

Panel Vision: Pipeline to Find Solar Feasible Properties

Unlock Solar Potential in Low-to-Moderate-Income Communities

Thibault Soubeste, Steve Zhou

Faculty Advisor: Allan Telio

PearlX Team: Tim Mac, Kristen Brown




Thibault Soubeste



Steve Zhou

Executive Summary

We built an end-to-end automated pipeline that compresses the 1-week manual rooftop solar screening process to 30 minutes, combining computer vision and optimization algorithms. Traditionally, this early-stage feasibility analysis required tedious work from four separate teams, but our AWS-based solution helps prioritize only the most promising sites, freeing team resources to focus where impact is maximal. This is especially critical for Low-to-Moderate-Income (LMI) communities, where thin margins have historically made solar development less feasible or scalable. With 11% higher solar efficiency from improved panel layout optimization, our pipeline empowers PearlX to identify more financially sustainable projects in LMI communities and expand equitable access to clean energy.



Tech Highlights

- Elevation + RGB rooftop segmentation
- Prompted image segmenting model
- Adaptive ViT-based obstacle detection
- Lightweight UI for expert intervention
- Per-segment solar yield modelling
- Modular, scalable AWS architecture

Scalability & Accuracy


Challenges for solar access in low-to-moderate-income (LMI) communities

- 7 days per property in current workflow
 - LMI projects have thin profit margin and must be evaluated carefully
- 439k+ LMI households in CA alone
 - Impossible to expand PearlX solution
 - Need scalability in screening LMI properties' solar feasibility
- Our solution
 - An AWS-deployed automated pipeline requiring minimal technical onboarding
 - Fast, scalable, accurate, and optimal

Our pipeline reduce per-project screening cost

- From **\$500** of labor to **\$0.86** of server, optionally \$11 manual quality check

Solar Screening Pipeline



Test: Dry Creek Apartments

PearlX's flagship project in 2024

- Home of 400+ families in Paso Robles, Central California
- Installed after latest Google Maps refresh (1-2 yr cycle)
- Panel-free imagery provides a "clean slate" for testing segmentation and obstacle detection algorithm

50+

Buildings with
distinct specs

<30

Minutes AWS
run-time

14%

**Higher system
efficiency**


9.4%

Higher ROI than
existing model

	System Size	System Efficiency
Current Assumption We Replaced	Unbounded	1,500 kWh/kW-yr
Google Maps Max Fit	3,536 kW	1,400 kWh/kW-yr
Our Pipeline Max Fit	1,099 kW	1,519 kWh/kW-yr
Our Pipeline Output Optimized	522kW	1,708 kWh/kW-yr

Above: Pipeline output, Layout efficiency comparison with Google Maps layout and Current Assumptions

Below: Part of Dry Creek Apartments satellite image, overlayed with optimized layout for human inspection




Computing Optimal Layout on Building

1. Building Identification and Preprocessing



We identify and isolate individual buildings from GeoTIFF Google Maps data. Based on the binary building mask, we crop and rotate each structure to ensure consistent axis-alignment for downstream roof analysis as well as efficient saving of intermediary results.

2. Roof Segmentation using Elevation




Using high-resolution Digital Surface Model (DSM) data, we apply the watershed algorithm on the Laplacian of DSM elevations to delineate distinct roof planes of each building. These preliminary elevation-based segments serve as informed prompts for segmentation models, capturing roof geometry critical to solar feasibility.

3. Segmentation Visual Enhancement




We strategically prompt the Segment Anything Model (SAM) with elevation-derived polygons to generate high-fidelity rooftop masks. This hybrid approach fuses 3D structural cues with state-of-the-art visual segmentation model, improving accuracy across varied roof textures and shade conditions.

4. Adaptive ViT-based Obstacle Detection




We extract patch-level DINOv2 visual features for each segment to identify rooftop obstacles such as vents and HVAC units. In the high-dimensional embedding space, we compute the distance of each patch to a dataset of expert-labelled obstacles, and if the distance is smaller than the tuned threshold, it's masked as obstacle. This approach enables adaptive learning based on human feedback in UI.

5. Max-Fit Layout Estimation



Given usable roof areas, we compute a "max-fit" layout of potential solar panels by accounting for panel size, spacing rules, and shading exclusions. Constrained by PearlX's internal engineering codes and regulations, this step provides an upper-bound estimation on solar capacity per roof segment in the real world.

6. Solar Layout Optimization



We query PVWatts API for accurate per-panel solar irradiance estimates and then using time-of-use hourly pricing, we solve a mixed-integer optimization problem to select the subset of panels that maximizes return of investment (ROI) for the entire project. This provides an evidence-based solar feasibility indicator for the multifamily communities for PearlX to target sales outreach.