

# SOLAR PANEL MAPPING VIA ORIENTED OBJECT DETECTION

Conor Wallace, Isaac Corley & Jonathan Lwowski

Zeitview Inc.

Santa Monica, CA 90401, USA

{conor.wallace, isaac.corley, jonathan.lwowski}@zeitview.com

## ABSTRACT

Maintaining the integrity of solar power plants is a vital component in dealing with the current climate crisis. This process begins with analysts creating a detailed map of a plant with the coordinates of every solar panel, making it possible to quickly locate and mitigate potential faulty solar panels. However, this task is extremely tedious and is not scalable for the ever increasing capacity of solar power across the globe. Therefore, we propose an end-to-end deep learning framework for detecting individual solar panels using a rotated object detection architecture. We evaluate our approach on a diverse dataset of solar power plants collected from across the United States and report a mAP score of 83.3%.

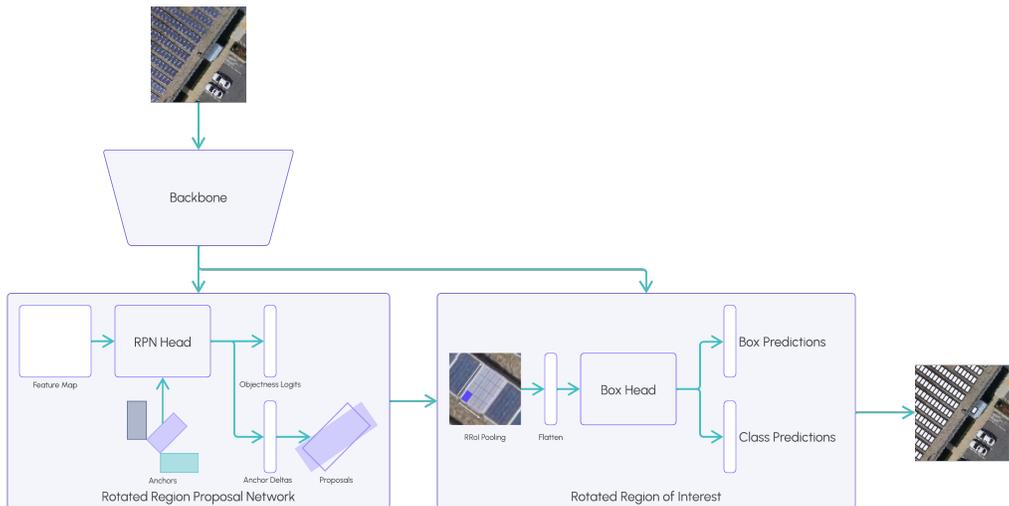


Figure 1: Oriented Panel Detection Architecture.

## 1 INTRODUCTION

The adoption of renewable energy resources is paramount to fighting the climate change crisis. Over the last decade renewable energy production has nearly quadrupled, 26% of which is contributed by solar energy (Rapier, 2020). To ensure that solar power plants are operating at maximum capacity, inspections utilizing remotely sensed imagery are important in identifying anomalies in damaged panels. However, with this rapid increase in solar production, it has become increasingly difficult to scale these inspections efficiently. The initial step in many inspection pipelines is to localize and georeference individual panels for downstream evaluation tasks. This can be prohibitively time-consuming for commercial sites, any of which can easily contain tens of thousands of solar panels. Therefore, this requires an automated approach for detecting and georeferencing individual solar

panels from large scale aerial imagery. To be specific, we wish to create a model that directly predicts the coordinates of the vertices of each solar panel. We also require the approach to generalize to images of solar arrays in any orientation across a wide variety of landscapes and environments.

In this paper we introduce a rotated object detection framework for localizing individual solar panels with arbitrary orientation. We preprocess large-scale orthomosaics into patches which are then processed in batches. The predictions are stitched back together and projected from pixel coordinates to geospatial coordinates, creating an accurate mapping of individual panels in each site. To the best of our knowledge this is the first study to model solar panel detection as a rotated object detection problem which allows us to efficiently map individual solar panels in an end-to-end fashion.

## 2 RELATED WORK

Most recent works build algorithms for mapping solar panels from aerial or satellite imagery (Yu et al., 2018; Camilo et al., 2017; Mujtaba & Wani, 2021; Golovko et al., 2021) with the goal of estimating surface area as a proxy for energy capacity, location density, and other potential insights. These studies have focused on both small-scale residential solar sites (Zhuang et al., 2020; Khomikov et al., 2022) as well as large-scale commercial sites (Yu et al., 2018). Most approaches use a two-stage architecture for mapping solar panels from aerial imagery (Hou et al., 2020; Parhar et al., 2021). The first stage consists of a convolutional neural network (CNN) classifier for predicting the probability of the given image containing solar panels or not. The second stage uses a segmentation model for segmenting out solar panels, contingent upon the first stage predicting the given image to contain solar panels. While these methods have shown promise in mapping the surface area of solar arrays across a diverse set of geographies, it remains nontrivial to extract the coordinates of individual solar panel vertices at scale.

## 3 METHODOLOGY

In this section, we describe the proposed framework’s architecture and the modeling of detection targets for solar panels.

### 3.1 ROTATED BOUNDING BOXES

Our key desideratum is to accurately localize individual solar panels, regardless of the orientation of the panels in the image, and to use this localization to extract the coordinates of the vertices of each panel. Accordingly, we model the problem as an arbitrarily oriented object detection task where the ground truth annotations of the solar panels are represented as rotated bounding boxes. Each box is comprised of the 5-dimensional tuple  $(x, y, w, h, \theta)$  where  $(x, y)$  is the centroid pixel coordinate of the box,  $w$  is the width defined as the shorter side of the box,  $h$  is the height defined as the longer side of the box, and  $\theta$  is the rotation angle in degrees required to obtain the rotated box. This bounding box formulation provides a compact representation which allows us to easily convert to the complete set of vertex coordinates  $(x_1, y_1, x_2, y_2, x_3, y_3, x_4, y_4)$  using equation 1 in Appendix B.

### 3.2 ROTATED ANCHORS

Following R-CNN style object detection architectures, we define a set of proposal anchors as a strong basis for regression to ground truth rotated bounding boxes. In addition to scale and aspect ratio, we also provide rotation angle anchors. We specify seven different rotation anchors:  $-90^\circ$ ,  $-60^\circ$ ,  $-30^\circ$ ,  $0^\circ$ ,  $30^\circ$ ,  $60^\circ$ , and  $90^\circ$ . We use standard anchors for scale and aspect ratio parameters which are defined as  $[32, 64, 128, 256, 512]$  and  $[1:2, 1:1, 2:1]$ , respectively. This produces a set of 105 anchors in total (7 rotation anchors, 5 scale anchors, and 3 aspect ratio anchors).

### 3.3 ROTATED OBJECT DETECTION

We adopt a general Faster R-CNN architecture (Ren et al., 2015) with a CNN backbone, region proposal network (RPN), classification and box regression head, and a region of interest (RoI) pooling layer. We use the oriented object detection framework from (Ma et al., 2018) in which the

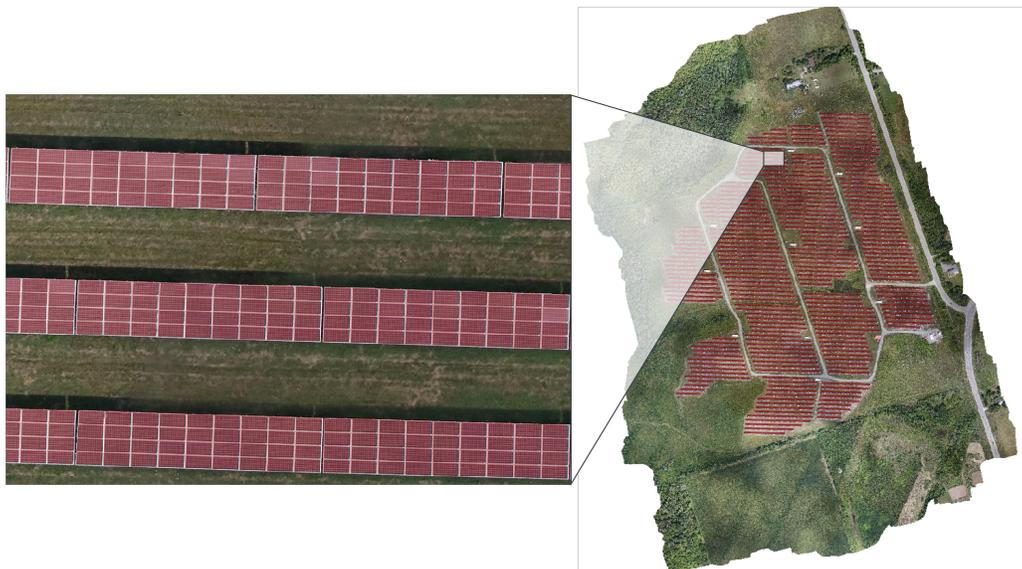


Figure 2: Sample Test Set Ortho Panel Predictions.

authors present rotated RPN (RRPN) and rotated RoI (RRoI) modifications. The RRPN performs two functions: it produces rotated bounding box proposals from the specified box anchors defined in Section 3.1 and regresses the box proposals to accurately localize panel instances for each feature map. From the RRPN, the proposals and feature maps are fed into the RRoI pooling layer which projects the arbitrarily oriented proposals onto each feature map and subsequently performs a max pooling operation. The pooled features from the RRoI layer are fed into a classification head and a box regression head which produce class labels, class scores, and box coordinates for each solar panel instance. The architecture can be seen in Figure. 1.

## 4 EXPERIMENTS

### 4.1 DATASET

To train the model, we create a diverse dataset of aerial imagery from various locations across North America. Figure 3 in Appendix A shows the exact locations where the data is collected. The dataset consists of 121 high-resolution stitched orthomosaics of large-scale solar farms captured using a fixed-wing aircraft with an RGB camera at approximately 2.5cm ground sample distance (GSD).

### 4.2 DATASET PREPROCESSING

We randomly split the 121 orthomosaics into train, validation, and test sets using an approximate split ratio of 80/10/10. We then preprocess each orthomosaic into 512x512 patches. Annotating large scale orthomosaics is time consuming and expensive. Additionally, since visual features of solar panels do not vary in comparison to natural images, we find it to be redundant to fully annotate and train on all patches from each site. Consequently, we randomly sample 10 unique foreground patches containing solar panels from each orthomosaic for annotating. We additionally sample 5 unique background patches from each orthomosaic to increase model precision on images of solar arrays with substantial portions of background information. We find that with proper augmentation, described in Appendix B.3, we achieve better results with a sampled dataset from a larger set of orthomosaics than we do with a smaller set of fully annotated orthomosaics. Finally, to test the efficacy of the proposed solar panel mapping framework, we create complete georeferenced maps of the 10 test orthomosaics containing the geospatial coordinates of the vertices of every solar panel. The spread of all three of our datasets is depicted in Table 2 in Appendix A.

Backbone	Validation				Test			
	AP	AR	AP <sub>75</sub>	AR <sub>75</sub>	AP	AR	AP <sub>75</sub>	AR <sub>75</sub>
ResNet-50-C4	80.9%	84.1%	92.5%	93.2%	84.8%	88.1%	97.9%	<b>98.4%</b>
ResNet-50-DC5	80.1%	84.1%	92.0%	93.2%	85.0%	88.3%	97.9%	98.3%
ResNet-50-FPN	<b>83.3%</b>	<b>88.4%</b>	<b>94.1%</b>	<b>96.9%</b>	84.3%	87.7%	97.8%	<b>98.4%</b>
ResNet-101-C4	82.1%	85.2%	92.4%	93.6%	84.9%	88.2%	<b>98.0%</b>	<b>98.4%</b>
ResNet-101-DC5	81.0%	84.5%	92.3%	94.0%	<b>85.5%</b>	<b>88.8%</b>	97.9%	<b>98.4%</b>
ResNet-101-FPN	79.5%	84.3%	92.9%	95.5%	82.0%	85.6%	97.5%	<b>98.4%</b>
ResNeXt-101-FPN	77.0%	82.0%	88.8%	92.4%	82.5%	86.2%	97.7%	98.2%

Table 1: Faster-RCNN Detection results.

### 4.3 METRICS

To evaluate our model, we report mean average precision (mAP) and mean average recall (mAR) as our primary performance metrics. We also report mAP and mAR with an IoU threshold of 0.75, notated as mAP<sub>75</sub> and mAR<sub>75</sub>, respectively. This is because we are interested in the performance of the model at high IoU thresholds as a proxy for evaluating how tight the localizations are.

### 4.4 RESULTS

We fine-tune the model using several different combinations of ResNet backbone networks, all of which are pretrained on the ImageNet dataset (Deng et al., 2009):

- C4: A conv4 backbone with a conv5 head as described in the original Faster R-CNN paper (Ren et al., 2015).
- DC5: A conv5 backbone with dilations in the conv5 layer and fully connected layers in the box regression head (Dai et al., 2017).
- FPN: A feature pyramid network backbone with fully connected layers in the box regression head (Lin et al., 2017).

We run an ablation study to investigate the effects of the choice of backbone on model performance. We don't run any hyperparameter tuning to keep things fair and we use the same hyperparameters for every model, with the exception that we reduce the batch size for deformable convnet backbones which require more memory during processing. The ResNet-50-FPN network comfortably achieves the best results on the validation set and comparable results on the test set.

### 4.5 DISCUSSION

The growth in cumulative solar capacity has made it difficult to scale up the inspection and maintenance of these valuable assets. In 2022, we estimate that 3.5 GW of potential solar power was lost due to faulty solar panels. By automating the initial mapping phase of the inspection process, we estimate that analysts will have a 43% increase in efficiency. This increase in analyst efficiency decreases the time that faulty solar panels are losing potential power. Finally, we estimate that our increase in inspection efficiency will lead to a potential retention of 1.4 GW of solar power.

## 5 CONCLUSION

In this paper we present a novel rotated object detection framework for end-to-end solar panel detection and mapping. With our framework, we can directly predict the coordinates of individual solar panels and project them from image space back to geographical space. We believe that this ability to automatically map the coordinates of every solar panel in each array will provide the efficiency required to monitor the overall health of large-scale sites and help continue the rapid growth of the solar power capacity worldwide.

## REFERENCES

- Joseph Camilo, Rui Wang, Leslie M. Collins, Kyle Bradbury, and Jordan M. Malof. Application of a semantic segmentation convolutional neural network for accurate automatic detection and mapping of solar photovoltaic arrays in aerial imagery. *IEEE Applied Imagery Pattern Recognition (AIPR) Workshop*, 2017.
- Jifeng Dai, Haozhi Qi, Yuwen Xiong, Yi Li, Guodong Zhang, Han Hu, and Yichen Wei. Deformable convolutional networks. *ICCV*, pp. 764–773, 2017.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. 2009.
- Vladimir Golovko, Alexander Kroshchanka, Egor Mikhno, and Myroslav Komar. Deep convolutional neural network for detection of solar panels. *Data-Centric Business and Applications*, pp. 371–389, 2021.
- Xin Hou, Biao Wang, Wanqi Hu, Lei Yin, Anbu Huang, and Haishan Wu. Solarnet: A deep learning framework to map solar plants in china from satellite imagery. *ICLR*, 2020.
- Maxim Khomiakov, Julius Radzikowski, Car Schmidt, Mathias Sørensen, Mads Andersen, Michael Andersen, and Jes Frellsen. Solardk: A high-resolution urban solar panel image classification and localization dataset. *NeurIPS*, 2022.
- Tsung-Yi Lin, Piotr Dollar, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. *CVPR*, pp. 2117–2125, 2017.
- Jianqi Ma, Weiyuan Shao, Hao Ye, Li Wang, Hong Wang, Yingbin Zheng, and Xiangyang Xue. Arbitrary-oriented scene text detection via rotation proposals. *IEEE Transactions on Multimedia*, 20:3111–3122, 2018.
- Tahir Mujtaba and M. Arif Wani. Segmentation of satellite images of solar panels using fast deep learning model. *International Journal of Renewable Energy Research (IJRER)*, 11 (1):31–45, 2021.
- Poonam Parhar, Ryan Sawasaki, Nathan Nusaputra, Felipe Vergara, Alberto Todeschini, and Hossain Vahabi. Hyperionsolarnet solar panel detection from aerial images. *NeurIPS*, 2021.
- Robert Rapier. Renewable energy growth continues at a blistering pace, 2020. URL <https://www.forbes.com/sites/rrapier/2020/08/02/renewable-energy-growth-continues-at-a-blistering-pace/?sh=2fcf9f3276b6>. Forbes Magazine.
- Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *NeurIPS*, 2015.
- Jiafan Yu, Zhecheng Wang, Arun Majumdar, and Ram Rajagopal. Deepsolar: A machine learning framework to efficiently construct a solar deployment database in the united states. *Joule*, 2: 2605–2617, 2018.
- Li Zhuang, Zijun Zhang, and Long Wang. The automatic segmentation of residential solar panels based on satellite images: A cross learning driven u-net method. *Applied Soft Computing Journal*, 92, 2020.
- Karel Zuiderveld. Contrast limited adaptive histogram equalization. *Graphics gems*, pp. 474–485, 1994.

## A DATASET DETAILS

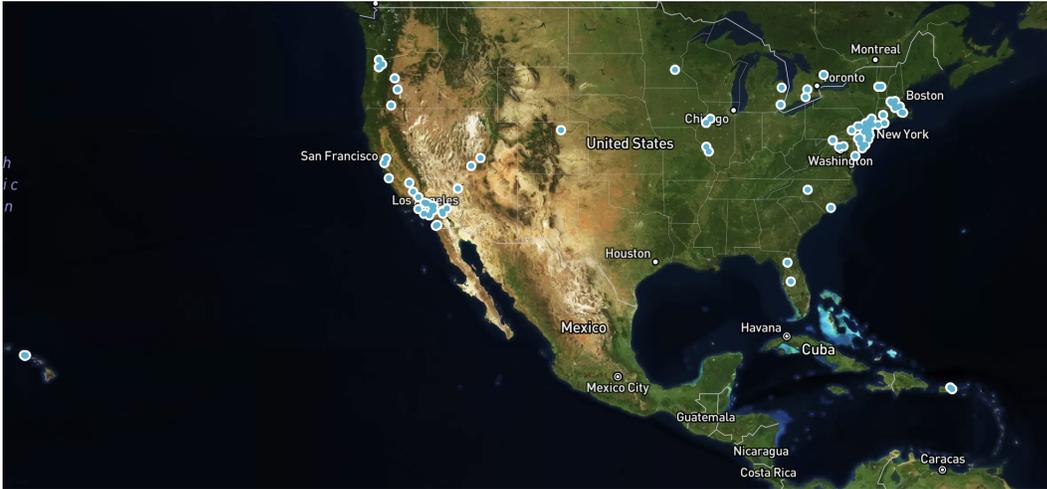


Figure 3: USA Solar farm locations.

We created a diverse dataset of aerial imagery collected from across North America. We collect 121 high resolution stitched orthomosaics from large scale solar farms using a fixed-wing aircraft with an RGB camera at roughly 2.5 cm ground sample distance (GSD). Of these orthomosaics, we reserve 10 specifically for testing.

Split	Background	Foreground	Total	Sampled Background	Sampled Foreground	Total Sampled
Train	57,323	19,151	76,474	566	795	1,361
Validation	3,815	1,133	4,948	88	216	304
Test	2,244	192	2,492	-	-	-

Table 2: Solar Panel Dataset.

## B IMPLEMENTATION DETAILS

### B.1 BOUNDING BOX CONVERSION

We use the following equation given  $\alpha$ , the orientation parameter  $\theta$  in radians to transform the compact oriented bounding box representation to the full set of vertex coordinates.

$$\begin{aligned}
 x_1 &= x - \frac{w}{2} \cos(\alpha) + \frac{h}{2} \sin(\alpha), & y_1 &= y + \frac{w}{2} \cos(\alpha) + \frac{h}{2} \sin(\alpha) \\
 x_2 &= x - \frac{w}{2} \cos(\alpha) - \frac{h}{2} \sin(\alpha), & y_2 &= y + \frac{w}{2} \cos(\alpha) - \frac{h}{2} \sin(\alpha) \\
 x_3 &= x + \frac{w}{2} \cos(\alpha) - \frac{h}{2} \sin(\alpha), & y_3 &= y - \frac{w}{2} \cos(\alpha) - \frac{h}{2} \sin(\alpha) \\
 x_4 &= x + \frac{w}{2} \cos(\alpha) + \frac{h}{2} \sin(\alpha), & y_4 &= y - \frac{w}{2} \cos(\alpha) + \frac{h}{2} \sin(\alpha)
 \end{aligned} \tag{1}$$

### B.2 TRAINING HYPERPARAMETERS

We trained each baseline model with the hyperparameters in Table 3 below. We train on 4 Tesla V100-SXM2-16GB GPUs which takes about 1 day to train for the non-FPN based models and about 4 days for the FPN models.

Hyperparameter	Value
Optimizer	Adam
Training Steps	25,000
Learning Rate	$1e - 4$
Linear Warmup	$1e - 3$
Batch Size	64
Grad Norm Clip	0.5
Angle Anchors	$[-90^\circ, -60^\circ, -30^\circ, 0^\circ, 30^\circ, 60^\circ, 90^\circ]$
Size Anchors	[32, 64, 128, 256, 512]
Aspect Ratio Anchors	[1:2, 1:1, 2:1]

Table 3: Solar Panel Detection Hyperparameters.

### B.3 DATA AUGMENTATION

During training we randomly augment the images and bounding boxes using the following methods:

- Random Rotation described in Ma et al. (2018) with probability  $p = 0.5$
- Color Jitter with brightness, contrast, saturation, and hue factors of  $\sigma = 0.1$  and probability  $p = 0.25$
- Contrast Limited Adaptive Histogram Equalization (CLAHE) Zuiderveld (1994) with kernel size  $k = 8$  and probability  $p = 0.01$
- Random Blur with kernel size  $k \in [3..7]$  and probability  $p = 0.01$

We find random rotation augmentation to be critical for accurate orientation prediction.