Frequency-Aware Gaussian Splatting Decomposition

Supplementary Material

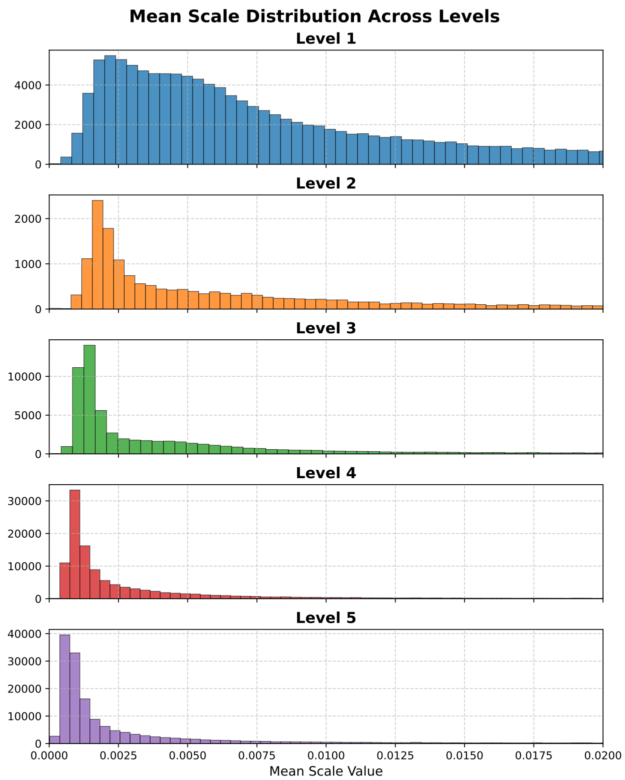

Figure 1. Mean scale distribution of Gaussians across different frequency levels. As the frequency level increases, the distribution shifts toward smaller values, indicating a higher concentration of small Gaussians. This trend highlights the hierarchical nature of the representation, where lower levels capture coarse structures, and higher levels represent finer details.

1. Supplementary Material

This supplementary document provides additional details, experiments, and qualitative results to complement our main paper. Below is an overview of the contents:

- **1.1 Analysis of Gaussian Scale Distribution Across Levels:** Distribution analysis of Gaussian scales across frequency levels.
- **1.2 Foveated Rendering:** Implementation details of our foveated rendering approach.
- **1.3 Ablation Studies:** Detailed discussion on key design choices, including residual color Gaussians, frequency magnitude loss, and image loss at lower levels.
- **1.4 Gaussian Count Dynamics:** Analysis of how the number of Gaussians evolves throughout training and across frequency levels.

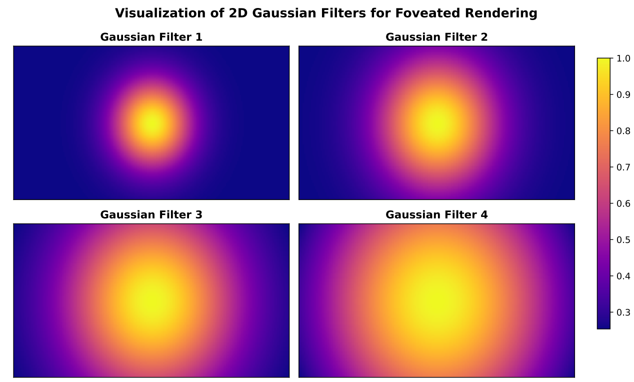

Figure 2. **Gaussian Filtering for Foveated Rendering:** Visualization of the hierarchical 2D Gaussian filters applied to projected 3D Gaussians at different frequency levels. Each filter progressively increases in standard deviation, ensuring that peripheral regions are rendered at lower fidelity while preserving details in the ROI.

- **1.5 Additional Examples of Level of Detail:** Visualization of level-of-detail decomposition in various scenes.
- **1.6 Additional Examples of Foveated Rendering:** Demonstrations of foveated rendering applied to different scenes.
- **1.7 Additional Examples of Promptable 3D Focus:** Illustration of selective object emphasis using our promptable 3D focus application in different scenes.
- **1.8 Additional 3D Artistic Filters:** Examples of artistic filters applied to different scenes using our method.

The following sections provide in-depth discussions and results for each of these topics.

1.1. Analysis of Gaussian Scale Distribution Across Levels

To gain deeper insights into the characteristics of each frequency level in our representation, we analyze the distribution of the mean scale of the Gaussians, computed as the average of the scale values along the x, y, and z axes for each Gaussian. Figure 1 presents this distribution across different levels.

From the figure, we observe a clear trend: as the frequency level increases, the distribution of mean scales shifts toward smaller values, indicating a higher concentration of small Gaussians. This suggests that higher frequency levels predominantly capture fine details in the scene, represented by numerous small Gaussians, whereas lower frequency levels retain larger-scale structures. The progressive shift in the distribution reinforces the hierarchical nature of our representation, where coarse details are preserved in the lower levels, and finer structures emerge at higher levels.

1.2. Foveated Rendering

To implement foveated rendering, we apply a set of spatially varying 2D Gaussian filters (centered at the fixation point (x, y)) to modulate Gaussians across our frequency levels. Each filter has an appropriately scaled standard deviation to retain peripheral lower-frequency components while prioritizing higher-frequency details in the ROI.

As depicted in Figure 2, these filters serve as soft importance weights for Gaussian opacities. Gaussians with contributions below a predefined threshold are discarded, effectively reducing computational load while maintaining visually coherent, high-detail renderings at fixation. This multi-scale filtering yields a notable 40% FPS improvement at 4K resolution, underscoring its effectiveness for real-time AR/VR applications requiring both efficiency and high visual quality.

1.3. Ablation Studies

Effect of Residual Color Gaussians. We evaluated the effectiveness of Residual Color Gaussians by training a model where all Gaussians retained colors in the standard $[0, 1]$ range, rather than using the proposed $[-1, 1]$ residual representation for higher-level Gaussians. This modification led to a significant increase in the total number of Gaussians, from 409K to 515K, representing a 20% growth in scene size. While this increase contributed to a slight improvement in PSNR (+0.2 dB), it also resulted in higher memory consumption and slower rendering times. These findings suggest that the residual color representation enables more compact and efficient scene encoding by allowing individual Gaussians to contribute more effectively to high-frequency details. The ability to add or remove color at different hierarchical levels grants each Gaussian greater expressive power, reducing the overall number of Gaussians required for accurate reconstruction.

Effect of Image Loss on Lower Frequency Levels. To assess the importance of applying the image loss at lower frequency levels, we trained a three-level model while omitting the image loss at lower levels. The resulting model achieved a comparable PSNR to our full method when evaluating the final rendered image; however, its intermediate frequency levels exhibited significant degradation. Specifically, when rendering downscaled resolutions using only the first level (corresponding to a $\times 4$ downscaled output) or the first two levels ($\times 2$ downscaled output), we observed a substantial loss of detail, leading to smeared and blurry reconstructions. Quantitatively, our full model achieved a PSNR of 29.65 dB for $\times 4$ downscaled rendering and 32.27 dB for $\times 2$ downscaled rendering. In contrast, the ablated model, trained without image loss at lower frequency levels, suffered a notable performance drop, with PSNR values of 25.97 dB and 29.36 dB for $\times 4$ and $\times 2$ downscaled renderings, respectively. These results highlight the crucial role of enforcing spatial consistency across frequency levels, ensuring that intermediate levels remain faithful to their expected resolution.

Effect of Frequency Magnitude Loss. We analyzed the impact of the frequency magnitude loss by training the same model without this term. Without frequency regulation, almost all scene details were modeled by the lowest-level Gaussians, while the higher levels contributed almost nothing. This disrupted the intended hierarchical decomposition, as the frequencies were not gradually distributed across levels but instead collapsed into the first level. Consequently, the model failed to leverage the benefits of multi-scale representation, relying predominantly on a single frequency band.

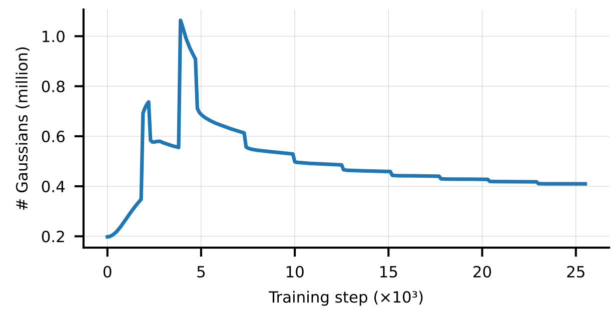

Figure 3. Evolution of the number of Gaussians during training. Each new frequency level increases the number of Gaussians, but adaptive density control quickly prunes low-utility ones, so the model converges to a compact, lightweight representation.

1.4. Gaussian Count Dynamics

Figure 3 illustrates the evolution of the number of Gaussians during training. Although each new frequency level initially doubles the number of Gaussians, the adaptive density control remains active throughout the entire training process, continuously removing low-utility components. As a result, the total number of Gaussians quickly stabilizes after each addition, resulting in a compact and efficient representation. This dynamic ensures that the final model remains lightweight despite the hierarchical structure.

1.5. Additional Examples of Level of Detail

In this section, we present qualitative results demonstrating the level of detail captured by our method across various scenes. Figure 4 showcases six different scenarios, each divided into multiple segments, representing different levels of detail in our method's representation. The images illustrate how our approach captures both coarse structural information and fine details, effectively adapting to the complexity of each scene. The selected examples include objects with varying textures, colors, and geometries, highlighting the versatility of our approach.

1.6. Additional Examples of Foveated Rendering

In this section, we showcase the application of our foveated rendering method across three different scenes: Counter, Playroom, and Train. Each scene contains two renderings for comparison, where the foveation is applied at different regions. The method selectively enhances details in areas of interest while reducing computational load in less critical regions, effectively balancing visual quality and performance. Fig. 5 demonstrates how our approach preserves important details in high-attention areas while applying aggressive downsampling elsewhere.

1.7. Additional Examples of Promptable 3D Focus

To further demonstrate the effectiveness of our promptable 3D focus application, we provide additional examples in Figure 6. These results highlight how our method enables selective object emphasis while maintaining 3D consistency. Each row in the figure presents different objects highlighted from a scene using text-based prompts, showcasing the robustness of our 3D voting mechanism. The approach effectively suppresses background details while preserving high-frequency structures for the selected objects, enabling seamless editing without introducing segmentation artifacts.

1.8. Additional 3D Artistic Filters

We present more examples of artistic filters applied to different scenes using our method, as shown in Figure 7. These effects are achieved by selectively modifying the attributes of Gaussians at different frequency levels.

For the **Brush** effect, we modify the 5-level trained model by keeping level 1 unchanged, while brightening levels 2 and 4 and darkening levels 3 and 5. Additionally, we introduce Gaussian noise to the center positions of Gaussians at levels 2 through 5, resulting in a painterly, textured appearance.

For the **X-ray** effect, we set the colors of Gaussians at levels 1 and 2 to a constant dark shade, simulating an underlying structure. Level 3 is brightened moderately, while levels 4 and 5 are brightened further, enhancing high-frequency details to create a glowing, skeletal appearance.

For the **Sharp** effect, we entirely remove level 2 and set the opacity of levels 3–5 to full, emphasizing fine details. To further enhance contrast, we darken the colors of Gaussians at level 5. This manipulation improves perceived sharpness while maintaining 3D consistency.

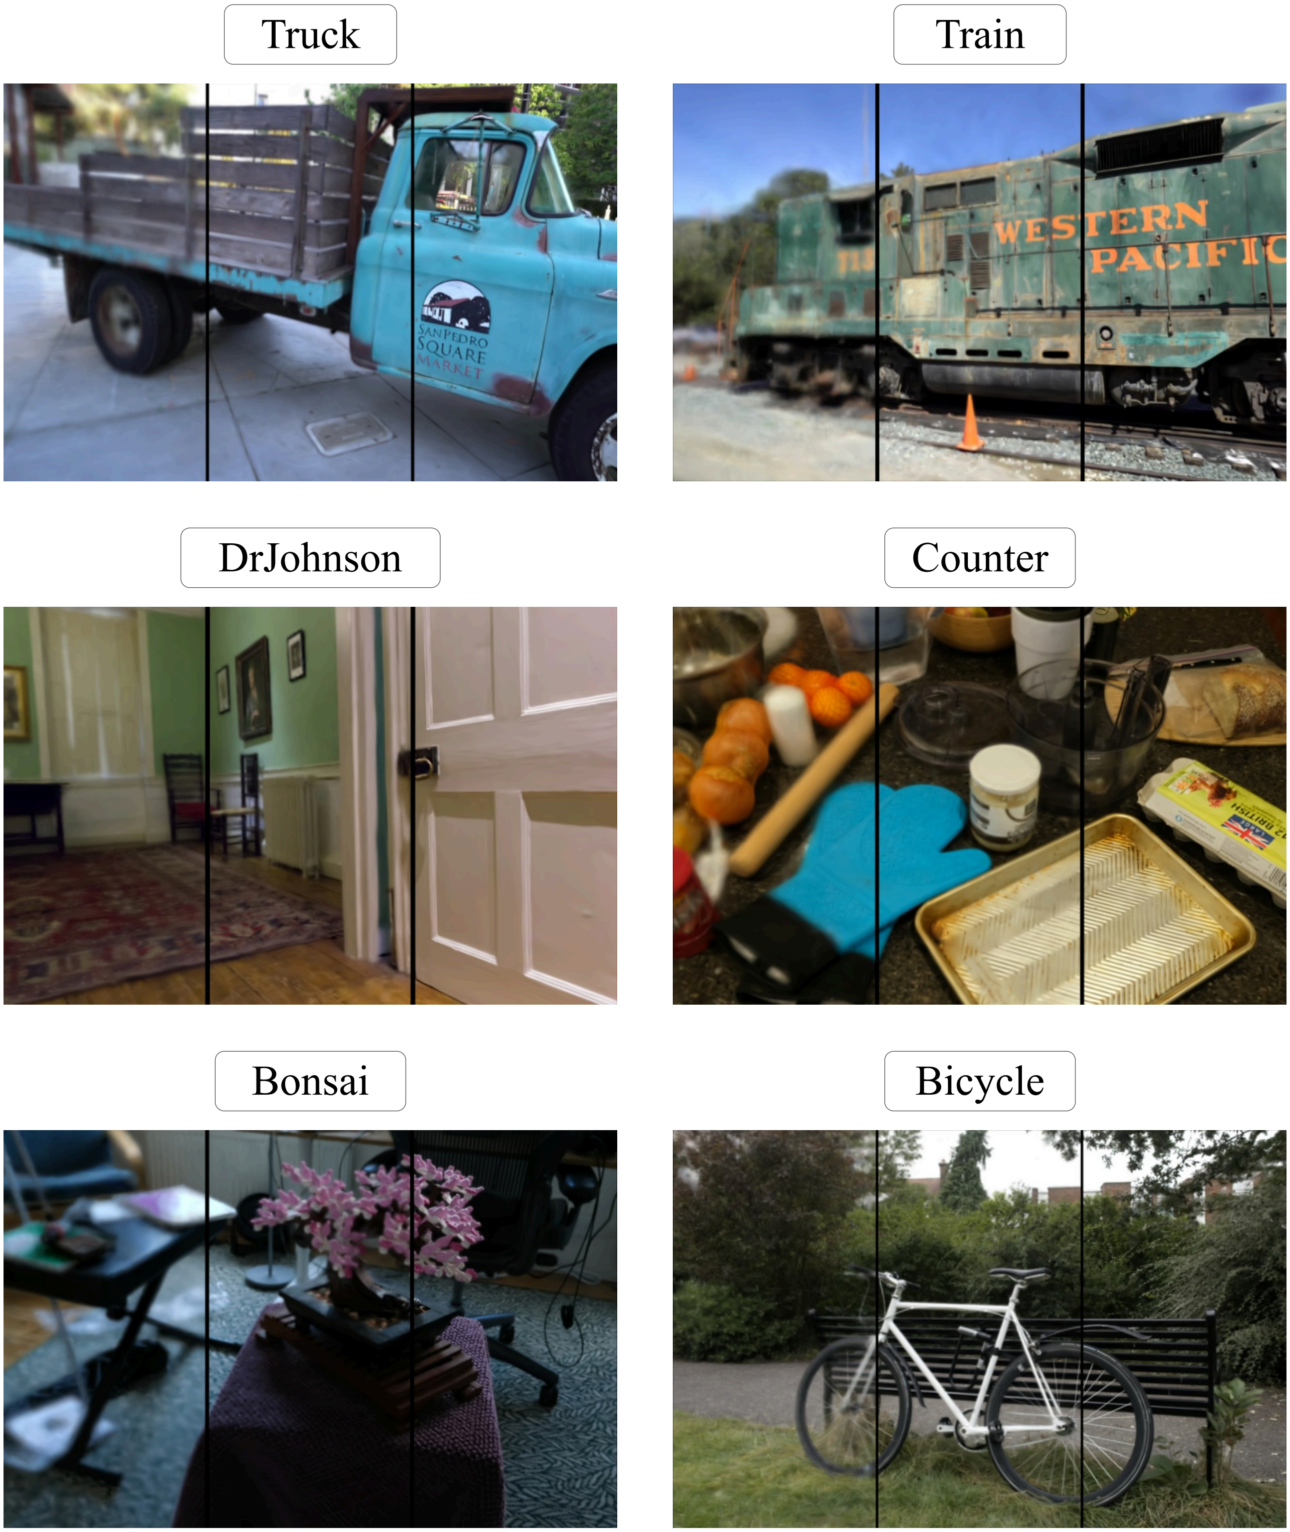

Figure 4. Visualization of the hierarchical detail representation across various scenes using our method. Each image is segmented to illustrate different levels of structural and textural detail captured by our approach. The selected examples—*Truck*, *Train*, *Dr. Johnson*, *Counter*, *Bonsai*, and *Bicycle*—demonstrate its effectiveness across diverse real-world scenarios. Our model is trained with 5 levels, and we render levels 1, 3, and 5 for comparison.

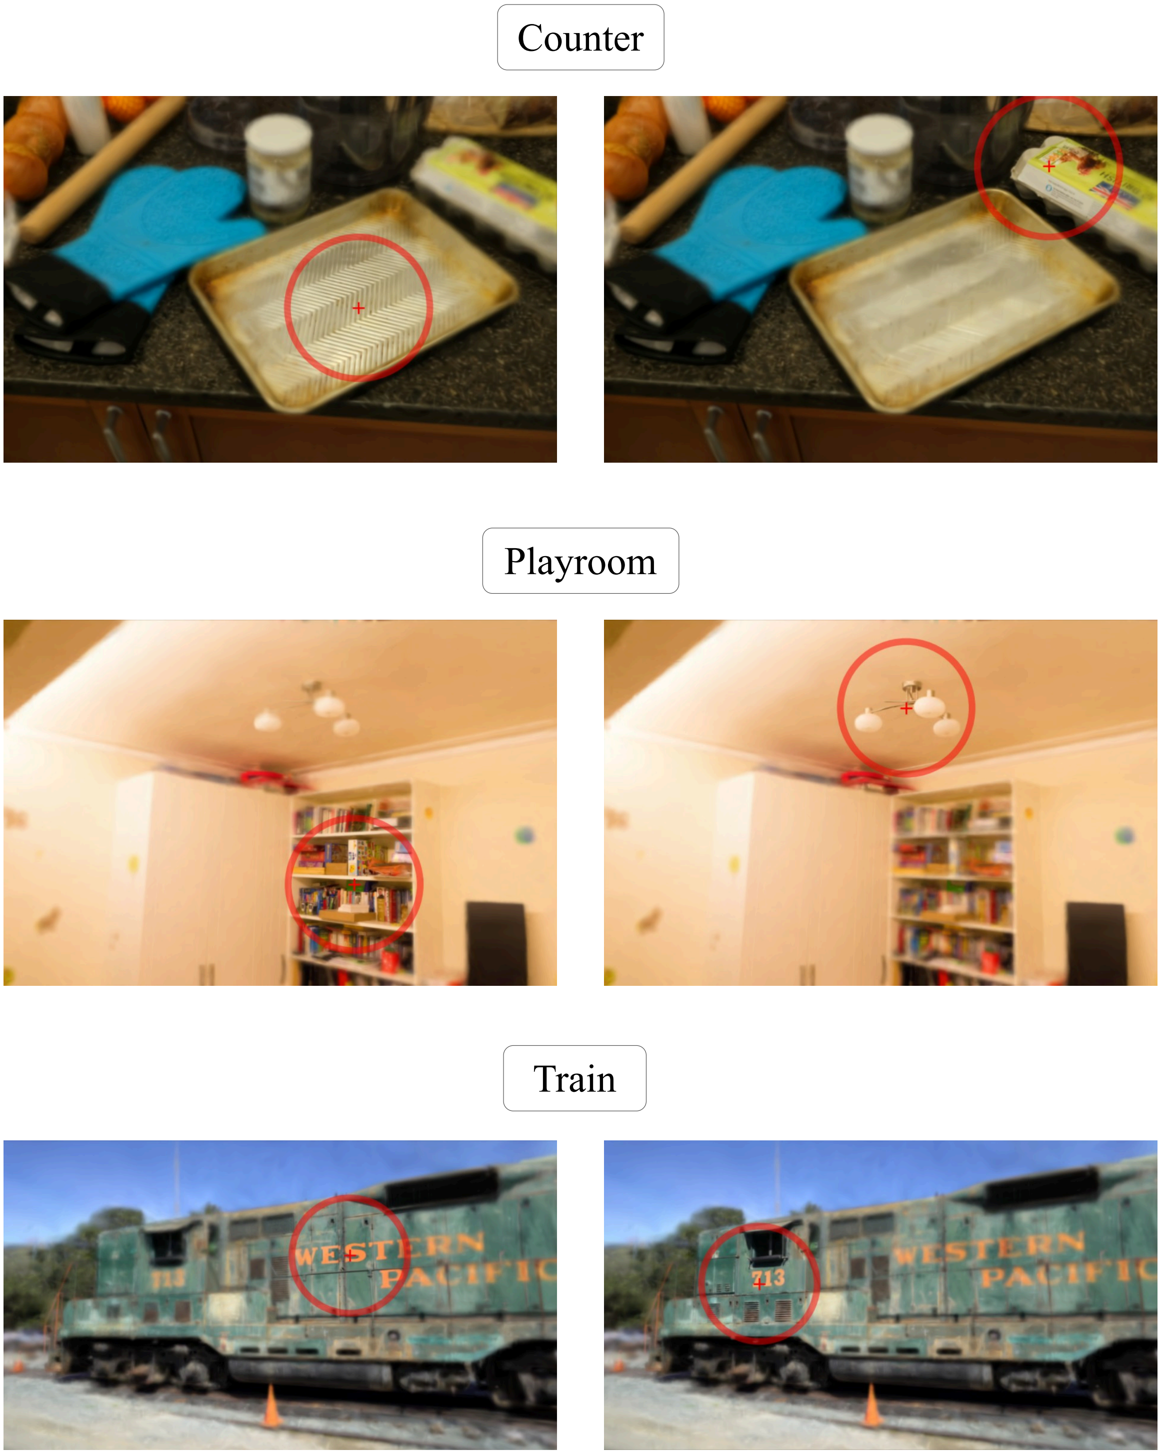

Figure 5. Foveated rendering applied to three different scenes: Counter, Playroom, and Train. Each scene contains two renderings with foveation applied to different areas (highlighted in red). The method prioritizes high-detail rendering in regions of interest while reducing detail elsewhere, optimizing rendering efficiency.

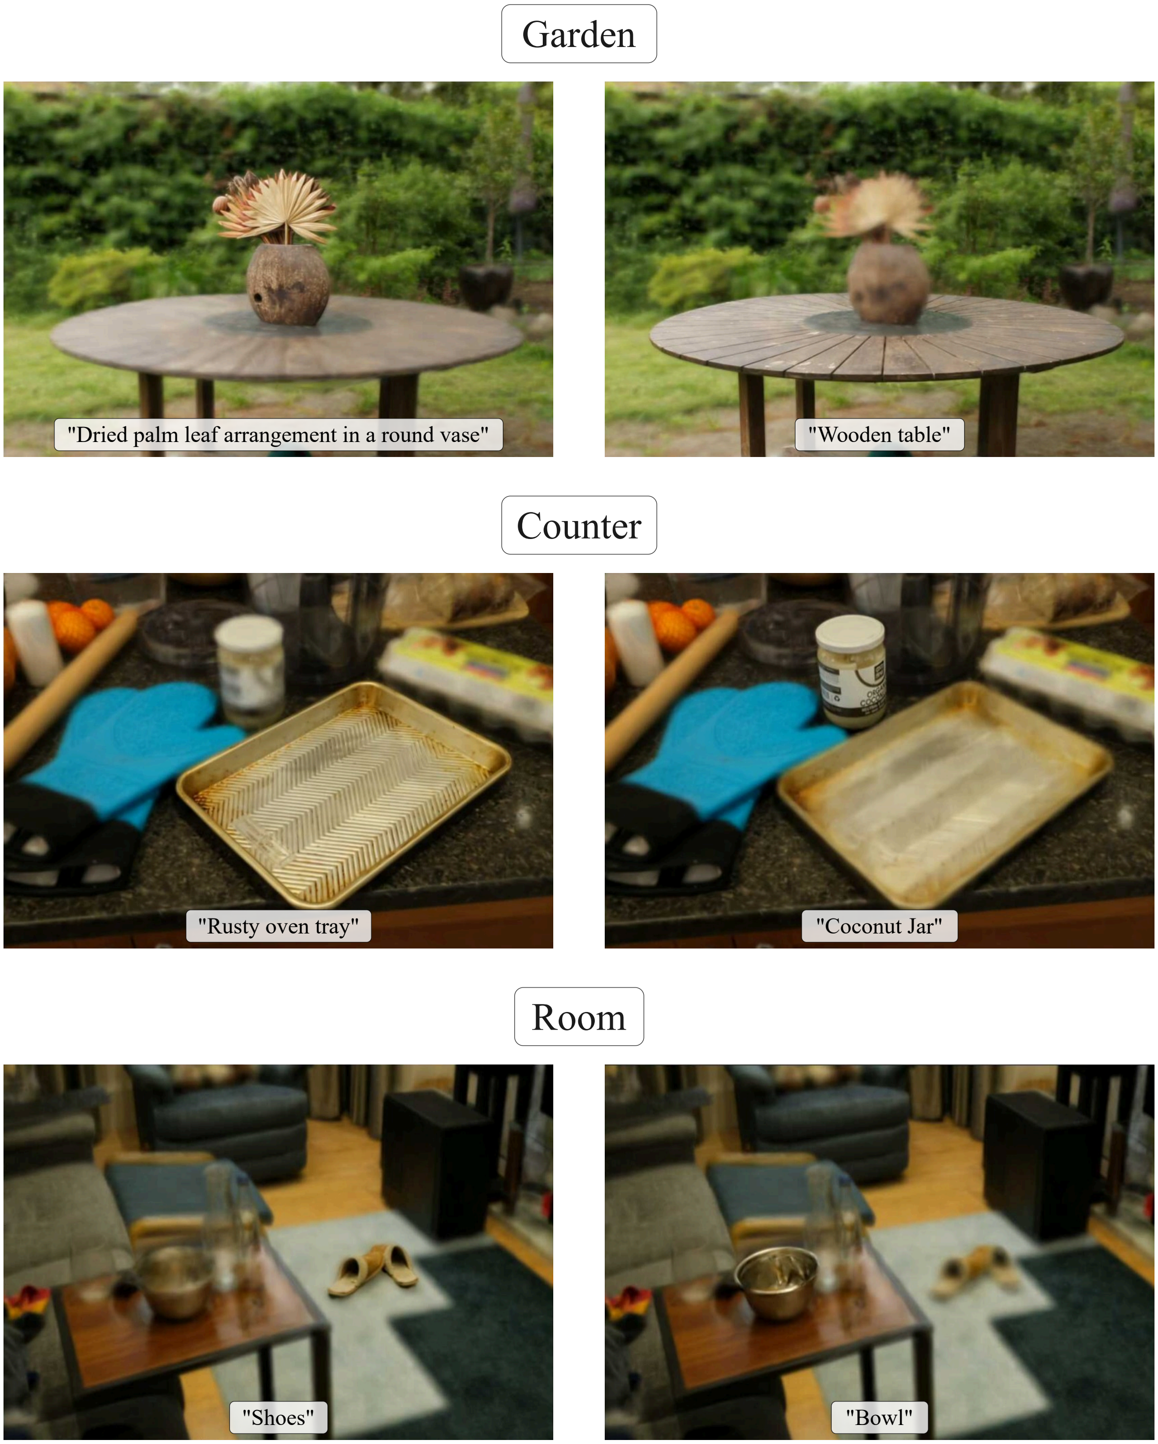

Figure 6. Additional results of our promptable 3D focus application. Each row shows a different object selected using a text prompt, demonstrating our method's ability to produce consistent 3D object emphasis.

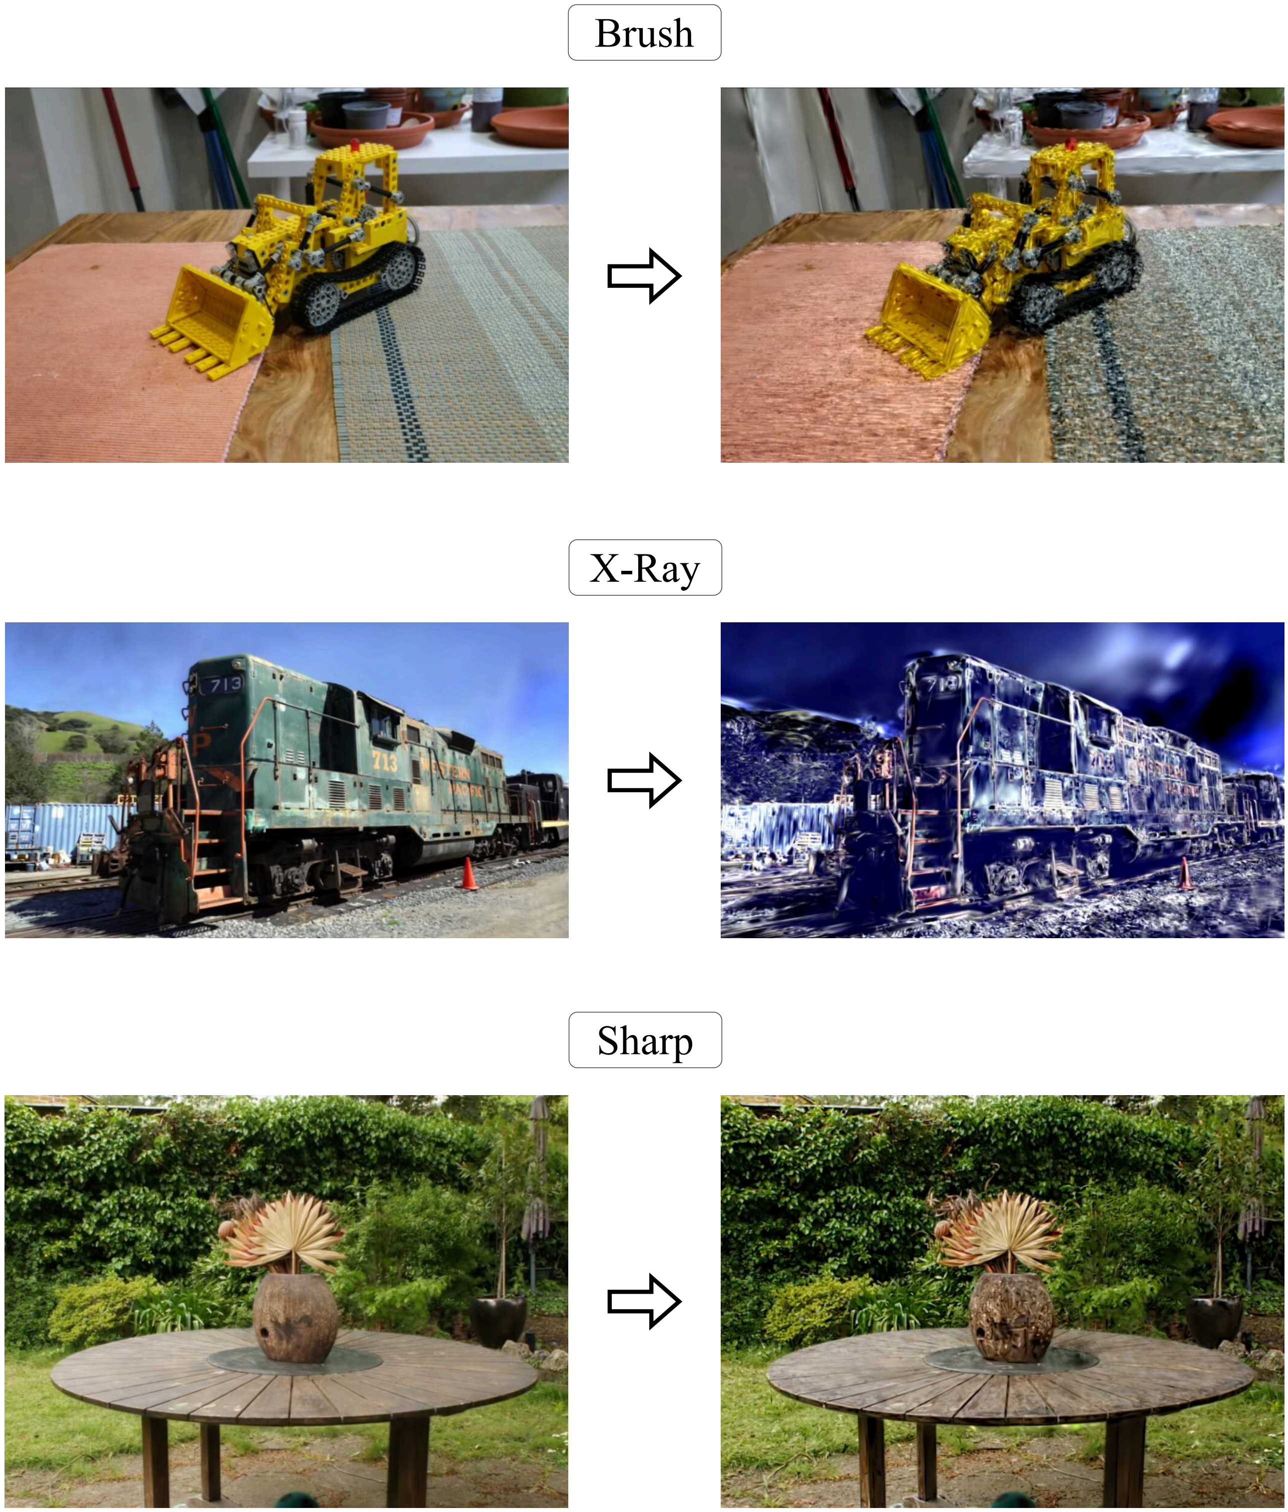

Figure 7. Three artistic effects applied to different scenes using our method. **Top:** The "Brush" effect introduces texture and painterly distortions by adjusting frequency bands and adding positional noise. **Middle:** The "X-ray" effect highlights high-frequency details while suppressing lower levels to create a glowing structural representation. **Bottom:** The "Sharp" effect enhances fine details by removing mid-frequency Gaussians and fully opacifying higher levels.