

Using Small-Object Detection to Track Solar Irrigation Scaling: Opportunities and Limitations

Sarath Chandra Koppolu¹, Linda Steinhuebel-Rasheed², Lucie Maruejols¹

¹University of Kiel, Germany

²International Food Policy Research Institute, Cairo, Egypt / University of Göttingen, Germany

Introduction

Scaling technologies that improve rural livelihoods is a central focus of many development strategies. For examples, innovations in irrigation, renewable energy, digital services, and mechanization are increasingly promoted as pathways to raise agricultural productivity, enhance resilience, and reduce poverty. Governments, donors, and development partners invest substantial resources in piloting these technologies, with the expectation that successful interventions will be replicated and scaled across regions and countries.

Yet monitoring the scaling process remains a persistent challenge. Many development interventions are implemented as pilots or small-scale programs dispersed over wide geographic areas, and reliable information on where, how fast, and to what extent technologies are adopted is often limited. This gap is particularly important for interventions with potential trade-offs alongside benefits. In Egypt, for example, small-scale solar-powered irrigation systems can improve farm-level water access and reduce energy costs, but widespread adoption without regulation may increase unsustainable groundwater extraction.

Traditional field-based monitoring approaches (e.g., surveys, site visits, and administrative reporting) are costly, time-consuming, and difficult to maintain over large areas or long-time horizons. Policymakers and practitioners therefore often lack timely, spatially explicit evidence on technology adoption and its interaction with environmental and socioeconomic outcomes.

Recent advances in Earth observation provide promising alternatives. High-resolution satellite imagery, combined with growing computational power and machine learning techniques, allows the remote monitoring of development interventions that leave a visible footprint on the landscape. For example, small-scale solar irrigation systems are detectable in principle, which would enable large-scale, near-real-time tracking at a fraction of the cost of traditional data collection. However, technical challenges remain. These solar systems are often small

relative to the spatial resolution of widely available imagery, and reliably detecting them requires specialized small-object detection techniques. Accuracy can be limited in heterogeneous rural environments, and false positives or missed detections are common when training data are scarce.

In this study, we pilot a workflow in Fayoum, Egypt, using freely available high-resolution imagery and an iteratively expanded, custom-labeled dataset, to explore whether small-object detection can feasibly track solar-powered irrigation adoption. If feasible, this approach can provide a low-cost, scalable foundation for evidence-based policy. Beyond mapping adoption, the method also has potential to link solar irrigation detection to environmental and agricultural outcomes, such as vegetation dynamics, cropping intensity, or water use efficiency.

Data and Methods

Data

Mapping small-scale solar panel installations in Egypt faces a key challenge: the absence of openly available, high-resolution annotated datasets. Most existing research focuses on large commercial solar farms and relies on proprietary imagery (e.g., Maxar WorldView) or curated datasets developed outside Egypt. While these datasets are effective for small-object detection, none cover Egypt, and commercial imagery remains prohibitively expensive for large-area applications. Freely available satellite imagery, such as Sentinel-2 or Landsat accessed through Google Earth Engine (GEE), has spatial resolutions that are too coarse to detect household- or farm-level solar installations. Prior studies show detection accuracy declines sharply when objects occupy fewer than 8–15 pixels, and spatial resolutions of 0.5 m or finer are typically required. This limitation is particularly relevant in rural Egypt, where installations are small, dispersed, and visually distinct from the large commercial facilities usually studied.

In this pilot study, we focused on Fayoum Governorate to explore feasibility and manage computational intensity. Fayoum also offered the unique advantage that georeferenced solar panel installation data from a previous IFPRI randomized controlled trial (RCT) was available and provided a reliable set of seed points for identifying small-scale solar panels. Using these points, we manually extracted high-resolution image tiles from Google Hybrid basemap imagery within QGIS, forming the initial dataset (Dataset 1) with 102 annotated solar panel instances (Table 1).

To expand the dataset, we adopted an iterative, model-assisted approach. The intent of such procedure was to improve the model's ability to detect the small-scale solar panels. An initial YOLO-based object detection model was trained on Dataset 1, applied to additional high resolution Google Hybrid basemap imagery from QGIS covering agricultural areas in Fayoum, and its predictions manually inspected by visually checking each image tile. Correct detections corresponding to visible solar panels were retained and added to the dataset while false predictions were discarded. This procedure produced Dataset 2 with 252 labeled instances and repeat similar process to generate Dataset 3 with 360 labeled instances (Table 1). By restricting the analysis to Fayoum for the moment, we tried to capture consistent environmental, architectural, and land-use conditions while keeping computational demands manageable.

We divided each dataset into training, validation, and test subsets using a 70% / 15% / 15% split to balance the need for a sufficiently large training set with the requirement for independent evaluation.

Because the number of original solar panel observations was limited, we applied a controlled data augmentation strategy to increase the effective size of the training set and reduce overfitting. We performed random horizontal and vertical flips, in-plane rotations, zoom transformations, additive noise, and Gaussian blur. Flips and rotations help the model learn invariance to object orientation, important for rooftop and farm-level photovoltaic systems that appear at arbitrary angles. Noise and blur simulate sensor noise, atmospheric haze, and minor focus variations, common in operational imagery and known to reduce small-object detection accuracy if not represented during training. Augmentation was applied only to training images; validation and test datasets were left unchanged to ensure unbiased performance evaluation.

Table 1. Description of augmentation and split of datasets 1 to 3.

Data	Train	Val	Test
Dataset1	102 (17+85)	4	4
Dataset2	252 (42+210)	9	9
Dataset3	360 (60+300)	12	12

Model selection and training

Our objective was to detect small-scale solar panel installations and extract their geographic coordinates with high spatial accuracy. We therefore focused on YOLO (You Only Look Once) one-stage object detection models, widely used for their balance of accuracy and computational efficiency.

We adopted YOLOv10 as implemented in the *Ultralytics* framework, using both the medium-capacity (YOLOv10m) and larger-capacity (YOLOv10b) variants. YOLOv10 introduces architectural improvements that enhance efficiency and localization accuracy, including streamlined multi-scale feature fusion and a dual-assignment training strategy that reduces reliance on heuristic post-processing.

We initialized both models using official *Ultralytics* weights pre-trained on large-scale object detection datasets and fine-tuned them on our Fayoum datasets. YOLOv10m was trained on Dataset 1 and Dataset 2, as this variant balances performance and computational cost and can be deployed on modest hardware. After the dataset reached sufficient size and diversity, we trained Dataset 3 using YOLOv10b to maximize detection accuracy.

Training configuration

Training was conducted using the *Ultralytics* YOLOv10 framework with minor adjustments to accommodate dataset size and image resolution. We used a batch size of 8 to preserve spatial detail and ensure that small solar panels remained clearly visible. Models were trained for 100 epochs, allowing losses to stabilize without evidence of substantial overfitting.

We largely followed default *Ultralytics* hyperparameters but adjusted confidence and intersection-over-union (IoU) thresholds to prioritize recall. In the context of rural energy mapping, missing small solar installations was considered more problematic than producing a limited number of false positives.

Performance metrics

We evaluated model performance using standard object detection metrics, including precision, recall, and F1-score, across different confidence thresholds. These metrics capture the trade-offs between correctly identifying solar panels and avoiding false detections.

We summarized overall detection performance using mean Average Precision (mAP). Following common practice, we report mAP@0.5 (mAP50), calculated at a single IoU threshold of 0.5, and mAP@0.5–0.95 (mAP50–95), which averages performance across IoU thresholds from 0.5 to 0.95 in increments of 0.05.

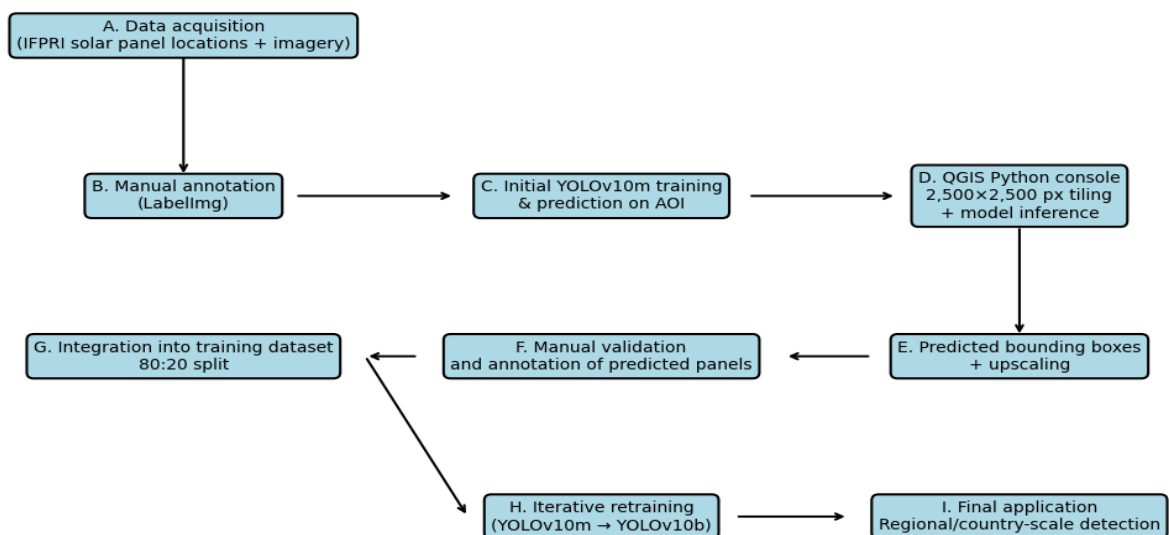


Figure 1. Iterative workflow for solar panel detection using YOLO and QGIS.

Workflow and deployment

We implemented the detection pipeline as an automated, end-to-end workflow that integrates high-resolution image extraction, model inference, validation, and geospatial output within QGIS (Figure 1). After training the final model, we applied it across tiled imagery covering Fayoum obtained from default Google Hybrid basemap and visually validated all detections to ensure geospatial reliability. This pilot deployment demonstrates that small-scale solar installations can be systematically identified using freely available basemap imagery and limited ground-truth data. While computationally intensive, the workflow is fully transferable and can be scaled to larger regions as additional computing resources and training data become available.

Results

Feasibility of Regional-Scale Photovoltaic Detection

Our study demonstrates the feasibility of detecting small-scale solar irrigation installations¹ across large geographic regions using freely available basemap imagery and an iteratively expanded, custom-labeled dataset. While object detection for photovoltaic infrastructure has been widely studied, applications focusing on very small installations and data upscaling detection across large rural areas using freely available basemap imagery remain limited in the existing literature. To our knowledge, this study represents an early attempt to work on such approach at regional scale. The primary aim was not to produce a production-ready, fully optimized detection model, but to illustrate a practical workflow for building an operational small-object detector in contexts where only limited labeled training data exist. Consequently, the results reflect both changes in detection accuracy across iterations and the utility of the proposed methodology for rapid, regional-scale photovoltaic mapping.

Applying the final YOLOv10b model to high-resolution imagery covering Fayoum, we generated a georeferenced layer of predicted solar panel locations (Figure 2). The final resulted in a few hundred detected installations across study area, clustered around agricultural zones, irrigation networks, and small farmsteads, reflecting the region's reliance on solar pumping systems. All detections were visually inspected one by one using individual image tiles to identify false positives; these false detections were primarily scattered over desert areas or water bodies. Overall, the map confirms that the pipeline can convert high-resolution imagery into actionable geospatial photovoltaic datasets. With expanded training data representing diverse rural and urban environments, the workflow could be extended to larger scale (governorate, national) photovoltaic infrastructure mapping, supporting energy planning and agricultural electrification assessments.

¹ Note however that we can not (yet) differentiate between solar pumps used to convey canal water and those used to extract groundwater.



Figure 3. Examples of prediction scenarios on test data with accuracy percentage.

Note however that some tiles produced no detections, while others included occasional false positives. These instances underscore the inherent challenge of detecting very small objects in heterogeneous rural environments. Nevertheless, the model consistently identifies true photovoltaic installations where present, validating the approach. Importantly, these results should not be interpreted as a final accuracy benchmark; instead, they demonstrate that regional-scale solar irrigation mapping is attainable once representative training data are incorporated, a goal achievable through the iterative workflow demonstrated here.

Conclusion

This study demonstrates that small-scale solar irrigation installations can be effectively detected across rural regions using high-resolution basemap imagery combined with an iteratively expanded, custom-labeled dataset. Our pilot in Fayoum shows that even with limited initial data, it is possible to develop a scalable workflow capable of converting satellite imagery into georeferenced layers of photovoltaic installations. This provides a practical pathway for monitoring the adoption of solar-powered irrigation technologies at regional scales, a critical step for supporting rural energy planning, agricultural electrification, and water resource management.

The approach highlights several opportunities for policymakers and researchers. First, it enables evidence-based tracking of solar irrigation expansion, identifying where smallholders are adopting photovoltaic pumps and how these systems cluster around agricultural and water infrastructure. Second, it lays the groundwork for integrating solar adoption data with multispectral or hydrological datasets, supporting analyses of crop productivity, water use efficiency, and spatial inequalities in irrigation access. Third, the modular and adaptable workflow can be extended to map other small-scale rural infrastructures, even in regions lacking commercial datasets.

At the same time, the study illustrates key limitations. Detection performance is currently constrained by the availability and diversity of labeled training data, computational requirements, and the resolution of freely available imagery. False positives in sparsely populated areas and occasional missed detections in heterogeneous landscapes underscore the need for careful validation. Scaling the workflow nationally will require additional high-quality data across diverse environments and seasonal conditions to en-

sure reliable mapping results. Furthermore, under the current approach, it is not yet possible to differentiate between solar pumps used to convey canal water and those used to extract groundwater—an important criterion for evaluating sustainability trade-offs.

In summary, small-object detection offers a promising tool for monitoring solar irrigation scaling and informing rural energy and agricultural policies. While technical and operational challenges remain, the methods demonstrated here provide a transparent, replicable framework that can be gradually expanded to generate actionable insights for policymakers, researchers, and development practitioners seeking to understand and support the diffusion of solar-powered irrigation technologies across Egypt and similar contexts.

ACKNOWLEDGMENTS

Funding for this work was provided by CGAIR's Scaling for Impact Science Program (S4I).

The Egypt Strategy Support Program (EGSSP) is managed by the International Food Policy Research Institute (IFPRI) and is financially supported by CGAIR's Scaling for Impact Science Program. This publication has been prepared as an output of Scaling for Impact Science Program and has not been independently peer reviewed. Any opinions expressed here belong to the author(s) and are not necessarily representative of or endorsed by IFPRI.

INTERNATIONAL FOOD POLICY RESEARCH INSTITUTE

A world free of hunger and malnutrition

IFPRI is a CGIAR Research Center

1201 Eye Street, NW, Washington, DC 20005 USA | T. +1-202-862-5600 | F. +1-202-862-5606 | Email: ifpri@cgiar.org | www.ifpri.org | www.ifpri.info

© 2026 International Food Policy Research Institute (IFPRI). This publication is licensed for use under a Creative Commons Attribution 4.0 International License (CC BY 4.0). To view this license, visit <https://creativecommons.org/licenses/by/4.0>.