Forma3D Tiny 3: continuation on mesh-derived Gaussian proxies

Lucas van de weerd | Lucasgitguy-co | 7 November 2026 | Version 1.0

Repository: Lucasgitguy-co/forma3d-tiny-3

Reviewed revision: 3469f527c5f6833d654f62653ee92e0f73854b33

Abstract

This report documents Forma3D Tiny 3, a 306,321-parameter image-to-Gaussian model. Published paired evaluation records show lower proxy reconstruction loss and higher occupancy IoU than its starting checkpoint on the same held-out distribution. The conclusion is limited to the recorded proxy task.

1. Scope and related work

This technical report reconstructs the release from pinned repository artifacts and the publisher's earlier portfolio report. It is an AI-assisted documentation synthesis, not an independent rerun or a peer-reviewed publication. Reported measurements belong to the original release records. [1, 2, 3, 4, 5]

The earlier portfolio report provides the organizing distinction between training, model modification, conversion and packaging. This report applies that distinction to one release and uses current records where later experiments supersede the historical overview.

2. Model and method

The model retains width 64 and a 16 cubed grid, with up to 4,096 Gaussian slots. It continues from Forma3D Tiny 2 with a fresh optimizer. Mesh surface samples form Gaussian targets; images are orthographic renders of these proxies on black. No pretrained 3D teacher or natural-photo training is established. 512 ABO objects with 2,048 orthographic proxy views. Group counts are 348 training, 87 validation and 77 test; test has 308 views. Shared Amazon item IDs and exact mesh/image/target hashes are grouped. The run completed and selected step 6,000. Checkpoint selection uses validation, with test excluded. [1, 2, 3, 4, 5]

3. Evidence and results

Both columns are measurements from evaluation.json on the same test set for this release. Loss is a composite cached-target regression objective; occupancy IoU is averaged per view at threshold 0.5. Results from different Forma3D releases use different test distributions and should not be ranked as if they were one common benchmark. [1, 2, 3, 4, 5]

Test metric	Starting checkpoint	Selected checkpoint
Mean proxy loss	0.690551	0.544487
Mean occupancy IoU	0.286866	0.339860

4. Limitations and intended use

Near duplicates can survive exact matching and grouping. Overlap with ancestor training datasets is not fully known. The measurements do not establish natural-photo reconstruction, novel-view photorealism, texture fidelity or mesh topology. There are no reported confidence intervals or component ablations. Use centered object images on black and retain dataset attributions. Publisher code/weights are offered under Apache-2.0; source assets retain their own terms.

5. Reproducibility and conclusion

The source snapshot above fixes what was reviewed. The linked source files and source-manifest.json record the evidence used, including file hashes. Consult the repository's usage instructions and recorded dependencies for execution. No new model training, inference benchmark, or claimed state-of-the-art result was introduced in preparing this report. Evidence status: recorded held-out proxy evaluation.

References

[1] Lucasgitguy-co. README.md at reviewed revision.

[2] Lucasgitguy-co. evaluation.json at reviewed revision.

[3] Lucasgitguy-co. training.json at reviewed revision.

[4] Lucasgitguy-co. config.json at reviewed revision.

[5] Lucasgitguy-co. preparation.json at reviewed revision.

[6] Lucasgitguy-co. Local Model Adaptation Across Vision Diagrams Geometry and Game Decisions. Historical portfolio report, 5 November 2026; its proposed v2 experiment is superseded by later release records.