

SUPERGAUSSIAN: Repurposing Video Models for 3D Super Resolution

Yuan Shen^{1,2}^{*}, Duygu Ceylan², Paul Guerrero², Zexiang Xu²,
Niloy J. Mitra^{2,3}, Shenlong Wang¹, and Anna Frühstück²

¹ University of Illinois at Urbana-Champaign

² Adobe Research

³ University College London

Abstract. We present a simple, modular, and generic method that up-samples coarse 3D models by adding geometric and appearance details. While generative 3D models now exist, they do not yet match the quality of their counterparts in image and video domains. We demonstrate that it is possible to directly repurpose existing (pre-trained) video models for 3D super-resolution and thus sidestep the problem of the shortage of large repositories of high-quality 3D training models. We describe how to repurpose video upsampling models – which are not 3D consistent – and combine them with 3D consolidation to produce 3D-consistent results. As output, we produce high-quality Gaussian Splat models, which are object-centric and effective. Our method is category-agnostic and can be easily incorporated into existing 3D workflows. We evaluate our proposed SUPERGAUSSIAN on a variety of 3D inputs, which are diverse both in terms of complexity and representation (e.g., Gaussian Splats or NeRFs), and demonstrate that our simple method significantly improves the fidelity of current generative 3D models.

Check our project website for details: supergaussian.github.io.

Keywords: 3D super resolution · video upsampling · category-agnostic · 3D generation · Gaussian splatting · 3D-consistent

1 Introduction

Generative 3D models have become a reality. Multiple methods [2, 25, 31, 46] have been developed to generate 3D models, optionally conditioned using text prompts or images. These methods use a combination of image and 3D data for supervision, are fast, and produce diverse results.

Unfortunately, the 3D models generated by the current methods still lack the level of detail and accuracy achieved by state-of-the-art generative models for images [26, 40] or videos [6]. Multiple challenges contribute to this limitation. First is the choice of 3D representation. While grid-based models are most popular as they do not need prior knowledge of the (generated) shape, their regular structure (e.g., volume grid [25, 34], triplanes [2, 5]) puts a limit on the fidelity

^{*} This project was done during Yuan’s internship at Adobe Research.

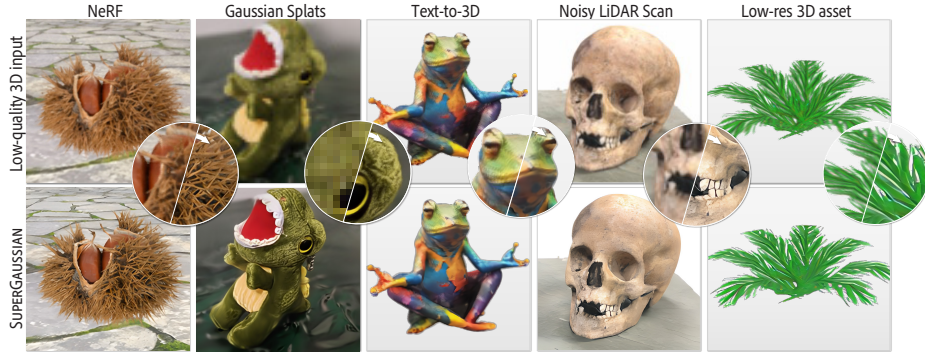


Fig. 1: We present SUPERGAUSSIAN, a novel method that repurposes existing video upsampling methods for the 3D superresolution task. SUPERGAUSSIAN can handle various input types such as NeRFs, Gaussian Splats, reconstructions obtained from noisy scans, models generated by recent text-to-3D methods [31], or low-poly meshes (e.g., assets used in Sim-on-Wheels [45]). SUPERGAUSSIAN generates high-resolution 3D outputs with rich geometric and texture details in the form of Gaussian Splats.

of the generation results. Secondly, acquiring large volumes of high-quality yet diverse 3D data continues to be difficult. While state-of-the-art image and video models are trained on several billion training samples, the most extensive 3D training datasets, at best, contain a few million objects.

We explore how to increase the fidelity of generated 3D objects. Starting from any generic coarse 3D representation, our goal is to ‘upsample’ the (coarse) 3D input model *without* requiring category-specific training. Our main observation is that any 3D representation can be rendered from multiple viewpoints along a smooth trajectory and mapped to an intermediate generic video representation. Hence, it is possible to repurpose existing video models [9, 69] to perform the 3D upsampling or super-resolution task. Such models are trained on large sets of video data, providing strong priors that can be applied in general scenarios. The critical challenge is to ensure 3D consistency; while being temporally smooth, video models, however, are not guaranteed to be 3D consistent. Note that our approach of using video, unlike those using image-based priors and treating each frame independently, significantly improves the (initial) consistency across time.

We address this challenge with a simple, modular, and generic approach that can be integrated into existing workflows. We proceed as follows: First, we render a video of the scene from the coarse 3D input given sampled view trajectories; second, we upsample the rendered video using a pretrained *video-based upsampler* that is optionally fine-tuned to handle domain-specific artifacts of the input modality. For 3D consolidation, we adopt Gaussian Splatting [27] as our output representation. Being an object-centric representation, Gaussian splats are ideally suited for encoding individual objects and are capable of capturing local details. Gaussian splats also strike a good balance between simplicity, fidelity of encoded models, and efficiency of rendering. Fig. 1 shows that our proposed

video-based 3D super-resolution framework can be successfully adapted to up-scale diverse sets of coarse 3D modalities.

We evaluate our algorithm on various low-resolution assets and report the quality and diversity of the generated high-resolution Gaussian splats. We compare our algorithm against possible alternatives and report an ablation study of our design choices. In summary, our main contributions include:

- (i) re-purposing pretrained video upsampling models for 3D super-resolution;
- (ii) finetuning video upsampling prior using generic 3D data to handle artifacts characteristic to low-res Gaussian splats; and
- (iii) proposing SUPERGAUSSIAN as an efficient, class-agnostic, and modular 3D super-resolution procedure that successfully upsamples diverse 3D scene representation, including Gaussian splats, NeRF, and others.

2 Related Work

Image Super-Resolution. Enhancing image details and resolution is a long-standing problem in computer vision. This is inherently an ill-posed problem, so generative methods are well suited to tackle the hallucination of missing information. For conciseness, we concentrate on recent generative approaches and refer the reader to a comprehensive review [11] of this subject area.

Early CNN-based approaches [14, 15, 28, 66] recognize the applicability of deep networks for super-resolution. Several approaches leverage powerful priors learned by image synthesis models to infer image details. Generative Adversarial Networks [17] have been extensively used for this task, either through optimization-based techniques [36] or conditional feed-forward network designs [3, 24, 29, 53–55]. Other approaches have explored normalizing flows [35] or transformer-based methods [12, 13]. The immensely popular denoising diffusion probabilistic models [22, 41] have also been recently explored for image upsampling [30, 42, 63].

Video Super-Resolution. Following the success of learning-based single-image super-resolution methods, there have been various efforts to extend this success to the video domain. Several strategies have been explored to make this extension. While some work [33] introduces temporal aggregation modules to aggregate information across time, recurrent neural networks [8, 60] with uni- and bi-directional information sharing have also been a choice of architecture. Deformable convolutions [48, 52] have been utilized to align per-frame features both for video upsampling and restoration tasks. More recently, diffusion-based video upsamplers have been proposed by injecting temporal layers into base diffusion model [69]. In our work, we leverage VideoGigaGAN [57], a state-of-the-art generative video upsampler. However, our method is agnostic to the video model we use and can further be improved in the future with advances in video upsampling techniques, e.g., diffusion-based video upsampler.

Image-Based 3D Super-Resolution. Several approaches have recently aimed at increasing the resolution of implicit radiance fields, in particular Neural Radiance Fields (NeRFs). This is an important and practical application for both capture and synthesis scenarios where higher detail and quality are desirable. Start-

ing with a low-res radiance field, at each iteration, this thread of work [19, 23, 49] renders images from the current representation. Next, a pre-trained image upsampler is used to upsample the renderings. Then the gradient flow coming from the difference between the current renderings and upsampled images, is used to optimize the 3D representation. Similarly, CROP [61] also leverages a single image upsampler but introduces an uncertainty volume to improve the multi-view consistency. Unlike these methods, however, we adopt a video upsampler instead of a single-image upsampling network. Thanks to better handling of temporal consistency, we show that using the video upsampler improves the visual quality of the final results. Moreover, unlike previous work mainly focusing on upsampling NeRF representations, SUPERGAUSSIAN is a generic approach, which we demonstrate on various 3D input types.

In the context of 3D-aware GANs, several works have been proposed [7, 16] that jointly learn to generate a triplane-based 3D representation as well as a 2D image upsampling module. The upsampling module is applicable only to the specific triplane features learned by the generator and the final output is not fully 3D consistent since each view is upsampled independently.

Feed-forward 3D Super-Resolution. Several 3D generation and reconstruction methods have proposed direct 3D superresolution modules. Smith et al. [47] represent a 3D shape as multi-view orthographic depth maps which are then upsampled. The recent diffusion-based 3D generation method LASDiffusion [68] presents a method to predict signed distance fields from an occupancy map to improve the resolution of the 3D shapes. Both methods, however, focus only on geometry refinement while our work jointly performs appearance and geometry upsampling. Rodin [51], a triplane-based diffusion method, presents a triplane upsampling strategy. This approach is applicable only to the specific triplane features learned. Finally, our approach does not need large-scale 3D datasets and instead leverages existing video models trained on video data.

Deep Implicit Representations. Recently, neural implicit representations such as NeRF [38], VolSDF [59], and NeuS [50] have been popularized as an efficient and compact representation for 3D shapes. Efficiently capturing high-resolution details, however, still remains a challenge for such representations due to the low-dimensional volumetric structure and the computationally intensive volumetric rendering process. Recently, Gaussian Splatting [27] has been proposed as an alternative that provides an object-centric representation with fast rendering capabilities. The scene is represented by a large number of 3D Gaussian blobs whose properties are optimized through differentiable rendering. In our work, we adapt Gaussian splats as the final 3D representation output due to the efficient optimization and rendering capabilities they offer.

3 Our Approach

3.1 Overview

Given a coarse 3D representation, our goal is to perform super-resolution to improve the fidelity of the 3D representation and capture more local details. Based

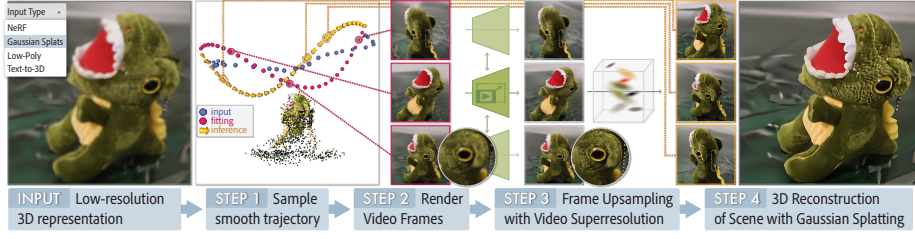


Fig. 2: SUPERGAUSSIAN pipeline. Given an input low-res 3D representation, which can be in various formats, we first sample a smooth camera trajectory and render an intermediate low-resolution video. We then upsample this video using existing video upsamplers and obtain a higher resolution 3D representation that has sharper and more vivid details. Our method, SUPERGAUSSIAN, produces a final 3D representation in the form of high-resolution Gaussian Splats.

on the observation that 3D content can be represented as a video depicting the 3D scene from multiple viewpoints, the main premise of our work is to leverage existing video upsampling priors for the task of 3D upsampling. As illustrated in Fig. 2, our method SUPERGAUSSIAN consists of two main steps. First, we run upsampling on the video rendered from the coarse 3D representation to gain resolution and obtain sharp results. Then, we perform 3D reconstruction to generate a consistent 3D representation. In addition to leveraging the prior of the video upsampler, which is trained on a large set of video data, we perform fine-tuning over domain-specific low-res videos (i.e., videos rendered from low-res 3D representations). Hence, SUPERGAUSSIAN is able to handle complicated degradation caused by various 3D capture and generation processes. Note that each component in our framework is highly modularized and could easily be replaced with other state-of-the-art video methods.

Next, we unfold this section by formulating the input and output of our method (Sec. 3.2), introducing the key video upsampling (Sec. 3.3) in our framework, and discussing the details of the 3D optimization (Sec. 3.4).

3.2 Problem Formulation

Our framework can handle a diverse set of coarse 3D representations of a static scene, which we denote as ψ_{low} . For example, ψ_{low} can be Gaussian Splats, NeRF, a low-poly mesh, a low-quality captured video, or a generated 3D object with recent text-to-3D methods [31]. Any such 3D representation, ψ_{low} , can be rendered from multiple viewpoints to yield a video, a common intermediate representation. In practice, we render each 3D input from N smooth trajectories with viewpoints $\xi_{1:T}^{1:N}$ resulting in a sequence of RGB images $I_{1:T}^{1:N}$. Here, the subscript indexes the viewpoint or the pose along each trajectory, and the superscript denotes the trajectory ID. We assume the camera movement between adjacent frames is sufficiently small such that a standard video upsampler can leverage enough temporal alignment. After we perform a sequence of video up-

sampling, as output, we generate a high-fidelity 3D representation ψ_{high} in the form of Gaussian Splats (note that camera views are known in our setup and do not need to be estimated). This final 3D optimization produces a true 3D output and, in the process, removes any remaining temporal inconsistencies in the refined video representations.

3.3 Initial Upsampling

First, we manually sample trajectories near the target scene in the empty region. Given the trajectory $I_{1...T} \in \{\mathbb{R}^{W \times H \times 3}\}$ describing the camera path for each individual video, the video upsampler outputs a trajectory with $\times r$ up-sampling ($r = 4$ in our experiments), $\hat{I}_{1...T} \in \mathbb{R}^{W \times rH \times 3}$. We assume the initial rendering resolution should be sufficiently high such that the rendered fidelity is bottlenecked by the coarse level of the input 3D representation.

Our framework can easily integrate any of the state-of-the-art pre-trained video upsamplers. For our case, we use VideoGigaGAN [57], a generative video upsampler. Please refer to Sec. 4.1 to see how we set up the video upsampler to have fair comparisons with baselines. In cases of handling input representations with severe domain bias (see Fig. 6), additional fine-tuning is desired. For example, with stripy or blob-like artifacts after zooming in, renderings from Gaussian Splats follow different degradation from standard augmentations deployed in state-of-the-art video upsamplers. Hence, to finetune the video upsampler, we need pairs of low and high-resolution videos that depict the specific degradation we would like to model. For this purpose, we use the multi-view dataset MVImgNet [62], which depicts a variety of 3D objects and scenes. First, we bilinearly downsample the original images in the dataset by a factor of 8, i.e., to 64×64 px resolution, to obtain a set of low-resolution images. We then fit low-resolution Gaussian Splats to these images [27] as described in more detail in Section 3.4. We render the optimized low-res Gaussians in the original camera trajectory provided by the dataset as input to the video upsampler. As the target ground truth, we use the original videos from the dataset, which are resized to be 4 times the resolution of the input. We finetune the model using the Charbonnier regression loss [4] for its robustness to outlier pixels, the LPIPS loss [64] for its perceptual level improvements, and GAN loss [18] for its generative behavior.

3.4 3D Optimization via Gaussian Splats

We use the official Gaussian Splatting codebase⁴ to perform 3D optimization to fit Gaussians to the upsampled videos. In particular, for our experiments, we follow the standard Gaussian Splatting optimization process by running 2k steps. Note that as we have perfect camera information, we directly provide them to the optimization, rather than estimating them using SfM. As for our loss function, following the practice of the original paper [27], we use L_1 and L_{SSIM} losses when optimizing the Gaussians. The advantages of adopting 3D Gaussian Splatting lie

⁴ <https://github.com/graphdeco-inria/gaussian-splatting>

in its being an object-centric representation and its efficiency in terms of training and rendering. Besides, it is able to capture view-dependent effects well for the upsampled frames. However, SUPERGAUSSIAN can also be easily integrated with other types of 3D representations, e.g., Neural Radiance Fields, which we use in our experiments on the Blender-synthetic dataset [38].

4 Experiments

4.1 Setup

Dataset: We finetune the video upsampler component of our method on the MV-ImageNet dataset [62], which consists of a rich number of real object-centric videos with high quality. This dataset is processed to fit low-resolution Gaussian Splats reconstructed from 64×64 px posed images as described in the previous section. In total, 200k scenes are used for training and 5k scenes for validation. A separate set of 500 scenes are used for quantitative evaluation. We obtain the ground-truth high-resolution images by center-cropping and resizing them to square images at 256×256 px resolution on MVImageNet. To demonstrate the generalizability of our approach, we also show results on the Blender-synthetic dataset, consisting of synthetic 3D objects. We use the original data split as prior work [32, 49] for evaluation.

Metrics: On the MVImageNet dataset, we utilize metrics focused on perceptual quality when performing quantitative comparisons, i.e., LPIPS [65], NIQE [39], Inception score [43] and FID [21]. For the FID computation, we use the real images provided by MVImageNet as the real distribution. On the NeRF Synthetic dataset, we report the same metrics, i.e., LPIPS, PSNR, and SSIM, analogous to FastSR-NeRF [32] and NeRF-SR [49]. Aside from quantitative metrics, we strongly encourage the readers to check our supplementary website for demo.

Baselines: While there are a recent few domain-specific 3D upsampling solutions [23, 49, 58], we focus on comparing against baselines that are applicable to generic 3D representations. Previous work for generic 3D upsampling follows a similar paradigm by leveraging upsampling priors to iteratively optimize the chosen 3D representation. The popular choice of the prior is image-based upsampling models, instead of our proposed video-based upsampler, one seemingly trivial design choice that profoundly impacts the upsampled quality. With that, we narrow down to three state-of-the-art methods: (1) Instruct-G2G (Instruct-NeRF2NeRF [20] with Gaussian Splats as 3D format), (2) Super-NeRF [19] with Gaussian Splats as 3D format, and (3) FastSR-NeRF [67]. For completeness, we include one domain-specific baseline, NeRF-SR [49]. Besides, we introduce a customized baseline, *Pre-hoc image*, as discussed below.

For a fair comparison, we use the same state-of-the-art image upsampler, GigaGAN [24] for all the baselines. Specifically for the GigaGAN checkpoint, we follow the same strategy proposed in the original paper to train a $4\times$ super-resolution model on LAION [44] that learns to upsample images from 64×64 to

Table 1: Benchmark of video and image upsampler priors. We report the performance of the upsampled frames for pre-hoc image and ours after applying the corresponding upsampler. Note: reconstruction is not performed for these results.

Model	LPIPS ↓	NIQE ↓	FID ↓	IS ↑
image upsampler	0.1533	7.06	14.67	11.67 ± 1.14
video upsampler	0.1511	6.04	14.61	12.45 ± 1.19

256 × 256px with great details. In SUPERGAUSSIAN, we use VideoGigaGAN [57] as the video upsampler, which mostly re-uses the former GigaGAN architecture except for additional BasicVSR++ layers for temporal feature extraction and propagation [9] that processes the low-res video input frames. The model size of our video prior turns out to be slightly smaller than the image prior after adjusting feature dimensions. Finally, we finetune both the image and video upsampler on MVImgNet until convergence. Benchmark of both priors are reported in Tab. 1. Both priors show generative behaviors after upsampling, and the video prior shows great temporal consistency across upsampled frames. Note SUPERGAUSSIAN can be easily integrated with a stronger video prior, e.g., Upscale-a-video [69] or Sora from OpenAI [1] by finetuning over low-res video as conditions. For Instruct-G2G, we first optimize 2k steps for the warm-up stage, which prepares the low-res representation, and then optimize another 2k steps with upsampling by upsampling every 20 steps. For SuperNeRF, since the authors did not release their codebase, we self-implemented their method with our best knowledge. We wrap our GigaGAN with the consistency enforcement module [3]. Similarly, we optimize 2k step for low-res representation and then optimize for another 2k steps. We set the learning rate of the CCLC embedding to be 5×10^{-5} , which is critical to be small to avoid blurry results. In our experiments, we observe that running an iterative optimization with an image upsampler, i.e., SuperNeRF [19] and Instruct-G2G, often results in more blurry results compared to upsampling all the frames individually first, in a *pre-hoc* manner, and then performing an optimization. Hence, we define another baseline, *Pre-hoc image*, where we individually upsample all the frames once and then optimize the 3D representation.

Finally, to remove the effect of representation on the experiment results, we modify our baselines to all use 3D Gaussian Splattings as their global representation for comparisons on MVImgNet. On the Blender synthetic dataset, Neural Radiance field is chosen as the global representation across methods to fairly compare with other baselines represented in NeRF.

Implementation details: During finetuning, all model parameters are finetuned with a learning rate of 5×10^{-5} . We weigh the LPIPS, Charbonnier regression losses, and adversarial loss with weights 15, 10, and 0.05, respectively. Following GigaGAN [24], we apply R_1 regularization [37] on the discriminator at a weight of 0.02048 and interval of 16. Batch size is 64 with 12 video frames in each. Our models are trained on 64 A100 GPUs, each with 80GB VRAM.

Table 2: Quantitative Comparisons on MVImgNet to upsample low-res Gaussian Splattings. We compare our methods against baselines using perceptual metrics. Our method, even though it is generic, consistently produces the best quantitative results. Instruct-G2G stands for Instruct-N2N with 3D Gaussian Splats as its 3D representation. We encourage the reader to inspect the visual results in our website, which highlights that the visual quality of our method surpasses the baselines.

Method	LPIPS ↓	NIQE ↓	FID ↓	IS ↑
Instruct-G2G [20]	0.1867	8.33	32.56	10.52 ± 1.06
Super-NeRF [19]	0.2204	8.84	37.54	10.40 ± 1.03
Pre-hoc image	0.1524	7.65	27.04	11.27 ± 0.99
Ours	0.1290	6.80	24.32	11.69 ± 1.08

Table 3: Quantitative Comparison on Blender synthetic dataset to upsample low-res posed RGB images. Here, we compare on $\times 4$ upsampling from 200×200 to 800×800 px. We compare our methods against baselines on the official test set using metrics reported in prior work. Our method produces on-par quantitative results. Besides, our results yield more generative details, which are not captured by the reference metrics. For a fair comparison, we use Neural Radiance Field, i.e., TensoRF [10], as our 3D representation. Other baseline results are directly taken from their paper.

Method	LPIPS ↓	PSNR ↑	SSIM ↑
FastSR-NeRF [32]	0.075	30.47	0.944
NeRF-SR [49]	0.076	28.46	0.921
Ours	0.067	28.44	0.923

For a set with 30 frames, the video upsampling takes ~ 1 second to upsample 64×64 inputs. High-resolution Gaussian splat optimization takes around 30 seconds for 2k steps to fit images of resolution 256×256 . On average, our pipeline takes 141.33 sec to complete, the most efficient compared with baselines. We note that our codebase can be further optimized for efficiency, e.g., in I/O.

4.2 Comparison Studies

Experiment protocol: On MVImgNet, for quantitative evaluations, given a low-res 3D representation, we prepare a smooth camera trajectory that passes through a set of camera poses perturbed from the original camera poses provided by the corresponding datasets. We use a video rendered from the resulting camera trajectory as input to the video upsampler. To perform pose perturbation, we randomly sample camera centers along the line segments between the scene center and each pose in the original camera trajectory. We fit a B-Spline curve to such sampled camera centers and resample camera positions on the resulting curve; finally, we add sinusoidal perturbation along the up direction. With the camera positions settled, we set the orientation of each camera by defining a look-at vector towards the scene center. The resulting perturbed camera trajectories we obtain differ from the original camera trajectories we use for finetuning



Fig. 3: Qualitative comparisons on MVImgNet to upsample low-res Gaussian Splatting. The low-resolution inputs are novel views rendered from fitted low-res Gaussian Splatting along the sampled trajectories. Our method produces the best upsampled *3D scenes* with generative texture and geometric level details preserved. Please check our supplementary website to interactively compare results in 3D.



Fig. 4: Qualitative results of SUPERGAUSSIAN on Blender synthetic dataset. Besides increased sharpness, SUPERGAUSSIAN can nicely lift *generative* details from video upsampling to 3D. The background image is our upsampled 3D NeRF. The blue, yellow, and red boxes indicate a zoom-in view of low-res, our upsampled results, and ground truth. The test pose IDs from left to right are `chair_40`, `hotdog_36`, `mic_60`, `lego_18`, which are chosen to align with Fig. 5 in NeRF-SR [49] for direct comparison.



Fig. 5: Qualitative evaluation of both video and image upsampling priors for 3D upsampling on Wild RGB-D dataset [56] without finetuning. The results visualize a 3D upsampled output from one test scene. Compared with using the image upsampler, adopting the video upsampler can distill more detail to 3D, as can be seen (for example) from the crisper characters of the book.

the video upsampler on the MVImgNet dataset. Once we obtain a high-res 3D representation, to evaluate the result quantitatively, we use the same perturbation strategy to sample a novel evaluation trajectory that is different than the trajectories used for rendering inputs to the video models. To prepare the corresponding ground truth on this evaluation trajectory, we first fit Gaussian splats to the original high-resolution images provided by the dataset and render views from the cameras on the evaluation trajectory. When testing ours on examples where we do not have access to ground truth, camera paths, or high-resolution images, we sample trajectories close to the region of interest with sufficiently high resolution that reflects the fidelity of the target 3D representation.

On the Blender synthetic dataset, we follow the same experiment protocol for $\times 4$ upsampling as prior work [32, 49]. Specifically, we set the target upsampled resolution fixed at 800×800 px and bilinearly downsample the target images to prepare low-res images to 200×200 px. For fair comparisons, we also use Neural Radiance Field for 3D representation, instead of Gaussian Splattings. Results are evaluated on the original held-out test set in Blender synthetic dataset.

Results and Analysis: We provide quantitative results in Tab. 2 and visual examples in Fig. 3 on MVImgNet. Using video as an intermediate representa-



Fig. 6: SUPERGAUSSIAN over severely degraded input modalities. By finetuning on 10% data over MVImgNet, SUPERGAUSSIAN is able to handle severe artifacts from the low-res modality. In particular, the Chinese characters in the NeRF case become legible after upsampling. The low-res Gaussian Splats is obtained by fitting with only 4096 Gaussian splats, and NeRF by optimizing Nerfacto for only 1000 steps.



Fig. 7: Upsampling real-world 3D Gaussian Splats. SUPERGAUSSIAN can remove Gaussian artifacts and increase sharpness with off-the-shelf video upsampler.

tion, as opposed to individual images, is effective for 3D tasks. We observe that image upsampling baselines tend to create blurry upsampled 3D content due to inconsistency across upsampled images, as can clearly be seen in the supplementary videos. However, our method, which utilizes video priors, provides the best quantitative metrics and high-quality visual results. Similarly, we achieve the best LPIPS on the Blender-synthetic dataset in Tab. 3 and Fig. 4. Our results yield more generative details, not captured by other reference-based metrics.

4.3 Ablation Studies

We perform detailed ablation studies to evaluate different design decisions. First, we assess whether an image upsampler or video upsampler is better suited for the task of 3D upsampling. Then, we demonstrate via qualitative examples the usefulness of finetuning to handle severely degraded input modality, including both Neural Radiance Field and 3D Gaussian Splatting. Finally, we investigate how trajectory sampling influences the final 3D upsampling quality by changing trajectories along normal and tangent directions of the target surface separately. **Video or image upsampling prior for 3D upsampling:** For fair comparison, as is described above, we carefully select two priors accordingly for image and



Fig. 8: Ablation on trajectory sampling if the upsampling path has a varying distance to the target during inference. A spiral-shape trajectory is sampled during 3D upsampling with different distances towards the target. With the upsampled 3D Gaussians, we evaluate the 3D upsampling quality at the same pose on the left.



Fig. 9: Ablation on trajectory sampling if evaluating over test trajectories with varying elevations relative to the upsampling path. We vary the elevation of the test trajectory while keeping distance-to-target constant. After upsampling at the gray path, columns 2-4 are the 25th frame at three test paths.

video, both of which use the similar GigaGAN asymmetric auto-encoder to do $\times 4$ upsampling and are finetuned on MVImgNet. As is shown in Tab. 2 and Fig. 3, our results clearly demonstrate that 3D upsampling with video prior produces sharper texture and geometric details. Meanwhile, image upsampling produces sharp but inconsistent results, hence causing blurriness after 3D reconstruction. For generalizability, we test both priors on another set, Wild RGB-D [56], by 3D upsampling over posed RGB images. An unseen scene type, **boat**, is chosen (Fig. 5). Without finetuning, LPIPS for prehoc img is 0.257 vs. **ours 0.242**.

Effect of finetuning on severely degraded low-res input: We prepare two sets of data for finetuning low-res NeRF and Gaussians. When preparing the low-res representation, instead of using full capacity to fit low-res images, we only use 4k Gaussian splats to intensify modality-induced degradation. For NeRF, we only optimize for 1k steps. As is shown in Fig. 6, after finetuning video upsampler over the above sets, SUPERGAUSSIAN can successfully upsample severely degraded 3D low-res inputs for diverse modalities.

Effect of trajectory sampling during inference: By varying the distance of the upsampling trajectory to the target, SUPERGAUSSIAN performs better if the trajectory is sampled closer to the target (Fig. 8). As is shown in Fig. 9, with

less occlusion, a test path closer to the upsampling path gives better performance.

Ablation on the output 3D representation: We ablate the output 3D format on ours with three output representations over the chestnut scene (check Fig. 1). The sampling trajectory is a circular path around the target, where we randomly sample training and test poses. Regarding LPIPS, Nerfacto-triplane gives 0.4669, Nerfacto-hashgrid gives 0.3495, and Gaussian Splats gives **0.1493**.

4.4 Additional Results

To demonstrate the generalizability of our method, we evaluate it on various types of input as shown in Fig. 1, Fig. 6 and 7. In particular, we show upsampling results over captured low-res NeRFs and low-res Gaussian splats, reconstructions obtained from noisy RGB-D sensors such as the iPhone, generated 3D results from Instant 3D [31], a low-resolution mesh, and one other real-world captured scene with dense 3D Gaussian Splats. Since the intermediate video representation can be rendered from any of these input types, SUPERGAUSSIAN provides a generic 3D superresolution framework. Additionally, we provide more qualitative results on other input types, i.e., indoor scene and more text-to-3D upsampling results, in our supplementary paper.

5 Conclusion

We have presented a simple, modular, and generic framework for improving 3D model fidelity by hallucinating geometric and textural details through an upsampling pipeline designed to repurpose existing video models. We demonstrated the universality of our approach, both in terms of handling multiple input modalities and *not* requiring any category-specific training. Given the simplicity and modular nature of our approach, we believe that SUPERGAUSSIAN can be easily integrated into most 3D content creation workflows.

Our method does have limitations: (i) Since we use pretrained video models, we cannot improve on their generalization and inference speed. While we cannot solve this by training on synthetic 3D data, we expect foundational video models to continue to improve, which can be easily used in our modular setup. (ii) We cannot recover from missing/occluded parts in the input or equivalently from insufficient viewpoint coverage. We would like to combine generative completion methods in our framework to use data priors (e.g., 3D priors or video priors) to complete and infill, as well as super-resolve the results. However, this can lead to the loss of our universality.

Acknowledgements

We thank Yiran Xu, Difan Liu, and Taesung Park for discussions related to VideoGigaGAN [57], whose support was essential to the development of Super-Gaussian. We also thank Nathan Carr for providing Gaussian Splat examples.

References

1. Sora: Creating video from text. <https://openai.com/sora>, accessed: 2024-03-05
2. Anciukevičius, T., Xu, Z., Fisher, M., Henderson, P., Bilen, H., Mitra, N.J., Guerrero, P.: Renderdiffusion: Image diffusion for 3d reconstruction, inpainting and generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2023)
3. Bahat, Y., Michaeli, T.: Explorable super resolution. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 2716–2725 (2020)
4. Barron, J.T.: A general and adaptive robust loss function. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
5. Bautista, M.A., Guo, P., Abnar, S., Talbott, W., Toshev, A., Chen, Z., Dinh, L., Zhai, S., Goh, H., Ulbricht, D., Dehghan, A., Susskind, J.: GAUDI: A neural architect for immersive 3d scene generation (2022)
6. Chai, W., Guo, X., Wang, G., Lu, Y.: Stablevideo: Text-driven consistency-aware diffusion video editing. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2023)
7. Chan, E.R., Lin, C.Z., Chan, M.A., Nagano, K., Pan, B., De Mello, S., Gallo, O., Guibas, L.J., Tremblay, J., Khamis, S., et al.: Efficient geometry-aware 3D generative adversarial networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 16123–16133 (2022)
8. Chan, K.C., Wang, X., Yu, K., Dong, C., Loy, C.C.: BasicVSR: The search for essential components in video super-resolution and beyond. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4947–4956 (2021)
9. Chan, K.C., Zhou, S., Xu, X., Loy, C.C.: BasicVSR++: Improving video super-resolution with enhanced propagation and alignment. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2022)
10. Chen, A., Xu, Z., Geiger, A., Yu, J., Su, H.: TensorRF: Tensorial radiance fields. In: Proceedings of the European Conference on Computer Vision (ECCV) (2022)
11. Chen, H., He, X., Qing, L., Wu, Y., Ren, C., Sheriff, R.E., Zhu, C.: Real-world single image super-resolution: A brief review. *Information Fusion* **79**(C), 124–145 (mar 2022)
12. Chen, X., Wang, X., Zhang, W., Kong, X., Qiao, Y., Zhou, J., Dong, C.: HAT: Hybrid attention transformer for image restoration. arXiv preprint arXiv:2309.05239 (2023)
13. Chen, X., Wang, X., Zhou, J., Qiao, Y., Dong, C.: Activating more pixels in image super-resolution transformer. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 22367–22377 (June 2023)
14. Dong, C., Loy, C.C., He, K., Tang, X.: Learning a deep convolutional network for image super-resolution. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 184–199 (2014)
15. Dong, C., Loy, C.C., Tang, X.: Accelerating the super-resolution convolutional neural network. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 391–407 (2016)
16. Dong, Z., Chen, X., Yang, J., J.Black, M., Hilliges, O., Geiger, A.: AG3D: Learning to generate 3D avatars from 2D image collections. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2023)

17. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: *Advances in Neural Information Processing Systems (NeurIPS)*. vol. 27 (2014)
18. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. *Advances in Neural Information Processing Systems (NeurIPS)* (2014)
19. Han, Y., Yu, T., Yu, X., Wang, Y., Dai, Q.: Super-NeRF: View-consistent detail generation for nerf super-resolution. *arXiv preprint arXiv:2304.13518* (2023)
20. Haque, A., Tancik, M., Efros, A., Holynski, A., Kanazawa, A.: Instruct-nerf2nerf: Editing 3d scenes with instructions. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* (2023)
21. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: GANs trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in Neural Information Processing Systems (NeurIPS)* (2017)
22. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. In: *Advances in Neural Information Processing Systems (NeurIPS)*. vol. 33, pp. 6840–6851 (2020)
23. Huang, X., Li, W., Hu, J., Chen, H., Wang, Y.: RefSR-NeRF: Towards high fidelity and super resolution view synthesis. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 8244–8253 (June 2023)
24. Kang, M., Zhu, J.Y., Zhang, R., Park, J., Shechtman, E., Paris, S., Park, T.: Scaling up GANs for text-to-image synthesis. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 10124–10134 (June 2023)
25. Karnewar, A., Vedaldi, A., Novotny, D., Mitra, N.J.: Holodiffusion: Training a 3D diffusion model using 2D images. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2023)
26. Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2019)
27. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3D gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics (ToG)* **42**(4), 1–14 (2023)
28. Kim, J., Lee, J.K., Lee, K.M.: Accurate image super-resolution using very deep convolutional networks. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 1646–1654 (2016)
29. Ledig, C., Theis, L., Huszár, F., Caballero, J., Cunningham, A., Acosta, A., Aitken, A., Tejani, A., Totz, J., Wang, Z., et al.: Photo-realistic single image super-resolution using a generative adversarial network. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 4681–4690 (2023)
30. Li, H., Yang, Y., Chang, M., Chen, S., Feng, H., Xu, Z., Li, Q., Chen, Y.: SRDiff: Single image super-resolution with diffusion probabilistic models. *Neurocomputing* **479**, 47–59 (2022)
31. Li, J., Tan, H., Zhang, K., Xu, Z., Luan, F., Xu, Y., Hong, Y., Sunkavalli, K., Shakhnarovich, G., Bi, S.: Instant3D: Fast text-to-3d with sparse-view generation and large reconstruction model. *arXiv preprint arXiv:2311.06214* (2023)
32. Lin, C.Y., Fu, Q., Merth, T., Yang, K., Ranjan, A.: FastSR-NeRF: Improving nerf efficiency on consumer devices with a simple super-resolution pipeline. In: *WACV* (2024)

33. Liu, D., Wang, Z., Fan, Y., Liu, X., Wang, Z., Chang, S., Huang, T.: Robust video super-resolution with learned temporal dynamics. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 2507–2515 (2017)
34. Liu, M., Xu, C., Jin, H., Chen, L., Xu, Z., Su, H., et al.: One-2-3-45: Any single image to 3d mesh in 45 seconds without per-shape optimization. arXiv preprint arXiv:2306.16928 (2023)
35. Lugmayr, A., Danelljan, M., Van Gool, L., Timofte, R.: SRFlow: Learning the super-resolution space with normalizing flow. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 715–732 (2020)
36. Menon, S., Damian, A., Hu, S., Ravi, N., Rudin, C.: PULSE: Self-supervised photo upsampling via latent space exploration of generative models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 2437–2445 (2020)
37. Mescheder, L., Geiger, A., Nowozin, S.: Which training methods for GANs do actually converge? In: Proceedings of the International Conference on Machine Learning (ICML). PMLR (2018)
38. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: NeRF: Representing scenes as neural radiance fields for view synthesis. Communications of the ACM **65**(1), 99–106 (2021)
39. Mittal, A., Soundararajan, R., Bovik, A.C.: Making a “completely blind” image quality analyzer. IEEE Signal Processing Letters (2013)
40. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2022)
41. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 10684–10695 (June 2022)
42. Saharia, C., Ho, J., Chan, W., Salimans, T., Fleet, D.J., Norouzi, M.: Image super-resolution via iterative refinement. IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI) **45**(4), 4713–4726 (2023)
43. Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., Chen, X.: Improved techniques for training GANs. Advances in Neural Information Processing Systems (NeurIPS) (2016)
44. Schuhmann, C., Beaumont, R., Vencu, R., Gordon, C.W., Wightman, R., Cherti, M., Coombes, T., Katta, A., Mullis, C., Wortsman, M., Schramowski, P., Kundurthy, S.R., Crowson, K., Schmidt, L., Kaczmarczyk, R., Jitsev, J.: LAION-5b: An open large-scale dataset for training next generation image-text models. In: Neural Information Processing Systems Datasets and Benchmarks Track (2022)
45. Shen, Y., Chandaka, B., Lin, Z.H., Zhai, A., Cui, H., Forsyth, D., Wang, S.: Sim-on-wheels: Physical world in the loop simulation for self-driving. IEEE Robotics and Automation Letters (2023)
46. Shen, Y., Ma, W.C., Wang, S.: SGAM: Building a virtual 3d world through simultaneous generation and mapping. In: Advances in Neural Information Processing Systems (NeurIPS) (2022)
47. Smith, E., Fujimoto, S., Meger, D.: Multi-view silhouette and depth decomposition for high resolution 3d object representation. In: Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., Garnett, R. (eds.) Advances in Neural Information Processing Systems (NeurIPS), pp. 6479–6489. Curran Associates, Inc. (2018)

48. Tian, Y., Zhang, Y., Fu, Y., Xu, C.: TDAN: Temporally-deformable alignment network for video super-resolution. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2020)
49. Wang, C., Wu, X., Guo, Y.C., Zhang, S.H., Tai, Y.W., Hu, S.M.: NeRF-SR: High quality neural radiance fields using supersampling. In: Proceedings of ACM Multimedia. pp. 6445–6454 (2022)
50. Wang, P., Liu, L., Liu, Y., Theobalt, C., Komura, T., Wang, W.: NeuS: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. Advances in Neural Information Processing Systems (NeurIPS) (2021)
51. Wang, T., Zhang, B., Zhang, T., Gu, S., Bao, J., Baltrušaitis, T., Shen, J., Chen, D., Wen, F., Chen, Q., Guo, B.: RODIN: A generative model for sculpting 3d digital avatars using diffusion. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4563–4573 (2022)
52. Wang, X., Chan, K.C., Yu, K., Dong, C., Loy, C.C.: EDVR: Video restoration with enhanced deformable convolutional networks. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops (June 2019)
53. Wang, X., Xie, L., Dong, C., Shan, Y.: Real-ESRGAN: Training real-world blind super-resolution with pure synthetic data. In: International Conference on Computer Vision Workshops (ICCVW) (2021)
54. Wang, X., Yu, K., Dong, C., Loy, C.C.: Recovering realistic texture in image super-resolution by deep spatial feature transform. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 606–615 (2018)
55. Wang, X., Yu, K., Wu, S., Gu, J., Liu, Y., Dong, C., Qiao, Y., Change Loy, C.: ESRGAN: Enhanced super-resolution generative adversarial networks. In: The European Conference on Computer Vision Workshops (ECCVW) (2018)
56. Xia, H., Fu, Y., Liu, S., Wang, X.: Rgb-d objects in the wild: Scaling real-world 3d object learning from rgb-d videos. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2024)
57. Xu, Y., Park, T., Zhang, R., Zhou, Y., Shechtman, E., Liu, F., Huang, J.B., Liu, D.: Videogigagan: Towards detail-rich video super-resolution (2024)
58. Yang, C., Li, S., Fang, J., Liang, R., Xie, L., Zhang, X., Shen, W., Tian, Q.: GaussianObject: Just taking four images to get a high-quality 3D object with gaussian splatting. arXiv preprint arXiv:2402.10259 (2024)
59. Yariv, L., Gu, J., Kasten, Y., Lipman, Y.: Volume rendering of neural implicit surfaces. Advances in Neural Information Processing Systems (NeurIPS) (2021)
60. Yi, P., Wang, Z., Jiang, K., Shao, Z., Ma, J.: Multi-temporal ultra dense memory network for video super-resolution. IEEE Transactions on Circuits and Systems for Video Technology (2019)
61. Yoon, Y., Yoon, K.J.: Cross-guided optimization of radiance fields with multi-view image super-resolution for high-resolution novel view synthesis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 12428–12438 (June 2023)
62. Yu, X., Xu, M., Zhang, Y., Liu, H., Ye, C., Wu, Y., Yan, Z., Liang, T., Chen, G., Cui, S., Han, X.: MVImgNet: A large-scale dataset of multi-view images. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2023)
63. Zhang, K., Gool, L.V., Timofte, R.: Deep unfolding network for image super-resolution. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 3217–3226 (June 2020)

64. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2018)
65. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2018)
66. Zhang, Y., Li, K., Li, K., Wang, L., Zhong, B., Fu, Y.: Image super-resolution using very deep residual channel attention networks. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 286–301 (2018)
67. Zhao, X., Ding, W., An, Y., Du, Y., Yu, T., Li, M., Tang, M., Wang, J.: Fast segment anything. arXiv preprint arXiv:2306.12156 (2023)
68. Zheng, X.Y., Pan, H., Wang, P.S., Tong, X., Liu, Y., Shum, H.Y.: Locally attentional SDF diffusion for controllable 3D shape generation. *ACM Transactions on Graphics (ToG)* **42**(4) (2023)
69. Zhou, S., Yang, P., Wang, J., Luo, Y., Loy, C.C.: Upscale-A-Video: Temporal-consistent diffusion model for real-world video super-resolution. arXiv preprint arXiv:2312.06640 (2023)