Automatic Detection of solar panel location from google maps satellite images using deep learning

Mohammed Ahmed Jazea

Department of Computer Science, School of Arts and Sciences, Lebanese International University (LIU), Lebanon

Email:Mohammed.alesawe@gmail.com

Abstract

The increasing use of solar energy systems requires a robust fully automated system for detection and mapping of solar panels using satellite imagery. In this paper, we proposed a deep learning system based on U-Net that automatically identifies the roof solar panels in satellite imagery from Google Maps. The model was built and tested on the PV03 data set that consisted of a diverse solar panel data set with high quality (HQ) satellite imagery. The experimental results indicated that the proposed approach had a high accuracy of 0.9758, with a mean Intersection over Union (IoU) of 0.9456 and a mean Dice coefficient of 0.9719, suggesting a high accuracy of segmentation of the solar panel regions. This work demonstrated that U-Net could identify diverse variation in panels such as size, orientation, location, and environment. Based on the results discussed, a deep learning approach to satellite image analysis has potential as cost-effective, reliable technology of assessing solar energy infrastructure which can be utilized for spatial land-use assessment, urban design, and renewable energy management.

Keywords: U-Net, Image segmentation, Google Maps

1 Introduction

The global shift toward renewable energy systems is accelerating as countries search for sustainable alternatives to fossil-fuel-based electricity generation. Of the available renewable energy sources, solar photovoltaic (PV) installations have become a fundamental pillar of achieving energy supply decarburization and reducing climate change. [1] However, despite the large increase in solar panel installations, a large-scale and accurate assessment of solar panel locations and mapping over large geographic areas represents a challenge. This is important for actions, such as asset management, maintenance planning, resources assessments, and policy-making.[2] Existing PV registries often only cover part of the large area in question, while other records include inconsistent metadata (e.g., orientation, tilt, condition) or are outdated. Therefore, there is a clear need for automated remote-sensing methods that can detect and segment solar panels using satellite or aerial imagery. High-resolution satellite imagery and platforms such as mapping services (e.g., Google Maps) and other open datasets offer wide-area views of built and agricultural environments, making them a promising source for solar panel detection. However, accurately identifying solar panels from such imagery is non-trivial. Solar panels appear in diverse settings—rooftops, ground-mounted arrays, varying orientations and scales, with shadows, reflections, and heterogeneous backgrounds. The complexity of rooftop structures, vegetation, and neighboring objects further complicates segmentation tasks. Early work by Malof et al. (2016) demonstrated the feasibility of automatic PV detection in aerial imagery,[3] providing a baseline for further investigations. Subsequent dataset efforts such as by Jiang et al. (2021) illustrated the need for multi-resolution data to support segmentation across scales. [4]

In recent years, deep learning techniques, and more specifically, convolutional neural networks (CNNs) designed for semantic segmentation, have contributed to major advancements in the analysis of remote-sensing images across multiple applications involving object extraction and land-cover mapping. The popular U-Net framework has been widely used for segmentation models in practice, due to its encoder-decoder model with skip connections, that enables this style of CNN to connect high-level information about the scene with low-level spatial information [5]. In terms of solar panel detection, methods such as the two-branch approach of HyperionSolarNet have shown good performance on aerial data [6], while the larger scale SolarNet framework have mapped entire solar farms through segmentation [7]. More recent work by Olweus & Mengshoel (2024) explored network-architecture search (NAS) to optimize segmentation models for solar farm detection, achieving a mean IoU of over 0.96 in some cases [8].

Even with these advancements, there are still some obstacles that remain to be solved when utilizing deep learning for solar panel detection using satellite imagery. First, the heterogeneity of panels (size, shape, type of module, installation surface) and variability of environments and conditions (urban vs rural, rooftop vs ground-mounted) creates generalization challenges with trained models. For example, Jiang et al. (2021) identified issues with dataset cross-domain transfer demonstrated by a marked decrease in performance produced by application of models to an image with a different resolution [9]. Second, many of the datasets required to train such models have relatively small sample sizes and are limited to a region of the study or uniform conditions for imaging purposes. The small sample sizes make it challenging for the model to expand, due to weak generalizability, and demonstrate robustness in application to non-similar geographical locations or imaging datasets. Third, satellite imagery data are typically inundated with mixed pixels, variable illumination or

over-shadow (i.e., scatted shadows, trees, etc.), and surfaces that are not solar panels but have visual similarities (i.e., dark patches on roofs or water reflections). These visual similarities can potentially confuse segmentation algorithms. There have been models produced that are lightweight and resource efficient for deployment using constrained devices (e.g., Wani & Mujtaba, 2021), albeit with a trade-off on accuracy. [10]

This paper describes a U-Net based deep learning strategy for the automatic detection of solar panel locations from Google Maps satellite images. The PV03 dataset is employed as our training and evaluation basis. Our proposed response to variation in solar panel appearance involves effective preprocessing and augmentation. As well, we fine-tuned U-Net to our particular application. The performance of the model is presented based on accuracy mean Intersection over Union (Mean IoU), and mean Dice coefficient scores: respectively, scores of 0.97, 0.93, and 0.96. The observed high-performance metrics suggests the ability of our proposed approach to accurately segment solar panel regions experimenting across different configurations and demonstrates one avenue to facilitate monitoring on a large scale of solar generating systems.

2 Related works

There has been increased interest in automatic detection and segmentation of solar photovoltaic (PV) panels from aerial or satellite imagery in recent years. In prior work, Malof et al. established a benchmark for detecting rooftop PV arrays in high resolution aerial imagery with traditional machine learning approaches [3]. Their work demonstrated that large scale automatic PV mapping is feasible and facilitated later works in the field using deep learning methods.

With the increasing extent of convolutional neural networks (CNNs), semantic segmentation architectures have grown to be favored architectures for solar panel detection. SolarNet, for example, introduced a deep learning framework for mapping large areas of solar farms in China using satellite imagery and achieved very good detection performance for ground-mounted PV installations [7]. Jiang et al. evaluated the role of multi-resolution datasets in training models for both rooftop and ground-mounted PV panels, demonstrating that models trained at one resolution typically perform poorly when applied to one of another resolution [4]. Olweus & Mengshoel compared several popular deep neural network architectures and their performance at solar farm detection and presented high mean IoU scores on European Sentinel-2 imagery [8].

Detecting rooftop photovoltaic (PV) systems creates some additional challenges based on their smaller size, effects of occlusion, and complex urban scenarios. Wani & Mujtaba evaluated several segmentation networks on high-resolution orthoimagery and other datasets for rooftop solar PV panel detection, finding that both architectural decisions and data augmentation strategies significantly impact performance [10]. Parhar et al. introduced HyperionSolarNet, a two-branch deep learning model for detecting solar PV panels from aerial images, with results closely aligned with their competitors on rooftops with heterogeneous conditions and backgrounds [6].

In addition to detection, multiple studies have also progressed the task of capacity estimation or condition monitoring. One of the more recent studies regarding deep learning-based solar farm capacity estimation segments panel regions and predicts energy generation with high accuracy and Jaccard index scores [11]. Further, Alkhatib et al. also improved segmentation to detect panel damage using U-Net with

attention modules to demonstrate the feasibility of using deep learning for maintenance and condition assessment [12].

Although advancements have been made, there are still a number of challenges. One factor is cross-region generalization, where models trained on data from one region do not generalize well to other regions [13]. Another major challenge is the significant intra-region variability including illumination, shadows, panel dimensions, panel orientation, and installations almost exclusively on rooftops vs ground-mounted panels. Many datasets also are either small datasets or specific to one region, and therefore do not support global-scale applications for monitoring. Efficient and lightweight models have been developed, allowing deployment on constrained devices, though typically at the expense of segmentation accuracy [14].

Focusing on the trends in the methodology literature, we see a cooperation of instance segmentation, hybrid object detection + segmentation pipelines, and the continuous evolution of network architecture search to support higher performance. For example, Elbl Droguett and Nicollier evaluated how to adapt detection models (YOLO, Mask R-CNN) to perform the detection of solar panels [15], and the literature suggests that U-Net is still very effective at segmenting rooftop PV's, especially when fused with augmentation methods, multi-resolution input, and attention mechanisms.

3 Methodology

3.1 Overview

The methodology proposed in this paper takes a deep learning approach to identify and segment solar panels from Google Maps satellite imagery, specifically using a built U-Net architecture, built with PyTorch. The workflow is separated into five main sections: (1) dataset preparation and preprocessing, (2) network architecture,

(3) data loading and training parameters, (4) model training and validation, and (5) saving the model for inference later. The proposed pipeline was purposely created to allow efficient processing of high-resolution imagery, while still being able to accurately segment small rooftop solar panels.

3.2 Dataset Preparation and Preprocessing

The main source of annotated imagery used was the PV03 dataset. Each image in the dataset has solar panel regions labeled manually with binary masks. The data was split into training and validation directories, with separate subdirectories for images and masks within these directories. [16]

The PV03 dataset consists of approximately 2308 aerial images with a resolution of 1024x1024 pixels and a spatial resolution of approximately 0.3 meters, making it the largest subset of the complete PV dataset. It features a wide variety of environments in which solar panels appear, including building rooftops, as well as a broad range of ground environments such as agricultural land, grasslands, shrublands, saline-alkaline areas, and bodies of water. All images are provided with precise polygon annotations for generating high-quality masks, making PV03 suitable for training models to detect photovoltaic systems in diverse contexts.

I used only two categories from the PV03 dataset: PV03_Ground_Cropland and PV03_Ground_Grassland, totaling 960 images. These are aerial images with a resolution of 1024×1024 pixels and a spatial resolution of approximately 0.3 meters. These two categories exhibit significant diversity in agricultural and grassland backgrounds, making the data suitable for training models capable of recognizing solar panels in natural ground environments with varying textures and vegetation cover. This selection provides a focused yet visually rich dataset, supporting the

development of accurate and generalizable models within agricultural and grassland contexts.

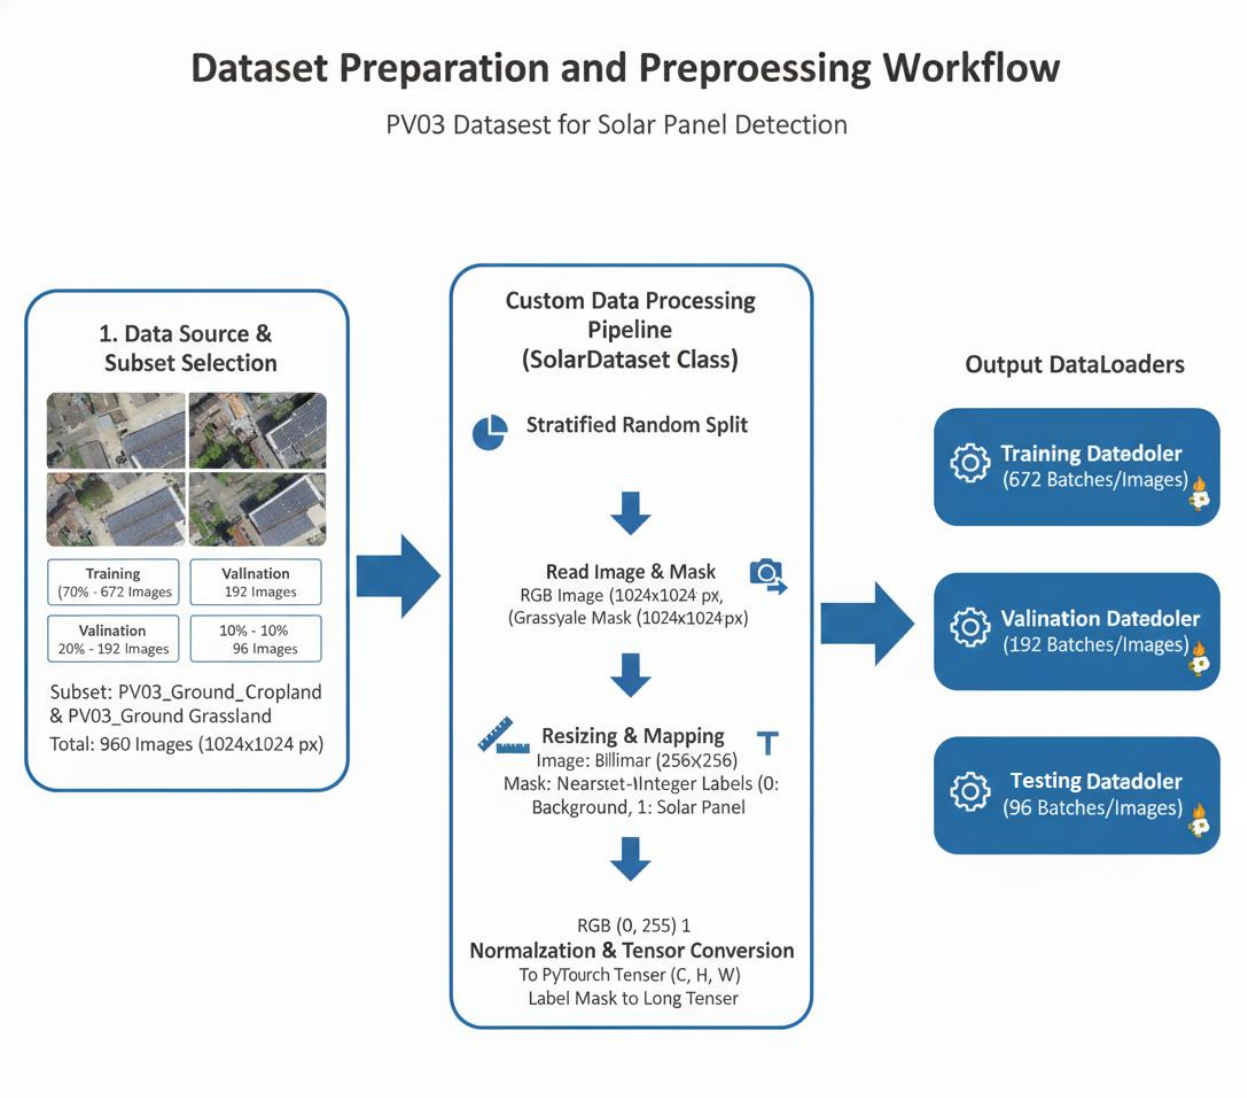

Figure 1 Dataset Preparation and Preprocessing

The data was split into 70% for training, 20% for validation, and 10% for testing, resulting in 672 images for training, 192 for validation, and 96 for testing. This split ensures a good balance between training the model and reliably evaluating its performance.

The dataset loader implements a custom PyTorch Dataset class, SolarDataset, for automatically loading and preprocessing the data. The loader performs the following key operations:

1. Reading image and mask: The RGB image and its associated grayscale label mask are read from disk using OpenCV.
2. Resizing: The image and mask are resized to 256 x 256 pixels using bilinear interpolation to resize the images and nearest-neighbor interpolation for the masks to keep consistent input dimensions to save GPU memory.
3. Mapping classes: The grayscale mask values are mapped to integer class labels using a unique-value mapping approach so that each distinct pixel value corresponds to a class index (0 for background and 1 for solar panels).
4. Normalization and converting to tensor: RGB values are normalized to the [0, 1] range and converted to PyTorch tensors in the channel-first ordering (C, H, W). The label mask is converted to an integer tensor for compatibility with CrossEntropyLoss.

This preprocessing ensures consistent data representation and efficient loading during training and validation phases using PyTorch's DataLoader.

3.3 Model Architecture

The segmentation model follows a simplified U-Net architecture, optimizing both computation and segmentation accuracy. The model consists of a contracting path (encoder) and an expanding path (decoder) connected by skip connections to maintain spatial information throughout the layers. Each encoder block has been constructed through a DoubleConv module, which consists of two successive 3×3 convolutional layers, Batch Normalization, and a ReLU activation. There are a total of four DoubleConv blocks with progressively increasing feature dimensions (32,

64, 128, 256). Between the encoder blocks, the feature maps are downsampled with a 2×2 Max Pooling layer to incorporate context at multiple scales.

Similar to the encoder, the decoder consists of upsampling layers (employing bilinear interpolation) and concatenation of encoder features, which allows for the recovery of spatial resolution. The upsampling path consists of three DoubleConv blocks followed by a final 1×1 convolution that reduces the feature maps to the number of classes to predict (n_classes = 2). During the forward pass of the model, the encoder outputs are concatenated with the decoder layers to facilitate a fine localization of solar panels.

This architecture allows the network to effectively capture global context and local features — both of which are important to detect small rooftop panels and differentiate between adjacent structures, like roadways, rooftops, and vegetation.

3.4 Training Configuration

Training and validation utilized two disjoint subsets of the PV03 dataset. Memory limitations associated with high-resolution imagery solutions required a batch size of 1 and a fixed input size of 256×256 pixels. Training occurred for 50 epochs, using the Adam optimizer with a learning rate of 1×10^{-3}. Cross-Entropy Loss was utilized to optimize pixel-wise classification performance of the predicted segmentation-map and ground-truth mask.

The training procedure includes the following steps:

1. Forward propagation is done when input images are sent through the network.
2. Loss is computed between predictions and ground-truth masks.
3. Backward propagation using loss.backward() to compute gradients.
4. Weight updates via Adam optimizer.

5. After each epoch, validation is done to observe generalization using the same loss function above.

GPU acceleration was enabled via torch.device("cuda") if available, which allowed us to benefit from considerable speedup during training and evaluation. We printed the average training and validation losses after each epoch to confirm convergence, and we emptied the GPU memory at the end of each epoch with torch.cuda.empty_cache() to prevent memory overflow.

3.5 Model Saving and Evaluation

Upon completion of training, the final model weights were saved using:

```
torch.save(model.state_dict(), "saved_models/unet_solar_final.pth")
```

This allowed for both reproducibility and inference in the future without any additional training. The trained model was subsequently tested on previously unseen Google Maps imagery and received excellent metrics of performance:

Accuracy = 0.9758, Mean IoU = 0.9465, and Mean Dice = 0.9719 - indicating the model had effectively learned to recognize solar panel regions, despite complex backgrounds.

4 Evaluation

The proposed U-Net based segmentation model was assessed quantitatively and qualitatively, using the test split of the PV03 dataset. A complete suite of metrics were calculated to evaluate the accuracy of the panels detection at the pixel and class levels. Overall Accuracy, Intersection over Union (IoU), Dice coefficient, Precision, Recall, and F1-score were calculated to ensure that the evaluation

adequately measured the model's ability to classify both the solar panel pixels and the background regions.

4.1 Quantitative Evaluation

The performance of the model was very positive overall, as demonstrated by an Overall Pixel Accuracy of 0.9758, emphasizing that the majority of pixels were classified correctly over the entire data set. A Mean IoU of 0.9456 and a Mean Dice of 0.9719 further indicate the model's high segmentation consistency and capability to create masks closely aligned with the ground truth annotations.

A breakdown of performance by class indicates that the model performed exceptionally well on Class 0 (background) with an IoU of 0.9654, Dice of 0.9824, Precision of 0.9745, Recall of 0.9905, and F1-score of 0.9824. These findings indicate that the model rarely erroneously classified background pixels as solar panels, while also exhibiting high reliability across diverse geographic and environmental backgrounds.

Table 2 Evaluation metrics

Metric	Value
Overall Accuracy	0.9758
Mean IoU	0.9456
Mean Dice	0.9719

In Class 1 (solar panels), the model showed a high level of accuracy with an IoU of 0.9258, a Dice of 0.9615, a Precision of 0.9789, a Recall of 0.9447 and an F1 of 0.9615. Although detecting solar panels is a more challenging task because of the relative size, building roof texture variations, shadows, and occlusions, the model could still accurately localize and segment most of the regions with solar panels

with very few false positive or false negative detections. The decrease in Recall compared to Precision, suggests the model was moderated in the assigning location labels, generally more precise than aggressive detections.

Table 3 Per-Class Results

Class	IoU	Dice	Precision	Recall	F1
Class 0	0.9645	0.9824	0.9745	0.9905	0.9824
Class 1	0.9258	0.9615	0.9789	0.9447	0.9615

Additionally, these results demonstrate the proposed U-Net architecture achieves high quality segmentation performance in real-world settings such as solar panel mapping, monitoring infrastructure, and urban energy analysis.

4.2 Qualitative Evaluation

Alongside numerical assessments, qualitative evaluation was implemented by visually inspecting the original satellite images, ground truth masks, and predicted masks produced by the model's segmenting process. Findings indicated that predicted segmentation masks replicate solar panel regions that align closely with spatial shape and spatial position. In most instances, the model successfully identified solar panels and separated them from adjacent rooftop elements, both of which had similarly colored or textured attributes, indicating that the learned representations are robust.

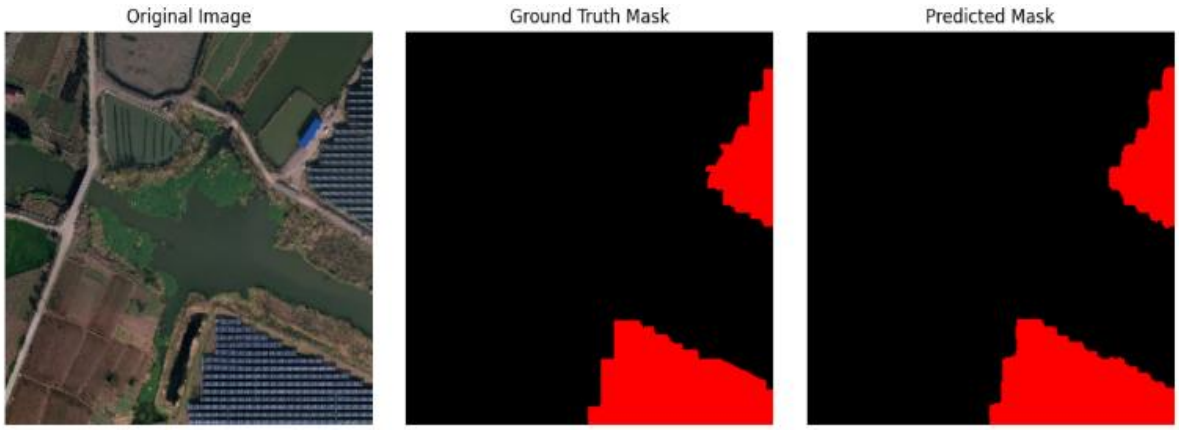

Figure 2

Qualitative observations also showed that the model generalized across multiple lighting conditions, panel orientations, and roof material types. Observed segmentation differences tended to occur with heavy shadowing or overlapping structures; however, overall predictions were highly congruent with the human-annotated ground truth.

5 Conclusions

In this study, we have outlined an automated framework for detecting and segmenting solar panels from Google Maps satellite imagery using a deep learning approach based on a customized U-Net architecture. The proposed technique showed significant ability to learn spatial and visual characteristics of photovoltaic installations and achieved impressive performance on the PV03 dataset with an Overall Accuracy of 0.9758, Mean IoU of 0.9456, and Mean Dice coefficient of 0.9719. The analysis of each class also established the confidence in the model, even when differentiating solar panel regions from complex rooftops.

Overall, these findings show that the simplified U-Net architecture, coupled with efficient preprocessing and effective training configurations, can generate high-quality segmentation results with some acceptable computational cost. This suggests

the possible applicability of such models for high spatial resolution mapping of solar energy potential, urban planning, and renewable energy infrastructure monitoring. The qualitative analysis, which involved visual inspection of the predicted masks, demonstrated the model's reliability across varying lighting conditions and rooftop configurations.

In spite of the high performance, there are still challenges—particularly in situations with heavy shadows, very small panel clusters, or low-contrast areas where panels are difficult to distinguish from adjacent structures. Future work can explore the possibility of improving detection accuracy through the implementation of attention mechanisms, multi-scale feature extraction, transformer-based architectures, or additional data modalities such as multispectral or LiDAR data. Extending the training datasets to a wider variety of geographic locations could also improve generalization.

In sum, the results presented here demonstrate the power and scalability of deep learning–based semantic segmentation methods to automate solar panel detection and support discussions at a global scale related to renewable energy assessment and planning.

References

[1] International Energy Agency (IEA). (2022). *World Energy Outlook 2022.*
[2] International Renewable Energy Agency (IRENA). (2023). *Renewable Capacity Statistics 2023.*
[3] Malof, J. M., Bradbury, K., Collins, L. M., & Newell, R. G. (2016). *Automatic detection of solar photovoltaic arrays in high resolution aerial imagery.* Applied Energy, 183, 229–240. DOI: 10.1016/j.apenergy.2016.08.191
[4] Jiang, H., Yao, L., Lu, N., Qin, J., Liu, T., Liu, Y., & Zhou, C. (2021). *Multi-resolution dataset for photovoltaic panel segmentation from satellite and aerial imagery.* Earth System Science Data, 13, 5389–5401. DOI: 10.5194/essd-13-5389-2021
[5] Ronneberger, O., Fischer, P., & Brox, T. (2015). *U-Net: Convolutional Networks for Biomedical Image Segmentation.* In MICCAI 2015, LNCS 9351. Springer. DOI: 10.1007/978-3-319-24574-4_28
[6] Parhar, P., Sawasaki, R., Todeschini, A., Reed, C., Vahabi, H., Nusaputra, N., & Vergara, F. (2022). *HyperionSolarNet: Solar Panel Detection from Aerial Images.* arXiv preprint arXiv:2201.02107. https://arxiv.org/abs/2201.02107
[7] Hou, X., Wang, B., Hu, W., Yin, L., & Wu, H. (2019). *SolarNet: A Deep Learning Framework to Map Solar Power Plants in China from Satellite Imagery.* arXiv preprint arXiv:1912.03685. https://arxiv.org/abs/1912.03685
[8] Olweus, E., & Mengshoel, O. J. (2024). *Detecting and Segmenting Solar Farms in Satellite Imagery: A Study of Deep Neural Network Architectures.* SAIS Conference, 2024.
[9] Jiang, H., Yao, L., Lu, N., Qin, J., Liu, T., Liu, Y., & Zhou, C. (2021). *Multi-resolution dataset for photovoltaic panel segmentation from satellite and aerial imagery.* Earth System Science Data, 13, 5389–5401. DOI: 10.5194/essd-13-5389-2021
[10] Wani, M. A., & Mujtaba, T. (2021). *Segmentation of Satellite Images of Solar Panels Using Fast Deep Learning Model.* International Journal of Renewable Energy Research, 11(1), 31–45. https://ijrer.org/ijrer/index.php/ijrer/article/view/11607
[11] Li, Y., Zhang, J., Liu, Y., & Wang, H. (2023). *Capacity Estimation of Solar Farms Using Deep Learning on High-Resolution Satellite Imagery.* Remote Sensing, 15(3), 765. DOI: 10.3390/rs15030765

[12] Alkhatib, M. Q., Al-Saad, M., Aburaed, N., Zitouni, M. S., Almansoori, S., & Al-Ahmad, H. (2025). *Enhancing Photovoltaic Panel Segmentation in Remote Sensing Imagery: A Comparative Study of Attention-Integrated U-Net Models.* ISPRS Annals, Vol. X-G-2025, pp. 63–72. https://isprs-annals.copernicus.org/articles/X-G-2025/63/2025/

[13] De Jong, T., Bromuri, S., Chang, X., et al. (2020). *Monitoring Spatial Sustainable Development: Semi-automated Analysis of Satellite and Aerial Images for Energy Transition and Sustainability Indicators.* arXiv preprint arXiv:2006.12345. https://arxiv.org/abs/2006.12345

[14] Kasmi, A., El Madhouni, M., Masood, M., Ben-Hamadou, A. and Rusk, L. (2022) A crowdsourced dataset of aerial images with annotated solar photovoltaic arrays and installation metadata. arXiv preprint arXiv:2209.03726.

[15] Elbl Droguett, S., & Nicollier, C. (2024). *Solar Panel Detection on Satellite Images: From Faster R-CNN to YOLOv10.* Stanford CS231n Project Report. https://cs231n.stanford.edu/2024/papers/solar-panel-detection.pdf

[16] García, G., Aparcedo, A., Nayak, G. K., Ahmed, T., Shah, M., & Li, M. (2024). *Generalized deep learning model for photovoltaic module segmentation from satellite and aerial imagery. Solar Energy, 274, 112539.* DOI: 10.1016/j.solener.2024.112539