SPOG: A GPU Rasterization Pipeline for Artifact-Reduced Photovoltaic Farm Orthomosaics

Udit Kandpal, Koustav Kumar Mondal, *Graduate Student Member, IEEE*, and Debasis Das, *Senior Member, IEEE*

***Abstract*—Orthomosaics generated from standard unmanned aerial vehicle (UAV) photogrammetric pipelines often contain rendering artifacts that degrade automated photovoltaic (PV) inspection by increasing false defect detections. Three failure modes are specific to repetitive planar PV surfaces: panel-surface waviness from structure-from-motion errors on specular targets, intertile seams from multiband blending, and phantom structures from posthoc sharpening. We propose solar panel optimized generator (SPOG), a graphics processing unit (GPU)-accelerated rasterization pipeline that addresses these failure modes through five domain-specific primitives operating at the rendering stage without altering upstream reconstruction geometry. Evaluated on a real utility-scale solar farm, SPOG reduces convolutional neural network (CNN) false positives by 63% and substantially improves grid-line fidelity relative to commercial and open-source baselines, demonstrating that rendering quality is a measurable and consequential factor in automated PV inspection. SPOG produces native-resolution 16-bit orthomosaics and integrates as a drop-in rendering stage for standard SfM/MVS reconstruction pipelines. The implementation is available at https://github.com/UditKandpal/SPOG**

***Index Terms*—Defect detection, graphics processing unit (GPU) rasterization, orthomosaic generation, solar farm inspection, unmanned aerial vehicle (UAV) photogrammetry.**

I. Introduction

UTILITY-SCALE photovoltaic (PV) installations now exceed 1 TW of global installed capacity [1], but cost-effective inspection depends on orthomosaics with radiometric fidelity and spatial accuracy that general photogrammetric pipelines fail to provide on specular PV surfaces. Standard pipelines [2], [3] rely on texture-atlas blending that introduces three systematic failure modes: *panel-surface waviness* from SfM errors on specular surfaces, *inter-tile seams* leaving discontinuities at panel edges where defect signatures are most informative, and *phantom structure* from aggressive posthoc sharpening that injects edge halos absent from the scene. Quantitatively, Agisoft detects 471 958 grid lines on a farm whose geometry supports at most 288 000 (derived from the documented panel count and per-module specifications; see Section IV), a 64% phantom inflation that substantially elevates false-positive rates in automated defect detection.

Existing orthomosaic generation methods do not address these rendering artifacts. Conventional image stitching [4], [5] fails under parallax, while PV inspection methods [6], [7] primarily focus on downstream defect analysis rather than rendering quality. We therefore propose solar panel optimized generator (SPOG), a graphics processing unit (GPU)-accelerated rasterization pipeline that reduces all three failure modes through five domain-specific primitives, operating solely at the rasterization stage and integrating seamlessly with any upstream SfM/MVS reconstruction (see Section III). Contributions include: 1) rendering-stage failure modes for specular PV surfaces; 2) GPU primitives addressing each without upstream modification; and 3) evaluation on a utility-scale farm.

TABLE I
UAV-Based Solar Farm Orthomosaic Generation: A Comparison

Method	Year	Rendering Focus	Defect Preserv.	Radiometric Consist.	Geometry Correct.	GPU Accel.	Bit Depth
Liao & Lu [11]	2021	Thermal + RGB	✓	✗	✗	✗	8-bit
Morando et al. [12]	2022	Thermal	✓	✗	✗	✗	8-bit
Akay et al. [13]	2024	Thermal	✓	✗	✗	✗	8-bit
Agisoft Metashape [2]	2024	General	✗†	✓*	✗	✓	8-bit
WebODM [3]	2024	General	✗‡	✓	✗	✓	8-bit
Vien et al. [14]	2025	LiDAR + Photo	N/A§	✓	✓	✓	16-bit¶
SPOG (Ours)	**2026**	**Solar**	✓	✓	✓	✓	**16-bit**

†Phantom grid (64% over max). ‡Over-smoothing (25% of physical max). *Gradient blur at tiles. §Non-solar (lagoon focus). ✓=Yes, ✓=Partial, ✗=No. ¶DEM elevation; orthomosaic RGB at 8-bit (Agisoft default).

II. Related Work

Traditional photogrammetric pipelines [2], [3] commonly employ multiband gradient blending [8] for texture-atlas composition. On solar farms, this introduces systematic artifacts, including phantom structures in Agisoft outputs, oversmoothing in WebODM outputs, and seam discontinuities on specular, low-texture surfaces. 2-D stitching approaches [4], [5] avoid 3-D reconstruction but fail under parallax around elevated structures. Table I summarizes key differences.

Domain-specific unmanned aerial vehicle (UAV) inspection frameworks [6], [7] and deep learning methods [9] build on orthorectified UAV imagery, focusing on downstream classification rather than rendering fidelity, underscoring the impact of rendering quality as evidenced by the 64% phantom inflation reported in Section I. Physics-based radiometric correction [10] addresses atmospheric compensation but not texture-atlas blending artifacts inherent to multiband atlas composition. SPOG targets the rendering stage through geometry-driven primitives (UV wrapping, confidence weighting, and hardware depth biasing) that improve radiometric consistency and spatial fidelity for downstream PV inspection and geospatial analysis.

Received 17 May 2026; revised 30 June 2026; accepted 27 July 2026. Date of publication 4 August 2026; date of current version 9 September 2026. This work was supported in part by the DST–NSF Project under Project DST/INT/USA/NSF-DST/Debasis/P-13/2024 and in part by the TIH IoT and IoE Project under Project TDP06-A-26. *(Corresponding author: Debasis Das.)*

Udit Kandpal and Debasis Das are with the Department of Computer Science and Engineering, IIT Jodhpur, Jodhpur 342030, India (e-mail: m24cse027@iitj.ac.in; debasis@iitj.ac.in).

Koustav Kumar Mondal is with the Center for AIoT and Applications, IIT Jodhpur, Jodhpur 342030, India (e-mail: mondal.4@iitj.ac.in).

Digital Object Identifier 10.1109/LGRS.2026.3720215

1558-0571 © 2026 IEEE. All rights reserved, including rights for text and data mining, and training of artificial intelligence and similar technologies. Personal use is permitted, but republication/redistribution requires IEEE permission.
See https://www.ieee.org/publications/rights/index.html for more information.

Authorized licensed use limited to: University of Melbourne. Downloaded on November 03,2026 at 13:55:42 UTC from IEEE Xplore. Restrictions apply.

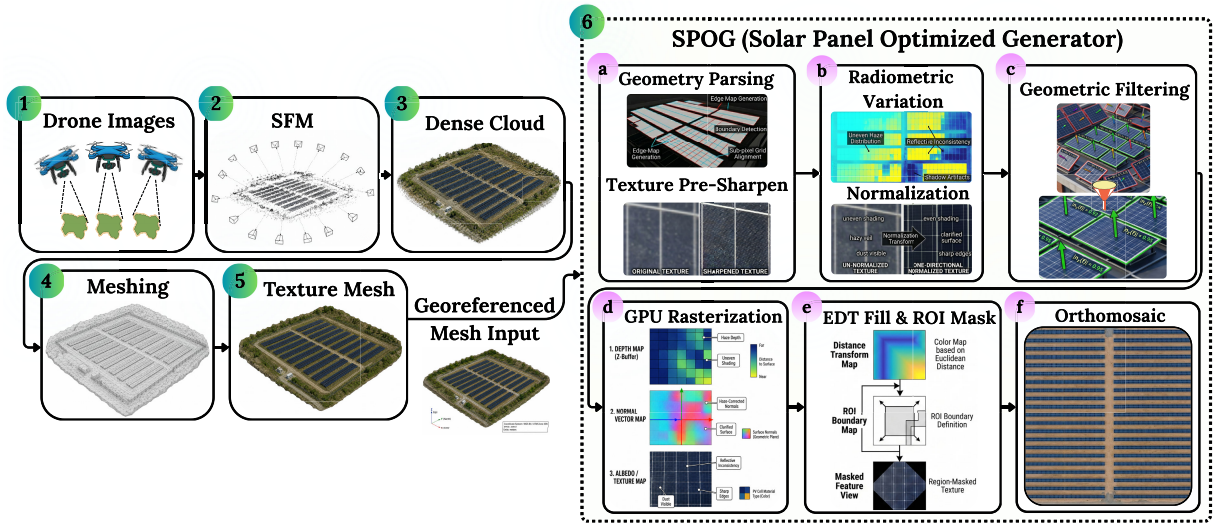


Fig. 1. SPOG end-to-end pipeline. After standard SFM/MVS reconstruction, the proposed five-stage module generates a 16-bit georeferenced orthomosaic optimized for solar panel analysis.

III. SPOG Methodology

This section describes the five-stage SPOG pipeline (see Fig. 1), detailing the geometry-driven primitives that distinguish it from texture-atlas methods. SPOG ingests an OBJ mesh and a GPS-calibrated texture atlas from any photogrammetry tool and produces a georeferenced 16-bit TIFF orthomosaic via five sequential stages: geometry parsing with texture presharpening, one-directional radiometric normalization, geometric filtering, hardware GPU rasterization, and EDT gap filling with ROI contouring (see Algorithm 1).

A. Stage 1–Geometry Parsing and Texture Presharpening

An ϵ-tolerant barycentric parser handles nonmanifold OBJ geometry. Each texture tile is presharpened using an SNR-adaptive unsharp mask to restore fine detail suppressed by MVS multiview averaging, avoiding the halo artifacts of posthoc sharpening.

B. Stage 2–One-Directional Radiometric Normalization

Global brightness variation across texture tiles causes spurious seam detections. SPOG applies asymmetric luminance lift in CIE-LAB space [see (1)], lifting only underexposed tiles halfway toward the scene median μ_{ref} while preserving high-contrast busbar and hotspot signatures

$$\text{Tile}_L \mathrel{+}= \text{clip}\left(0.5\,(\mu_{\text{ref}} - \mu_{\text{tile}}),\ 0,\ 255\right) \quad \text{if } \mu_{\text{tile}} < \mu_{\text{ref}}. \tag{1}$$

C. Stage 3–Geometric Filtering

Nonpanel surfaces (terrain, mounting structures) are excluded by retaining only mesh faces satisfying $|n_z(f)| \geq 0.25$ [see (2)], corresponding to maximum tilt $\approx 75°$ from nadir, ensuring rasterization operates exclusively on panel top-surfaces

$$|n_z(f)| \geq 0.25. \tag{2}$$

Algorithm 1 SPOG Algorithm

Require: OBJ Model (M), Texture Atlas (T), Target Resolution (S_{img})
Ensure: Normalized Orthomosaic (I_{out}), ROI Mask (M_{roi})

procedure GenerateOrthomosaic (M, T, S_{img})
  Stage 1: Data Parsing and Presharpening
  $Vertices, Faces, UVs \leftarrow$ ParseOBJ(M)
  $T \leftarrow$ TexturePreSharpen(T) ▷Restore signatures before blending
  Stage 2: One-Directional Radiometric Normalization
  $\mu_{\text{ref}} \leftarrow$ median$(\{\mu_{\text{tile}}$ for each tile $\in T\})$
  for each $Tile \in T$ **do**
    if $\mu_{\text{tile}} < \mu_{\text{ref}}$ **then** ▷Only lift darker tiles
      $shift \leftarrow 0.5 \cdot (\mu_{\text{ref}} - \mu_{\text{tile}})$
      $Tile_L \leftarrow \text{clip}(Tile_L + shift, 0, 255)$ ▷ Modify L-channel only
    end if
  end for
  Stage 3: Geometric Filtering and VBO Preparation
  $Faces_{\text{active}} \leftarrow \{f \in Faces \mid |n_z| \geq 0.25\}$ ▷Cull nonpanel surfaces
  UploadToGPU$(Vertices, Faces_{\text{active}}, UVs)$
  Stage 4: Hardware-Accelerated Rendering (GPU)
  Vertex Shader: Apply orthographic projection
  Fragment Shader:
    $u', v' \leftarrow \text{fract}(u, v)$ ▷Explicit wrapping (Eq. 3)
    $color \leftarrow$ BilinearSample(T, u', v')
    $conf \leftarrow |n_z|$ ▷Confidence from nadir alignment
    Output:$(conf \cdot color, conf)$ to float32 FBO
  Depth Test: `glPolygonOffset(1.0, 1.0)` ▷ Slope-dependent bias
  Stage 5: Reconstruction and ROI Contouring
  $I_{\text{raw}} \leftarrow \sum(conf \cdot color) / \sum(conf)$ ▷Normalize by weights
  $M_{\text{pre}} \leftarrow$ BinaryMask$(I_{\text{raw}} > 8)$ ▷Capture mesh footprint
  $I_{\text{fill}} \leftarrow$ UniversalEDTFill(I_{raw}) ▷Fill gaps via nearest neighbor
  $M_{\text{roi}} \leftarrow$ MorphClosing$(M_{\text{pre}}, 80\text{px})$
  $I_{\text{out}} \leftarrow I_{\text{fill}} \cdot M_{\text{roi}}$ ▷Apply hard-boundary mask
  $I_{\text{sharp}} \leftarrow$ DefectAwareSharpen$(I_{\text{out}}, \alpha = 0.3, \tau = 3)$
  return $I_{\text{sharp}}, M_{\text{roi}}$
end procedure

D. Stage 4–GPU Rasterization

Active faces are rendered via an OpenGL 3.3 Core headless context (NVIDIA RTX A4000 GPU, EGL backend), with three

Authorized licensed use limited to: University of Melbourne. Downloaded on November 03,2026 at 13:55:42 UTC from IEEE Xplore. Restrictions apply.

hardware-level innovations addressing panel-specific problems. First, explicit UV wrapping [see (3)] enforces panel texture periodicity at the fragment shader level, preventing subpixel UV errors from bleeding adjacent tile colors

$$(u', v') = (\text{fract}(u),\ \text{fract}(v)). \tag{3}$$

SPOG processes each OBJ material group with its own texture in a separate draw call; `fract(v_uv)` corrects subpixel UV overshoot at material-boundary seam vertices. Second, confidence-weighted accumulation [see (4)] weights each fragment by its nadir-alignment score in a float32 framebuffer so that oblique-view fragments contribute less

$$I_{\text{out}}(x, y) = \frac{\sum_f |n_z(f)| \cdot c(f)}{\sum_f |n_z(f)|}. \tag{4}$$

Third, slope-dependent hardware depth bias via `glPolygonOffset(1.0, 1.0)` substantially reduces Z-fighting on coplanar panel surfaces, providing consistent depth priority without per-face sorting (prohibitively slow for 300K+ face meshes).

E. Stage 5–EDT Fill, ROI Mask, Defect-Aware Sharpening

Internal rendering holes are filled using Universal Euclidean Distance Transform [15], propagating the nearest valid pixel color without Gaussian diffusion blur. The panel ROI mask is derived geometrically via morphological closing (80 px structural element), ensuring all quality metrics are computed strictly over active panel area. A threshold-gated unsharp mask [see (5)] is applied as the final step

$$I_{\text{sharp}} = I + \alpha \cdot \text{clip}\left(|\nabla^2 I| - \tau,\ 0,\ \infty\right) \cdot \nabla^2 I \tag{5}$$

with $\alpha = 0.3$, $\tau = 3$ DN, $\sigma = 0.8$ px. The threshold τ gates noise-floor gradients; $\tau = 3$ DN is grounded in the nominal sensor read-noise floor ($\sigma_{\text{read}} \approx 1$ –2 DN), below which gradients are indistinguishable from quantization noise, selectively enhancing genuine defect edges (microcracks, hotspot boundaries) while leaving uniform panel regions untouched.

F. Parameter Selection

Constants were selected on a 20-tile held-out validation subset: $\lambda = 0.5$ minimizes seam visibility; $|n_z| \geq 0.25$ corresponds to a 75° tilt limit consistent with fixed-axis PV geometry; ($\alpha = 0.3$, $\tau = 3$ DN, and $\sigma = 0.8$ px) maximize edge density while suppressing noise amplification; and the 80 px closing element matches typical interrow gaps at 2.5 cm/px GSD.

IV. Experimental Evaluation and Results

Experiments assess SPOG's ability to reduce the three failure modes identified in Section I across two datasets of increasing scale, using four complementary metric dimensions and an independent convolutional neural network (CNN)-based defect detector.

A. Datasets and Evaluation Protocol

1) Datasets: Dataset 1 is a real solar farm spanning 459.55 × 252.03 m, reconstructed from 200 DJI Phantom 4 RTK [16] drone images (RGB, nadir-to-15° oblique) using COLMAP [17] for sparse reconstruction and OpenMVS [18] for dense reconstruction, producing a 339 627-face OBJ mesh.

TABLE II

Comprehensive Evaluation: 200-Image Solar Farm. ↑ Higher is Better, ↓ Lower is Better. Bold = Best; † = Artifact-Inflated; ‡ = Sharpening Artifact; ∗ = Waviness-Masking

Metric	SPOG	WebODM	Agisoft
Output Specification			
Resolution (MP)	**147.2**	43.8	255.8
Bit depth	**16**	8	8
Render time (s)	**142.4**	1280	158
Pattern Preservation			
Edge Density ↑	**0.0254**	0.0309†	0.0163
Grid Lines Detected ↑	**262,143**	73,092	471,958‡
HFE ($\times 10^{12}$) ↑	**5.20**	1.60	13.64‡
Texture Uniformity ↓	**0.081**	0.062†	0.084
Image Quality			
Image SNR ↑	**4.071**	4.007	2.899
BRISQUE-NSS ↓	0.644	**0.467**	0.643
Radiometric Fidelity			
Seam ΔL (CIE-LAB) ↓	1.400	0.785	**0.479**
Intra-panel CV ↓	0.199	0.226	0.136∗
CNN Defect Detection (EfficientNet-B0, elpv)			
Average Precision ↑	**0.9505**	0.9274	0.9487
F1-Score ↑	**0.8754**	0.8377	0.8644
Precision ↑	**0.9489**	0.8716	0.8726
False Positives (count) ↓	**7**	19	20
Detect. Reliability Score ↑	**0.9389**	0.9151	0.9393
Z-Score Statistical Detection (detector-agnostic)			
Hotspot AP ↑	**0.8377**	0.8325	0.6318
Hotspot F1 ↑	**0.7619**	0.7536	0.6781
Micro-crack AP ↑	**0.4106**	0.4061	0.4048
Cross-Method Generalization (Train SPOG → Test Agisoft)			
AP ↑	**0.9433** (99.4% of Agisoft within-domain)		

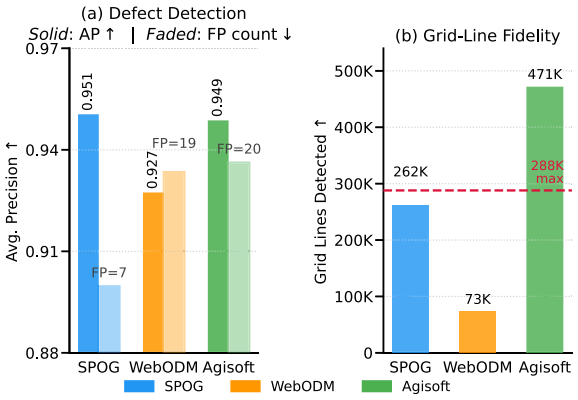


Fig. 2. Quantitative evaluation (200-image dataset). (a) SPOG achieves AP = 0.9505 with 63% fewer false positives (7 versus 19/20). (b) SPOG's grid-line count (262K) is within 9% of the physical maximum (≈288K); Agisoft inflates by 64% (phantom structure), WebODM recovers only 25%.

SPOG was run on an NVIDIA RTX A4000 GPU (EGL headless, OpenGL 3.3 Core) under Ubuntu 22.04; output: 16 384 × 8985 px (147.2 MP), 16-bit TIFF, 142.4 s. Agisoft Metashape Professional 2.1 and WebODM (OpenDroneMap) were run with default high-quality orthomosaic settings on the same host machine. Dataset 2 spans 1176 × 389 m (0.46 km^2) with 351 806 faces; SPOG output: 24 000 × 7945 px (190.7 MP), 16-bit TIFF, 175.6 s. The 500-image (Dataset 2) comparison in Fig. 3 is qualitative only, all quantitative results in Table II correspond to the 200-image Dataset 1.

2) Baselines and Metrics: WebODM [3] and Agisoft Metashape [2] receive the identical mesh and GPS metadata as

Authorized licensed use limited to: University of Melbourne. Downloaded on November 03,2026 at 13:55:42 UTC from IEEE Xplore. Restrictions apply.

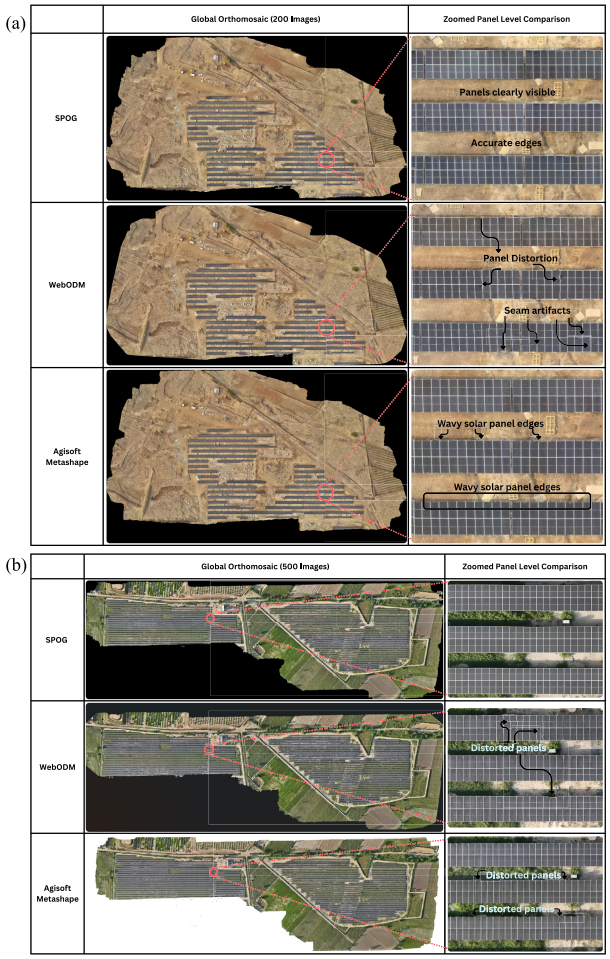


Fig. 3. Full-scene orthomosaic comparison. (Top to bottom) SPOG, WebODM, Agisoft. (a) 200-image dataset (459.55 × 252.03 m). (b) 500-image dataset (qualitative). SPOG output is uniform; WebODM shows lower resolution; Agisoft exhibits farm-scale waviness and sharpening artifacts.

SPOG. We acknowledge a key limitation of the comparison: the three pipelines produce outputs at different resolutions (SPOG: 147.2 MP/16-bit, WebODM: 43.8 MP/8-bit, Agisoft: 255.8 MP/8-bit) and bit depths. These differences may directly influence edge density, grid-line count, texture uniformity, SNR, BRISQUE-NSS, and detection performance. All image-quality and pattern-preservation metrics are computed over identical geographic ROIs (the common panel extent visible in all three outputs, extracted via the SPOG ROI mask projected onto each baseline) to minimize area effects; however, a strictly controlled comparison at a matched ground sample distance and bit depth is left as future work. Hough and Canny detection parameters are held constant across all three methods: Canny thresholds (50, 150) with a 3 × 3 aperture; probabilistic Hough transform with ρ = 1 px, θ = 1°, vote threshold = 50, minLineLength = 30 px, maxLineGap = 10 px. Evaluation spans four dimensions aligned with the failure modes from Section I: *pattern preservation* via Canny edge density, Hough grid-line count, and high-frequency energy; *image quality* via SNR and BRISQUE-NSS [19]; *radiometric fidelity* via CIE-LAB seam ΔL and the intrapanel coefficient of variation; and *defect detection* via EfficientNet-B0 [20] trained on elpv-dataset [9] morphologies with synthetic 5%–12% intensity darkening applied to UAV orthomosaic tile crops to simulate panel degradation in the UAV imaging domain. The elpv-dataset (2624 EL images; 1508 healthy, 1116 defective) provides spatial templates; masks are split 70/30 before extraction (781 train, 335 val masks). Each method trains an independent EfficientNet-B0 (AdamW, CosineAnnealingWR) on 1280 patches (150 × 150 px, 5%–12% intensity darkening) and is evaluated on 320 held-out patches.

B. Quantitative Results and Analysis

1) Defect Detection: SPOG achieves the highest CNN Average Precision (0.9505) and $F1$ -Score (0.8754), with only seven false positives versus 19 (WebODM) and 20 (Agisoft), a 63% reduction reflecting, as shown in Fig. 2(a), the preservation of genuine defect boundaries with substantially reduced blur and sharpening halos. Fisher's exact test (Bonferroni-corrected α = 0.0167) shows that this reduction is statistically significant versus Agisoft (p = 0.0145, OR= 0.320) and borderline versus WebODM (p = 0.0228; one-sided p = 0.0114); SPOG's precision Wilson CI [0.898, 0.975] does not overlap with WebODM [0.808, 0.916] or Agisoft [0.811, 0.916]. SPOG also achieves the second-lowest false alarm rate (FAR = 0.90%; WebODM: 1.33%, Agisoft: 0.53%), measured independently by a detector-agnostic z-score analysis without classifier involvement; achieving the highest CNN AP alongside a competitive FAR is consistent with SPOG encoding physically meaningful panel surface information rather than rendering artifacts. Cross-method generalization (AP = 0.9433 on Agisoft imagery without retraining, 99.4% within-domain retention) suggests that SPOG may encode genuine panel surface properties rather than pipeline-specific rendering patterns, though definitive validation requires geo-referenced ground-truth defect labels, such as coregistered EL imagery or multitemporal acquisition.

2) Context for Agisoft and BRISQUE-NSS Metrics: Agisoft achieves better seam ΔL (0.479 versus 1.400) and intra-panel CV (0.136 versus 0.199), and WebODM achieves a lower BRISQUE-NSS (0.467 versus 0.644)—these metrics do not uniformly favor SPOG. Agisoft's low seam ΔL reflects gradient-domain blending that smooths tile boundaries but attenuates defect signatures spanning them; its low CV reflects waviness bowing that makes local patches appear uniform while the global surface is distorted. WebODM's lower BRISQUE-NSS reflects its smoother output, which suppresses the high-frequency defect content quantified by edge density and grid-line metrics.

3) Pattern Fidelity and Operational Implications: SPOG's grid-line count tracks the physical module geometry within 9%, whereas WebODM underrecovers by 75% and Agisoft overrecovers by 64% [Fig. 2(b)]. Operationally, the 63% false-positive reduction reduces unnecessary maintenance dispatches. Georeferencing accuracy is not evaluated here; RTK or GCP validation is needed before deployment in multiepoch change-detection workflows.

4) Throughput and Visual Fidelity: SPOG renders in 142.4 s (versus WebODM's 1280 s), scaling near-linearly with mesh face count, while achieving an SNR of 4.07 versus Agisoft's 2.90 (+40%), showing that throughput gains do not compromise image quality. Fig. 3 shows farm-scale visual trends across both datasets, and Fig. 4 shows cell-level detail where SPOG's 2.81 cm/px GSD preserves panel boundaries blurred at WebODM's 5.06 cm/px output. The detector-agnostic z-score analysis corroborates visual fidelity: SPOG achieves hotspot AP = 0.8377 versus Agisoft's 0.6318,

Authorized licensed use limited to: University of Melbourne. Downloaded on November 03,2026 at 13:55:42 UTC from IEEE Xplore. Restrictions apply.

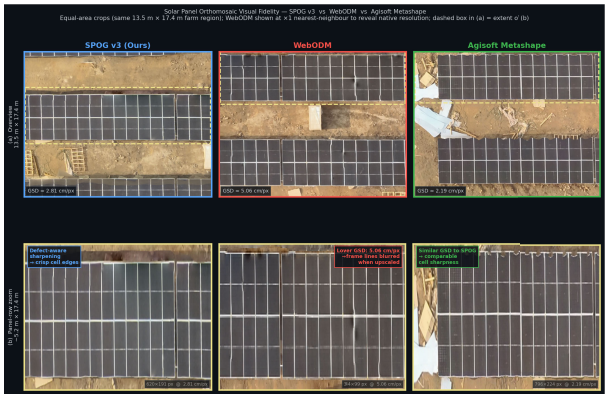


Fig. 4. Cell-level comparison (200-image dataset). (Left to right) SPOG, WebODM, Agisoft.

TABLE III

ABLATION–GLOBAL-METRIC CONTRIBUTIONS PER STAGE. Z-DET: Z-SCORE ANOMALY BLOB COUNT ($|z| > 2.5\sigma$); *PHANTOM DETECTIONS (HIGHER IS WORSE); †NOISE-INFLATED ED

Variant	ED↑	GL↑	Z-Det	t (s)
SPOG-Full	**0.2078**	**203,559**	**1,884**	**46**
w/o Pre-Sharp	0.1927	191,161	1,849	30
w/o Defect-Sharp	0.2034	199,993	1,868	45
w/o EDT-Fill	0.2077	202,454	1,982*	43
CPU baseline	0.2426†	194,955	3,541*	426

consistent with phantom structure suppression without classifier involvement.

5) Ablation Study: Table III quantifies SPOG stages on the 200-image dataset (8192×4492 px). Presharpening shows the largest global effect: removing it reduces real grid lines by 12 398, and 35 genuine z-score detections fall below threshold; SNR rising from 8.19 to 9.14 when disabled is consistent with structure recovery rather than edge amplification. EDT fill prevents 98 phantom detections from unfilled black-pixel gaps. The CPU baseline's higher ED (0.2426†) reflects aggressive postcomposition noise sharpening—GL is 4.4% lower and phantom detections are 88% higher than SPOG-Full at 9.2× the run time. Depth bias, geo-cull, and confidence weighting address localized artifacts (0.07%–0.71% of pixels, max 203 DN) that are diluted in farm-scale metrics; radiometric normalization reduces tile brightness discontinuities at seam lines, as confirmed visually when it is disabled.

V. CONCLUSION

We presented SPOG, a GPU rasterization pipeline that substantially reduces panel-surface waviness, intertile seams, and phantom structure through five domain-specific primitives applied at the rasterization stage, without altering upstream SfM/MVS geometry. Waviness is substantially reduced through nadir-alignment confidence weighting and slope-dependent depth biasing; seam discontinuities are substantially reduced through one-directional radiometric normalization (SNR 4.07 versus Agisoft's 2.90, + 40%); phantom structure is suppressed through threshold-gated sharpening, bringing the grid-line count within 9% of the physical maximum against Agisoft's 64% inflation. The combined effect is a 63% reduction in CNN false positives (7 versus 19/20) and cross-method generalization at AP = 0.9433, achieved 9× faster than WebODM while producing 16-bit output, establishing rendering quality as a measurable and consequential factor in automated PV inspection. Limitations include unmatched GSD and bit depth across baselines and reliance on synthetic defect labels; open directions include a matched-GSD evaluation, centimeter-level georeferencing via RTK or LiDAR [14], and extension to thin-film panels, tracker-mounted arrays, and thermal channels.

ACKNOWLEDGMENT

The authors gratefully acknowledge the financial support provided by these projects.

REFERENCES

[1] *Renewables 2024: Analysis and Forecasts to 2030*, International Energy Agency (IEA), Paris, France, 2024.
[2] Agisoft LLC.(2024). *Agisoft Metashape: Professional*. [Online]. Available: https://www.agisoft.com
[3] OpenDroneMap. (2024). *WebODM: Open-Source Photogrammetry and Mapping Platform*. [Online]. Available: https://www.opendronemap.org
[4] W. Qian, Y. Yang, Y. Xiao, K. Xu, S. Xie, and Y. Xie, "Enhanced incremental image stitching for low-altitude UAV imagery with depth estimation," *IEEE Geosci. Remote Sens. Lett.*, vol. 21, pp. 1–5, 2024.
[5] S. Bu, Y. Zhao, G. Wan, and Z. Liu, "Map2DFusion: Real-time incremental UAV image mosaicing based on monocular SLAM," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, Oct. 2016, pp. 4564–4571.
[6] M. Aghaei, F. Grimaccia, C. A. Gonano, and S. Leva, "Innovative automated control system for UAV missions in photovoltaic plant inspection," *IEEE Trans. Energy Convers.*, vol. 30, no. 2, pp. 534–543, Jun.2015.
[7] Y. Zefri, I. Sebari, H. Hajji, and G. Aniba, "Developing a deep learning-based layer-3 solution for thermal infrared large-scale photovoltaic module inspection from orthorectified big UAV imagery data," *Int. J. Appl. Earth Observ. Geoinf.*, vol. 106, Feb.2021, Art. no. 102652.
[8] P. J. Burtand E. H. Adelson, "A multiresolution spline with application to image mosaics," *ACM Trans. Graph.*, vol. 2, no. 4, pp. 217–236, Oct.1983.
[9] S. Deitsch et al., "Automatic classification of defective photovoltaic module cells in electroluminescence images," *Sol. Energy*, vol. 185, pp. 455–468, Jul.2019.
[10] R. A. Schowengerdt, *Remote Sensing: Models Methods for Image Processing*. Amsterdam, The Netherlands: Elsevier, 2006.
[11] K.-C. Liaoand J.-H. Lu, "Using UAV to detect solar module fault conditions of a solar power farm with IR and visual image analysis," *Appl. Sci.*, vol. 11, no. 4, p. 1835, Feb.2021.
[12] L. Morando, C. T. Recchiuto, J. Calla, P. Scuteri, and A. Sgorbissa, "Thermal and visual tracking of photovoltaic plants for autonomous UAV inspection," *Drones*, vol. 6, no. 11, p. 347, Nov.2022.
[13] S. S. Akay, O. Özcan, O. Özcan, and Ö. Yetemen, "Efficiency analysis of solar farms by UAV-based thermal monitoring," *Eng. Sci. Technol., Int. J.*, vol. 53, May2024, Art. no. 101688.
[14] B. S. Vien, T. Kuen, L. R. F. Rose, and W. K. Chiu, "Structural health monitoring of anaerobic lagoon floating covers using UAV-based LiDAR and photogrammetry," *Remote Sens.*, vol. 17, no. 20, p. 3401, Oct.2025.
[15] P. F. Felzenszwalband D. P. Huttenlocher, "Distance transforms of sampled functions," *Theory Comput.*, vol. 8, no. 1, pp. 415–428, 2012.
[16] *Phantom 4 RTK User Manual*, DJI Enterprise (SZ DJI Technol. Co., Ltd.), 2024.
[17] J. L. Schönbergerand J.-M. Frahm, "Structure-from-motion revisited," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Las Vegas, NV, USA, Jun. 2016, pp. 4104–4113.
[18] (2023). *OpenMVS: Open Multi-View Stereo Reconstruction Library*. [Online]. Available: https://github.com/cdcseacave/openMVS
[19] A. Mittal, A. K. Moorthy, and A. C. Bovik, "No-reference image quality assessment in the spatial domain," *IEEE Trans. Image Process.*, vol. 21, no. 12, pp. 4695–4708, Dec.2012.
[20] M. Tanand Q. Le, "EfficientNet: Rethinking model scaling for convolutional neural networks," in *Proc. Int. Conf. Mach. Learn.*, 2019, pp. 6105–6114.

Authorized licensed use limited to: University of Melbourne. Downloaded on November 03,2026 at 13:55:42 UTC from IEEE Xplore. Restrictions apply.