
Satellite Sunroof: High-res Digital Surface Models and Roof Segmentation for Global Solar Mapping

Vishal Batchu^{*1} Alex Wilson^{*1} Betty Peng¹ Carl Elkin¹ Umangi Jain²
Christopher Van Arsdale¹ Ross Goroshin¹ Varun Gulshan¹

¹Google ²University of Toronto[†]

{vishalbatchu, alexwilson, bettypeng, celkin}@google.com

{umangi.jain}@mail.utoronto.ca

{cvanarsdale, goroshin, varungulshan}@google.com

* Equal contribution from these authors.

[†] Work done while at Google.

Abstract

The transition to renewable energy, particularly solar, is key to mitigating climate change. Google’s Solar API aids this transition by estimating solar potential from aerial imagery, but its impact is constrained by geographical coverage. This paper proposes expanding the API’s reach using satellite imagery, enabling global solar potential assessment. We tackle challenges involved in building a Digital Surface Model (DSM) and roof instance segmentation from lower resolution and single oblique views using deep learning models. Our models, trained on aligned satellite and aerial datasets, produce 25cm DSMs and roof segments. With ~1m DSM MAE on buildings, ~5° roof pitch error and ~56% IOU on roof segmentation, they significantly enhance the Solar API’s potential to promote solar adoption.

1 Introduction

Google Maps Platform Solar API [Google, 2024] aims to increase the speed and depth of rooftop solar photo-voltaic roll-out by accurately estimating the solar potential of all suitable buildings worldwide using high quality aerial imagery. Since its release, we estimate that Solar API has been used in over 1 million residential solar projects in US, Europe and Japan. Previous work [Goroshin et al., 2023] demonstrated the potential of such mapping with lower quality aerial imagery in the US and parts of Europe increasing data coverage by 10x over high quality aerial coverage. With satellite imagery, Solar API has the potential to increase its coverage by a further 1B buildings, focusing on the global south (20+ countries across South America, Asia, Africa, Australia and Europe). This expansion represents a further 10x increase in potential area coverage over high quality aerial imagery and enables a higher refresh rate to assess changes in solar potential over time.

Accurate rooftop geometry and shading analysis are crucial for the Solar API pipeline, which relies on precise Digital Surface Models (DSMs) and roof segmentation. We propose training deep learning models to predict DSMs and roof segments from satellite imagery, using high-quality aerial imagery as labels. Our method addresses challenges inherent in satellite data, such as lower resolution, oblique angles, and temporal discrepancies compared to aerial labels. Building on prior work [Goroshin et al., 2023], we demonstrate the potential of Solar API to enhance global solar potential assessment, advancing the transition to sustainable energy.

Recent advances in DSM estimation such as [Panagiotou et al., 2020], [Stucker and Schindler, 2022], [Kunwar, 2019], have made significant strides, but often overlook fine-grained roof geometry detail which is crucial for solar design. While NeRF-based approaches [Marí et al., 2022] offer compelling 3D reconstructions, their multi-view requirement limits scalability. Similarly, existing




Figure 1: Overview of inputs and outputs from our ML models. Both DSMs are visualized with a hillshade visualizer. A few sides of buildings are highlighted in the off-nadir satellite RGB with red ovals to emphasize the off-nadir nature of the image. Location: Ankara, Turkey.

roof segmentation techniques such as [Chen et al., 2019], while effective in certain contexts, often fall short when dealing with satellite inputs. This paper introduces solutions aimed at enhancing DSM estimation and achieving precise instance-level roof segmentation from satellite imagery, thereby facilitating effective solar design.

2 Data

2.1 Inputs, labels and pre-processing

Inputs comprise off-nadir satellite (Pleiades Neo [ESA, 2024]) RGB at 30cm resolution along with optional photogrammetry derived DSMs and DTMs (plane sweep stereo [Collins, 1996] with graph cut optimization [Boykov et al., 2001]) computed from a stack of satellite imagery wherever available. These photogrammetry-derived DSMs (even when present) often lack the necessary roof detail required, so we propose enhancements on top.

Labels comprise high quality nadir aerial RGB and corresponding DSMs + DTMs computed via photogrammetry for the DSM estimation task. We compute building instances from the imagery using the a high-quality building detection model [Sirko et al., 2021]. We then run graph cut [Boykov et al., 2001] on the DSMs within each building to produce (somewhat noisy) roof segment labels [Goroshin et al., 2023] for the roof segment prediction task. We then use simple geometry to reproject the labels into the view frame of each satellite image (see section A.2). Lastly we compute a few masks: Building presence/consistency masks are computed by comparing satellite and label imagery detected buildings. Similarly, roof segment consistency masks are generated by comparing label roof segments and building instances. These masks are used to filter areas of disagreement. (see section B.1).

2.2 Dataset

The final dataset is constructed by pairing processed inputs with two sets of labels: off-nadir labels aligned with the input view and nadir labels. This results in $\sim 1.1\text{M}$ datapoints in total, with $\sim 275\text{k}$ containing satellite DSMs/DTMs as well (referred to as "RGB+DSM" here onwards), and the remaining $\sim 860\text{k}$ without (referred to as "RGB only" here onwards). All the inputs + labels are re-sampled to a 25cm resolution in UTM projection. Each datapoint consists of images of 1024×1024 pixels. Spatial distributions of the dataset splits are visualized in section A.3. We generate 60:20:20 train:val:test splits at the city level (i.e. the splits consist of different cities) and then sub-sample the validation and test splits to ensure an equal per-country sampling of imagery.

In addition we collect a human annotated dataset of 1647 tiles where roof segments were annotated in the off-nadir RGB satellite imagery (see section A.1). We use these labels for additional validation to complement our noisier graph cut segmentation.




Figure 2: Sample inputs and outputs from the Satellite Solar API pipeline. All outputs are nadir. Location: Brasilia, Brazil.

3 Methods

We train base and refinement models which take in off-nadir satellite imagery as inputs and produce nadir outputs which are consumed by the Satellite Solar API pipeline.

3.1 Base model

The base model processes off-nadir satellite RGB, optionally incorporating a photogrammetry derived height map (DSM minus DTM) and satellite viewing angles (elevation and azimuth) to generate enhanced off-nadir height maps and roof segment instances. Subsequently, we reproject the off-nadir satellite imagery (see section A.2), enhanced height maps and roof segments to a nadir view using the enhanced height maps.

Similar to [Goroshin et al., 2023], we use a U-Net [Ronneberger et al., 2015] styled architecture employing a Swin Transformer [Liu et al., 2021] encoder for feature extraction, followed by a convolutional up-sampling decoder (see section B.2 for details and hyper-parameters). Two prediction heads are added on top of the decoder: one dedicated to height map regression, the other to affinity mask prediction for roof segment delineation as described in Goroshin et al. [2023].

The model is trained using an L1 loss to minimize height map discrepancies, a Sobel gradient [Kanopoulos et al., 1988] loss to achieve smooth gradients and capture finer roof details (crucial for panel placement later on), and an affinity mask loss for roof instance segmentation [Goroshin et al., 2023]. Performance is evaluated using L1 height map error, roof segment intersection over union, and pitch/azimuth errors computed from per-pixel surface normals (obtained from the height maps) averaged within each segment.

To enhance the accuracy and alignment of labels used in losses and metrics with respect to the satellite inputs, we employ various masking techniques (discussed in detail in section B.1).

Table 1: Validation and test results for the combined base + refinement model. Columns left to right: Split, dataset, input channels used by the model, overall height map MAE (mean absolute error), height map MAE over buildings only, graph-cut (GC) based roof pitch/azimuth errors, graph-cut based IoU (intersection over union), human-labels based IoU.

Split	Dataset	Input	Overall MAE	Building MAE	GC pitch error	GC azimuth error	GC segment IOU	Human labels segment IOU
Val	RGB+DSM	RGBH	1.51m	1.17m	4.73°	13.2°	54.2%	56.0%
	RGB+DSM	RGB	1.65m	1.33m	4.81°	13.7°	53.6%	56.2%
	RGB only	RGB	1.41m	1.07m	5.08°	12.3°	52.3%	-
Test	RGB+DSM	RGBH	1.24m	1.03m	4.83°	13.1°	52.4%	-
	RGB+DSM	RGB	1.33m	1.16m	4.97°	14.0°	51.8%	-
	RGB only	RGB	1.26m	0.92m	5.13°	10.2°	54.7%	-

3.2 Refinement model

This stage takes as input nadir RGB imagery and height maps (containing occlusions due to reprojection) from the base model, and produces refined nadir RGB, height maps, and roof segments devoid of such artifacts.

The model architecture largely mirrors that of the base model, with the key addition of an RGB prediction head alongside the height map and roof segment heads. An L2 loss is employed between aerial RGB and the model’s predicted RGB. Model predicted RGB is then utilized to fill any occlusions present in the original nadir satellite RGB from the base model (more details in section B.3).

To derive the final refined DSMs, we combine model-predicted height maps with available low-quality satellite DTMs. In the absence of such DTMs, we incorporate re-sampled NASA DEMs [Crippen et al., 2016] originally at 30m resolution.

3.3 Satellite Solar API pipeline

The Solar API pipeline, as detailed in Goroshin et al. [2023], estimates solar potential and panel layouts (visualized in Figure 2) by utilizing nadir RGB imagery, DSMs, and roof segments as input. It runs model inference on overlapping tiles, and stitches the output seamlessly using weighted kernels. It then uses ray-tracing to generate building-level solar flux estimates and proposes optimal solar panel placement configurations.

We assess end-to-end performance through metrics such as Mean Absolute Percentage Error (MAPE) and MAPE@5kW [Goroshin et al., 2023] on the Solar API pipeline outputs in section C.2.

4 Results and future work

Quantitative evaluation of our base + refinement models combined, using metrics defined in the preceding section, is presented in Table 1. We divide our metrics into results from the "RGB only" and "RGB+DSM" datasets and further split the results for the latter into RGBH (RGB + height map) inputs and RGB (no heightmap) inputs to probe the performance impact of losing the input height map on a consistent dataset. These results suggest that while height map inputs improve overall DSM accuracy, they are not essential for achieving high-quality roof geometry and segmentation results which are of higher importance for solar potential estimation.

In addition to these top-level metrics, in the appendix we include insights into per-country performance (section C.3, where aerial data was available), end-to-end metrics from the Solar API pipeline (section C.2) and ablations and sensitivity analyses studying the impact of key factors (section C.1, including dataset size and masking). Overall, our models demonstrate strong generalization, achieving consistent performance across different countries.

Future work will focus on refining solar potential estimates by addressing challenges such as roof obstacle detection, roof material/type detection, existing solar panel identification and roof segment fine-tuning with human labels.

Acknowledgments and Disclosure of Funding

We would like to thank John C. Platt, Artem Zholus, Christopher Schmidt, Jordan Raisher, Saleem V. Groenou, Dana Kurnaiwan, Juliet Rothenberg and Courtney Maimon for their valuable insights and suggestions.

References

- Y. Boykov, O. Veksler, and R. Zabih. Fast approximate energy minimization via graph cuts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(11):1222–1239, 2001. doi: 10.1109/34.969114.
- Qi Chen, Lei Wang, Yifan Wu, Guangming Wu, Zhiling Guo, and Steven L. Waslander. Temporary removal: Aerial imagery for roof segmentation: A large-scale dataset towards automatic mapping of buildings. *ISPRS Journal of Photogrammetry and Remote Sensing*, 147:42–55, 2019. ISSN 0924-2716. doi: <https://doi.org/10.1016/j.isprsjprs.2018.11.011>. URL <https://www.sciencedirect.com/science/article/pii/S0924271618303083>.
- R.T. Collins. A space-sweep approach to true multi-image matching. In *Proceedings CVPR IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 358–363, 1996. doi: 10.1109/CVPR.1996.517097.
- R Crippen, S Buckley, P Agram, E Belz, E Gurrola, S Hensley, M Kobrick, M Lavalle, J Martin, M Neumann, et al. Nasadem global elevation model: Methods and progress. *isprs-international archives of the photogrammetry, remote sensing and spatial information sciences*. xli-b4. 125-128. 10.5194/isprs-archives-xli. 2016.
- ESA. Pleiades neo. <https://earth.esa.int/eogateway/missions/pleiades-neo>, 2024. Accessed: 2024-08-07.
- Google. Google solar api. <https://developers.google.com/maps/documentation/solar/overview>, 2024. Accessed: 2024-08-07.
- Ross Goroshin, Alex Wilson, Andrew Lamb, Betty Peng, Brandon Ewonus, Cornelius Ratsch, Jordan Raisher, Marisa Leung, Max Burq, Thomas Colthurst, et al. Estimating residential solar potential using aerial data. *arXiv preprint arXiv:2306.13564*, 2023.
- N. Kanopoulos, N. Vasanthavada, and R.L. Baker. Design of an image edge detection filter using the sobel operator. *IEEE Journal of Solid-State Circuits*, 23(2):358–367, 1988. doi: 10.1109/4.996.
- Günter Klambauer, Thomas Unterthiner, Andreas Mayr, and Sepp Hochreiter. Self-normalizing neural networks. *Advances in neural information processing systems*, 30, 2017.
- Saket Kunwar. U-net ensemble for semantic and height estimation using coarse-map initialization. In *IGARSS 2019 - 2019 IEEE International Geoscience and Remote Sensing Symposium*, pages 4959–4962, 2019. doi: 10.1109/IGARSS.2019.8899861.
- Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021.
- 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. doi: 10.1109/CVPRW56347.2022.00137.
- Emmanouil Panagiotou, Georgios Chochlakis, Lazaros Grammatikopoulos, and Eleni Charou. Generating elevation surface from a single rgb remotely sensed image using deep learning. *Remote Sensing*, 12(12), 2020. ISSN 2072-4292. doi: 10.3390/rs12122002. URL <https://www.mdpi.com/2072-4292/12/12/2002>.
- 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.

- Wojciech Sirko, Sergii Kashubin, Marvin Ritter, Abigail Annkah, Yasser Salah Eddine Bouchareb, Yann Dauphin, Daniel Keysers, Maxim Neumann, Moustapha Cisse, and John Quinn. Continental-scale building detection from high resolution satellite imagery. *arXiv preprint arXiv:2107.12283*, 2021.
- Corinne Stucker and Konrad Schindler. Resdepth: A deep residual prior for 3d reconstruction from high-resolution satellite images. *ISPRS Journal of Photogrammetry and Remote Sensing*, 183:560–580, 2022. ISSN 0924-2716. doi: <https://doi.org/10.1016/j.isprsjprs.2021.11.009>. URL <https://www.sciencedirect.com/science/article/pii/S0924271621003075>.
- Sherrie Wang, François Waldner, and David B Lobell. Unlocking large-scale crop field delineation in smallholder farming systems with transfer learning and weak supervision. *Remote Sensing*, 14(22):5738, 2022.

A Data

A.1 Human labeled roof segments



Figure 3: Human labeled roof segments visualization.

Recognizing the inherent noise in our graph cut roof segment labels, we add human-annotated labels to our evaluation process to produce more reliable roof segmentation metrics. Each annotator is presented with a satellite RGB scene and is tasked with labeling roof segments. Due to the high building density in each scene, we ask annotators to label as many buildings as possible in a contiguous area in the center of each image [Wang et al., 2022]. To aid in discerning finer details, annotators can reference corresponding aerial RGB and DSM data alongside the satellite RGB. An example of human annotated labels are shown in Figure 3.

A.2 Geometry based reprojection




Figure 4: Geometry-based reprojection.

We opt for a simplified geometric reprojection approach that is efficient and suits our use case. Predicated on the assumptions of an infinitely distant satellite and parallel ground rays, we derive the following set of reprojection equations. Sample reprojected outputs are shown in Figure 4.

Given input satellite angles (elevation, azimuth), we obtain y and x angles for reprojection as,

$$\begin{aligned} angle_y &= \arctan(\cos(azimuth)/\tan(elevation)) \\ angle_x &= \arctan(\sin(azimuth)/\tan(elevation)) \end{aligned}$$

We reproject each pixel individually using the angles computed above, in ascending order of original height. If multiple pixels map to the same target location, we retain the value of the highest original pixel to handle occlusion.

$$\begin{aligned} reprojected_y_i &= \text{round}(y_i + (height_i/spatial_resolution) * \tan angle_y) \\ reprojected_x_i &= \text{round}(x_i + (height_i/spatial_resolution) * \tan angle_x) \end{aligned}$$

When reprojecting height maps from nadir to off-nadir in particular, we make a small modification to reproject entire sides of buildings and not just leave them as masked out pixels. For each input pixel, we get the neighbouring pixel with the smallest height (h_{base_i}) and then reproject the current pixel multiple times starting from h_{base_i} till h_i at 1m intervals. This produces smooth gradients on sides of buildings which are useful approximations for training models.

A.3 Distribution of train/val/test dataset splits

Figure 5 illustrates the geographic distributions of our train, validation, and test splits. To ensure greater geographic diversity in the validation and test sets, we sub-sample the data to maintain roughly equal representation from each country. The training set utilizes all available data.

Note that RGB+DSM and RGB only data are mutually exclusive.

B Modeling

B.1 Masking techniques for loss/metrics

The quality of training data significantly impacts the performance of our models. To mitigate the influence of noisy or erroneous data during training (loss computation) and evaluation (metrics), we implement the following masking schemes:

1. Temporal building mismatch masking: Temporal discrepancies between aerial DSM labels and satellite inputs can lead to inconsistencies due to building demolitions or new constructions. We mask out pixels where the high-quality building masks [Sirko et al., 2021] derived from satellite and aerial RGB imagery disagree with each other, thereby excluding potentially changed areas.
2. Roof segment masking: Graph cut roof segment labels from aerial imagery can be noisy and may not cover buildings fully. Utilizing high-quality building instances from satellite imagery, we mask out buildings where less than 50% of their area is covered by graph cut roof segments. This helps filter out potential graph cut errors, as roof segments are generally expected to encompass a majority of the building area.

B.2 Model details + Hyper-parameters

Our model adopts a U-Net [Ronneberger et al., 2015] style architecture, featuring a SWIN-B [Liu et al., 2021] encoder for feature extraction and a convolutional up-sampling decoder. The decoder comprises three up-sampling stages, each consisting of two convolutional blocks (each block consisting of convolution, normalization and activation layers) with skip connections. We attach two prediction heads to the decoder output, each comprising of a single convolution layer. We employ SELU activations [Klambauer et al., 2017] throughout the model.

We use the following hyper-parameters:

- Learning rate: 0.0003
- Weight decay: 1e-07
- Training batch size: 1024
- Input/label size: 1024x1024 (randomly cropped to 512x512 during training)
- Evaluation tile size: 1024x1024
- Total training steps: 75000




Figure 5: Geographical distribution of train, validation and test splits (in order). Red represents RGB+DSM data and blue represents RGB only data.

Table 2: Masking ablation results on the RGB+DSM validation split with RGBH inputs.

Method	Overall MAE	Building MAE	GC pitch error	GC azimuth error	GC segment IOU	Human labels segment IOU
Baseline	1.42m	1.38m	4.36°	18.6°	53.7%	55.4%
No temporal building mismatch masking	1.46m	1.44m	4.5°	18.4°	53.2%	53.9%
No roof segment masking	1.47m	1.51m	4.38°	18.9°	53.3%	53.5%
No masking	1.47m	1.47m	4.55°	18.2°	52.5%	51.1%

- Learning rate schedule: Warm-up + cosine decay (warm-up steps = 7,500)

The following loss weighting schemes are also applied:

- Affinity mask losses: Up-weighted by 5x to align with the scale of DSM losses.
- DSM and gradient losses within buildings: Up-weighted by 5x to prioritize rooftop geometry.
- Vegetation: Implicitly down-weighted to mitigate the impact of noisy vegetation labels due to seasonal variations and temporal gaps between labels and input data.

B.3 Refinement RGB post-processing




Figure 6: Generating the final infilled nadir RGB output.

Occlusions in geometry reprojected satellite imagery appear as distracting black regions. We address this by in-filling these areas with a blurry model-predicted RGB as outlined in Figure 6. This maintains visual clarity without obscuring the distinction between real and synthetic data.

C Results

C.1 Ablations

To gain a deeper understanding of the contributions of various modeling components, we conduct two sets of ablation studies.

1. Masking ablation: We drop each of our masking techniques one at a time to assess its impact on the overall performance of our models in Table 2.
2. Input dataset size ablation: We progressively decrease the size of the input dataset (logarithmically) to examine its effect on model performance in Table 3.

Note that both these sets of ablation results are on the base model only (opted for faster experimentation since refinement models take longer to train). So these are not comparable with the main results in Table 1 directly.

We observe (with the marginal exception of azimuth error) that the best model performance is obtained when combining both masking strategies during training, and that the model performance continues to improve with increasing dataset size: with training with 1M datapoints resulting in the best performance.

Table 3: Dataset size ablation results on the RGB+DSM validation split with RGBH inputs.

Dataset size	Overall MAE	Building MAE	GC pitch error	GC azimuth error	GC segment IOU
1M (baseline)	1.42m	1.38m	4.36°	18.6°	53.7%
360k	1.45m	1.39m	4.49°	19.4°	53.7%
200k	1.45m	1.49m	4.52°	19.4°	53.6%
60k	1.65m	1.9m	4.62°	20.6°	53.0%
20k	1.66m	1.9m	4.91°	21.4°	51.8%
6k	1.88m	2.55m	5.12°	23.9°	50.5%

Table 4: End-to-end performance metrics - MAPE@5kW and MAPE computed across randomly sub-sampled validation and test splits where we include one region from each country.

Split	MAPE@5kw	MAPE
Validation	2.6%	19.2%
Test	2.5%	18.3%

C.2 End-to-end MAPE metrics

End-to-end performance metrics are presented in Table 4. Test outputs are obtained by feeding the model-predicted nadir RGBs, DSMs, and roof segments into the Solar API pipeline. Ground truth outputs are derived by directly running the Solar API pipeline on high-quality aerial imagery of the same regions wherever overlap exists. Predicted fluxes and panel placements are then compared against each other to compute performance metrics (as outlined in [Goroshin et al., 2023]).

C.3 Country level analysis

Figures 7 and 8 show a per-country breakdown of the height error and roof segmentation performance for our RGB-only dataset (chosen as it has the greatest number of countries). We observe that the error variation between countries is small (with the exception of Chile and the Philippines, for which our ground-truth aerial imagery is atypically noisy - likely explaining their outlier status) and conclude that the model is able to adapt well to different regions and styles of housing.

D Qualitative results

Figures 9 and 10 showcase sample model results. Figure 9 compares ground truth nadir aerial data with model predictions from off-nadir satellite imagery, highlighting accurate DSM and roof segment inference. Figure 10 further demonstrates performance on diverse off-nadir inputs, showcasing predicted nadir RGB, DSM, and roof segments.




Figure 7: Per-country (RGB-only) building height error distribution




Figure 8: Per-country (RGB-only) roof segmentation IoU distribution




Figure 9: Comparison of sample model outputs with ground truth (GT) data. Location: Bloemfontein, South Africa




Figure 10: Qualitative visualizations of model outputs from various geographies.