

C3DAG: Controlled 3D Animal Generation using 3D pose guidance

Sandeep Mishra*
University of Texas at Austin
Austin
sandy.mishra@utexas.edu

Oindrila Saha*
Univeristy of Massachusetts
Amherst
osaha@umass.edu

Alan C. Bovik
University of Texas at Austin
Austin
bovik@ece.utexas.edu

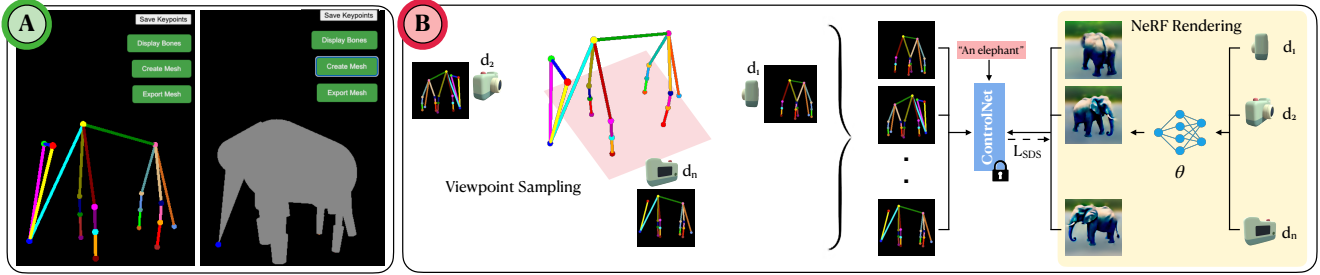


Figure 1. **C3DAG : A pipeline to generate highly detailed and anatomically accurate 3D Animals.** (A) A web-based tool that lets the user modify keypoints in 3D space and generates balloon animals using simple geometric shapes such as spheres, cylinders, and cones with a single button click. The figure depicts the 3D pose and skeleton of an elephant and its generated balloon animal mesh. (B) We illustrate our method to create a pose consistent animal by sampling 2D pose control images from a 3D pose based on camera parameters. The sampled 2D poses are used as input to our tetrapod-pose ControlNet along with a text prompt and a NeRF rendered image to backpropagate Score Distillation Sampling (SDS) gradients to the learnable NeRF parameters θ .

Abstract

Recent advancements in text-to-3D generation have demonstrated the ability to generate high quality 3D assets. However while generating animals these methods underperform, often portraying inaccurate anatomy and geometry. Towards ameliorating this defect, we present C3DAG, a novel pose-Controlled text-to-3D Animal Generation framework which generates a high quality 3D animal consistent with a given pose. We also introduce an automatic 3D shape creator tool, that allows dynamic pose generation and modification via a web-based tool, and that generates a 3D balloon animal using simple geometries. A NeRF is then initialized using this 3D shape using depth-controlled SDS. In the next stage, the pre-trained NeRF is fine-tuned using quadruped-pose-controlled SDS. The pipeline that we have developed not only produces geometrically and anatomically consistent results, but also renders highly controlled 3D animals, unlike prior methods which do not allow fine-grained pose control.

1. Introduction

The past decade has observed dramatic improvements in the field of computer vision, especially on 3D generative

tasks. The introduction of denoising diffusion models and implicit representations for 3D reconstruction has dramatically improved the quality of generative art. DreamFusion [23] showed that knowledge of 2D image generative models can be utilized to guide generation of 3D objects. These concepts alleviate the need for substantial datasets of 3D ground truth, by lifting 2D images and videos into 3D shapes by distillation. With this in mind, we focus on the controlled generation of 3D animals guided by text and pose, with the specific goal of generating anatomically plausible 3D animals.

3D animal generation has been previously studied, yet current methods still suffer by the creation of problematic anatomical distortions such as multiple heads or limbs, or displaced or missing body parts. Attempts at tackling these problems, such as PerpNeg [2], have improved the quality of generation but only in selective cases. Another inherent issue with distillation based methods is the amount of time required to conduct test-time optimization. Although it is not possible to avoid test-time optimization when generating highly detailed 3D assets, efforts to reduce the time required for optimization is an important consideration.

Another approach to generate 3D animals involves 3D Morphable Models (3DMMs) which are parametric models. While this approach can generate high quality 3D shapes, due to the time-consuming data collection process of 3D

*These authors contributed equally to this work

scanning, the diversity of generated animals is quite limited. Some methods utilize 2D images and videos to create banks of such models to expand the range of animals they can generate, but out-of-domain generation remains an unsolved problem with this approach. Methods such as MagicPony [31], which use images and videos to learn parametric models, are unable to capture fine details and end up completely omitting tails and other shape details resulting in low quality 3D meshes.

DreamWaltz [9] demonstrates the use of SMPL [15] based initialization and projected 2D poses to generate highly controlled 3D human-like avatars, while usually reducing inaccuracies in the generated asset. Animals take on a tremendous number of different shapes and body proportions, so starting from an initialization using 3DMM based models like those generated by SMPL is not viable, as they have limited diversity. Towards remediating these limitations, we make the following contributions:

- An automatic 3D shape creator tool that generates a naive 3D shape and pose which can be used to initialize a NeRF.
- A tetrapod-pose guided ControlNet trained on a diverse dataset including mammals, reptiles, birds and amphibians to provide highly controlled guidance via Score Distillation Sampling (SDS).
- We combine (a) and (b) in an efficient two-stage pipeline that is able to generate high quality and anatomically consistent 3D animal assets. Our method, which we call C3DAG, is efficient, and produces a high quality animal much faster than current SOTA methods.

2. Related Work

Generating 3D assets without specific conditions requires understanding the diverse distributions of 3D data. There are two primary strategies: explicit and implicit. Structured representations such as point clouds [1, 17], voxel grids [13, 28], and mesh models [34] are categorized as explicit methods. Implicit techniques generally rely on abstract representations, such as signed distance functions (SDFs) [6, 7, 20], tri-planes [4], the parameters of multi-layer perceptrons (MLPs) [8], and radiance fields [16].

The recent surge in text-to-3D asset generation research has been significantly fueled by the availability of vast datasets of text-image pairs and the success of text-to-image generative models. These innovations have paved the way for the adaptation of pre-trained text-to-image models into the 3D domain, primarily those that are CLIP-guided, or are 2D diffusion-guided. CLIP-guided methods such as [10, 21] leverage cross-modal matching models like CLIP [24] for text-to-3D conversion, whereas diffusion-guided strategies utilize text-to-image diffusion models [5, 35], like Imagen [26] and Stable Diffusion [25], to generate 3D assets from textual descriptions. The latter has been noted to deliver superior text-to-3D generation performance, em-

ploying techniques such as Score Distillation Sampling (SDS) [23], which refines noise in images captured from NeRFs [19], and Score-Jacobian-Chaining [29], which aggregates image gradients into 3D asset gradients.

The field of 3D animal generation has made significant progress thanks to groundbreaking studies that have provided new methods and insights for modeling the complex structures and movements of animals in 3D. SMAL [36] introduced a method to fit a parametric 3D shape model, derived from 3D scans, to animal images using 2D keypoints and segmentation masks, with extensions to multi-view images [37]. Subsequent efforts, such as LASSIE [32], have focused on deriving 3D shapes directly from smaller image collections by identifying self-supervised semantic correspondences to discover 3D parts. Despite these advances, the limited diversity of currently available 3D assets confines most research [11, 30, 31] to the production of 3D models that are both class-specific and small scale.

3. Method

We propose C3DAG: Controlled 3D Animal Generation, which is an efficient two-stage approach to generate high-quality 3D assets visually representative of posed tetrapods, including birds, animals and reptiles. The first stage initializes a NeRF to a generic balloon animal created using our automatic 3D shape creator tool. The second stage refines the initialization by using occlusion aware quadruped-3D-pose consistent Score Distillation Sampling (SDS) to generate high quality 3D animal assets. Unlike DreamWaltz [9], we avoid using 3DMM models, since they are always limited to a fixed set of animals, thus enabling C3DAG to generate a more diverse set of fauna including quadruped land animals, reptiles and birds. We also introduce control scale and guidance scale annealing, ensuring that the 3D optimization is able to fully utilize the diversity and quality of Stable Diffusion. See Supplementary Sec. 6 for more details about the preliminary concepts used in this article. Hereafter we describe each aspect in detail.

3.1. 2D Pose guided ControlNet

Following the human-pose ControlNet introduced in the original paper, we train a ControlNet that can produce various animal images including quadrupeds, reptiles and birds using the same set of 18 keypoints. We introduce data augmentations such as random rotations, translations, and scaling, so that the model is robust against various factors such as occluded poses and different scales of animals. This is crucial in the 3D generation stage as the randomly sampled camera views of a 3D skeleton will often include heavy occlusions. We provide more details in the section on Implementation Details.




Figure 2. **Comparison with text-to-3D generation guided by 2D diffusion.** Given the prompt as shown, we qualitatively compare with the open-source version of Dreamfusion and HiFA using their default settings. For both of these we append “, full body” to the prompt. It is clearly observed that both prior work suffer from various inconsistencies producing multiple heads or limbs. Stable Dreamfusion usually produces lower details in textures. HiFA produces high-quality textures but almost always produces anatomically incorrect animals. We provide more results in supplementary.

3.2. NeRF initialization

To conduct NeRF pre-training using depth ControlNet, DreamWaltz uses the neutral SMPL mesh. However initialization using a single generic model is ineffective for generating tetrapods, as they vary vastly in shape and size. Existing parametric models such as SMAL [36] also cannot be used, since they represent only a small set of animals that does not include any birds or reptiles. Another parametric model called 3DFauna [12], has wider diversity, but also fails to represent elephants, and appendages like tails in certain animals, as evidenced in the results in their original paper and website. To remediate these issues, we instead resort to an automatic method of generating an initial 3D mesh given a 3D skeleton, as described in Sec. 3.3. The effect of using this pre-training is shown in supplementary.

3.3. Generating 3D pose and shape

We employ a set of 18 3D keypoints and 18 bone connections to represent a generic body pose, as visualized in Fig. 1 A. To enable the modification of the 3D pose we created a simple THREE.js web tool that provides a clean UI to move the 3D keypoints around. The same tool also allows the user to initialize the 3D shape using spherical, cylindrical and conical meshes to form the head, limbs, tails, and nose, respectively. The parameters of the constructive components are tunable, enabling the creation of balloon animals of different shapes and sizes. The combined components are then re-meshed into a single mesh, which is then used in the pre-training stage using SDS, with depth ControlNet as the guidance.

3.4. 3D Aware Score Distillation Sampling

Given the input 3D keypoints we project 2D poses onto the camera coordinate system, then use the projections as control images when generating pose-based-guidance signals from our trained Tetrapod-Pose guided ControlNet. Since the 3D pose is fixed, this makes it possible to produce 3D consistent animal views, ensuring a high quality reconstruction. We illustrate this in Fig. 1 B. It is important to note that C3DAG uses no 3D ground-truth data unlike MV-Dream [27] or Zero-1-2-3 [14], instead purely lifting 2D poses to 3D, guided only by training on 2D image datasets. We show that without pose control, other 2D diffusion SDS guided methods fail to maintain geometric and anatomic consistency (Fig. 2). However, since ControlNet is trained using a limited number of images from datasets that contain annotated poses, exploiting the quality and diversity learned by stable diffusion is important. Animal pose datasets inherently consider occlusion, as only visible keypoints are annotated. In light of this, we deployed a more generic variation of occlusion culling than that in DreamWaltz, to control the visibility of each individual part of the head (left eye, right eye, and nose) depending on the view-description. We utilized randomly acquired camera parameters to determine the view description, including: front, left-side, back, right-side, top and bottom.

4. Implementation details

ControlNet Training: We used annotated poses from the AWA-pose [3] and Animal Kingdom [22] datasets to train ControlNet in a similar way as the original, which uses sta-

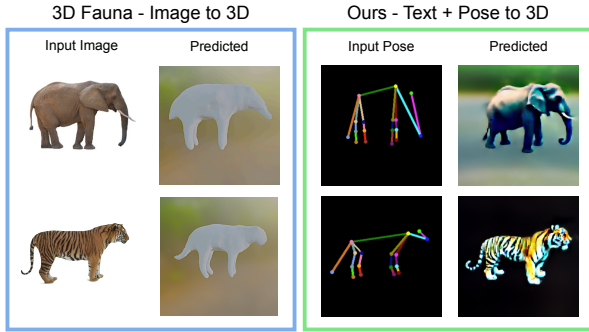


Figure 3. **Comparison with parametric model based method.** Given the input image 3DFauna fails to capture high-frequency details and follow the input image (see tail), whereas our method produces a highly detailed animal given input pose and text, which closely follows the input pose control.

ble diffusion version 1.5. AwA-pose consists of 10k annotated images covering 35 quadruped animal classes, while Animal Kingdom provides 33k annotated images spanning 850 species, including mammals, reptiles, birds, amphibians, fishes and insects. From a combined set of 43k samples, we carefully selected a subset including only mammals, reptiles, birds, and amphibians. We also eliminated any sample having less than 30% of its keypoints annotated. The curated dataset consists of 13k annotated samples. To increase diversity in learning, and to improve test-time generation at any scale and transformations, we used a combination of data augmentation strategies consisting of random rotations, translations, and scaling while training. The model was trained over 229k iterations with a batch size of 12, a constant learning rate of $1e^{-5}$, on a single Nvidia RTX 6000. The model converged after around 120k iterations and would not overfit even up to 200k iterations, owing in part to the augmentation strategy.

3D Pose and Shape generation: We used the following 18 keypoints to represent every quadruped: left eye, right eye, nose, neck end, $4 \times$ thighs, $4 \times$ knees, $4 \times$ paws, back end, and tail end. For the upper limbs of birds, i.e. wings, their front - thighs, knees, and paws are defined in accordance with how their upper limbs move. We began with an initial pose of a tiger from SMAL, and modify its keypoints using the balloon animal creator tool. An example of this process and the tool UI is depicted in Fig. 1A. After appropriate modification of the pose the user can press the button to *create mesh* around bones. This button press invokes calls to various functions defined to create each body part, based on their natural appearances using simple mesh components such as ellipses, cylinders, and cones. The combined mesh and the corresponding keypoints can be downloaded by another button click.

Mesh depth guided NeRF initialization: The mesh downloaded in the previous step was used to provide depth

maps to the pre-trained depth guided ControlNet, which produces the gradient loss by SDS, which in turn is used to pre-train the NeRF. The pre-training helps achieve a reasonable initial state for the NeRF weights, which can then be refined in the final pose-guided training stage. The diffusion model was pre-trained for 10,000 iterations using the Adam optimizer with a learning rate of $1e^{-3}$ and a batch size of 1. During training, the camera positions were randomly sampled in spherical coordinates, where the radius, azimuthal angle, and polar angle of camera position were sampled from $[1.0, 2.0]$, $[0, 360]$, and $[60, 120]$.

Pose-guided SDS for NeRF fine-tuning: Finally, we fine-tune the NeRF using the pre-trained ControlNet to provide 2D pose guidance to SDS. The gradients computed using the noise residual from SDS was weighted in a similar manner as DreamFusion, where $w(t) = \sigma_t^2$ and t was annealed using $t = t_{max} - (t_{max} - t_{min})\sqrt{\frac{iter}{total_iters}}$. We set t_{max} to be 0.98, t_{min} to be 0.4. Similar to the previous stage, we trained the model over $total_iters = 10,000$ using the same settings for the optimizer. Using cosine annealing, we reduced the control scale from an initial value of 1 to a final value of 0.25, while updating guidance scale as $\omega = \omega_{init} \cdot (1 + \frac{iter}{total_iters})$, where $\omega_{init} = 50$. These settings helped reduce the impact of ControlNet gradually over the training process, while improving quality by gradually increasing ω of classifier-free guidance. The camera positions were randomly sampled as in stage 1, as were the radius, azimuthal angle, and polar angle of the camera. The 3D avatar representation renders “latent images” in the latent space of $\mathbb{R}^{64 \times 64 \times 4}$ following LatentNeRF [18], where the “latent images” can be decoded into RGB images of $\mathbb{R}^{512 \times 512 \times 4}$ using the VAE decoder of Stable Diffusion [25].

5. Discussion and Conclusion

We showed our method performs consistently well (Fig 2) while Stable DreamFusion and HiFA usually produce physically implausible or low-quality assets. It is important to note that HiFA, which is a current SOTA model, requires about 7 hours on a 80GB A100 GPU, while ours needs only about 20 minutes on the same setting. Fig 3 shows that our generated animal contains much more detail than the one produced by 3DFauna [12]. Our method currently requires some human effort to re-position 3D keypoints in our balloon animal creator tool. It is a direction of future work to fully automate this process. To summarize, we present C3DAG – an efficient method to generate anatomically and geometrically consistent 3D animals. We contribute 1) a 2D Tetrapod-pose ControlNet trained on mammals, amphibians, reptiles and birds, 2) a tool to create 3D balloon animals automatically given 3D keypoints, and 3) a method to pipeline 1) and 2) to create high quality 3D animals.

References

- [1] Panos Achlioptas, Olga Diamanti, Ioannis Mitliagkas, and Leonidas Guibas. Learning representations and generative models for 3d point clouds. In *International Conference on Machine Learning*, pages 40–49. PMLR, 2018. 2
- [2] Mohammadreza Armandpour, Ali Sadeghian, Huangjie Zheng, Amir Sadeghian, and Mingyuan Zhou. Re-imagine the negative prompt algorithm: Transform 2d diffusion into 3d, alleviate janus problem and beyond. *arXiv preprint arXiv:2304.04968*, 2023. 1
- [3] Prianka Banik, Lin Li, and Xishuang Dong. A novel dataset for keypoint detection of quadruped animals from images. *arXiv preprint arXiv:2108.13958*, 2021. 3
- [4] Hansheng Chen, Jiatao Gu, Anpei Chen, Wei Tian, Zhuowen Tu, Lingjie Liu, and Hao Su. Single-stage diffusion nerf: A unified approach to 3d generation and reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2416–2425, 2023. 2
- [5] Rui Chen, Yongwei Chen, Ningxin Jiao, and Kui Jia. Fantasia3d: Disentangling geometry and appearance for high-quality text-to-3d content creation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 22246–22256, 2023. 2
- [6] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5939–5948, 2019. 2
- [7] Yen-Chi Cheng, Hsin-Ying Lee, Sergey Tulyakov, Alexander G Schwing, and Liang-Yan Gui. Sdfusion: Multimodal 3d shape completion, reconstruction, and generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4456–4465, 2023. 2
- [8] Ziya Erkoç, Fangchang Ma, Qi Shan, Matthias Nießner, and Angela Dai. Hyperdiffusion: Generating implicit neural fields with weight-space diffusion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14300–14310, 2023. 2
- [9] Yukun Huang, Jianan Wang, Ailing Zeng, He Cao, Xianbiao Qi, Yukai Shi, Zheng-Jun Zha, and Lei Zhang. Dreamwaltz: Make a scene with complex 3d animatable avatars. *Advances in Neural Information Processing Systems*, 36, 2024. 2
- [10] Ajay Jain, Ben Mildenhall, Jonathan T Barron, Pieter Abbeel, and Ben Poole. Zero-shot text-guided object generation with dream fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 867–876, 2022. 2
- [11] Tomas Jakab, Ruining Li, Shangzhe Wu, Christian Rupprecht, and Andrea Vedaldi. Farm3d: Learning articulated 3d animals by distilling 2d diffusion. *arXiv preprint arXiv:2304.10535*, 2023. 2
- [12] Zizhang Li, Dor Litvak, Ruining Li, Yunzhi Zhang, Tomas Jakab, Christian Rupprecht, Shangzhe Wu, Andrea Vedaldi, and Jiajun Wu. Learning the 3d fauna of the web. *arXiv preprint arXiv:2401.02400*, 2024. 3, 4
- [13] Chieh Hubert Lin, Hsin-Ying Lee, Willi Menapace, Menglei Chai, Aliaksandr Siarohin, Ming-Hsuan Yang, and Sergey Tulyakov. Infinitcity: Infinite-scale city synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 22808–22818, 2023. 2
- [14] Ruoshi Liu, Rundi Wu, Basile Van Hoorick, Pavel Tokmakov, Sergey Zakharov, and Carl Vondrick. Zero-1-to-3: Zero-shot one image to 3d object. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9298–9309, 2023. 3
- [15] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. Smpl: A skinned multi-person linear model. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*, pages 851–866, 2023. 2
- [16] Jonathan Lorraine, Kevin Xie, Xiaohui Zeng, Chen-Hsuan Lin, Towaki Takikawa, Nicholas Sharp, Tsung-Yi Lin, Ming-Yu Liu, Sanja Fidler, and James Lucas. Att3d: Amortized text-to-3d object synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 17946–17956, 2023. 2
- [17] Shitong Luo and Wei Hu. Diffusion probabilistic models for 3d point cloud generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2837–2845, 2021. 2
- [18] Gal Metzer, Elad Richardson, Or Patashnik, Raja Giryes, and Daniel Cohen-Or. Latent-nerf for shape-guided generation of 3d shapes and textures. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12663–12673, 2023. 4
- [19] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021. 2
- [20] Paritosh Mittal, Yen-Chi Cheng, Maneesh Singh, and Shubham Tulsiani. Autosdf: Shape priors for 3d completion, reconstruction and generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 306–315, 2022. 2
- [21] Nasir Mohammad Khalid, Tianhao Xie, Eugene Belilovsky, and Tiberiu Popa. Clip-mesh: Generating textured meshes from text using pretrained image-text models. In *SIGGRAPH Asia 2022 Conference papers*, pages 1–8, 2022. 2
- [22] Xun Long Ng, Kian Eng Ong, Qichen Zheng, Yun Ni, Si Yong Yeo, and Jun Liu. Animal kingdom: A large and diverse dataset for animal behavior understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19023–19034, 2022. 3
- [23] Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. *arXiv preprint arXiv:2209.14988*, 2022. 1, 2
- [24] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021. 2
- [25] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of*

- the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022. 2, 4
- [26] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in Neural Information Processing Systems*, 35:36479–36494, 2022. 2
- [27] Yichun Shi, Peng Wang, Jiandong Ye, Long Mai, Kejie Li, and Xiao Yang. Mvdream: Multi-view diffusion for 3d generation. In *The Twelfth International Conference on Learning Representations*, 2023. 3
- [28] Edward J Smith and David Meger. Improved adversarial systems for 3d object generation and reconstruction. In *Conference on Robot Learning*, pages 87–96. PMLR, 2017. 2
- [29] Haochen Wang, Xiaodan Du, Jiahao Li, Raymond A Yeh, and Greg Shakhnarovich. Score jacobian chaining: Lifting pretrained 2d diffusion models for 3d generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12619–12629, 2023. 2
- [30] Shangzhe Wu, Tomas Jakab, Christian Rupprecht, and Andrea Vedaldi. Dove: Learning deformable 3d objects by watching videos. *International Journal of Computer Vision*, 131(10):2623–2634, 2023. 2
- [31] Shangzhe Wu, Ruining Li, Tomas Jakab, Christian Rupprecht, and Andrea Vedaldi. Magicpony: Learning articulated 3d animals in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8792–8802, 2023. 2
- [32] Chun-Han Yao, Wei-Chih Hung, Yuanzhen Li, Michael Rubinstein, Ming-Hsuan Yang, and Varun Jampani. Lassie: Learning articulated shapes from sparse image ensemble via 3d part discovery. *Advances in Neural Information Processing Systems*, 35:15296–15308, 2022. 2
- [33] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3836–3847, 2023. 1
- [34] Song-Hai Zhang, Yuan-Chen Guo, and Qing-Wen Gu. Sketch2model: View-aware 3d modeling from single free-hand sketches. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6012–6021, 2021. 2
- [35] Junzhe Zhu, Peiye Zhuang, and Sanmi Koyejo. Hifa: High-fidelity text-to-3d generation with advanced diffusion guidance. In *The Twelfth International Conference on Learning Representations*, 2023. 2
- [36] Silvia Zuffi, Angjoo Kanazawa, David Jacobs, and Michael J. Black. 3D menagerie: Modeling the 3D shape and pose of animals. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2017. 2, 3
- [37] Silvia Zuffi, Angjoo Kanazawa, and Michael J Black. Lions and tigers and bears: Capturing non-rigid, 3d, articulated shape from images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3955–3963, 2018. 2

C3DAG: Controlled 3D Animal Generation using 3D pose guidance

Supplementary Material

6. Preliminaries

Recent works have shown impressive results in 2D and 3D content generation using a combination of the following:

- **Diffusion Model:**

Generative pre-training on extensive image-text datasets using denoising diffusion models serves as effective priors for text-to-3D generation. These models undergo two primary processes:

Forward Process: This involves incrementally adding noise to the data $x \sim p(x)$, transforming it towards a Gaussian distribution over T steps. The noise-augmented data at step t is described by:

$$z_t = \sqrt{\bar{\alpha}_t}x + \sqrt{1 - \bar{\alpha}_t}\epsilon, \quad \epsilon \sim \mathcal{N}(0, I),$$

where $\bar{\alpha}_t = \prod_{s=1}^t \alpha_s$ and $\alpha_t \in (0, 1)$ as part of the pre-defined noising schedule.

Backward Process: In this phase, the model learns to reverse the noise addition, aiming to reconstruct the original data from its noised version. The model learns to estimate the noise by minimizing

$$L_t = \mathbb{E}_{x, \epsilon \sim \mathcal{N}(0, I)} [\|\epsilon - \epsilon_\phi(z_t, t)\|^2].$$

This dual process enables the diffusion model to accurately estimate and reconstruct x from noisy observations, effectively leveraging clean data from noised counterparts.

- **Neural Radiance Fields (NeRF):**

Widely embraced for text-to-3D generation tasks, NeRFs utilize a trainable Multilayer Perceptron (MLP) parameterized by θ as their core 3D representation. When rendering, a set of rays $r(k) = o + kd$ are sampled, where o represents the camera position and d the direction, calculated on a per-pixel basis. The MLP processes each ray $r(k)$, outputting the density τ and color c for each sampled point. The final color C of a pixel is then determined by approximating the volume rendering integral using numerical quadrature:

$$\hat{C}(r) = \sum_{i=1}^{N_c} \Omega_i \cdot (1 - \exp(-\tau_i \delta_i)) c_i,$$

where N_c is the number of points sampled along a ray, $\Omega_i = \exp\left(-\sum_{j=1}^{i-1} \tau_j \delta_j\right)$ is the accumulated transmittance, and δ_i is the distance between adjacent sample points.

- **Score Distillation Sampling (SDS):**

Introduced by DreamFusion and further utilized in various studies, SDS is a method designed to transfer the




Figure 4. **Effect of NeRF pre-training using shape generated by the automatic 3D shape creator tool.** The top row shows the final generated 3D without any initial pre-training stage. This has visible shortcomings in form of body occlusion as well as extra volume in the foreground which does not belong to the body of the animal. In contrast, fine-tuning after a pre-training stage generates a clean result.

knowledge from a pre-trained diffusion model ϵ_ϕ into a differentiable 3D representation, enhancing a NeRF model’s parameters θ . The rendered output x from a NeRF model can be obtained via $x = g(\theta)$, where g signifies a differentiable rendering function. The core of SDS lies in computing the gradients of the NeRF parameters θ , formulated as:

$$\nabla_\theta L_{\text{SDS}}(\phi, x) = \mathbb{E}_{t, \epsilon} \left[w(t) (\epsilon_\phi(x_t; y, t) - \epsilon) \frac{\partial z_t}{\partial x} \frac{\partial x}{\partial \theta} \right],$$

where $w(t)$ is a timestep-dependent weighting function and y is the input text prompt.

- **ControlNet guided SDS:**

A novel approach that extends the principles of Score Distillation Sampling (SDS) to include conditioning on additional image-based information, facilitating enhanced control when generating 3D models. Like SDS, ControlNet [33] leverages a pre-trained diffusion model ϵ_ϕ but also introduces a conditioning image c , which could be a skeleton/pose map, depth map, normal map, or a combination thereof, to guide the generation process more precisely. For a model parameterized by θ , and its output x , ControlNet refines the gradient calculation of the parameters θ by incorporating c , as shown in:

$$\nabla_\theta L_{\text{SDS}}(\phi, x) = \mathbb{E}_{t, \epsilon} \left[w(t) (\epsilon_\phi(x_t; y, t, c) - \epsilon) \frac{\partial z_t}{\partial x} \frac{\partial x}{\partial \theta} \right],$$

where $w(t)$ is a weighting function dependent on the timestep t , y denotes the given text prompt, and c is

the conditioning image that significantly influences the generation process. This approach benefits by ControlNet's innovative method of integrating additional contextual cues into the generative framework, thereby enriching the model's capacity to compute detailed and accurate 3D reconstructions using complex conditioning information.

7. Impact of NeRF Pre-training

Fig. 4 shows the impact of the NeRF pre-training stage using depth ControlNet as described in Sec 3.2 using the 3D balloon animal generated by the automatic 3D shape creator tool. This shows that using the pre-training stage results in visibly cleaner results without any body occlusions.