

Figure 1: Visual comparison of OSEDiff with and without fine-tuning the VAE encoder.




Figure 2: Visual comparison of outputs by OSEDiff, StableSR (2 samples) and SeeSR (2 samples) using the same LQ image.




Figure 3: Comparison of distributions: pre-trained SD (red), OSEDiff (green), and OSEDiff w/o VSD loss (orange).




Figure 4: Visual comparison of results using null prompts and tag-style prompts.




Figure 5: Visual comparison with and without VSD loss.