INFORMATION
Volume 29, Number 1, pp.117-122

ISSN 1343-4500 eISSN 1344-8994

Automated Generation of Synthetic 3D Gaussian Splatting Datasets for Shape and Texture Descriptor Analysis

Motofumi Suzuki *

* *The Open University of Japan, Dept. of Informatics, Chiba-shi 261-8586, Japan,*
E-mail: motofumi@ouj.ac.jp

Abstract

This paper presents an automated framework for generating synthetic three-dimensional Gaussian splatting (3DGS) data to support the quantitative analysis of shape and texture descriptors. The system integrates procedural shape modeling, texture mapping, multiview photorealistic rendering, and structure-from-motion reconstruction to produce 3DGS datasets enriched with geometric and appearance-based attributes, including spherical harmonics (SH) coefficients. These coefficients capture lighting and material interactions, enabling the rotation-invariant analysis of the surface appearance. Experiments using superellipsoids combined with various materials demonstrated that the framework can systematically control the geometry and texture, producing consistent and reproducible datasets. Analyses based on SH-derived principal component analysis, opacity correlation, spatial density, and entropy confirmed that the generated data effectively revealed how surface roughness, reflectance, and scale contributed to the internal structure of 3DGS scenes. The results highlight the potential of synthetic 3DGS data as a controlled testbed for studying shape–appearance descriptors and for advancing data-driven 3D representation research.

Keywords: 3D Gaussian splatting (3DGS); Synthetic dataset generation; Spherical harmonics; Texture descriptors; Material descriptors; Superellipsoids

1. Introduction

Three-dimensional Gaussian splatting (3DGS) [1] has emerged as a powerful alternative to traditional 3D representation, such as polygonal meshes, point clouds, and voxel grids. Unlike volumetric neural rendering methods such as neural radiance fields, 3DGS represents a scene as a collection of explicit anisotropic Gaussian primitives, enabling real-time rendering with efficient memory usage. This capability supports a wide range of applications, including virtual and augmented reality, cultural heritage preservation, e-commerce, and scientific visualization. Traditional 3D model retrieval systems have focused primarily on shape-based descriptors, often employing handcrafted features or deep neural representations. However, the structure of 3DGS requires a more comprehensive approach that captures not only the geometric shape but also the color, material, and appearance under varying illumination conditions. This paper presents a system for the synthetic generation of 3DGS datasets and the computation of descriptors that integrate both geometric and appearance-related features, supporting tasks such as classification, retrieval, and pattern analysis.

2. Methodology

The 3DGS data are typically stored in formats such as PLY, SPZ, or SPLAT, which encode Gaussian primitives with attributes such as position, scale, rotation, color, opacity, and spherical harmonics (SH) coefficients. Among these, SH coefficients are particularly important as they encode directional lighting and surface appearance. SH functions provide a mathematical framework for representing functions on a sphere, and the power spectrum derived from these coefficients is inherently rotation-invariant. This property has long been utilized in fields such as chemistry and drug discovery for molecular shape comparison. Extending this concept to 3DGS enables robust similarity searches across models with different orientations.

To systematically generate 3DGS data, we designed an automated pipeline that integrates shape modeling, texture mapping, rendering, and reconstruction. The pipeline begins with a database of 3D shapes, either parametric such as superellipsoids or mesh-based, and a library of textures, including materials such as rock, wood, metal, and glass. The rendering was performed using the Persistence of Vision Raytracer (POV-Ray), a widely used ray tracer for generating photorealistic computer graphics images. Multiple cameras are positioned around each object to capture two-dimensional images from different viewpoints. The intrinsic and extrinsic camera parameters were recorded for each image to enable an accurate 3D reconstruction.

We then used structure-from-motion (SfM) via COLMAP [2] [3] to reconstruct the 3D geometry from the rendered images, producing both a sparse point cloud and corresponding camera poses. The reconstructed data were converted into Gaussian primitives, and the SH coefficients and other rendering attributes were estimated. Although 3DGS data can be generated directly from original 3D models, the SfM-based approach was chosen because it captures realistic lighting and appearance interactions that are difficult to model analytically.

3. System Architecture

The system consists of multiple modules, including input management, scene generation, rendering, camera layout, SfM processing, 3DGS conversion, feature extraction, and data export. The workflow proceeds as follows:

1. Input of shape and texture data
2. Scene generation and multiview rendering using POV-Ray
3. Recording of camera metadata
4. 3D reconstruction using COLMAP and conversion into Gaussian primitives, including the estimation of SH coefficients via the official graphdeco-inria/Gaussian-splatting [1] implementation
5. Feature extraction and export as CSV descriptors

This modular design enables the efficient batch processing of large datasets and parameter sweeps across lighting, texture, and shape variations. Each module operates independently, but contributes to a unified data pipeline that converts 3D assets into structured 3DGS datasets suitable for quantitative analysis. At the end of the pipeline, three main tables describe the reconstructed 3DGS data. Table 1 lists all raw attributes exported from the PLY file, such as the position, normal, scale, rotation, opacity, and SH coefficients. Table 2 summarizes the derived metrics computed for each Gaussian point, including mean scale, anisotropy, SH norms, and positional magnitude. Table 3 lists the global and statistical descriptors, such as neighborhood density, principal component projections, and spatial entropy. Together, these tables capture both the local and global properties of the Gaussian primitives, forming a comprehensive dataset for analyzing the geometry, texture, and appearance.

Table 1. Raw attributes (PLY fields)

Field	Meaning
`x, y, z`	3D coordinates of the Gaussian center.
`nx, ny, nz`	Normal vector at the Gaussian center (if present).
`f_dc_*`	Spherical Harmonic DC (base) color coefficients.
`f_rest_*`	Higher-order SH color coefficients.
`opacity`	Opacity (0–1 range).
`scale_0..2`	Scaling factors of the Gaussian along principal axes.
`rot_0..3`	Rotation quaternion parameters.

Table 2. Derived per-Gaussian metrics

Field	Meaning
`point_index`	Row index of the Gaussian point (0-based).
`mean_scale`	Mean of the scale components ((scale□ + scale□ + scale□) / 3).
`scale_anisotropy`	Ratio of scale standard deviation to mean scale
`sh_norm_rgb`	L2 norm over all SH RGB coefficients (color magnitude).
`sh_norm_lab`	Luminance-weighted SH norm, using perceptual weights (Lab-like)
`position_norm`	Euclidean distance from the origin ($\sqrt{(x^2 + y^2 + z^2)}$).

Table 3. Global and statistical descriptors

Field	Meaning
`pca_x, pca_y, pca_z`	PCA projection of the position into the three principal axes.
`knn_mean_dist`	Mean distance to the eight nearest neighbors (spatial locality).
`knn_density`	Local density estimate ($1 / (1e^{-9}$ + knn_mean_dist)).
`bbox_volume`	Global bounding box volume, repeated per Gaussian.
`entropy_global`	Global spatial entropy, repeated per Gaussian.

4. Experimental Results

To evaluate the system performance, experiments were conducted using superellipsoids, which provide precise control over geometric parameters. Superellipsoids are parametric surfaces that extend traditional ellipsoids by introducing two shape exponents that independently define the curvature along orthogonal axes. This property enables the smooth interpolation between spherical and cuboidal forms, making superellipsoids ideal for systematically exploring geometric variations. Their analytical formulation also ensures consistent surface sampling and a well-defined curvature, which are essential for accurate reconstruction and descriptor evaluation.

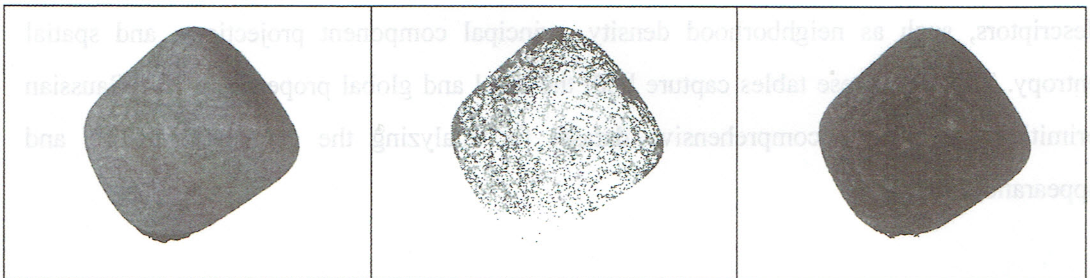

(a) 3DGS | (b) Initial point cloud | (c) 3DGS (scale = 0.50)

Figure 1: 3D Gaussian splatting (rust texture; 80,185 points; 30,000 iterations)

Five base shapes were paired with rock, wood, metal, and glass textures to produce 80 objects. These textures represent the diffuse, specular, and refractive materials commonly used in graphics. Each object was rendered from multiple viewpoints to ensure comprehensive spatial coverage.

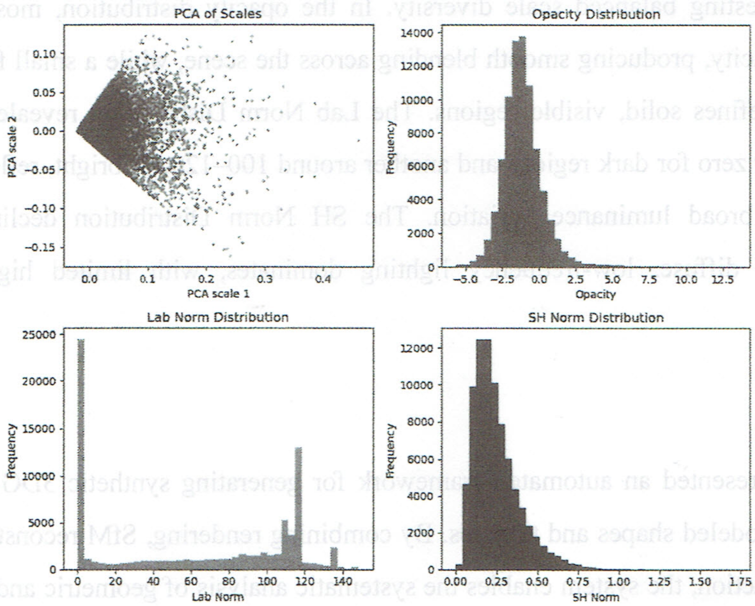

Figure 2: Statistical summary of Gaussian splatting data.

As an example, one superellipsoid defined by the shape parameters $e1 = 0.8$ and $e2 = 0.6$ was rendered with a stone texture labeled "Rust." The object was placed at the center of the scene, and 72 virtual cameras were arranged along a spiral orbit around it. Images were captured from each viewpoint at a resolution of 1024 × 768 pixels, providing sufficient detail for SfM reconstruction. The resulting images were then converted for compatibility with the INRIA 3DGS model and used for training. As the synthesized objects were the primary targets, the system employed mask images to isolate each object from the background. In addition, accurate camera positions and orientations were passed to the training process to ensure correct spatial alignment. Using the Rust texture provided in the "textures.inc" file as an example, the total training time was approximately 30 min on a PC equipped with a Ryzen 9 5950X CPU and a GeForce RTX 4060 Ti GPU, resulting in 3DGS data of approximately 80,185 Gaussian center points. These Gaussian primitives collectively encoded the spatial structure, color, opacity, and lighting, resulting in a compact and expressive representation compatible with the SIBR [4] viewer system.

Figure 2 provides a concise summary of the structure and lighting characteristics of the Gaussian splatting dataset. Principal component analysis (PCA) of the scale plot showed that most Gaussians clustered near the origin, indicating small, nearly isotropic shapes with minimal variation. A few outliers with higher PCA values corresponded to larger structural elements, suggesting balanced scale diversity. In the opacity distribution, most Gaussians exhibit low opacity, producing smooth blending across the scene, while a small fraction with high opacity defines solid, visible regions. The Lab Norm Distribution revealed two main peaks: one near zero for dark regions and another around 100–120 for bright, reflective areas, demonstrating broad luminance variation. The SH Norm Distribution declines steeply, indicating that diffuse, low-frequency lighting dominates, with limited high-frequency highlights.

5. Conclusion

This paper presented an automated framework for generating synthetic 3DGS data from procedurally modeled shapes and textures. By combining rendering, SfM reconstruction, and descriptor extraction, the system enables the systematic analysis of geometric and appearance features. Experiments with superellipsoids and materials such as rust texture confirmed the ability of the system to capture variations in roughness, reflectance, and spatial structure. Future work will focus on the similarity retrieval of 3DGS objects and expanding the dataset with real materials for broader validation.

References

[1] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ACM Transactions on Graphics,* 42(4), July 2023.

[2] Johannes Lutz Schönberger and Jan-Michael Frahm. Structure-from-Motion revisited. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR),* 2016.

[3] Johannes Lutz Schönberger, Enliang Zheng, Marc Pollefeys, and Jan-Michael Frahm. Pixelwise View Selection for Unstructured Multi-View Stereo. In *Proceedings of the European Conference on Computer Vision (ECCV),* 2016.

[4] Sebastien Bonopera, Peter Hedman, Jerome Esnault, Siddhant Prakash, Simon Rodriguez, Theo Thonat, Mehdi Benadel, Gaurav Chaurasia, Julien Philip, and George Drettakis. SIBR: A System for Image-Based Rendering, 2020.