GeoGS: Geospatial Gaussian Splatting for Robust 3D Reconstruction from Sparse Satellite Imagery

Supplementary Material

A. Robustness of the Prior Extraction Pipeline

A potential concern regarding our prior extraction module is whether the heuristics are overfitted to the specific characteristics of the JAX and IARPA datasets (WorldView-3). To demonstrate the robustness and zero-shot generalizability of our pipeline, Fig. S1 visualizes the intermediate outputs across highly diverse urban morphologies and cross-sensor datasets, including SkySat imagery of Seville and Jamsil.

Crucially, all masks presented were generated using a single, fixed set of hyperparameters without any per-scene or per-sensor tuning. The successful disentanglement of static flat surfaces and dynamic shadows across these radically different domains strictly validates that our preprocessing pipeline learns generalized physical and semantic concepts.

B. Detailed Qualitative Results

Due to space constraints in the main paper, we provide comprehensive per-scene visual comparisons for all baselines and ablation variants across different view settings in this supplementary document.

B.1. Full-View 3D Reconstruction Comparisons

Figure S2 expands upon the results presented in the main text by providing a detailed qualitative comparison across all individual AOIs in the JAX and IARPA datasets under the full-view setting. The visual results clearly demonstrate the consistent superiority of our proposed GeoGS over existing NeRF and 3DGS baselines in reconstructing clean, high-fidelity surfaces and preserving structural boundaries without floater artifacts.

B.2. Comprehensive Sparse-View Comparisons

Figure S3 extends the visual analysis provided in the main text by presenting the complete set of qualitative comparisons across all baseline methods—including EO-NeRF, Sat-DN, Skyfall-GS, and EOGS—under extreme data-scarce conditions ($N \in \{3, 5, 7\}$).

This comprehensive breakdown further confirms that GeoGS is the only framework capable of preserving sharp building footprints and stable terrain structures, effectively preventing the catastrophic collapses and massive floater artifacts that plague existing approaches when the number of available views is extremely limited.

B.3. Comprehensive Ablation Study Visualizations

Figure S4 visually details the impact of each core component of our framework—specifically, the Depth Prior, Physics-based Surface Regularization, and Adaptive Curriculum—under both full-view and sparse-view settings. The severe geometric collapse and noisy surfaces observed when removing the depth prior or surface regularization, particularly in data-scarce conditions, rigorously validate their critical roles as geometric anchors.

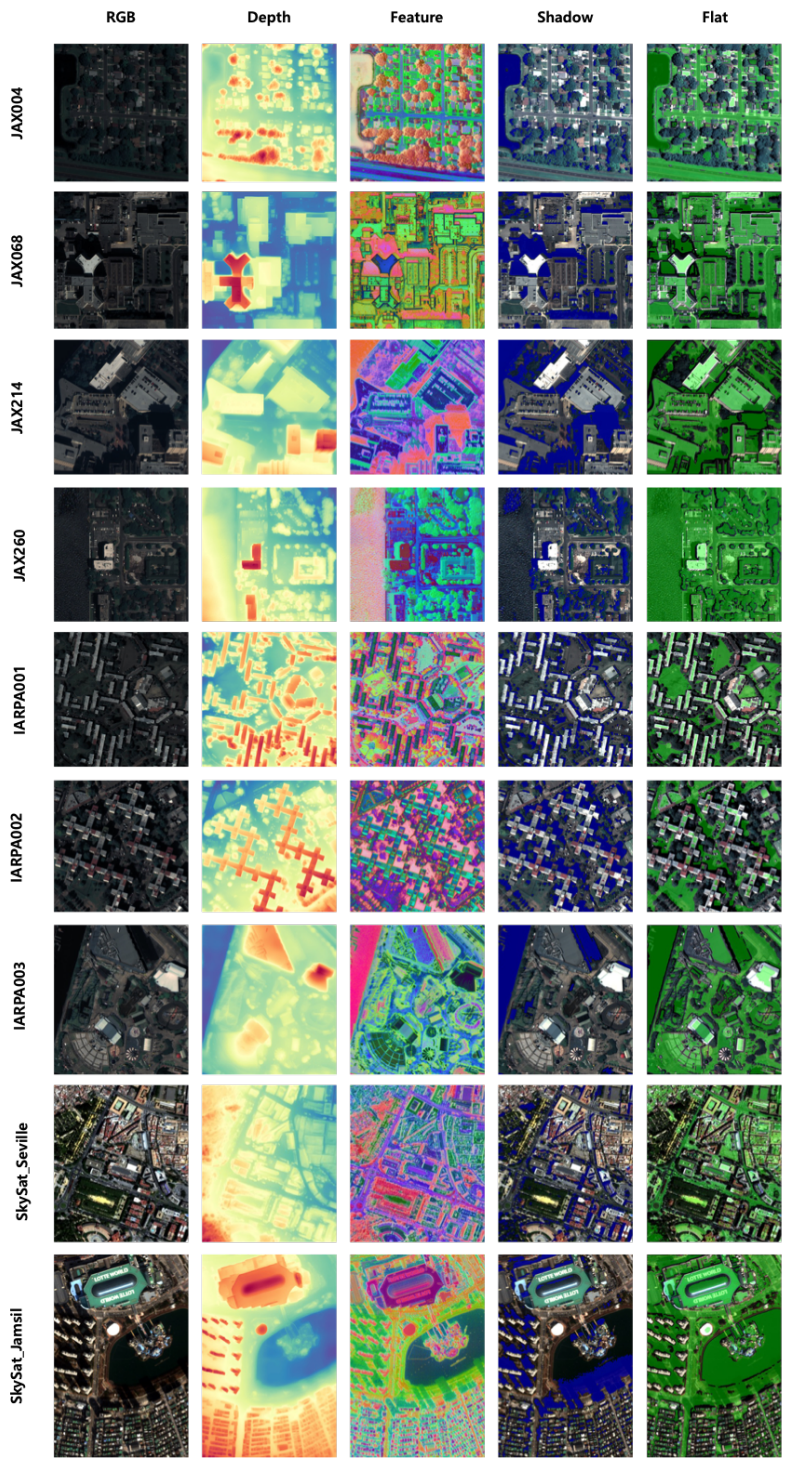

Figure S1. Robustness of Prior Masks across Diverse Domains. Visualization of the prior extraction pipeline on JAX/IARPA (WorldView-3) and Seville/Jamsil (SkySat). The consistent quality of the masks across different sensors—achieved with zero parameter tuning—demonstrates strong generalizability. Note on Shadow Mask (Blue): The shadow extraction intentionally adopts a conservative, high-recall strategy, occasionally encompassing low-reflectance surfaces such as water bodies. In 3DGS optimization, missing actual shadows causes severe geometric holes, whereas false positives (e.g., water) merely apply appearance conditioning without harming the structure. Any semantic ambiguity is systematically resolved downstream, as GeoGS independently extracts a text-prompted water mask to enforce explicit planarity on water surfaces. Note that the Feature column visualizes the 1024-dimensional DINOv3 semantic features reduced to RGB via PCA.

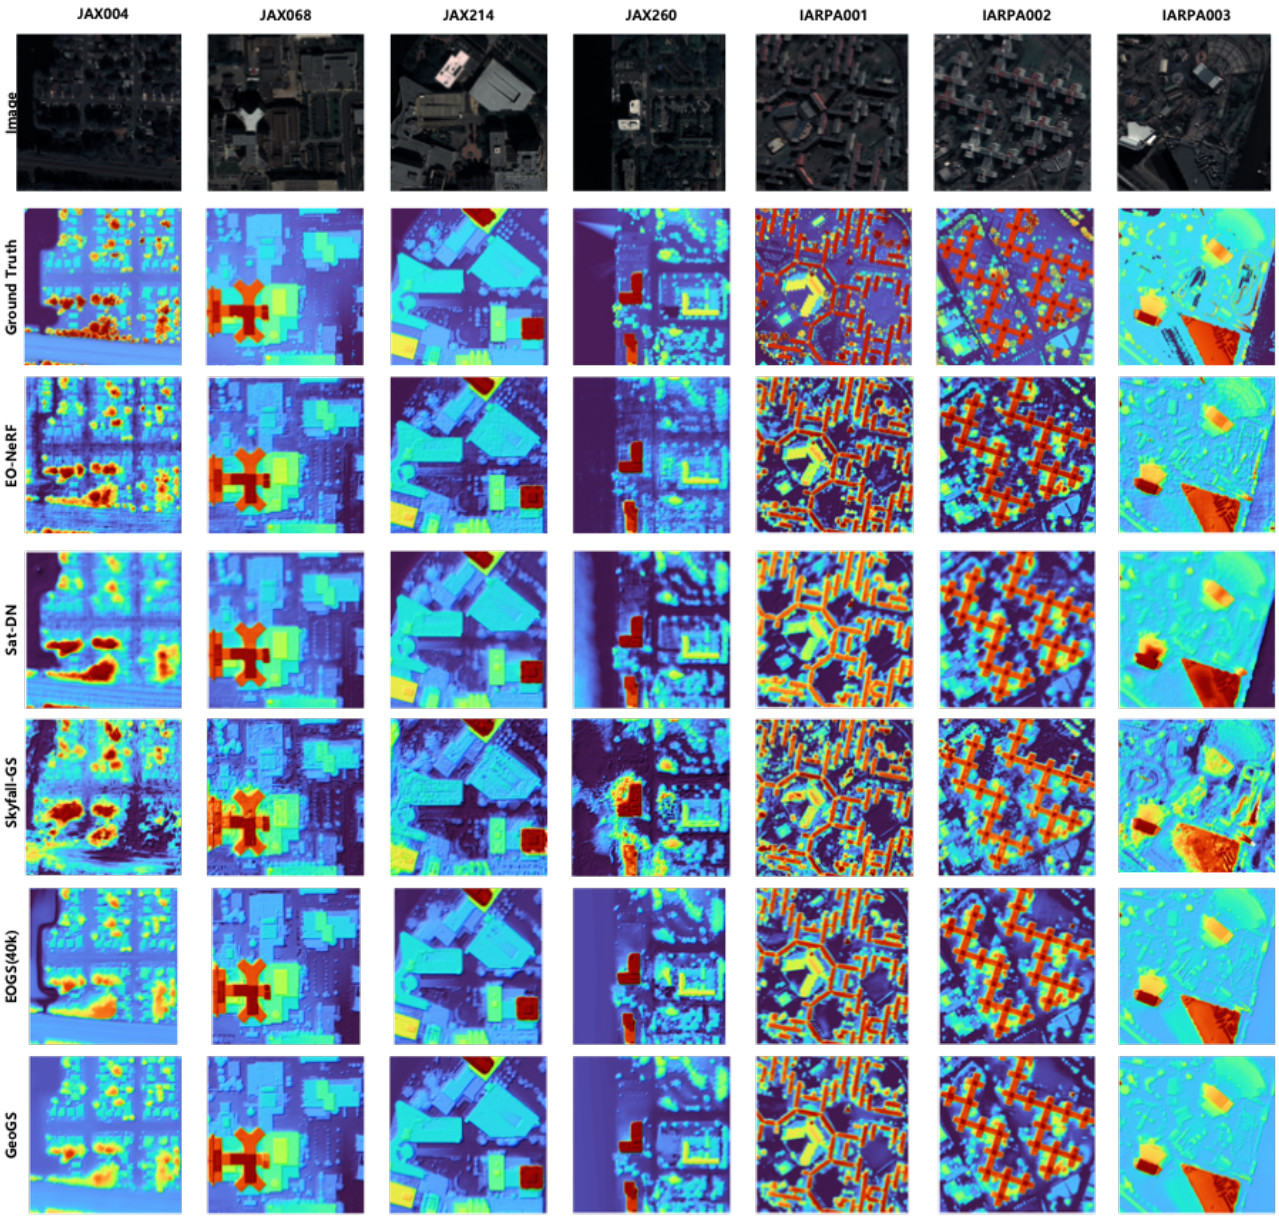

Figure S2. Detailed Per-Scene Qualitative Results for Full-View 3D Reconstruction. Comprehensive visual evaluation across all individual AOIs. Compared to baselines, GeoGS consistently produces cleaner DSMs with sharper building footprints and fewer geometric artifacts.

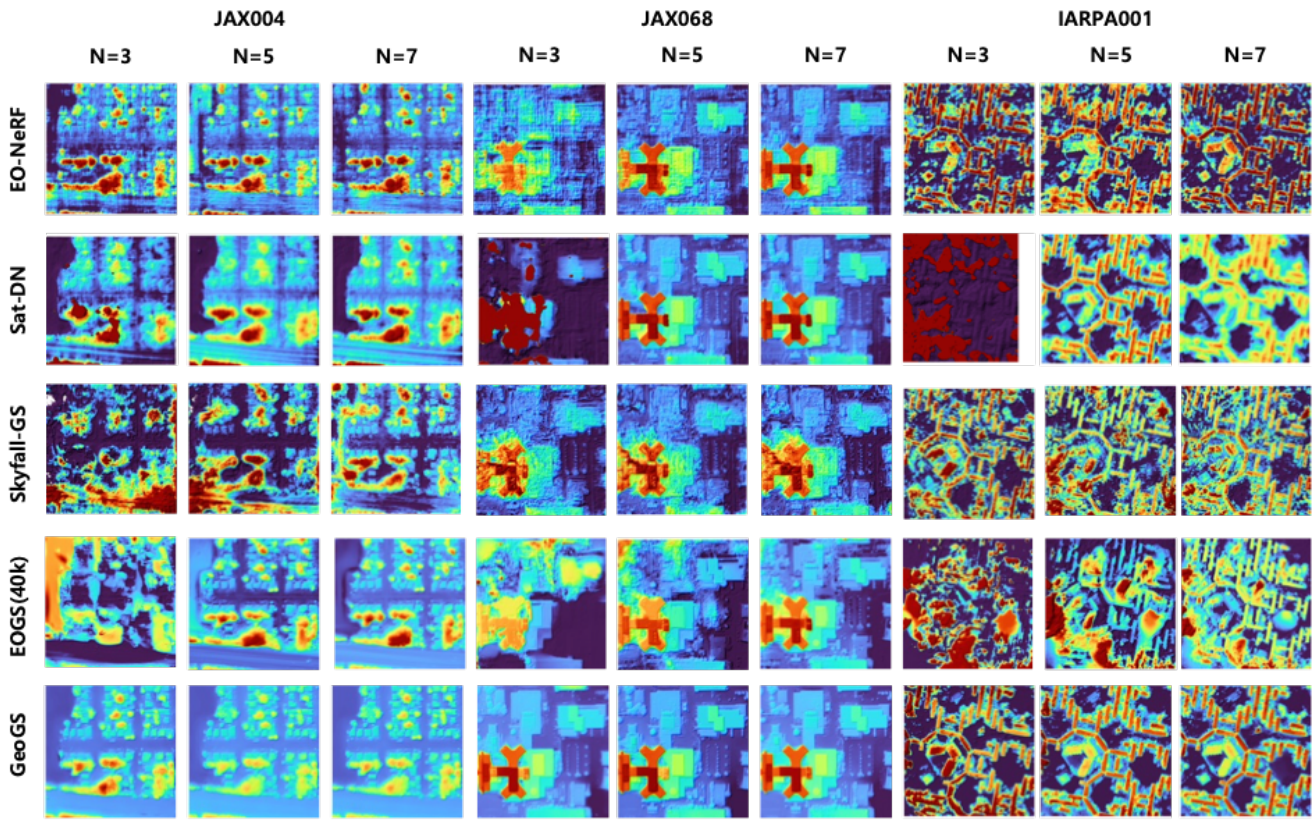

Figure S3. Comprehensive Qualitative Comparison for Sparse-View Reconstruction. Extended visual results including all evaluated baseline methods and our proposed GeoGS across varying sparse-view settings. The results clearly highlight the extreme vulnerability of existing methods to data scarcity, whereas GeoGS demonstrates unmatched geometric resilience and structural integrity.

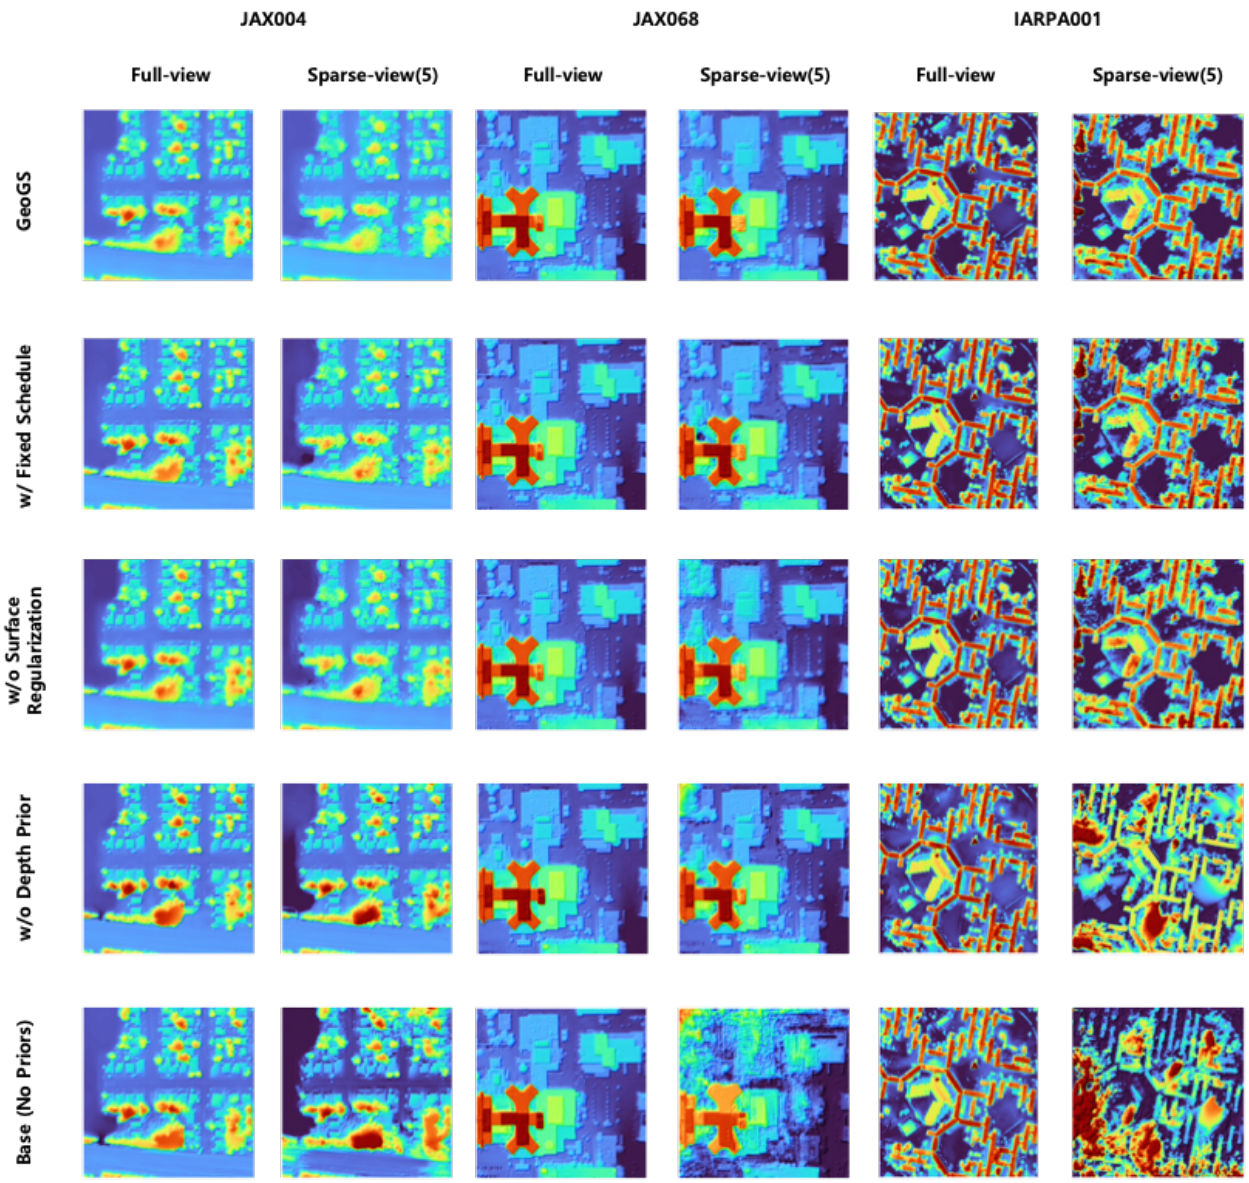

Figure S4. Comprehensive Qualitative Ablation Results. Per-scene visual evaluation demonstrating the critical contribution of each proposed module. The baseline without priors suffers from severe geometric degradation under extreme view scarcity, whereas GeoGS maintains robust structural integrity.