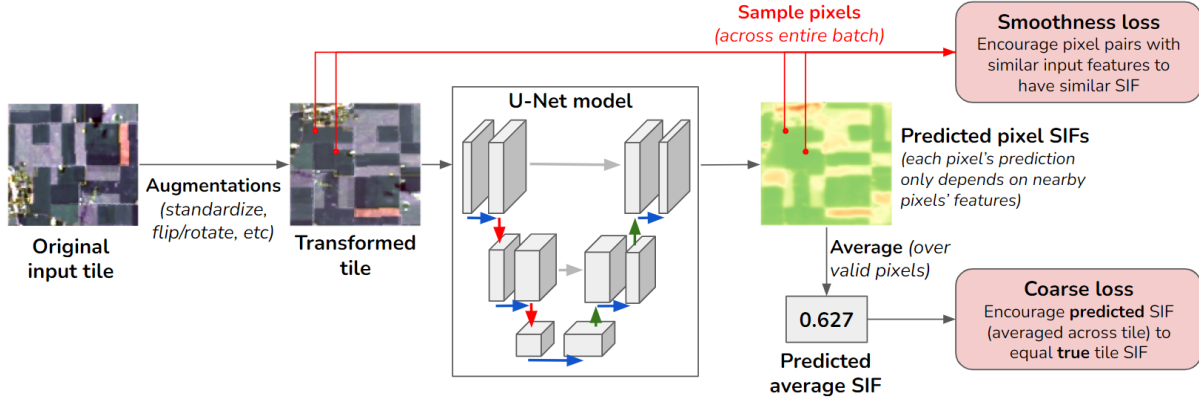


Figure 2: Coarsely-Supervised Smooth U-Net (CS-SUNet). At each iteration, we first apply some augmentations to the original input tile (satellite images, etc). Then we pass it through the U-Net to obtain pixel SIF predictions. Finally, we compute the coarsely-supervised loss and the smoothness regularization loss, and backpropagate to improve the model.

in the train set. Then, we select hyperparameters based on NRMSE on the *fine-resolution CFIS validation set*.

## 4 Methods

We propose CS-SUNet, a novel approach for coarsely-supervised regression tasks. CS-SUNet utilizes the localization abilities of the U-Net architecture, adapts it to our coarsely-supervised regression setting, and employs regularization techniques informed by prior knowledge, such as a smoothness loss and early stopping.

### 4.1 U-Net Architecture

CS-SUNet makes use of the U-Net architecture [Ronneberger *et al.*, 2015], which was originally designed for segmenting biomedical images using relatively little training data. The U-Net is a fully convolutional neural network which can take in input images of any size. It consists of a contracting path which aggregates information from local regions (left side), and an expanding path which combines broader large-area information with precise local characteristics (right side). The final output is a map, with a prediction for each pixel.

In our implementation, we used 2 upsampling and 2 downsampling blocks, with  $\{64, 128\}$  hidden units. For half of the convolutional layers, we decreased the filter size from 3 to 1 to ensure that each pixel’s prediction only depends on nearby pixels’ features. The output is a single scalar per pixel. We removed the batch normalization layers as they caused training to be unstable; they do not appear to be a good fit for regression tasks like ours that depend on the absolute magnitude of the inputs (see Appendix G). Further model details can be found in Appendix J.

### 4.2 Coarsely Supervised Training

CS-SUNet is trained in a coarsely-supervised fashion. Our training set consists of pairs  $\{(\mathbf{X}_i, y_i)\}_{i=1}^N$ :  $\mathbf{X}_i \in \mathbb{R}^{C \times H \times W}$  is a large tile “image” with  $C$  channels, height  $H$ , and width  $W$ , and  $y_i \in \mathbb{R}$  is the ground-truth average (coarse) SIF for tile  $i$ .

We train a U-Net model  $f_\theta$  that takes in large tile  $\mathbf{X}_i$  and predicts SIF for each pixel,  $\mathbf{Z}_i \in \mathbb{R}^{H \times W}$ . Let  $z_i^{(p)} \in \mathbb{R}$  be the predicted SIF of pixel  $p$ . During training, we do not have pixel-level SIF labels. Instead, we train the model by optimizing a combination of two loss functions: (1) a coarse loss that encourages the average predicted SIF of the entire tile to equal the true SIF of the tile, and (2) a smoothness loss that encourages pixel pairs with similar input features to map to similar SIF predictions.

The coarse loss is defined as the Mean Squared Error between the true and predicted coarse tile SIFs; this is similar to [Zhang *et al.*, 2020]. We average predictions over  $P$ , the set of valid fine pixels that contributed to the tile’s coarse SIF:

$$\mathcal{L}_{coarse}(y_i, \mathbf{Z}_i) = \left( y_i - \frac{1}{|P|} \sum_{p \in P} z_i^{(p)} \right)^2$$

However, achieving a low coarse loss does not guarantee that the fine-resolution predictions will be accurate. It is possible for the U-Net to output extreme pixel predictions that still produce the correct coarse-tile average. We find that this occurs frequently in practice (see Appendix B). To reduce this unique form of overfitting, we introduce a smoothness loss [Kotzias *et al.*, 2015], which encourages pixels with similar input features to map to similar SIFs. To compute this loss, for each batch we randomly sample a set of pixels  $S$ ; in our experiments we sample 1000 pixels per batch. For each pair of pixels in  $S$ , we compute a similarity score between their input feature vectors ( $\mathbf{x}^{(j)}$ ,  $\mathbf{x}^{(k)}$ ), and multiply that by the squared difference between their predicted SIFs ( $z^{(j)}$ ,  $z^{(k)}$ ):

$$\mathcal{L}_{smooth}(\mathbf{X}, \mathbf{Z}) = \sum_{j \in S} \sum_{k \in S} \text{sim}(\mathbf{x}^{(j)}, \mathbf{x}^{(k)}) \cdot (z^{(j)} - z^{(k)})^2$$

This places a large penalty if two pixels with high input similarity have very different SIF, as this represents overfitting. The similarity function can be customized based on one’s prior knowledge of the problem. We set the similarity to 0 if pixels  $j$  and  $k$  are of different land cover types. Otherwise,

if they are of the same land cover type, we define their input similarity as

$$\text{sim}(\mathbf{x}^{(j)}, \mathbf{x}^{(k)}) = \exp\left(-\tau \|\mathbf{x}_{ref}^{(j)} - \mathbf{x}_{ref}^{(k)}\|_2^2\right)$$

Here  $\mathbf{x}_{ref}^{(j)}$  is the reflectance feature vector for pixel  $j$ . Note that if  $\tau$  is positive, the similarities are bounded between 0 and 1.  $\tau$  is a hyperparameter that controls how quickly the similarities decay towards 0; we set this to 0.5. Appendix C discusses the impact of the smoothness loss, and Appendix D motivates the design of this loss.

The final loss is a weighted combination of the coarse loss and the smoothness loss. For a batch,

$$\mathcal{L}(\mathbf{X}, \mathbf{Z}, \mathbf{y}) = \frac{1}{B} \left[ \sum_{i=1}^B \mathcal{L}_{coarse}(y_i, \mathbf{Z}_i) \right] + \lambda \mathcal{L}_{smooth}(\mathbf{X}, \mathbf{Z})$$

where  $\lambda$  is a hyperparameter that balances the loss terms. We optimize this loss via backpropagation. Figure 2 depicts the training process.

### 4.3 Early Stopping

CS-SUNet does not look at any fine-resolution SIF data during training. However, during validation, CS-SUNet does peek at the model’s performance on the *fine-resolution validation set* to select which epoch’s model to use. Like the smoothness loss, early stopping encourages “input smoothness” (pixels with similar input features should map to similar SIF) [Rosca *et al.*, 2020], which is crucial for regularizing our model. The fine-resolution data is only used to determine how much regularization is necessary, so it may not need to be very similar to the evaluation setting. If we do not have any fine-resolution data at all, it is difficult to determine the optimal amount of regularization, but our smoothness loss can still prevent overfitting on later epochs (see Appendix C).

## 5 Experiments

We evaluate our pixel-level SIF predictions against airborne measurements from CFIS, which are available at a 30-meter resolution in limited regions. To reduce measurement noise, we only evaluate on pixels for which CFIS had at least 30 observations and where the SIF value was at least 0.1. We use Normalized Root Mean Squared Error (NRMSE) and  $R^2$  as evaluation metrics. (NRMSE is Root Mean Squared Error, normalized by the average SIF of the training dataset.)

### 5.1 Training Details

The CS-SUNet approach is compared against existing “statistical downscaling” baselines that have been used in papers such as [Yu *et al.*, 2019; Li and Xiao, 2019]. These involve averaging each input feature (channel) over all valid pixels in the tile, and training a model to predict SIF from the feature averages. We consider 4 types of models: Ridge Regression, Gradient Boosting Regressor, Random Forest, and a fully-connected artificial neural network (ANN).

We also include a few novel baselines. The “Pixel NN” approach involves feeding each pixel through the same neural network, averaging all pixels’ predictions across a tile, and

Method	NRMSE ( $\downarrow$ )	$R^2$ ( $\uparrow$ )
<i>Predict coarse</i>	0.248	0.372
Ridge Regression	0.203	0.582
Gradient Boosting	0.212	0.541
Random Forest	0.220 $\pm$ 0.002	0.507 $\pm$ 0.008
ANN	0.215 $\pm$ 0.005	0.528 $\pm$ 0.020
Pixel NN	0.208 $\pm$ 0.004	0.561 $\pm$ 0.016
Vanilla U-Net	0.223 $\pm$ 0.005	0.492 $\pm$ 0.021
CS-SUNet	<b>0.184 <math>\pm</math> 0.005</b>	<b>0.656 <math>\pm</math> 0.019</b>

Table 1: Results on interpolating CFIS to 30m resolution, on train tiles where only *coarse-resolution* SIF was known. The “predict coarse” row is predicting the coarse SIF for each pixel. “ANN” is artificial neural network. For methods that involve randomness, we report the average  $\pm$  the standard deviation over 3 runs.

Method	NRMSE ( $\downarrow$ )	$R^2$ ( $\uparrow$ )
Ridge Regression	0.205	0.361
Gradient Boosting	0.216	0.289
Random Forest	0.229 $\pm$ 0.001	0.199 $\pm$ 0.006
ANN	0.215 $\pm$ 0.002	0.300 $\pm$ 0.016
Pixel NN	0.207 $\pm$ 0.002	0.348 $\pm$ 0.013
Vanilla U-Net	0.241 $\pm$ 0.011	0.113 $\pm$ 0.081
CS-SUNet	<b>0.185 <math>\pm</math> 0.002</b>	<b>0.481 <math>\pm</math> 0.013</b>
<i>U-Net fine supervision</i>	0.161 $\pm$ 0.003	0.605 $\pm$ 0.014

Table 2: Results on extrapolating CFIS to 30m resolution, on test tiles where even coarse-resolution SIF is unknown.

then training using coarse loss (tile-level supervision); this substitutes the U-Net architecture with an per-pixel architecture where each pixel’s prediction is only influenced by that pixel’s features. The “vanilla U-Net” approach involves training a U-Net with only the coarse loss, but without regularization (early stopping or smoothness loss). Finally, “U-Net fine supervision” trains a U-Net using fine-resolution labels, and represents an upper limit on our model performance if we were actually able to train on fine-resolution labels.

We train all deep models with the AdamW optimizer [Loshchilov and Hutter, 2018]. For each method, a grid-search was performed over hyperparameter values (see Appendix I for details); we selected the hyperparameters that performed best on the *fine-resolution* validation set. For CS-SUNet, we used a batch size of 128, learning rate  $3e-4$ , weight decay  $1e-4$ ,  $\tau = 0.5$ , and  $\lambda = 0.5$ . For methods that involve randomness, we took the average of 3 runs using different random seeds. For tiles where the coarse-resolution SIF label is observed during training, we also include a trivial “predict coarse” baseline where we predict the SIF of every pixel to be the same as the SIF of the entire tile.

### 5.2 Results

Table 1 presents the results for fine pixels in train tiles (i.e. the coarse-resolution SIF of the entire tile was known during training, but not the fine-resolution SIF of the individual pixels). The model is tasked with interpolating the fine-resolution SIF from the coarse-resolution labels. Since we

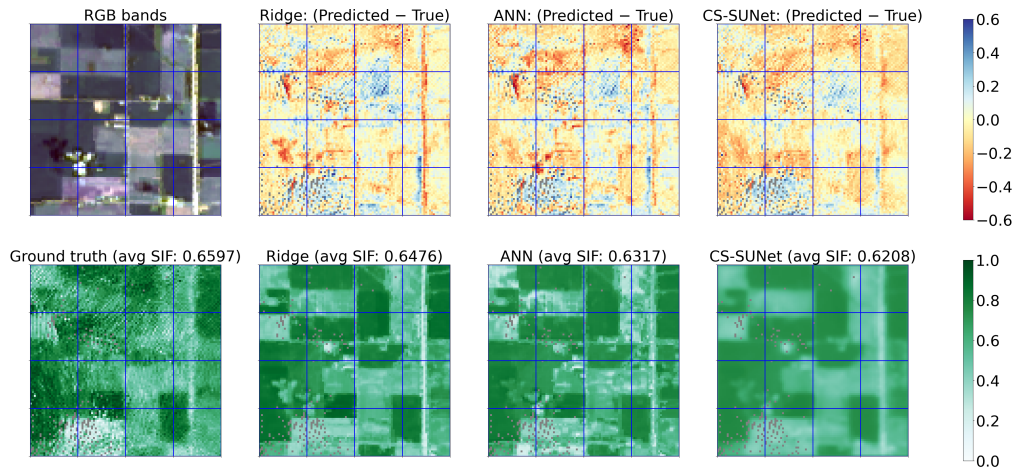


Figure 3: Example result. **Top-left:** RGB bands, used as input features. **Bottom-left:** ground-truth CFIS SIF (gray pixels are missing data). **Other bottom images:** predictions with different methods (CS-SUNet is rightmost). **Other top images:** error maps (yellow is accurate, blue is over-prediction, red is under-prediction). CS-SUNet is generally more accurate than other methods – for example, there are fewer dark red or blue pixels (which indicate major errors).

know the coarse-resolution SIF of these tiles, we multiply each tile’s pixel predictions by a per-tile constant, such that the average of the pixel predictions exactly equals the true coarse SIF. Table 2 presents results for fine pixels where not even coarse-resolution SIF was seen during training; it tests the model’s ability to generalize to “gaps” where SIF observations were not directly available. (Note that the relatively lower  $R^2$  in Table 2 is because there is less total variation in these tiles; the absolute error magnitudes (NRMSE) are similar to Table 1.)

According to both metrics, CS-SUNet outperforms standard “statistical downscaling” methods in both settings, reducing NRMSE by 9-10% and increasing  $R^2$  by 13-33% (relative) over the best baseline, Ridge Regression. This improvement occurs across multiple land cover types and resolutions (see Appendix A). CS-SUNet outperforms the “Pixel NN” architecture, which shows that the prediction accuracy for a pixel can be improved if you look at neighboring pixels’ features, not just that pixel’s features. Notably, the “Vanilla U-Net” actually does poorly – simply applying a U-Net out of the box does not work, due to the “overfitting to coarse labels” problem described earlier. In order for the U-Net to work, it needs to be highly regularized, which we achieve in this paper by early stopping and a novel smoothness loss.

While an NRMSE of 18-19% is not perfect, note that even a model trained with fine-resolution labels can only achieve an NRMSE of 16% on the test tiles, so our performance is not far off. In any case, increasing the resolution of SIF measurements by 100x is an extremely challenging task, and our method provides significant improvements.

**Qualitative evaluation.** Figure 3 presents an example of the predictions outputted by various measurements. All methods provide useful results; even the basic Ridge Regression method is able to identify fields with higher and lower SIF. However, the error maps on the top row show that CS-SUNet produces more reliable predictions, as there are fewer dark

blue or red pixels (which indicate large errors). In addition, CS-SUNet takes the spatial context of pixels into consideration, so the predictions outputted by CS-SUNet are much more smooth and consistent across local regions.

**Limitations.** The evaluation of our models is limited by the fact that there is significant noise in both SIF measurements and Landsat reflectance (and there may be a few days’ gap between them). Our approach for SIF prediction requires Landsat reflectance data, yet Landsat data is unavailable or noisy in some places/times due to cloud cover. This could be alleviated in the future as more satellites are deployed.

## 6 Conclusion

We present Coarsely-Supervised Smooth U-Net (CS-SUNet), which is capable of predicting SIF at a very fine resolution (30m), even when only coarse-resolution (3km) SIF measurements are available. Due to its localization properties, CS-SUNet is able to figure out which farms within a large tile had higher and lower SIF. CS-SUNet harnesses the expressive power of deep convolutional networks, yet avoids overfitting due to its smoothness loss and early stopping.

Although CS-SUNet’s predictions are not perfect, they clearly outperform existing state-of-the-art methods, and can provide valuable vegetation information. Even noisy fine-resolution SIF estimates may facilitate improvements in crop yield prediction and monitoring, as SIF contains detailed information about plant photosynthesis that is not captured in vegetation indices such as NDVI [Peng *et al.*, 2020], which are simple combinations of a few spectral bands. Moreover, in addition to SIF, CS-SUNet could also potentially be applied to make fine-resolution predictions of any numerical variable that is only available at coarse resolution, if there exists additional fine-resolution data that is correlated with the variable of interest. Such applications could include predicting soil moisture, precipitation, disease prevalence, and species distributions at fine resolutions.

## Acknowledgements

This research was supported by NSF award CCF-1522054 (Expeditions in Computing) and the NSF NRT Digital Plant Science Fellowship (1922551). We thank Junwen Bai for proofreading and providing valuable suggestions.

## References

- [Duveiller *et al.*, 2020] G. Duveiller, F. Filippini, S. Walther, P. Köhler, C. Frankenberg, L. Guanter, and A. Cescatti. A spatially downscaled sun-induced fluorescence global product for enhanced monitoring of vegetation productivity. *Earth System Science Data*, 12(2):1101–1116, 2020.
- [Kira and Sun, 2020] O. Kira and Y. Sun. Extraction of sub-pixel c3/c4 emissions of solar-induced chlorophyll fluorescence (sif) using artificial neural network. *ISPRS Journal of Photogrammetry and Remote Sensing*, 161:135–146, 2020.
- [Kotzias *et al.*, 2015] D. Kotzias, M. Denil, N. De Freitas, and P. Smyth. From group to individual labels using deep features. In *Proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining*, pages 597–606, 2015.
- [Li and Xiao, 2019] X. Li and J. Xiao. A global, 0.05-degree product of solar-induced chlorophyll fluorescence derived from oco-2, modis, and reanalysis data. *Remote Sensing*, 11(5):517, 2019.
- [Liu *et al.*, 2020] Y. Liu, A. R. Ganguly, and J. Dy. Climate downscaling using ynet: A deep convolutional network with skip connections and fusion. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 3145–3153, 2020.
- [Loshchilov and Hutter, 2018] I. Loshchilov and F. Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2018.
- [Musicant *et al.*, 2007] D. R. Musicant, J. M. Christensen, and J. F. Olson. Supervised learning by training on aggregate outputs. In *Seventh IEEE International Conference on Data Mining (ICDM 2007)*, pages 252–261. IEEE, 2007.
- [Ortiz-Bobea *et al.*, 2021] A. Ortiz-Bobea, T. R. Ault, C. M. Carrillo, R. G. Chambers, and D. B. Lobell. Anthropogenic climate change has slowed global agricultural productivity growth. *Nature Climate Change*, 11(4):306–312, 2021.
- [Peng *et al.*, 2017] J. Peng, A. Loew, O. Merlin, and N. E. C. Verhoest. A review of spatial downscaling of satellite remotely sensed soil moisture. *Reviews of Geophysics*, 55(2):341–366, 2017.
- [Peng *et al.*, 2020] B. Peng, K. Guan, W. Zhou, C. Jiang, C. Frankenberg, Y. Sun, L. He, and P. Köhler. Assessing the benefit of satellite-based solar-induced chlorophyll fluorescence in crop yield prediction. *International Journal of Applied Earth Observation and Geoinformation*, 90:102126, 2020.
- [Ronneberger *et al.*, 2015] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [Rosca *et al.*, 2020] M. Rosca, T. Weber, A. Gretton, and S. Mohamed. A case for new neural networks smoothness constraints. In *“I Can’t Believe It’s Not Better!” NeurIPS 2020 workshop*, 2020.
- [Song *et al.*, 2020] Y. Song, J. Wang, and L. Wang. Satellite solar-induced chlorophyll fluorescence reveals heat stress impacts on wheat yield in india. *Remote Sensing*, 12(20):3277, 2020.
- [Sturrock *et al.*, 2014] H. J. W. Sturrock, J. M. Cohen, P. Keil, A. J. Tatem, A. Le Menach, N. E. Ntshalintshali, M. S. Hsiang, and R. D. Gosling. Fine-scale malaria risk mapping from routine aggregated case data. *Malaria journal*, 13(1):1–9, 2014.
- [Sun *et al.*, 2017] Y. Sun, C. Frankenberg, J. D. Wood, D. S. Schimel, M. Jung, L. Guanter, D. T. Drewry, M. Verma, A. Porcar-Castell, T. J. Griffis, et al. Oco-2 advances photosynthesis observation from space via solar-induced chlorophyll fluorescence. *Science*, 358(6360):eaam5747, 2017.
- [Tanaka *et al.*, 2019] Y. Tanaka, T. Iwata, T. Tanaka, T. Kurashima, M. Okawa, and H. Toda. Refining coarse-grained spatial data using auxiliary spatial data sets with various granularities. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 5091–5099, 2019.
- [Wang and Yang, 2018] Z. Wang and J. Yang. Diabetic retinopathy detection via deep convolutional networks for discriminative localization and visual explanation. In *Workshops at the thirty-second AAAI conference on artificial intelligence*, 2018.
- [Wang *et al.*, 2020] S. Wang, W. Chen, S. M. Xie, G. Azzari, and D. B. Lobell. Weakly supervised deep learning for segmentation of remote sensing imagery. *Remote Sensing*, 12(2):207, 2020.
- [Wen *et al.*, 2020] J. Wen, P. Köhler, G. Duveiller, N. C. Parazoo, T. S. Magney, G. Hooker, L. Yu, C. Y. Chang, and Y. Sun. A framework for harmonizing multiple satellite instruments to generate a long-term global high spatial-resolution solar-induced chlorophyll fluorescence (sif). *Remote Sensing of Environment*, 239:111644, 2020.
- [Wood *et al.*, 2004] A. W. Wood, L. R. Leung, V. Sridhar, and D. P. Lettenmaier. Hydrologic implications of dynamical and statistical approaches to downscaling climate model outputs. *Climatic change*, 62(1):189–216, 2004.
- [Yu *et al.*, 2019] L. Yu, J. Wen, C. Y. Chang, C. Frankenberg, and Y. Sun. High-resolution global contiguous sif of oco-2. *Geophysical Research Letters*, 46(3):1449–1458, 2019.
- [Zhang *et al.*, 2020] Y. Zhang, N. Charoenphakdee, Z. Wu, and M. Sugiyama. Learning from aggregate observations. *Advances in Neural Information Processing Systems*, 33:7993–8005, 2020.