Edge Computing in Photogrammetry: Performance and Trade-offs of OpenDroneMap on NVIDIA Jetson Nano

Carlos Yan Matos Ferreira, Gustavo F. Freitas Sampaio, Antonio Alberi de Lima Júnior, Murilo Oliveira Bispo, Antonio Wendell de O. Rodrigues

[1]Laboratory of Innovation and Technology (LIT)
Federal Institute of Ceará (IFCE) – Fortaleza, CE, Brazil

{carlos.yan13, gustavo.fernandes62}@aluno.ifce.edu.br,

{alberi.junior07, murilo.oliveira.bispo03}@aluno.ifce.edu.br,

wendell@ifce.edu.br

Abstract. *Digital photogrammetry via OpenDroneMap (ODM) traditionally demands high-performance infrastructure, restricting geospatial analysis to centralized data centers. This work investigates whether NVIDIA Jetson Nano (4 GB RAM) enables on-site photogrammetric processing with acceptable quality. We benchmark ODM on 46 aerial images (181 MB) comparing Jetson Nano versus Intel Core i5 desktop and reveal: (i) 8.3x slowdown in Structure-from-Motion with 77% point cloud density reduction (1.9M vs. 8.3M); (ii) GPU utilization under 2% during costly matching phases, indicating CPU-GPU algorithmic misalignment; (iii) RAM saturation (80%) is the primary constraint, not thermal. Jetson Nano is viable for inspecting small areas (under 50 images, overlap over 85%), but scalability requires heterogeneous acceleration or larger memory capacity.*

Resumo. *Fotogrametria digital via OpenDroneMap (ODM) exige infraestrutura de alto desempenho, restringindo análise geoespacial a centros centralizados. Este trabalho investiga se NVIDIA Jetson Nano (4 GB RAM) viabiliza processamento on-site com qualidade aceitável. Benchmark de ODM em 46 imagens aéreas (181 MB) comparando Jetson Nano versus desktop Intel Core i5 mostra: (i) slowdown 8,3x em Structure-from-Motion com degradação de 77% em densidade de pontos (1,9M vs. 8,3M); (ii) utilização GPU menor que 2% em fases custosas, indicando desalinhamento algorítmico; (iii) saturação RAM (80%) é constrangimento primário, não térmico. Jetson Nano é viável para inspeção de pequenas áreas (menos de 50 imagens, overlap maior que 85%), mas escalabilidade exige aceleração heterogênea ou maior memória.*

1. Introduction

Digital photogrammetry reconstructs 3D scenes from 2D aerial imagery, enabling precision geospatial analysis essential for infrastructure monitoring, cadastral mapping, and environmental surveys [Westoby et al. 2012]. The OpenDroneMap ecosystem provides open-source photogrammetry pipelines [OpenDroneMap Authors 2026], reducing vendor lock-in compared to proprietary solutions (Pix4D, DroneDeploy).

However, ODM workflows impose severe computational demands: Structure-from-Motion (SfM) stages (feature extraction, matching, sparse reconstruction) and Multi-View Stereo (MVS) densification typically require 32+ GB RAM and GPU acceleration, restricting processing to fixed data centers. This centralization creates latency bottlenecks, dependency on network connectivity, and operational inflexibility in remote field settings.

Research Question: Can constrained edge devices (Jetson Nano, 4 GB RAM) execute photogrammetric workflows while preserving acceptable quality levels, potentially enabling on-site analysis and reducing dependence on centralized infrastructure?

This work evaluates the technical feasibility and performance trade-offs of ODM on NVIDIA Jetson Nano, a compact single-board computer (SBC) designed for embedded AI [NVIDIA Corp. 2026]. We establish a rigorous performance baseline and identify algorithmic bottlenecks that inform future edge photogrammetry research.

2. Related Work and Motivation

Modern photogrammetry relies on Structure-from-Motion (SfM) via OpenCV SIFT/ORB descriptors and Multi-View Stereo (MVS) densification. OpenDroneMap (ODM) integrates these into a unified Docker-based pipeline [OpenDroneMap Authors 2026], but assumes hardware with sufficient memory and GPU support. Prior work explores GPU acceleration in SfM [Lowe 2004], yet ODM's default CPU-centric workflow leaves GPU largely idle.

The Jetson Nano represents a practical edge alternative: quad-core ARM Cortex-A57, 4 GB shared LPDDR4, and Maxwell GPU (128 CUDA cores) [NVIDIA Corp. 2026]. Unlike datacenter approaches, edge execution reduces bandwidth and enables real-time field decisions [Shi et al. 2016]. However, literature lacks rigorous performance evaluation of photogrammetry on constrained edge platforms [Mittal and Vetter 2015].

This study contributes to this gap by providing an end-to-end benchmark of ODM on Jetson Nano, offering insights into speed, memory, and quality trade-offs. Our findings aim to support decision-making for practitioners and suggest directions for future optimizations (GPU-accelerated matching, memory-efficient MVS).

3. Materials and Methods

3.1. Dataset and Experimental Setup

Test dataset: 46 aerial images (181 MB total), native resolution 4608×3456 px, captured by Canon PowerShot A3300 IS. Images exhibit 85% forward overlap and nadir viewing angle, typical of precision agriculture and inspection workflows.

ODM Configuration:

- `--min-num-features 10000`
- `--feature-quality high`
- `--cog`
- `--orthophoto-resolution 3.0`
- Dense point target: ˜1.2M points

Platforms:

- **Edge (Target):** Jetson Nano 4 GB, 5W power mode, swap disabled [NVIDIA Corp. 2026].
- **Baseline:** IdeaPad laptop, Intel Core i5 (12th gen), 16 GB RAM, Intel Iris Xe GPU.

3.2. Profiling and Measurement

Docker container deployment: official `OpenDroneMap/odm` image [ODM Team 2026]. Metrics collected via `jtop` [Bonghi 2026] (1 Hz sampling): CPU/GPU frequency, RAM usage, core utilization, and temperature. Wall-clock times extracted from `stats.json` per pipeline stage.

4. Results

4.1. Performance Comparison: Speedup and Slowdown

The desktop completed the processing in 19m 50s, while the Jetson Nano took 2h 31m 42s (7.6× overall slowdown).

Table 1. SfM execution time (seconds) and speedup ratio

Stage	Nano (s)	Desktop (s)	Slowdown (x)
Feature Extraction	296.72	25.72	11.5
Features Matching	713.58	95.12	7.5
Sparse Reconstruction	254.02	30.06	8.4
SfM Total	**1284.71**	**153.61**	**8.3**

4.2. Hardware Utilization and Bottleneck Analysis

CPU: All 4 cores sustained 100% utilization throughout the pipeline.

GPU: Remained with $< 2\%$ utilization during 95% of the execution. Sparse activation (2–31%) appeared only during densification. This underutilization reveals that ODM's SfM path lacks GPU kernels for matching.

RAM: Peak consumption was 3194 MB (79.9%). Swap was disabled, forcing the densification algorithm to reduce quality to prevent OOM (Out of Memory).

Thermal: Temperature remained between 47–61°C. Computational saturation (CPU + RAM), and not thermal throttling, was the limiting factor.

4.3. Quality Degradation Due to Memory Constraint

The desktop generated 8.3M points; the Nano produced 1.9M (77% reduction). RAM saturation forced ODM to reduce densification recursion, an automatic and irreversible trade-off. Figure 1 illustrates this density difference visually.

5. Discussion and Insights

The GPU remains idle because ODM lacks GPU kernels for descriptors [Lowe 2004]. Quality loss (77% fewer points) is the cost of operating within 4 GB of RAM. Viability depends on application tolerance: structural inspection might accept 1.9M points, but cadastral mapping likely will not [ISPRS Authors 2025]. Upgrading to a Jetson Orin Nano (16 GB RAM) is recommended for scalability [Mittal and Vetter 2015].

Desktop: 8.3M points **Jetson Nano: 1.9M points**

Figure 1. Point cloud density comparison: Desktop (left) vs. Jetson Nano (right).

6. Conclusion and Future Work

The Jetson Nano (4 GB) is technically stable for small datasets (< 50 images) but limited by RAM. Future work includes porting ORB descriptors to cuDNN, benchmarking the Jetson Orin Nano, and evaluating memory-efficient MVS libraries like Colmap.

Acknowledgments

Work supported by IFCE and Polo de Inovação Embrapii.

References

Bonghi, R. (2026). jetson_stats: A control dashboard and python library for monitoring nvidia jetson devices. Accessed: May 2, 2026.

ISPRS Authors (2025). Validation of opendronemap as a leading open-source tool for photogrammetric processing. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLVIII-1-W4-2025:65–72.

Lowe, D. G. (2004). Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 60(2):91–110.

Mittal, S. and Vetter, J. S. (2015). A survey of cpu-gpu heterogeneous computing management notes. *ACM Computing Surveys*, 47(4):1–34.

NVIDIA Corp. (2026). Jetson nano developer kit technical specifications. Accessed: May 2, 2026.

ODM Team (2026). Docker hub repository: opendronemap/odm. Accessed: May 2, 2026.

OpenDroneMap Authors (2026). Opendronemap: The open source ecosystem for photogrammetry. Accessed: May 2, 2026.

Shi, W., Cao, J., Zhang, Q., Li, Y., and Xu, L. (2016). Edge computing: Vision and challenges. *IEEE Internet of Things Journal*, 3(5):637–646.

Westoby, M. J., Brasington, J., Glasser, N. F., Hambrey, M. J., and Reynolds, J. M. (2012). Structure-from-motion photogrammetry: A low-cost, effective tool for geoscience applications. *Geomorphology*, 179:300–314.