

## ENERGY EFFICIENCY IN PHOTOGRAMMETRY: A COMPARATIVE ANALYSIS OF DATASET, HARDWARE, AND RESOLUTION EFFECTS IN AGISOFT METASHAPE

Branimir Kolarek<sup>1</sup>, Davor Davidović<sup>1</sup>, Marko Maričević<sup>2</sup>

<sup>1</sup>Ruđer Bošković Institute, Bijenička cesta 54, 10000 Zagreb, Croatia

<sup>2</sup>University of Zagreb, The Faculty of Graphic Arts, Getaldićeva 2, 10000 Zagreb, Croatia

### ABSTRACT

Photogrammetry enables detailed 3D reconstruction from images but often involves computationally intensive workflows. While processing time is frequently analysed, energy consumption is a critical but less quantified factor, particularly concerning variations in datasets, hardware, and processing parameters. This study quantitatively evaluates the impact of dataset type (uniform Single-Source vs. heterogeneous Multi-Source), hardware platform (energy-efficient Apple Mac Mini M4 vs. high-performance Windows/NVIDIA RTX 4090), and input image resolution (4000px, 6000px, Native) on processing time and total energy consumption using Agisoft Metashape Professional. Experiments involved processing two 335-image datasets across the hardware and resolution configurations using a standard workflow with consistent 'Medium' quality settings, while monitoring time and energy usage. Results indicate the high-performance system was significantly faster (average 1.74x) but consumed substantially more energy (average 4.65x) than the energy-efficient platform. Energy consumption scaled approximately linearly with the total number of aligned pixels processed under constant settings. The Single-Source dataset demonstrated greater robustness, succeeding where the Multi-Source dataset failed in one low-resource scenario, although the Multi-Source dataset showed competitive or better efficiency in some successful tests on the energy-efficient hardware. These findings highlight crucial trade-offs between speed, energy efficiency, and robustness, providing empirical data to inform workflow optimisation based on specific project constraints.

**Keywords:** photogrammetry, Agisoft Metashape, energy consumption, processing time, hardware comparison, dataset quality, image resolution, Structure from Motion (SfM), energy efficiency.

### 1. INTRODUCTION

Image-based 3D reconstruction, commonly known through Structure from Motion (SfM) and Multi-View Stereo (MVS) techniques collectively referred to as photogrammetry, has become an indispensable tool across diverse fields such as cultural heritage documentation, topographic surveying, infrastructure inspection, visual effects, and game development (1). Software solutions like Agisoft Metashape Professional have emerged as industry standards, enabling users to generate detailed 3D models, orthomosaics, and digital elevation models (DEMs) from sets of overlapping photographs (2).

While these photogrammetric workflows offer powerful capabilities, they are often computationally intensive, demanding significant processing time and hardware resources (3). Traditionally, research and end-user focus have centred on optimising processing speed and geometric accuracy. However, with the growing emphasis on sustainable computing practices, rising energy costs, and the potential expansion of photogrammetry to resource-constrained or battery-powered platforms, understanding and quantifying the

energy consumption of these workflows is becoming increasingly critical (4). Factors such as the characteristics of the input dataset, the choice of processing hardware, and the selected input image resolution are known to influence processing time, but their specific impact on total energy usage within a defined workflow is less comprehensively documented.

Previous studies have often benchmarked photogrammetry performance primarily in terms of speed, particularly highlighting the advantages of utilising Graphics Processing Units (GPUs) for specific calculation stages (3,4). The influence of dataset characteristics, such as image count or overlap, on processing time and accuracy has also been investigated (5). However, a gap exists in quantitative, comparative data regarding the energy expenditure associated with varying fundamental aspects of the workflow: the inherent complexity and consistency of the input dataset (e.g., imagery from a single, uniform source versus imagery from multiple, mixed sources), the hardware profile (e.g., energy-efficient System-on-Chip vs. high-power discrete CPU/GPU combinations), and the resolution of the input images processed.

This study aims to address this gap by quantitatively evaluating the processing time and, crucially, the total energy consumption of a standard photogrammetric workflow in Agisoft Metashape Professional (Version 2.2.0 build 19890) under varying conditions. Specifically, the objectives are:

1. To compare the processing time and energy consumption required for a Single-Source Dataset (SSD), captured with uniform camera settings, versus a Multi-Source Dataset (MSD), composed of images from multiple cameras, lenses, and resolutions.
2. To compare the processing time and energy consumption between an energy-efficient hardware platform (Apple Mac Mini M4 with integrated GPU) and a high-performance workstation (Windows PC with Intel i9 CPU and NVIDIA RTX 4090 GPU).
3. To evaluate the effect of input image resolution (downscaled to approx. 4000px and 6000px wide, and native camera resolution) on processing time and energy consumption for both datasets and hardware platforms.
4. To analyse the interplay between these factors (dataset type, hardware, resolution) and their combined effect on the overall time and energy budget.

The contribution of this paper lies in providing empirical data on energy consumption for these common photogrammetry scenarios. These findings offer practical insights for users and researchers seeking to optimise workflows not only for speed and accuracy but also for energy efficiency, enabling more informed decisions regarding data acquisition strategies, hardware selection, and processing parameter choices.

## **2. METHODOLOGY**

This study employed an experimental approach to compare processing time and energy consumption across different dataset types, hardware platforms, and input resolutions within a standardized photogrammetric workflow using Agisoft Metashape Professional.

### **2.1. Datasets**

Two distinct datasets, each comprising 335 images of the same physical object but captured under different conditions, were used for this analysis.

- **Single-Source Dataset (SSD):** This dataset represents a controlled capture scenario. It consists of 335 images acquired using a single camera body (Sony ILCE-7RM3) equipped with a single prime lens (FE 35mm F2.8 ZA). All images share the same native resolution of 7952 x 5304 pixels (~42 MPix) and intrinsic camera parameters, simplifying the calibration process within Metashape. This dataset corresponds to the reports prefixed with "SSD-".
- **Multi-Source Dataset (MSD):** This dataset simulates a less controlled or more op-

opportunistic capture scenario, incorporating heterogeneity in the source data. It also consists of 335 images but utilizes a combination of two camera bodies (Sony ILCE-7RM3 and Canon EOS RP) and multiple lenses (Sony FE 24-105mm F4 G OSS zoom, Sony FE 35mm F2.8 ZA prime, and Canon RF 24-70mm range). Images within this dataset were captured at various focal lengths using the zoom lenses and possess different native resolutions, primarily 7952 x 5304 pixels (~42 MPix) and 6240 x 4160 pixels (~26 MPix), including both landscape and portrait orientations. This inherent variability necessitates the estimation of multiple distinct camera calibration groups by Metashape during processing. This dataset corresponds to the reports prefixed with “MSD-”.

## 2.2. Hardware Platforms

Two distinct computing systems were utilized to represent different hardware profiles:

- System 1 (Energy-Efficient): An Apple Mac Mini featuring an Apple M4 SoC (System on Chip) with integrated graphics and 16 GB of unified RAM, running Mac OS Sequoia 15.3.2 64 bit.
- System 2 (High-Performance): A custom-built Windows 11 PC equipped with an Intel Core(TM) i9-14900K CPU, 64 GB of DDR5 RAM, and utilizing an NVIDIA GeForce RTX 4090 discrete GPU alongside integrated Intel(R) UHD Graphics 770, though primary acceleration relies on the NVIDIA GPU, running Windows 64 bit.

## 2.3. Software

All processing was conducted using Agisoft Metashape Professional, version 2.2.0 (build 19890). GPU acceleration was enabled in Metashape preferences for all tests, utilizing the available capabilities of each hardware platform (Metal on Mac M4, CUDA on NVIDIA RTX 4090).

## 2.4. Experimental Design

A full factorial experimental design was implemented, crossing the two dataset types (SSD, MSD), the two hardware platforms (Mac M4, Win/RTX4090), and three input resolution levels:

- 4000px: Images downsampled prior to processing so the longer dimension was approximately 4000 pixels (e.g., 4000x2668 or 4000x2666).
- 6000px: Images downsampled prior to processing so the longer dimension was approximately 6000 pixels (e.g., 6000x4002 or 6000x4000).
- Native: Images processed at their original captured resolution. For the SSD, this was 7952 x 5304 pixels. For the MSD, this involved the mix of native resolutions (primarily 7952 x 5304 and 6240 x 4160 pixels) present in that dataset.

—> Datasets:

<https://e.pcloud.link/publink/show?code=kZhqudZCjKAksIqMmkS6dMTiJyxsBIWiLYV>  
This resulted in 12 distinct test runs, corresponding to the reports analysed: MSD-MAC-4000PX, MSD-MAC-6000PX, MSD-MAC-NATIVE, MSD-RTX4090-4000PX, MSD-RTX4090-6000PX, MSD-RTX4090-NATIVE, SSD-MAC-4000PX, SSD-MAC-6000PX, SSD-MAC-NATIVE, SSD-RTX4090-4000PX, SSD-RTX4090-6000PX, SSD-RTX4090-NATIVE.

—> Reports:

<https://e.pcloud.link/publink/show?code=kZnqudZkunLlcRIQ1pNLiVcWW9ajjLTzi0X>  
It is noted that the MSD-MAC MINI M4-4000PX run failed to produce a usable 3D model and is therefore excluded from comparative performance analysis where applicable. The circumstances of the failure are considered in the Discussion section.

## 2.5. Processing Workflow and Parameters

A standard photogrammetric workflow was executed in batch (all steps are run one after the other automatically, without interaction of the user.) for each test run within Metashape, consisting of the following steps:

1. Load Images
2. Align Photos
3. Build Depth Maps
4. Build Dense Cloud
5. Build Mesh (from Depth Maps)
6. Build Texture

Key processing parameters were consistent across all runs:

- Align Photos: Accuracy set to Medium. Generic Preselection: Yes. Key point limit: 40,000. Tie point limit: 4,000.
- Build Depth Maps: Quality set to Medium. Filtering mode: Moderate.
- Build Dense Cloud: Derived directly from Depth Maps (Source data: Depth maps).
- Build Mesh: Source data: Depth Maps. Surface type: Arbitrary. Interpolation: Enabled.
- Build Texture: Mapping mode: Generic. Blending mode: Mosaic. Texture size: 8,192 x 8,192.

### 2.6 Data Acquisition

- Processing Time: Timestamps for the completion of major processing stages (Matching, Alignment, Depth Maps, Dense Cloud, Mesh Reconstruction, UV Mapping, Texturing) were extracted directly from the “Processing Parameters” section of the generated Metashape PDF reports. Total processing time was calculated by summing the durations of these key stages.
- Energy Consumption: Total energy consumed for the entire Metashape processing workflow (from the start of Align Photos to the end of Build Texture) was measured using an external Commel digital power meter, recording cumulative Watt-hours (Wh).
- Quality Metrics: Geometric accuracy indicators (RMS Reprojection Error in pixels) and output model complexity metrics (Dense Cloud points, Mesh faces) were extracted from the reports for comparative analysis.
- Input Data Size: Total input pixels for each dataset at each resolution level were calculated based on the image dimensions and the number of aligned images reported for each specific run.

## 3. RESULTS

The processing time, total energy consumption, and key quality metrics were recorded for each test run. The MSD-MAC MINI M4-4000PX configuration, utilizing the Multi-Source Dataset (MSD) on the Mac M4 at the lowest resolution, failed during the alignment stage and did not produce a usable 3D model; its energy consumption up to the point of failure (14 Wh) is noted, but it is excluded from time-based performance comparisons. The primary results for the 11 successful test configurations, plus the energy data for the failed run, are summarised in Table 1. Total processing time represents the sum of durations for Alignment (Matching + Alignment), Depth Maps, Dense Cloud, Mesh (from Depth Maps), and Texture (UV Mapping + Blending) stages as reported by Metashape. Total Aligned Pixels were calculated based on the resolution and number of successfully aligned images for each run. Energy consumption was measured using a Commel digital power meter logger.

**Table 1: Summary of Processing Time, Energy Consumption, and Output Metrics**

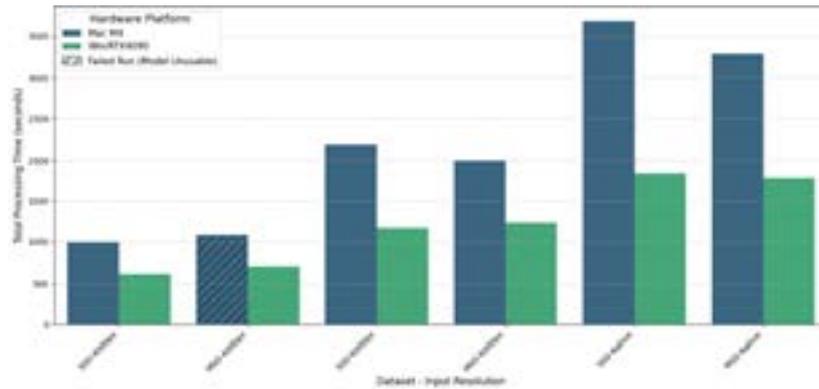
Test ID	Aligned Images	Total Aligned Pixels (Gpix)	RMS Error (pix)	Dense Points (Millions)	Mesh Faces (Millions)	Total Time (H:M:S)	Total Energy (Wh)
SSD-MAC-4000PX	335	3.58	2.09	2.75	0.12	0:16:36	12
SSD-MAC-6000PX	335	8.04	2.58	5.17	0.28	0:36:20	28
SSD-MAC-NATIVE	335	14.11	3.17	9.74	0.45	1:01:24	47
SSD-RTX4090-4000PX	335	3.58	2.10	2.81	0.12	0:10:19	63
SSD-RTX4090-6000PX	335	8.04	2.59	5.26	0.26	0:19:28	123
SSD-RTX4090-NATIVE	335	14.11	3.19	9.12	0.40	0:30:42	194
MSD-MAC MINI M4-4000PX*	329*	3.52*	1.70*	19.62*	0.52*	0:18:02*	14*
MSD-MAC-6000PX	331	7.96	1.95	10.36	0.67	0:33:17	26
MSD-MAC-NATIVE	330	12.81	2.35	14.79	0.82	0:54:46	43
MSD-RTX4090-4000PX	330	3.54	1.79	4.71	0.27	0:11:52	76
MSD-RTX4090-6000PX	330	7.96	1.95	9.61	0.61	0:20:33	130
MSD-RTX4090-NATIVE	330	12.81	2.35	18.24	0.80	0:29:44	193

\*Processing completed according to report, but the resulting 3D model was unusable/broken as per experimental observation (Figure 8.).

### 3.1. Impact of Hardware Platform on Speed and Energy Consumption

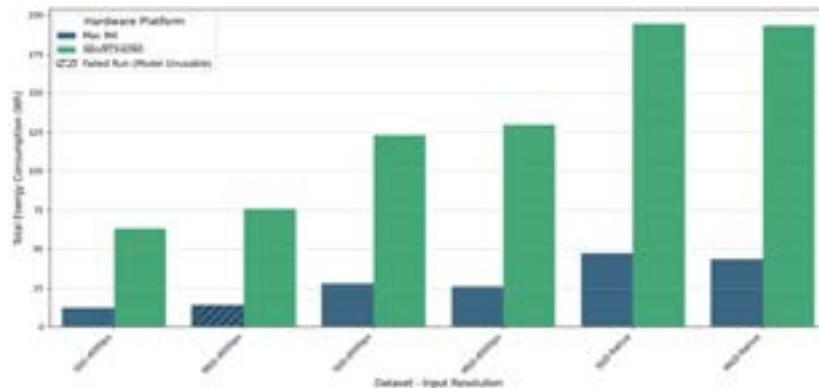
Significant differences in processing speed and energy usage were observed between the energy-efficient Mac M4 and the high-performance Win/RTX4090 systems. The Win/RTX4090 platform consistently completed the photogrammetric workflow faster than the Mac M4 across all comparable successful runs (Figure 1). The average speedup was substantial, with the Win/RTX4090 system being approximately 1.74 times faster on average. The largest speedup (2.15x) was observed for the SSD at Native resolution, while the

smallest (1.20x) occurred for the MSD at 6000px resolution.



**Figure 1: Processing Time Comparison by Hardware.** This figure shows the total processing time for each configuration, grouped by Dataset/Resolution, with bars colored by Hardware, includes the failed run (marked with hatching).

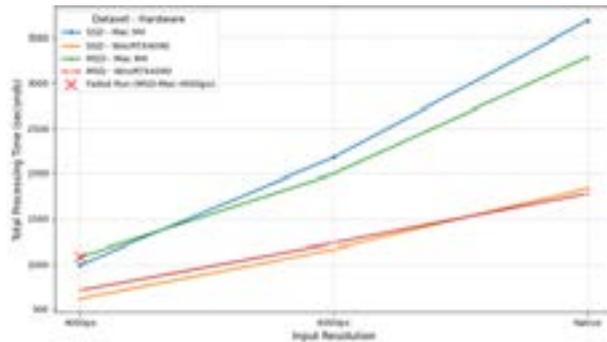
Conversely, the enhanced speed of the Win/RTX4090 system demanded significantly higher energy consumption (Figure 2). Across all comparable successful runs, the high-performance system used, on average, 4.65 times more energy (Wh) than the Mac M4. The energy increase factor ranged from 4.13x for the SSD-Native configuration to 5.25x for the SSD-4000px configuration. Figure 2 also highlights the energy consumed (14 Wh) by the MSD-MAC MINI M4-4000PX run before it failed to produce a usable model, indicated by the hatched bar.



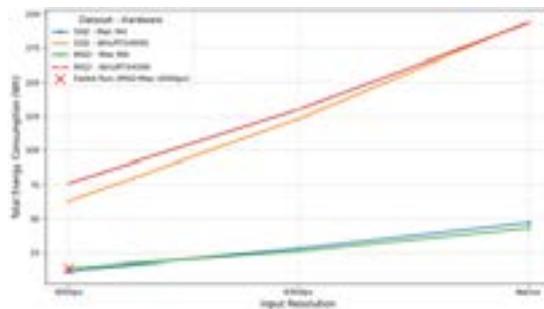
**Figure 2: Energy Consumption Comparison by Hardware.** This figure shows total energy consumption, similar to Figure 1, includes the failed run (marked with hatching).

### 3.2. Impact of Input Resolution on Efficiency

Processing time and energy consumption generally increased with higher input image resolution for both hardware platforms and dataset types, as illustrated in Figure 3 and Figure 4. On the Mac M4 with the SSD, increasing the resolution from 4000px to 6000px resulted in a ~50% time increase and a ~133% energy increase. The subsequent step from 6000px to Native resolution led to a further ~120% increase in time and ~68% increase in energy. Similar upward trends were observed for the Win/RTX4090 system and the MSD, although the specific scaling factors varied. Figure 4 includes the energy data point for the failed MSD-MAC MINI M4-4000PX run, marked with a red 'x', showing the energy consumed despite the lack of a corresponding completion time in Figure 3. Higher resolutions consistently yielded denser point clouds and meshes (Table 1).



**Figure 3: Processing Time vs. Input Resolution.** This figure uses line plots to show how processing time changes with resolution for each Dataset/Hardware combination and includes a marker for the failed run’s energy usage.

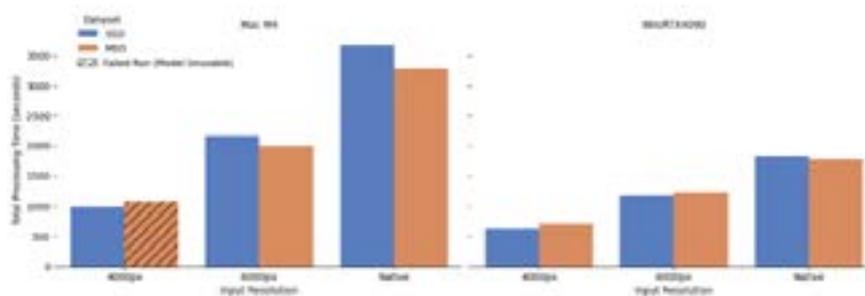


**Figure 4: Energy Consumption vs. Input Resolution.** Similar to Figure 3, but for energy consumption, and includes a marker for the failed run’s energy usage.

### 3.3. Impact of Dataset Type (SSD vs. MSD) on Processing Efficiency

The comparison between the Single-Source Dataset (SSD) and the Multi-Source Dataset (MSD) revealed differences in robustness, processing time, energy use, and output quality, often interacting with the hardware platform.

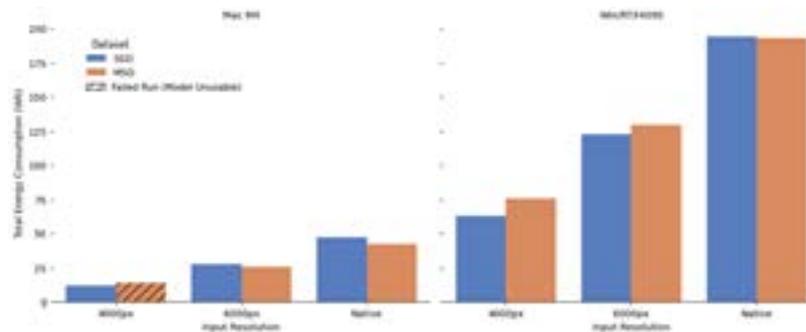
As noted previously, the MSD failed to process at 4000px on the Mac M4, whereas the SSD succeeded, demonstrating superior robustness for the simpler dataset under these conditions (Figure 5, annotation on Mac M4 subplot).



**Figure 5: Processing Time: SSD vs. MSD Comparison.** Uses facets (subplots) to compare SSD and MSD processing times side-by-side for each hardware platform and including the bar for the failed run (with hatching).

For successful runs on the Mac M4 (6000px and Native), the MSD processed faster than the SSD (Figure 5). At 6000px, MSD was ~15% faster, and at Native, ~9% faster. This time advantage translated to lower energy consumption for the MSD on the Mac M4 (Figure 6), using ~7% less energy at 6000px and ~9% less at Native resolution.

On the Win/RTX4090 platform, the trend reversed slightly. The SSD processed marginally faster than the MSD at all resolutions (Figure 5). Correspondingly, the SSD consumed less energy at 4000px (~17% less) and 6000px (~5% less), while energy consumption was nearly identical at Native resolution (Figure 6).



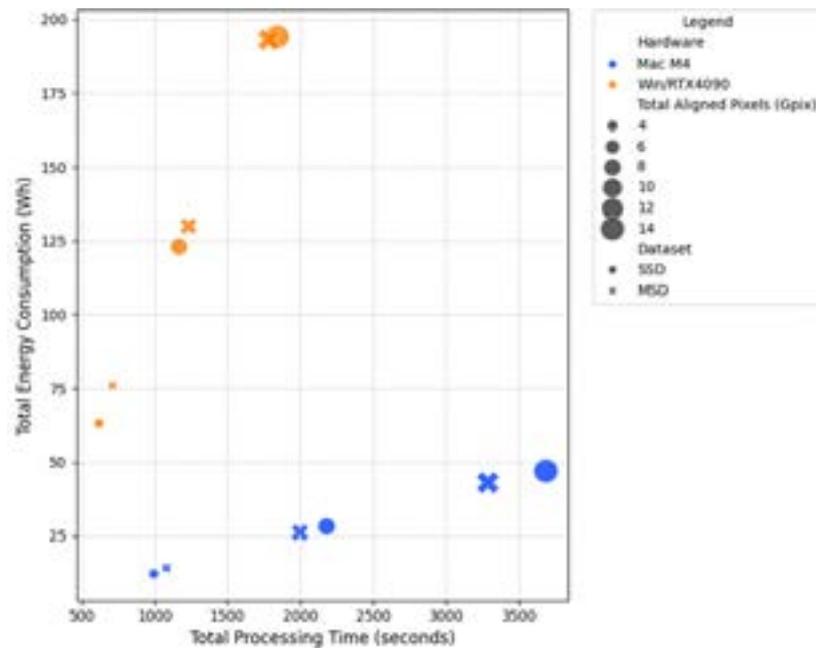
**Figure 6: Energy Consumption: SSD vs. MSD Comparison. Similar faceted plot as Figure 5, but showing energy consumption and including the bar for the failed run (with hatching).**

Despite sometimes being less efficient in terms of time and energy (especially on the Win/RTX4090), the MSD consistently produced lower RMS reprojection errors (Table 1), indicating a better geometric alignment relative to the image data. It also generated significantly denser point clouds and meshes compared to the SSD at equivalent settings (Table 1).

### 3.4. Relationship Between Processing Time and Energy Consumption

A strong positive correlation exists between total processing time and total energy consumption across all successful runs, as depicted in the scatter plot (Figure 7). The runs are clustered primarily by hardware platform, with the Win/RTX4090 runs occupying the upper-left region (shorter time, much higher energy) and the Mac M4 runs occupying the lower and lower right region (longer time, much lower energy). Within each hardware cluster, increasing resolution generally moves points diagonally up and to the right, indi-

cating increases in both energy and time.



**Figure 7: Processing Time vs. Energy Consumption.** A scatter plot showing the relationship between time and energy, with points colored by hardware and potentially styled by dataset or resolution.

#### 4. DISCUSSION

This study aimed to quantitatively evaluate the impact of dataset type (Single-Source vs. Multi-Source), hardware platform (energy-efficient vs. high-performance), and input image resolution (4000px, 6000px, Native) on the processing time and total energy consumption of a standard photogrammetric workflow in Agisoft Metashape Professional. The results presented provide several key insights into the complex interplay between these factors.

##### 4.1. Dataset Type: Robustness vs. Nuanced Efficiency

The comparison between the Single-Source Dataset (SSD) and the Multi-Source Dataset (MSD) yielded perhaps the most nuanced findings. As hypothesized, the structural simplicity of the SSD translated into greater robustness, particularly under challenging conditions. The complete failure of the MSD alignment on the Mac M4 at 4000px resolution (Figure 8), while the SSD succeeded, underscores this point (Table 1). This failure likely stems from the combination of lower feature distinctiveness at reduced resolution and the increased complexity of simultaneously optimizing multiple, potentially poorly constrained, camera calibration groups inherent to the MSD (6). From an efficiency standpoint, successful completion is paramount; thus, the SSD proved more efficient in this

specific scenario.



**Figure 8: The complete failure of the MSD alignment on the Mac M4 at 4000px resolution**



**Figure 9: Successfully processed model from MSD on the Mac M4 at Native resolution**

Counter-intuitively, however, when both datasets were processed successfully on the Mac M4 (at 6000px and Native resolutions), the MSD demonstrated slightly better efficiency, completing faster and consuming less energy than the SSD (Figure 5, Figure 6). This suggests that factors beyond mere structural complexity influence overall efficiency. The consistently lower RMS reprojection error achieved with the MSD (Table 1) implies a better geometric solution was found, potentially due to the richer geometric information provided by the varied focal lengths and camera positions. This superior geometric fit might enable more efficient processing in subsequent stages like depth map calculation or texturing, potentially outweighing the increased overhead during the initial alignment phase on the M4 platform, where computational bottlenecks may differ from the high-performance system. Furthermore, the slightly fewer aligned images for MSD at Native resolution (330 vs 335) contributed to processing less pixel data, directly impacting energy and time (Table 1).

On the high-performance Win/RTX4090 system, the results aligned more closely with initial expectations. The SSD consistently held a slight advantage in processing time and consumed notably less energy at 4000px and 6000px resolutions (Figure 5, Figure 6). This suggests that when computational power is abundant, the overhead associated with managing the MSD's multiple calibration groups becomes a more dominant factor relative to any potential benefits from its geometric diversity, making the simpler SSD structure more efficient overall. The near-identical energy consumption at Native resolution on this platform indicates that at very high data volumes, the sheer pixel processing load might overshadow the subtle differences in dataset structural overhead.

## **4.2. Hardware Platform: The Speed vs. Energy Trade-off**

The comparison between the Mac M4 and the Win/RTX4090 platform starkly illustrates the well-established trade-off between processing speed and energy consumption in high-performance workstation (7). The Win/RTX4090 system, leveraging its powerful CPU and particularly its high-end discrete GPU, achieved significant reductions in processing time, averaging a 1.74x speedup over the Mac M4 (Figure 1). As indicated by stage-specific timings in the reports, this advantage is primarily driven by the massive processing capabilities of the RTX 4090, accelerating GPU-dependent tasks within Metashape, such as depth map calculation and texture blending (8). Stages less amenable to GPU scaling or potentially more CPU-bound, like mesh reconstruction from depth maps in this workflow, showed less dramatic or even negligible speed differences.

However, this performance gain came at a steep energy cost. The Win/RTX4090 system consumed, on average, 4.65 times more energy than the Mac M4 to complete the same tasks (Figure 2). This disproportionate increase highlights that the higher clock speeds, larger core counts, and especially the power-hungry discrete GPU, while reducing wall-clock time, operate at a much lower energy efficiency (Joules per calculation) compared to the integrated SoC design of the Apple M4. This finding is crucial for users operating under energy constraints, such as those using battery power, facing high electricity costs, or aiming for more sustainable processing practices. The choice of hardware presents a clear decision point between optimizing for time or for energy efficiency.

## **4.3. Input Resolution and Data Size: A Near-Linear Cost**

The results demonstrate a strong, near-linear positive correlation between the amount of processed pixel data (Total Aligned Pixels in Gpix) and both total processing time and total energy consumption, particularly when viewed within a specific hardware/dataset configuration (Figure 3, Figure 4, Figure 7). The energy consumption, in particular, showed remarkable linearity with the GigaPixel workload for each hardware platform, irrespective of the dataset type (SSD vs MSD), suggesting an approximate energy cost per GigaPixel processed (roughly 3.4 Wh/Gpix for Mac M4, 12.5 Wh/Gpix for Win/RTX4090 under these 'Medium' settings). Processing time also scaled positively with data size, though with slightly more variability, potentially influenced by factors like alignment convergence or variations in feature density across resolutions.

It is critical to note that this observed linearity is contingent on the consistent use of "Medium" quality settings for alignment and depth map generation across all runs. These settings involve internal image downscaling by Metashape[2]. Therefore, while the relationship between aligned input pixels and resource use appears linear here, changing the quality settings (e.g., to High or Ultra High, which use less or no downscaling) would alter the number of effectively processed pixels and likely result in a non-linear increase in time and energy relative to the raw input pixel count. Nonetheless, for consistent processing parameters, the total aligned pixel count serves as a strong predictor of the computational workload and associated resource consumption.

## **4.4. Limitations**

This study provides valuable quantitative data but has several limitations. The findings are based on two specific datasets representing particular types of object geometry and capture strategies; results may differ for scenes with significantly different characteristics (e.g., highly complex geometry, poor texture, aerial vs. terrestrial). Only two hardware platforms were compared, representing specific points on the performance/efficiency spectrum. All processing was performed using a single version of Agisoft Metashape with predominantly "Medium" quality settings; different software or settings (e.g., High/Ultra High quality, different filtering modes) would likely yield different absolute and

relative results. The energy measurement method captured total system energy during the workflow, potentially including minor background OS activity, rather than isolating Metashape's consumption perfectly. Finally, the single failure of the MSD on the Mac M4 at 4000px limits direct comparison for that specific configuration.

#### **4.5. Implications and Future Work**

The results offer practical guidance for optimizing photogrammetry workflows based on specific constraints. For maximum robustness, especially on lower-power hardware or with lower-resolution imagery, using a consistent Single-Source Dataset is advisable. When processing speed is paramount and energy costs are secondary, a high-performance system with a powerful discrete GPU offers significant time savings, albeit at a high energy premium. Conversely, platforms like the Apple Silicon Macs provide remarkable energy efficiency, completing tasks with substantially less power, making them suitable for energy-constrained environments. The near-linear scaling of energy with processed pixels (under constant settings) allows for rough estimation of energy budgets based on dataset size.

Future work could expand this analysis across a wider range of hardware (including different GPUs, CPUs, and mobile platforms), investigate the impact of different Metashape quality settings (Low, High, Ultra High) on the time/energy/quality trade-offs, and measure energy consumption on a per-stage basis to better understand bottlenecks. Comparing energy profiles across different photogrammetry software packages for the same tasks would also provide valuable context. Further investigation into the factors causing the MSD to outperform the SSD on the Mac M4 platform could also yield interesting insights into software-hardware interactions.

### **5. CONCLUSION**

This study provided a quantitative comparative analysis of processing time and energy consumption in Agisoft Metashape Professional, evaluating the effects of dataset type (Single-Source vs. Multi-Source), hardware platform (energy-efficient Mac M4 vs. high-performance Win/RTX4090), and input image resolution (4000px, 6000px, Native).

The results demonstrate the significant impact of hardware choice on photogrammetric processing efficiency. While the high-performance Win/RTX4090 system offered substantial speed improvements, roughly halving processing time compared to the Mac M4, this came at the cost of consuming over four times the energy. This highlights a critical trade-off for users selecting hardware based on priorities for speed versus energy conservation or cost.

Input data size, measured as total aligned pixels processed, showed a strong, near-linear relationship with both time and energy consumption when processing settings were held constant. This suggests that the volume of pixel data is a primary driver of resource usage within a given hardware and software configuration.

The comparison between dataset types yielded nuanced results. The Single-Source Dataset (SSD) proved more robust, successfully processing under conditions where the Multi-Source Dataset (MSD) failed. On the high-performance system, the SSD generally offered slightly better time and energy efficiency. However, on the energy-efficient Mac M4, the MSD unexpectedly demonstrated superior efficiency in successful runs, potentially due to achieving better geometric alignment, indicating that dataset complexity interacts significantly with hardware capabilities and requires further investigation across more diverse scenarios to draw fully coherent conclusions.

Overall, this research underscores the importance of considering dataset characteristics, hardware selection, and processing resolution not only for optimizing processing speed and output quality but also for managing energy consumption in photogrammetric work-

flows. While high-performance systems offer speed, energy-efficient platforms present a preferred alternative when power usage is a primary concern, and simpler, consistent datasets generally offer greater processing reliability.

Acknowledgements:

This research was supported by the project EUROCC2 funded by the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 101101903 and the HybridScale project of the Croatian Science Foundation under grant agreement number HRZZ-UIP-2020-02-4559.

## **6. REFERENCES**

- [1] Westoby M, Brasington J, Glasser N, Hambrey M, Reynolds J. Structure-from-Motion photogrammetry: a novel, low-cost tool for geomorphological applications. 2012 Apr 1;936.
- [2] Agisoft LLC. Agisoft Metashape User Manual: Professional Edition, Version 2.2. [Software Documentation]. St. Petersburg: Agisoft LLC; 2024. [Internet]. [cited 2025 Apr 1]. Available from: [https://www.agisoft.com/pdf/metashape-pro\\_2\\_2\\_en.pdf](https://www.agisoft.com/pdf/metashape-pro_2_2_en.pdf)
- [3] Kolarek B, Gamulin L, Davidović D. Cultural Heritage on HPC - Creating High Resolution 3D Models Using Photogrammetry. In: 2024 47th MIPRO ICT and Electronics Convention (MIPRO) [Internet]. 2024 [cited 2025 Apr 1]. p. 1121–7. Available from: <https://ieeexplore.ieee.org/abstract/document/10569784>
- [4] ResearchGate [Internet]. [cited 2025 Apr 1]. Introduction to High Performance Computing for Scientists and Engineers | Request PDF. Available from: [https://www.researchgate.net/publication/215739126\\_Introduction\\_to\\_High\\_Performance\\_Computing\\_for\\_Scientists\\_and\\_Engineers](https://www.researchgate.net/publication/215739126_Introduction_to_High_Performance_Computing_for_Scientists_and_Engineers)
- [5] Remondino F, El-Hakim S. Image-based 3D Modelling: A Review. Photogramm Rec [Internet]. 2006 [cited 2025 Apr 1];21(115):269–91. Available from: <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1477-9730.2006.00383.x>
- [6] Agarwal S, Snavely N, Simon I, Seitz S, Szeliski R. Building Rome in a Day. Vol. 2009, Proceedings of the IEEE International Conference on Computer Vision. 2009. 72 p.
- [7] Huang Y, Li H, Lu D, Zhang Z, Tie W. Intelligence Evaluating Computational Power: A Multi-Factor Method. IEEE Access [Internet]. 2025 [cited 2025 Apr 1];13:27398–415. Available from: <https://ieeexplore.ieee.org/abstract/document/10872901>
- [8] Wang Q, Zhan Z, He J, Tu Z, Zhu X, Yuan J. High-Quality Spatial Reconstruction and Orthoimage Generation Using Efficient 2D Gaussian Splatting [Internet]. arXiv; 2025 [cited 2025 Apr 1]. Available from: <http://arxiv.org/abs/2503.19703>