

AI-Guided Optimization for 3D Gaussian Splatting on Drone Images

1 INTRODUCTION/ PROBLEM

3D Gaussian Splatting (3DGS) can produce high-quality 3D reconstructions from drone imagery. However, reconstruction quality depends on hyperparameter choice. Manual tuning is time-consuming because drone projects differ in overlap, viewing geometry, scene content, vegetation, texture, and acquisition conditions. This can lead to inconsistent results and limits efficient use of 3DGS across diverse projects.



2 GOAL

Develop and test a learning-based workflow that recommends suitable 3DGS hyperparameter multipliers for new drone-image projects

- ◆ Learn from previous 3DGS reconstruction runs.
- ◆ Use project descriptors, tested multipliers, and quality scores.
- ◆ Recommend geometry, appearance, and densification multipliers.
- ◆ Reduce trial-and-error while maintaining or improving reconstruction quality.

3 PROJECT DESCRIPTORS & Multipliers

10 standardized project descriptors


	Acquisition descriptors GSD, Image Overlap Proxy, Coverage Spread, Camera Angle, Heading Consistency
	Image descriptors Texture Density, Blur / Motion Risk
	Scene descriptors Terrain Roughness, Vegetation Complexity, Vegetation Cover

Multiplier Group, Hyperparameters

Geometry (bound 0.2 -5)	Position, Scale, Rotation learning rates
Appearance (bound 0.2 -5)	Feature, Opacity learning rates, DSSIM loss balance
Densification (bound 0.7 -1.43)	Densify grad threshold, Opacity threshold

Note: Project descriptors describe the conditions of each drone project and are used as model inputs. Multipliers define the controlled changes applied to 3DGS settings, which the model learns to recommend. Bounds keep the applied values within a practical range.

4 METHODOLOGY



Note: Baseline runs use default hyperparameters. Exploration runs apply sampled multiplier groups to the default hyperparameters before 3DGS training.

DATA & LEARNING MODELS

58	46	12	653
Drone Projects	Model-development Projects	Test Projects	Training Rows

Note: Each row represents one completed exploration run. The baseline run is used only to calculate the relative quality score, which summarizes whether the run improved or degraded quality compared with that project's baseline. Later Training Rows are used to train the models where descriptors and multipliers are used as input and relative quality score as a target.

MODEL TESTING

Input
10 project descriptors + tested multiplier combination

47 input values:
multipliers, squared multipliers, descriptors, descriptor-multipliers. Both models were tested using 30 candidate multiplier combinations.

Models


- Ridge Regression**
Main interpretable scoring model. It is simple, regularized, and stable for a limited data.
- Multilayer Perceptron (MLP)**
Nonlinear comparison model, compact 47 → 16 → 1 structure

Prediction & Selection

30 separate candidate multiplier groups used to predict score


Multiplier group with highest predicted quality score is selected and used to run 3DGS exploration run.

5 RESULTS



6 OUTPUTS & CONTRIBUTION

Open-source Bimba3D-re platform



It connects project creation, 3DGS processing, model training, testing, and result comparison in one workflow.

<https://github.com/geomatupen/bimba3d-re>

Results platform



Contribution

- ◆ Learning-based 3DGS hyperparameter recommendation
- ◆ Training data from descriptors, multipliers, and quality scores
- ◆ Comparison of Ridge Regression and MLP with limited training data
- ◆ Open workflow/platform for training, testing and result comparison

7 CONCLUSION

The work demonstrates that previous drone-image 3DGS runs can support more informed hyperparameter selection for unseen projects. Ridge Regression was the more stable and reliable model under the tested conditions, completing more projects normally and producing fewer hard-cap cases than the MLP. The recommended settings can reduce manual trial-and-error, but they do not guarantee faster or smaller reconstructions.

Visit Website



Note: Results platform is linked here