

# S-EO: A Large-Scale Dataset for Geometry-Aware Shadow Detection in Remote Sensing Applications

Elías Masquil<sup>1,2</sup> Roger Martí<sup>3</sup> Thibaud Ehret<sup>4</sup> Enric Meinhardt-Llopis<sup>5</sup>  
Pablo Musé<sup>1,5</sup> Gabriele Facciolo<sup>5</sup>

<sup>1</sup>IIE, Facultad de Ingeniería, Universidad de la República, Uruguay

<sup>2</sup>Digital Sense, Uruguay

<sup>3</sup>Eurecat, Centre Tecnològic de Catalunya, Multimedia Technologies, Barcelona, Spain

<sup>4</sup>AMIAD, Pôle Recherche, France

<sup>5</sup>Université Paris-Saclay, ENS Paris-Saclay, CNRS, Centre Borelli, 91190, Gif-sur-Yvette, France

## Abstract

We introduce the S-EO dataset: a large-scale, high-resolution dataset, designed to advance geometry-aware shadow detection. Collected from diverse public-domain sources, including challenge datasets and government providers such as USGS, our dataset comprises 702 georeferenced tiles across the USA, each covering  $500 \times 500$  m. Each tile includes multi-date, multi-angle WorldView-3 pansharpened RGB images, panchromatic images, and a ground-truth DSM of the area obtained from LiDAR scans. For each image, we provide a shadow mask derived from geometry and sun position, a vegetation mask based on the NDVI index, and a bundle-adjusted RPC model. With approximately 20,000 images, the S-EO dataset establishes a new public resource for shadow detection in remote sensing imagery and its applications to 3D reconstruction. To demonstrate the dataset’s impact, we train and evaluate a shadow detector, showcasing its ability to generalize, even to aerial images. Finally, we extend EO-NeRF — a state-of-the-art NeRF approach for satellite imagery — to leverage our shadow predictions for improved 3D reconstructions.

## 1. Introduction

Shadows play a crucial role in visual perception, providing key insights into depth, contours, textures, and lighting within a 3D scene. Although they can hide scene details by darkening regions where light is obstructed, they also encode valuable information about object shapes [19, 25] and spatial relationships, enabling scene interpretation without requiring additional data sources. Whether the goal is to remove shadows to recover hidden details or to leverage them as geometric cues, accurate shadow detection is an essential




Figure 1. Example multi-date satellite images and shadow masks of an area from the S-EO dataset. Shadows and their variations are critical cues for Earth Observation algorithms focused on geometry estimation and scene understanding.

task for image-based scene understanding.

As shadows are present in nearly all satellite images [30], the development of accurate shadow detection methods is particularly important to create efficient systems dedicated to Earth Observation (EO). As in many other domains, state-of-the-art approaches primarily rely on deep learning techniques [16], which require large, high-quality annotated datasets for effective training. While remote sensing data availability continues to grow, dedicated shadow datasets remain scarce. To our knowledge, only one prior work has publicly released a small dataset [30], with approximately 500 manually annotated images.

To address this limitation, we introduce a large-scale, high-resolution dataset designed to advance research in shadow detection and 3D reconstruction: the S-EO dataset (named after Shadow-aware Earth Observation). Our dataset, derived from the IARPA CORE3D program data [5, 12, 18, 24], consists of extensive multi-date, multi-angle

WorldView-3 imagery at 30 cm resolution, along with georeferenced shadow masks and aligned Digital Surface Models (DSMs) derived from USGS LiDAR scans (Map services and data available from U.S. Geological Survey, National Geospatial Program). Shadow masks, depicted in Figure 1, were automatically generated from the surface models using a shadow-casting algorithm and the sun position. All images and data are delivered in a ready-to-use format: [https://centreborelli.github.io/shadow-  
eo](https://centreborelli.github.io/shadow-eo).

With more than 700 georeferenced and annotated tiles of  $500 \times 500$  m, totaling approximately 20,000 images, and the associated DSMs, the S-EO dataset provides a new public resource for both shadow detection and shadow-aware 3D modeling from remote sensing data. The S-EO dataset includes a wide variety of samples, covering diverse cities with different urban layouts, vegetation types and climate, observed at different seasons and times of the day. For instance, it captures snow-covered landscapes, which introduce additional challenges for shadow detection models.

To demonstrate the relevance of this novel dataset, we train a shadow detection model and show that it generalizes effectively, even to aerial datasets. Our experiments find that the U-Net architecture [37] remains a strong baseline, performing on par with specialized shadow detection models. Furthermore, we extend the state-of-the-art multi-view 3D modeling framework EO-NeRF [34] to incorporate shadow mask supervision using the detector predictions. Our results show that shadow supervision consistently improves 3D reconstruction performance, boosting both the accuracy of altitude values and the quality of output shapes.

In summary our contributions consist of:

- **A large-scale, high-resolution shadow dataset (S-EO)**, built from public-domain WorldView-3 imagery and USGS LiDAR-derived DSMs.
- **A shadow detector for remote sensing images**, trained on the S-EO data. We evaluate its performance and demonstrate its generalization ability. With minimal fine-tuning, our model surpasses the state-of-the-art in shadow segmentation on previously unseen aerial imagery.
- **A demonstration of shadow-supervised 3D reconstruction from satellite imagery**. We incorporate shadow-based supervision into EO-NeRF and show that it improves 3D reconstruction quality.

Along with the dataset, we release the complete data processing pipeline, our shadow detection model, and our shadow-supervised implementation of EO-NeRF.

## 2. Related work

### 2.1. Deep supervised shadow detection

A comprehensive overview of the state-of-the-art in deep shadow detection is provided by [16]. Here, we highlight

key advances in the field, focusing on both model architectures and datasets.

**Models.** Supervised shadow detection has evolved alongside broader trends in computer vision. Early deep learning approaches [22, 23] replaced hand-crafted features with neural networks, which were later integrated into optimization models such as Conditional Random Fields (CRFs) [48] for refinement. Over time, the focus shifted to end-to-end models, such as GAN-based approaches [36], where the shadow detector operates as the generator of a conditional GAN.

Subsequent advancements introduced multi-scale architectures [14] to better capture fine details, while recent state-of-the-art models leverage popular backbones such as ResNeXt [45], U-Net variants [37], or EfficientNet [41]. Transformers are also being used for shadow detection [20], to efficiently capture contextual relationships with smaller models and achieve faster inference times.

However, as noted in [16], comparisons between existing methods often suffer from inconsistencies in input sizes, evaluation metrics, datasets, and implementation details, making it difficult to determine whether newer architectures truly outperform earlier ones. Since not all models in the literature are publicly available, this work relies on the implementations listed in [16]. The FSDNet [15] is selected as a baseline due to its balance between performance and efficiency. The FSDNet architecture is characterized by Direction-aware Spatial Context (DSC) modules [14] and a MobileNet V2 backbone [40].

**Datasets.** Popular benchmarks for shadow detection typically contain thousands of real-world images from diverse domains. SBU [13], one of the most widely used datasets, provides 4,087 training images and 638 testing images. However, as noted by [46], shadow annotations are often noisy and inconsistent. To address this, a method for self-supervised label refinement is introduced for training alongside expert-verified annotations for the test partition.

ISTD [44] introduces a dataset designed for both shadow detection and removal, providing triplets of shadow images, shadow masks, and shadow-free images. It contains 1,330 training images, 540 testing images, and 135 distinct backgrounds. CUHK-Shadow [15] presents a large-scale dataset of 10,500 shadow images, categorized into different shadow types, including building shadows, Google Maps images, and shadows cast by people and objects.

### 2.2. Shadow detection in remote sensing

AISD [30] was the first publicly available remote sensing dataset for shadow detection. It is derived from the Inria Aerial Image Labeling Dataset [31] and consists of 514 images with manually annotated shadow masks. The image sizes vary between  $256 \times 256$  and  $1688 \times 1688$  pixels. The dataset covers regions in both the United States (Austin,




Figure 2. Top to bottom: S-EO dataset images from three sites— in Jacksonville, Omaha, and San Diego. For each site, left to right: the pansharpened RGB image; the DSM computed from the minimum elevation per grid cell (DSM Min) and its derived shadow mask; and the DSM computed from the maximum elevation per grid cell (DSM Max) with its respective shadow mask. DSM Min yields a cleaner shadow mask by filtering out transient elements (e.g., trees), although it may erode building edges and produce smaller shadows. In contrast, DSM Max preserves more building details and generates larger shadows but with increased noise in the shadow mask.

Chicago) and Austria (Tyrol, Vienna, Innsbruck). It is split into 412 training images, while the validation and test sets each contain 51 images with their corresponding masks.

In the same work, the authors introduce DSSDNet, a deep learning architecture for remote sensing shadow detection. The model comprises two key components: (1) an encoder-decoder residual (EDR) structure, which extracts multi-level and discriminative shadow features, and (2) a deeply supervised progressive fusion (DSPF) module, which enhances detection by progressively refining feature maps during training. Unfortunately, the model weights are not publicly available, and the codebase is implemented in MATLAB, making it difficult to use in modern deep learning frameworks.

In [43], the authors address the domain gap in shadow detection models between ground-level images and aerial imagery, as well as the scarcity of annotated satellite image datasets. They propose a pipeline for automatically generating shadow masks using dense 3D point clouds reconstructed from city-scale Wide Area Motion Imagery (WAMI) sequences and the sun position. While their approach shares similarities with ours, there are some key differences. Unlike our method, which leverages large-scale

ground-truth 3D models and satellite imagery, their approach relies on 3D point clouds reconstructed from aerial imagery. Their contribution is limited to preliminary experiments and does not include a publicly released dataset.

### 2.3. Shadows in advanced 3D modeling

In recent years, multi-view 3D scene modeling has reached unprecedented levels of photorealism, driven by advanced volumetric representations such as NeRF [35] and Gaussian Splatting [21]. The variants of these methods adapted to satellite imagery have emphasized the importance of shadows in accurately interpreting both scene appearance and geometry. Notable examples include Shadow-NeRF [9], Sat-NeRF [33], EO-NeRF [34], SUNDIAL [3], Sat-NGP [4], and EOGS [1], among others. These works highlight the importance of incorporating shadow estimation into the optimization process to produce more accurate geometric reconstructions, as shadows and scene geometry are intrinsically linked. Early models predicted shadows directly using embedding vectors based on the sun position, whereas more recent approaches use additional constraints, such as mathematical modeling of light transmittance for analytical shadow computation [34] and irradiance models

of higher complexity [3].

Shadow masks have been successfully used for geometry recovery in NeRF-based methods for conventional imagery [28, 42]. With the emergence of new datasets and improved shadow detection techniques for remote sensing images, shadow masks could also be integrated into the optimization of these volumetric representations for remote sensing applications.

### 3. Method

#### 3.1. Shadow-aware Earth Observation Dataset

The dataset consists of 702 georeferenced tiles of  $500 \times 500$  m each across the cities of Jacksonville, Omaha, and San Diego [5, 12, 18, 24] as shown in Figure 2. For each tile, we provide multiple:

- Panchromatic (PAN) WorldView-3 images (30 cm)
- Pansharpened RGB images (30 cm);
- Bundle-adjusted RPC models for all PAN and pansharpened images;
- DSMs obtained from LiDAR data from USGS with a 50 cm grid resolution. Two versions: min and max aggregation of altitude values;
- Vegetation masks based on the Normalized Difference Vegetation Index (NDVI);
- Shadow masks. Two versions corresponding to min and max DSMs.

With multiple images per tile, captured on different dates and from various angles, we obtain a total of 19,162 images and masks of approximately  $1500 \times 1500$  pixels. Note that the exact image size varies depending on the viewing angle.

**Shadow mask generation.** Manually annotating shadow masks is a time-consuming task that becomes infeasible at the scale of the dataset presented here. In [30], the authors report having spent four months annotating just 514 images.

We use a shadow simulation algorithm to automatically generate shadow masks at scale from existing data. For each image in the dataset, the shadow mask is computed using the sun position at the time of capture, the aligned DSM for the area of interest, and the image camera models provided by the RPCs. Our algorithm consists of two steps: shadow casting, where shadows are simulated over the DSM based on the sun position and geometry, and shadow projection, where these shadows are projected onto each image using the corresponding RPCs. To make sure that the shadow projection step produces sufficiently dense maps, we generate higher resolution masks by upscaling the DSM by a factor of 4 before shadow casting.

The shadow casting algorithm simulates terrain shadows on the raster DSM by casting rays in the direction of the sun and marking occluded pixels as shadow (Algorithm 1). Discrete ray trajectories are computed for each pixel using a modified Bresenham algorithm [6]. Along each ray, oc-

---

#### Algorithm 1 Shadow casting algorithm

---

**Require:** A DSM, *i.e.* an altitude map of size  $H \times W$ , and the Sun position given by the azimuth  $\alpha$  and elevation  $\beta$  angles.

**Ensure:** A shadow mask  $S_{DSM} \in \{0, 1\}^{H \times W}$ , where 1 indicates a shadow.

```

1:  $p \leftarrow -\sin(\alpha) \times \cos(\beta)$       ▷ x-component of sun dir.
2:  $q \leftarrow \cos(\alpha) \times \cos(\beta)$     ▷ y-component of sun dir.
3:  $a \leftarrow \tan(\beta)$                 ▷ slope of sun rays
4:  $M \leftarrow 0$                       ▷ Initialize all pixels as illuminated
5:  $Paths \leftarrow \text{ComputeBresenhamPaths}(W, H, p, q)$ 
6: for each path  $\Pi$  in  $Paths$  do
7:    $\ell \leftarrow \Pi[0]$               ▷ Index of the first pixel of this path
8:   for  $j = 1 \dots |\Pi| - 1$  do
9:      $(x_\ell, y_\ell) \leftarrow \text{SubPixelCoord}(\ell)$ 
10:     $(x_j, y_j) \leftarrow \text{SubPixelCoord}(\Pi[j])$ 
11:     $d \leftarrow \sqrt{(x_j - x_\ell)^2 + (y_j - y_\ell)^2}$ 
12:     $Z_\ell \leftarrow \text{BilinearInterp}(DSM, x_\ell, y_\ell)$ 
13:     $Z_j \leftarrow \text{BilinearInterp}(DSM, x_j, y_j)$ 
14:     $l = (Z_\ell - Z_j)/a$                 ▷ Equation (1)
15:    if  $d < l$  then
16:       $M[\Pi[j]] \leftarrow 1$           ▷ Current pixel → shadow
17:    else
18:       $\ell \leftarrow \Pi[j]$           ▷ Current pixel → new occluder

```

---

clusion testing is performed using subpixel-accurate sampling of the DSM and a threshold defined by the sun’s elevation  $\beta$ .

Given the first pixel along a ray with elevation  $Z_{occluder}$ , by default this pixel is considered illuminated and defined as the first occluder. For each subsequent pixel along the ray with elevation  $Z_{current}$  and located at a horizontal distance  $d$  from the occluder, the horizontal length of the shadow casted by the occluder is

$$l = \frac{Z_{occluder} - Z_{current}}{\tan \beta}. \quad (1)$$

If the horizontal distance between the pixels is smaller than the shadow length  $d < l$ : the current pixel is marked as a shadow; otherwise the pixel is considered illuminated and becomes the new occluder. The result is a shadow cast map that we denote  $S_{DSM}$ .

After casting shadows on the DSM, the next step is to project them into the image coordinate system (Algorithm 2). The DSM pixel coordinates are localized to obtain the world coordinates, and then re-projected into the image coordinates using the RPC. Because of the projection, several DSM points may project onto the same image pixel. This is resolved with the use of a z-buffer, storing, for each image pixel, a pair  $(Z_{max}, S_{max})$  containing the elevation of the highest projected DSM point  $Z_{max}$ , and its corresponding value in  $S_{DSM}$  *i.e.*,  $S_{max}$ . The shadow value for the image pixel is then given by  $S_{img} = S_{max}$ .

---

**Algorithm 2** Shadow projection from a DSM to an image

---

**Require:** A DSM, a shadow mask  $S_{\text{DSM}}$ , the image size  $(H, W)$ , RPC parameters associated to the image

**Ensure:** A projected shadow mask  $S_{\text{img}}$ , an uncertainty mask  $U$

```
1:  $S_{\text{img}} \leftarrow \mathbf{0}^{H \times W}$            ▷ Initialize with non-shadow
2:  $U \leftarrow \mathbf{1}^{H \times W}$            ▷ Initialize uncertainty mask
3:  $Z_{\text{buffer}} \leftarrow -\infty^{H \times W}$    ▷ Initialize Z-buffer
4: for  $(x, y) \in \text{DSM}$  do
5:    $(X, Y, Z) \leftarrow \text{WorldCoords}(x, y)$ 
6:    $I \leftarrow \text{RPCProjection}(X, Y, Z)$ 
7:   if  $I$  is within image bounds and  $Z > Z_{\text{buffer}}(I)$  then
8:      $Z_{\text{buffer}}(I) \leftarrow Z$            ▷ Update Z-buffer
9:      $S_{\text{img}}(I) \leftarrow S_{\text{DSM}}(x, y)$    ▷ Update shadow
10:     $U(I) \leftarrow 0$                    ▷ Mark as certain
11:  $S_{\text{img}} \leftarrow \text{RemoveSmallRegions}(S_{\text{img}})$ 
```

---

During the projection step, a small number of image pixels may not receive any corresponding projection from the DSM. These pixels are assigned a default non-shadow value and stored in an uncertainty mask to indicate areas without valid information. Finally, a post-processing step is applied to remove small spurious shadow regions using connected component analysis.

While our shadow annotation method is well-suited for large-scale applications, it has certain limitations. First, DSMs and satellite images are often captured at different times—sometimes years apart—leading to discrepancies between the two (e.g., buildings appearing in one but not the other), which can result in missing or phantom shadows. Additionally, although our method primarily aims to label shadows of buildings and other static structures, transient objects such as large vehicles or seasonal vegetation changes can introduce spurious shadow artifacts. These artifacts are not always fully removed by post-processing or filtered out by vegetation masks.

Another limitation is the inability to handle hollow structures such as power lines or bridges. Since our method relies on elevation data to cast shadows, all objects are treated as solid walls, leading to inaccuracies in these cases. Furthermore, because our shadow masks are entirely derived from the DSM, any errors or inconsistencies in the DSM directly propagate into the generated shadows.

Figure 3 illustrates some of these challenges, along with predictions from a model trained on S-EO. Despite the noise in annotations, as detailed in the experiments section, the scale and overall quality of our dataset enable the training of robust models that ultimately generate even more accurate predictions than the provided labels.

**Satellite data processing pipeline.** The S-EO dataset primarily relies on two data sources: WorldView-3 imagery and corresponding metadata collected as part of the IARPA




Figure 3. Limitations of our shadow annotation method. Left to right: pansharpened RGB image, shadow annotations, model predictions (S-EO-trained), and an overlay (magenta: ground truth, cyan: predictions, orange: matches). DSM holes in the cross-shaped building cause false positives, which the model corrects.

CORE3D program [5, 12, 18, 24], and LiDAR-derived elevation and terrain data from the USGS 3D Elevation Program. The IARPA CORE3D data include respectively 26, 43 and 35 WorldView-3 images over Jacksonville, Omaha, and San Diego, captured between 2014 and 2016. Each WorldView-3 product contains panchromatic (PAN) images at 30 cm resolution and the corresponding multi-spectral (MSI) bands at 120 cm.

The S-EO processing pipeline begins with the definition of  $500 \times 500$  m square tiles, covering the footprints of all available images in each city. Tiles with more than 60% water coverage, determined using the SRTM Water Body Data, are discarded. For each remaining tile, we extract the corresponding MSI and PAN crops from all images. The PAN images undergo radiometric correction. The MSI images first undergo a top-of-atmosphere correction and then are pansharpened. For that, each band is first independently aligned to the PAN using ORB feature matching [39]. Lastly, a vegetation mask is computed by thresholding the NDVI index at 0.0, allowing the removal of vegetation and tree shadows from further analyses.

To generate the DSMs, we use LiDAR data from the USGS 3D Elevation Program, primarily relying on the *FL\_Peninsular\_FDEM\_Duval\_2018*, *IA\_FullState*, and *CA\_SanDiegoQL2\_2014* datasets. When data from these campaigns is unavailable, we use the most recent data available for the respective region. DSMs are computed at a grid resolution of 50 cm per pixel.

To minimize the inclusion of shadows casted by thin or transient structures such as trees, light poles, and transmission lines, we aggregate the point cloud using the minimum elevation value in each cell (DSM Min). However, experiments (Section 4.1) revealed that DSM Min introduces a systematic bias, causing buildings to appear thinner and their shadows smaller than expected. To address this issue, we also compute a maximum-based DSM (DSM Max) and generate an alternative set of shadow masks (see Figure 2).

Lastly, we ensure precise alignment and georeferencing

between the images and the DSMs. Our approach follows the same conceptual pipeline as [7], but differs in specific implementation details, such as the choice of stereo models. First, we perform bundle adjustment [32] on the PAN images to correct localization errors and discrepancies across the RPCs of different cameras. This step ensures internal consistency among cameras but does not provide absolute georeferencing with respect to the DSMs. After bundle adjustment, we generate a robust photogrammetric DSM by computing the median of 10 pairwise DSMs obtained with S2P [2]. The image pairs for stereoscopic reconstruction are selected based on the heuristics of [10, 11]. Once the photogrammetric DSM is generated, we use it to estimate the shift between the images and the LiDAR data, ensuring full alignment between the ground-truth DSM and the S-EO imagery.

### 3.2. Shadow Detection Network

We train a shadow detector for remote sensing images on the S-EO dataset, leveraging both the shadow masks and auxiliary masks. We experiment with two architectures: a U-Net [37] and FSDNet [15], a fast, state-of-the-art, shadow detection model. To improve model generalization, we incorporate data augmentation strategies proposed in [46].

Our models are initialized with pretrained weights from a different domain (ground-level images). To better align our dataset’s color distribution—enhancing contrast and visual consistency—we apply a two-step color correction process. First, we perform a channel-wise quantile clip of intensity values [26] to align the color bands, followed by a histogram equalization [29] to enhance the contrast.

For training supervision, we employ the focal loss [27] and explore different masking strategies. We exclude certain pixels from the loss computation using an uncertainty mask (marking pixels not projected from the DSM, see Algorithm 2) and a vegetation mask.

**Min-Max shadow masks.** Our training strategy leverages the systematic differences between DSM Min and DSM Max. DSM Min tends to erode buildings, shrinking their structures and shadows, while DSM Max dilates them, expanding both. This creates a margin between the two, which we use to mitigate bias. As shown in Figure 4, we supervise only in regions where both shadow masks agree, ignoring the rest. This reduces shadow overlap with buildings and improves the accuracy of predicted shadows (see Table 1).

## 4. Experiments

We conduct a series of shadow detection experiments to showcase the impact of our dataset and demonstrate its applicability beyond our evaluation setup. The S-EO dataset primarily functions as a large-scale training resource rather than a validation set, due to the inherent noise in its annota-



Figure 4. Impact of bias reduction when training with DSM Min and DSM Max shadow masks. The image shows shadow predictions for a San Diego tile: magenta (DSM Min-only training), cyan (Min-Max training), and orange (overlap, where both models agree). Magenta regions extend onto buildings, while cyan aligns better with actual shadows.

tions (Figure 3). Our experiments show that it enables solutions to various shadow-related challenges across different domains.

### 4.1. Shadow detection on the S-EO dataset

During training, we randomly sample  $512 \times 512$  patches to generate training batches. We use a rolling window approach to ensure full coverage of validation images. We set the batch size to 32 and use a learning rate of  $5 \times 10^{-3}$  with the AdamWScheduleFree optimizer [8]. The training data consists of almost all Jacksonville sites, along with 50 sites from San Diego. The remaining data is reserved for validation and testing. The excluded Jacksonville sites correspond to tiles that Jacksonville sites JAX\_341, JAX\_342, JAX\_335, and JAX\_334 are also excluded from training as they overlap with EO-NeRF areas used later in Section 4.2.

Our initial experiments revealed no significant performance differences between FSDNet, initialized with weights from [16], and a general-purpose U-Net pretrained on ImageNet. Given that U-Net is a widely used, well-understood architecture with broad applicability, we opted to conduct all subsequent experiments using the U-Net implementation from [17] with the usage of squeeze and excitation blocks [38].

Table 1 reports Balanced Error Rates (BER, Positive BER, Negative BER) and F-Score—standard shadow detection metrics [16]—for both the Min-Max trained model and the baseline trained with Min-only masks. To mitigate annotation noise, we use uncertainty masks and evaluate under two conditions: (1) the full image and (2) only pixels that are certain and where both Min and Max shadow masks agree. We test on nine diverse areas with varied structures across the three cities, where we manually verified the accuracy of shadow annotations. Qualitative results are shown in Figure 5.

Thanks to the use of fully convolutional models, we can process entire images at test time, allowing predictions to be context-aware. As shown in Figure 6, having access to the full scene can significantly impact shadow detection. When

Metric	Shadow detector		Baseline (min shadows)	
	All	Filt.	All	Filt.
BER ( $\downarrow$ )	30.81%	21.79%	33.59%	28.54%
Pos. BER ( $\downarrow$ )	59.76%	43.17%	66.47%	56.86%
Neg. BER ( $\downarrow$ )	1.86%	0.42%	0.71%	0.22%
F-Score ( $\uparrow$ )	48.47%	70.13%	46.76%	59.13%

Table 1. Evaluation metrics for shadow detection across diverse areas of the S-EO dataset: JAX\_341, JAX\_342, JAX\_335, JAX\_334, UCSD\_353, UCSD\_573, OMA\_93, OMA\_930, OMA\_967. Shadow detector is the final model trained with the min-max masks, baseline is only supervised with min shadows (Section 3.2)




Figure 5. Qualitative results on the S-EO dataset. Top to bottom: Jacksonville (JAX\_334.0), Omaha (OMA\_93.15), San Diego (UCSD\_353.10). Left to right: input image, ground-truth shadow mask, and model prediction.

only a patch is provided, a dark region may not be labeled as a shadow if the object casting it is not visible. However, feeding the entire image into the model allows it to correctly associate shadows with their sources, demonstrating its ability to leverage contextual cues beyond pixel intensity.

**Shadow detection generalization to AISD.** We assess the generalization ability of our shadow detector on unseen remote sensing imagery. To this end, we use the AISD dataset [30], which, to our knowledge, is the only available for this task. A noticeable domain shift exists between S-EO and AISD due to radiometric differences, as AISD is derived from aerial imagery. These discrepancies persist even after applying our color correction pre-processing, leading to suboptimal zero-shot generalization.

To bridge the gap between satellite and aerial domains, we fine-tune our model on a small subset of AISD’s train-




Figure 6. Impact of contextual information on shadow prediction. Left to right: input image, shadow mask from crop-based inference, and shadow mask from full-image inference. The red region highlights a crop with a misclassified shadow due to missing context in the cropped input.

Model	AUC-ROC	F-Score
DSSDNet [30]	0.985	<b>91.79%</b>
Fine-tuned U-Net (Ours)	<b>0.990</b>	91.52%
U-Net from scratch	0.976	89.02%

Table 2. Comparison of shadow detection performance on the AISD dataset. Higher values indicate better performance.

ing data, using only 10% (41 images). Despite the limited supervision, this adaptation significantly enhances performance. Our fine-tuned model surpasses the best model reported in the benchmark [30] in AUC-ROC and achieves a comparable F-score, as shown in Table 2.

These results highlight the advantages of leveraging a large-scale pretraining dataset. To further validate this, we also train a U-Net from scratch using the full AISD training partition and find that our pretrained model outperforms it. This demonstrates that, in this case, pretraining on a large-scale dataset is more effective than training solely on task-specific data, even when the whole, but small, target dataset is available. Moreover, our findings highlight the U-Net’s strong capability for shadow detection, showing that it can surpass concurrent task-specific architectures when provided with sufficient data. Table 2 reports the quantitative evaluation, and Figure 7 presents some visual results.

## 4.2. Extending EO-NeRF with shadow supervision

Lastly, we extended the EO-NeRF framework [34] to use the shadow masks predicted by our segmentation model as auxiliary inputs, alongside the satellite views.

EO-NeRF is a NeRF variant designed for multi-date satellite imagery that achieves state-of-the-art geometry estimates thanks to a geometrically consistent shadow rendering approach. The EO-NeRF geometry is optimized by penalizing pixel-wise color differences between sets of rendered and actual pixels in the input views. Notably, the rendered color results from a combination of physical variables, including a shadow component. As a result, shadow masks can be easily integrated into the optimization process to help the shadow component align with the input masks.




Figure 7. Qualitative results on the AISD dataset. Top to bottom: images from Austin, Chicago, and Innsbruck. Left to right: input image, ground-truth shadow mask, and model prediction.


Given the shadow values estimated by EO-NeRF, denoted  $\hat{\mathcal{M}}$ , and the GT shadow values in the input masks, denoted  $\mathcal{M}$ , we add the following term to the loss function:

$$\mathcal{L}_S = \lambda \mathcal{M}(\hat{\mathcal{M}} - \mathcal{M})^2, \quad (2)$$

where  $\lambda$  is a weight that we set to  $\sum^N \mathcal{M}/N$ , *i.e.*, the percentage of GT shadow pixels in a batch of  $N$  rays. Note that (2) penalizes only false negatives, *i.e.*, rays not rendered as shadows that are shadows in  $\mathcal{M}$ . We also experimented with a binary cross-entropy term to penalize both false negatives and false positives, but it yielded poorer performance. Allowing the flexibility to incorporate new shadows, even those absent from the ground-truth mask, resulted in the best performance.

Figure 8 shows the evolution of the altitude MAE with respect to the number of views across the *DFC2019* and *IARPA2016* areas of interest reported in EO-NeRF [34], with and without shadow mask supervision. In all experiments we used the bundle-adjusted RPC models provided in [33, 34]. No view selection was performed; for each  $K$ -views subset, the first  $K$  images from the training list were used. Experiments with fewer than 5 views are not included, as few-shot NeRF-based geometry becomes excessively noisy in the absence of depth priors [47].

The experiments that account for shadow masks using the loss (2) consistently achieve lower altitude MAE (Figure 8) and produce sharper, less noisy geometry compared to the concurrent experiments. Figure 9 shows a detail of the output geometry for one of the target areas [34].



Altitude MAE [m]	all views	9 views	7 views	5 views
Shadow masks, $\mathcal{L}_S$ (2)	1.40	1.90	2.44	3.41
No shadow masks	1.57	2.11	2.71	3.66

Figure 8. Average EO-NeRF altitude MAE (in meters) across all areas as a function of the number of input satellite views. Shadow supervision consistently improves altitude accuracy, as the error distribution of blue boxes is lower than that of the corresponding red boxes. Average MAE values are listed in the associated table.




Figure 9. Left to right: JAX\_214, 7 input views: EO-NeRF DSM detail without and with shadow supervision vs. LiDAR DSM.

## 5. Conclusion

This work presented S-EO, a novel large-scale, high-resolution dataset for geometry-aware shadow detection in satellite imagery. The dataset comprises WorldView-3 images, shadow masks, and ground-truth DSMs. All shadow masks in the dataset were automatically annotated using the sun position and the geometry of the associated DSMs.

The significant potential of the S-EO dataset is showcased by training a shadow detector model and demonstrating its capacity to generalize to unseen aerial imagery. The predicted shadow masks are subsequently used to incorporate shadow supervision into the state-of-the-art multi-view reconstruction method EO-NeRF. The geometry estimates from the shadow-supervised variant consistently outperform the geometry reconstructions of the original method. Note that the proposed shadow-based supervision is also compatible with any other 3D modeling method that incorporates shadow modeling, including recent advancements in Gaussian Splatting [1].

## Acknowledgments

This work was partially supported by Agencia Nacional de Investigación e Innovación (ANII, Uruguay) under the graduate scholarship POS\_NAC\_2023\_1\_177798. It was performed using HPC resources from GENCI-IDRIS (grant 2024-AD011012453R4). Centre Borelli is also with Université Paris Cité, SSA and INSERM.

## References

- [1] Luca Savant Aira, Gabriele Facciolo, and Thibaud Ehret. Gaussian splatting for earth observation. In *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025. 3, 8
- [2] Tristan Amadei, Enric Meinhardt-Llopis, Carlo de Franchis, Jeremy Anger, Thibaud Ehret, and Gabriele Facciolo. s2p-hd: Gpu-accelerated binocular stereo pipeline for large-scale same-date stereo. In *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2025. 6
- [3] Nikhil Behari, Akshat Dave, Kushagra Tiwary, William Yang, and Ramesh Raskar. SUNDIAL: 3D satellite understanding through direct ambient and complex lighting decomposition. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 522–532, 2024. 3, 4
- [4] Camille Billouard, Dawa Derksen, Emmanuelle Sarrazin, and Bruno Vallet. SAT-NGP: Unleashing neural graphics primitives for fast relightable transient-free 3d reconstruction from satellite imagery. In *2024 IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, pages 8749–8753, 2024. 3
- [5] Marc Bosch, Kevin Foster, Gordon Christie, Sean Wang, Gregory D. Hager, and Myron Brown. Semantic stereo for incidental satellite images. In *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1524–1532, 2019. 1, 4, 5
- [6] Jack E Bresenham. Algorithm for computer control of a digital plotter. In *Seminal graphics: pioneering efforts that shaped the field*, pages 1–6. 1998. 4
- [7] Myriam Cournet, Emmanuelle Sarrazin, Loïc Dumas, Julien Michel, Jonathan Guinet, David Youssefi, Véronique Defonte, and Quentin Fardet. Ground truth generation and disparity estimation for optical satellite imagery. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 43:127–134, 2020. 6
- [8] Aaron Defazio, Xingyu Yang, Ahmed Khaled, Konstantin Mishchenko, Harsh Mehta, and Ashok Cutkosky. The road less scheduled. *Advances in Neural Information Processing Systems*, 37:9974–10007, 2025. 6
- [9] Dawa Derksen and Dario Izzo. Shadow neural radiance fields for multi-view satellite photogrammetry. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 1152–1161, 2021. 3
- [10] Gabriele Facciolo, Carlo De Franchis, and Enric Meinhardt-Llopis. Automatic 3d reconstruction from multi-date satellite images. In *2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 1542–1551, 2017. 6
- [11] Alvaro Gómez, Gregory Randall, Gabriele Facciolo, and Rafael Grompone von Gioi. Improving the pair selection and the model fusion steps of satellite multi-view stereo pipelines. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 6344–6353, 2023. 6
- [12] S. Hagstrom, H. W. Pak, S. Ku, S. Wang, G. Hager, and M. Brown. Cumulative assessment for urban 3d modeling. In *2021 IEEE Geoscience and Remote Sensing Symposium*, 2021. 1, 4, 5
- [13] Le Hou, Tomás F Yago Vicente, Minh Hoai, and Dimitris Samaras. Large scale shadow annotation and detection using lazy annotation and stacked CNNs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(4):1337–1351, 2019. 2
- [14] Xiaowei Hu, Lei Zhu, Chi-Wing Fu, Jing Qin, and Pheng-Ann Heng. Direction-aware spatial context features for shadow detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7454–7462, 2018. 2
- [15] Xiaowei Hu, Tianyu Wang, Chi-Wing Fu, Yitong Jiang, Qiong Wang, and Pheng-Ann Heng. Revisiting shadow detection: A new benchmark dataset for complex world. *IEEE Transactions on Image Processing*, 30:1925–1934, 2021. 2, 6
- [16] Xiaowei Hu, Zhenghao Xing, Tianyu Wang, Chi-Wing Fu, and Pheng-Ann Heng. Unveiling deep shadows: A survey on image and video shadow detection, removal, and generation in the era of deep learning. *arXiv preprint arXiv:2409.02108*, 2024. 1, 2, 6
- [17] Pavel Iakubovskii. Segmentation models pytorch. [https://github.com/qubvel/segmentation\\_models.pytorch](https://github.com/qubvel/segmentation_models.pytorch), 2019. 6
- [18] Intelligence Advanced Research Projects Activity (IARPA). Creation of Operationally Realistic 3D Environment (CORE3D). <https://www.iarpa.gov/research-programs/core3d>. Accessed: 2025-03-01. 1, 4, 5
- [19] R Bruce Irvin and David M McKeown. Methods for exploiting the relationship between buildings and their shadows in aerial imagery. *IEEE Transactions on Systems, Man, and Cybernetics*, 19(6):1564–1575, 1989. 1
- [20] Leiping Jie and Hui Zhang. A fast and efficient network for single image shadow detection. In *2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2634–2638, 2022. 2
- [21] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4):139–1, 2023. 3
- [22] Salman Hameed Khan, Mohammed Bennamoun, Ferdous Sohel, and Roberto Togneri. Automatic feature learning for robust shadow detection. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1939–1946, 2014. 2

- [23] Salman H Khan, Mohammed Bennamoun, Ferdous Sohel, and Roberto Togneri. Automatic shadow detection and removal from a single image. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(3):431–446, 2015. 2
- [24] Bertrand Le Saux, Naoto Yokoya, Ronny Hänsch, and Myron Brown. Data Fusion Contest 2019 (DFC2019), 2019. 1, 4, 5
- [25] Gregoris Liasis and Stavros Stavrou. Satellite images analysis for shadow detection and building height estimation. *ISPRS Journal of Photogrammetry and Remote Sensing*, 119: 437–450, 2016. 1
- [26] Nicolas Limare, Jose-Luis Lisani, Jean-Michel Morel, Ana Belén Petro, and Catalina Sbert. Simplest color balance. *Image Processing On Line*, 1:297–315, 2011. 6
- [27] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2980–2988, 2017. 6
- [28] Jingwang Ling, Zhibo Wang, and Feng Xu. ShadowNeuS: Neural SDF reconstruction by shadow ray supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 175–185, 2023. 4
- [29] Jose-Luis Lisani, Ana-Belen Petro, and Catalina Sbert. Color and contrast enhancement by controlled piecewise affine histogram equalization. *Image Processing On Line*, 2:243–265, 2012. 6
- [30] Shuang Luo, Huifang Li, and Huanfeng Shen. Deeply supervised convolutional neural network for shadow detection based on a novel aerial shadow imagery dataset. *ISPRS Journal of Photogrammetry and Remote Sensing*, 167:443–457, 2020. 1, 2, 4, 7
- [31] Emmanuel Maggiori, Yuliya Tarabalka, Guillaume Charpiat, and Pierre Alliez. Can semantic labeling methods generalize to any city? The Inria aerial image labeling benchmark. In *2017 IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, pages 3226–3229, 2017. 2
- [32] Roger Marí, Carlo de Franchis, Enric Meinhardt-Llopis, Jérémy Anger, and Gabriele Facciolo. A generic bundle adjustment methodology for indirect RPC model refinement of satellite imagery. *Image Processing On Line*, 11:344–373, 2021. 6
- [33] Roger Marí, Gabriele Facciolo, and Thibaud Ehret. Sat-NeRF: Learning multi-view satellite photogrammetry with transient objects and shadow modeling using RPC cameras. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 1310–1320, 2022. 3, 8
- [34] Roger Marí, Gabriele Facciolo, and Thibaud Ehret. Multi-date earth observation NeRF: The detail is in the shadows. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 2035–2045, 2023. 2, 3, 7, 8
- [35] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. NeRF: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021. 3
- [36] Vu Nguyen, Tomas F Yago Vicente, Maozheng Zhao, Minh Hoai, and Dimitris Samaras. Shadow detection with conditional generative adversarial networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4510–4518, 2017. 2
- [37] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical image computing and computer-assisted intervention—MICCAI 2015: 18th international conference, Munich, Germany, October 5–9, 2015, proceedings, part III 18*, pages 234–241. Springer, 2015. 2, 6
- [38] Abhijit Guha Roy, Nassir Navab, and Christian Wachinger. Recalibrating fully convolutional networks with spatial and channel “squeeze and excitation” blocks. *IEEE Transactions on Medical Imaging*, 38(2):540–549, 2018. 6
- [39] Ethan Rublee, Vincent Rabaud, Kurt Konolige, and Gary Bradski. ORB: An efficient alternative to SIFT or SURF. In *2011 International Conference on Computer Vision*, pages 2564–2571, 2011. 5
- [40] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. MobileNetV2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4510–4520, 2018. 2
- [41] Mingxing Tan and Quoc Le. EfficientNet: Rethinking model scaling for convolutional neural networks. In *International Conference on Machine Learning*, pages 6105–6114, 2019. 2
- [42] Kushagra Tiwary, Tzofi Klinghoffer, and Ramesh Raskar. Towards learning neural representations from shadows. In *European Conference on Computer Vision*, pages 300–316, 2022. 4
- [43] Deniz Kavzak Ufuktepe, Jaired Collins, Ekinan Ufuktepe, Joshua Fraser, Timothy Krock, and Kannappan Palaniappan. Learning-based shadow detection in aerial imagery using automatic training supervision from 3d point clouds. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3926–3935, 2021. 3
- [44] Jifeng Wang, Xiang Li, and Jian Yang. Stacked conditional generative adversarial networks for jointly learning shadow detection and shadow removal. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1788–1797, 2018. 2
- [45] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1492–1500, 2017. 2
- [46] Han Yang, Tianyu Wang, Xiaowei Hu, and Chi-Wing Fu. SILT: Shadow-aware iterative label tuning for learning to detect shadows from noisy labels. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12687–12698, 2023. 2, 6
- [47] Lulin Zhang and Ewelina Rupnik. SparseSat-NeRF: Dense depth supervised neural radiance fields for sparse satellite images. *ISPRS Annals*, 2023. 8
- [48] Jiejie Zhu, Kegan GG Samuel, Syed Z Masood, and Marshall F Tappen. Learning to recognize shadows in monochro-

matic natural images. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 223–230, 2010. [2](#)