InsideOut: Integrated RGB-Radiative Gaussian Splatting for Comprehensive 3D Object Representation

Supplementary Material

A. Effect of Hyperparameters.

Loss Balancing. Tab. A shows that the values of the optimal parameters are $\lambda_s = 0.2$ and $\lambda_z = 0.005$. The value of $\lambda_s = 0.2$ is commonly used as a standard in the field, and our experimental results confirm its effectiveness. For λ_z, values that are too low introduce excessive noise, while values that are too high cause distortion in RGB Gaussian splats and fading in radiative Gaussian splats. Our experiments reveal that $\lambda_z = 0.005$ provides the ideal balance, effectively flattening noise while clearly distinguishing between layers.

	RGB Gaussian Splatting PSNR↑	RGB Gaussian Splatting SSIM↑	Radiative Gaussian Splatting PSNR↑	Radiative Gaussian Splatting SSIM↑
$\lambda_s = 0.0$, $\lambda_z = 0.005$	34.66	0.92	27.42	0.97
$\lambda_s = 0.1$, $\lambda_z = 0.005$	37.71	0.94	27.43	0.97
$\boldsymbol{\lambda_s = 0.2}$, $\boldsymbol{\lambda_z = 0.005}$	37.74	0.94	28.14	0.97
$\lambda_s = 0.3$, $\lambda_z = 0.005$	37.02	0.94	28.07	0.97
$\lambda_s = 0.2$, $\lambda_z = 0.000$	35.54	0.92	27.41	0.96
$\lambda_s = 0.2$, $\lambda_z = 0.001$	37.26	0.94	27.45	0.97
$\boldsymbol{\lambda_s = 0.2}$, $\boldsymbol{\lambda_z = 0.005}$	37.74	0.94	28.14	0.97
$\lambda_s = 0.2$,$\lambda_z = 0.010$	36.87	0.93	27.61	0.96

Table A. **Analysis of weighting parameters.** The optimal values are highlighted in bold.

Fine Gaussian Selection. As shown in Fig. A, setting $p = 95$ retains sufficient fine Gaussian splats to cover the full surface. Lower p values result in incomplete surface reconstruction, whereas higher values introduce redundant elements without significant improvement. Thus, $p = 95$ provides an optimal balance between surface fidelity and structural efficiency.

B. Computational Efficiency.

Our proposed modules—hierarchical fitting and visual alignment—are highly efficient, requiring only 9.83 minutes (46,129 MB) and 2 seconds (697 MB), respectively, compared to 32.66 minutes (12,332 MB) for rendering on an NVIDIA A6000 GPU. The training of our algorithm takes approximately 45 minutes—slightly longer than rendering X-ray and RGB images separately—yet it maintains the inference speed of baseline methods.

C. Real-world Applicability.

Our method enables multi-domain applications using only RGB and X-ray data without calibration. The elimination of complex calibration procedures significantly reduces setup

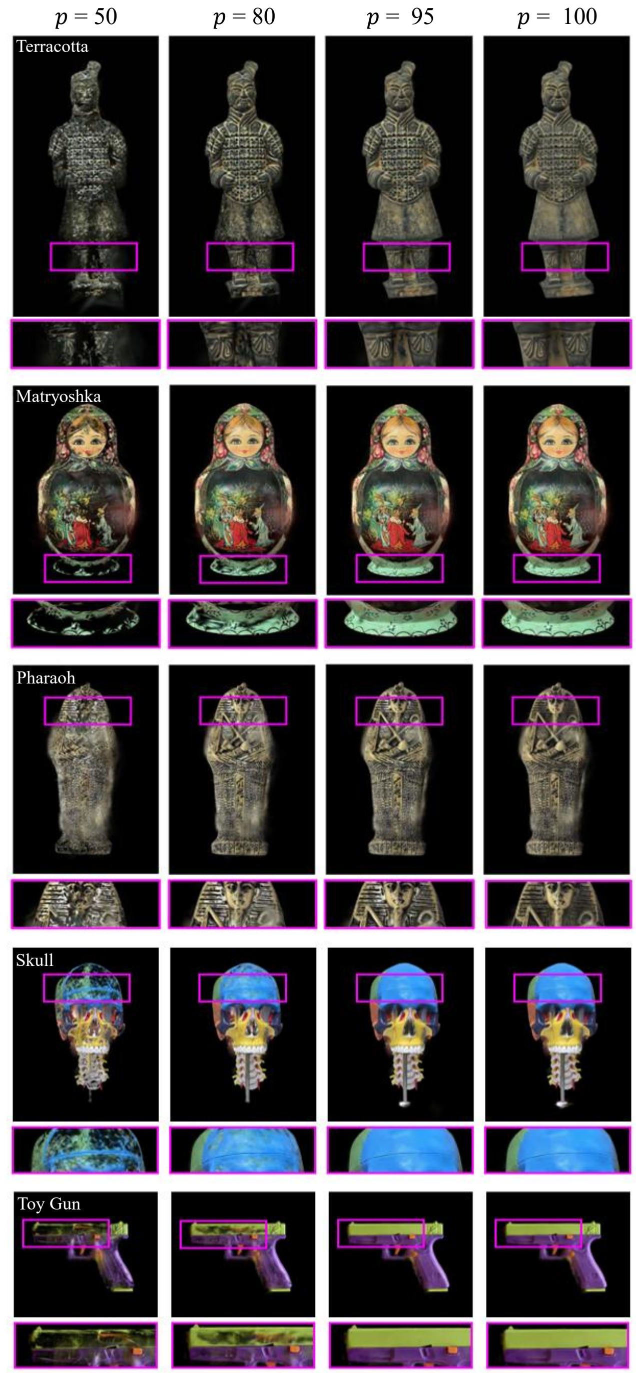

Figure A. **Optimal selection threshold p for efficient surface representation.** Pink boxes indicate zoomed regions.

time and makes the technology more accessible across different environments. In surgical planning, we provide visualization of underlying fractures beneath bruising, helping surgeons optimize incisions and employ real-time navigation. For cultural heritage preservation, we expose internal cracks connecting to the surface and provide museum visitors with augmented reality experiences to observe artifacts' interior and exterior simultaneously. In industrial defect detection, we examine both external appearance and internal structure to identify hidden defects, improving correlation analysis and root cause identification in production.