Image3D Student: compact cached-target reconstruction prototype

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

Repository: Lucasgitguy-co/image3d-student

Reviewed revision: 1738cdfaf1bb940bf576a96d11b742ce014b9139

Abstract

This report documents a compact image-to-Gaussian research prototype and its synthetic demonstration checkpoint. Its evidence concerns an implemented training/export pipeline, not pretrained natural-photo reconstruction or verified execution of multiple external teachers.

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]

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

A four-stage convolutional image encoder produces a shared feature vector. A coordinate MLP predicts Gaussian properties and auxiliary geometry occupancy on a fixed 3D grid. A 16 cubed grid supports up to 4,096 candidate Gaussians. Cached-target losses supervise occupancy, position, color, log-scale, sign-invariant rotation, opacity and optional geometry occupancy. Imported teacher outputs must share the correct image framing and coordinate system. Missing teacher supervision is masked. [1, 2, 3, 4]

3. Evidence and results

The included checkpoint learned procedural ellipsoids. The repository records CPU/Metal pipeline verification and synthetic smoke-test artifacts. It does not establish real-image accuracy. The interfaces can import Idol Gaussian PLY and aligned InstantMesh/Hi3D meshes; their presence is not evidence that those teachers produced training data for this checkpoint. [1, 2, 3, 4]

4. Limitations and intended use

The baseline has no differentiable-rendering loss or binary-weight student. Full pretrained Idol execution on the target Mac was unverified; local InstantMesh and Hi3D runners were not implemented. Cell compression is lossy. Input framing must match the training contract, and output is Gaussian PLY rather than a watertight mesh. Apache-2.0 applies to publisher source/weights with separate third-party 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: synthetic prototype verification.

References

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

[2] Lucasgitguy-co. VERIFICATION.md at reviewed revision.

[3] Lucasgitguy-co. src/image3d_student/model.py at reviewed revision.

[4] Lucasgitguy-co. THIRD_PARTY.md at reviewed revision.

[5] 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.