Learning Continuous 3D Words for Text-to-Image Generation

Ta-Ying Cheng[1]* Matheus Gadelha[2] Thibault Groueix[2] Matthew Fisher[2]

Radomír Měch[2] Andrew Markham[1] Niki Trigoni[1]
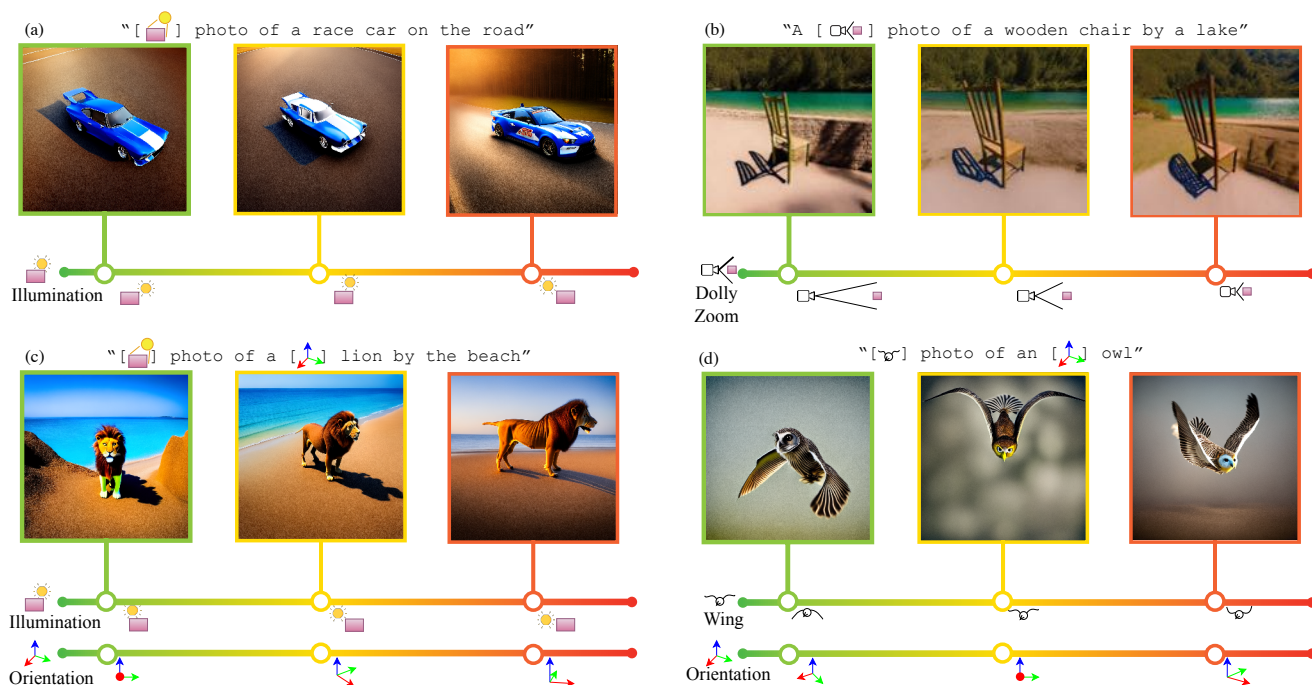
[1]University of Oxford [2]Adobe Research

Figure 1. We introduce **Continuous 3D Words** – special tokens in text-to-image models that allow users to have fine-grained control over several attributes like illumination [▦] (a and c), non-rigid shape change [⌁] (d), orientation [⤢] (c and d), and camera parameters [◁▣] (b). Our approach can be trained using a *single 3D mesh* and a rendering engine while incurring into negligible runtime and memory costs.

Abstract

*Current controls over diffusion models (e.g., through text or ControlNet) for image generation fall short in recognizing abstract, continuous attributes like illumination direction or non-rigid shape change. In this paper, we present an approach for allowing users of text-to-image models to have fine-grained control of several attributes in an image. We do this by engineering special sets of input tokens that can be transformed in a continuous manner – we call them **Continuous 3D Words**. These attributes can, for example, be represented as sliders and applied jointly with text prompts for fine-grained control over image generation. Given only a single mesh and a rendering engine, we show that our approach can be adopted to provide continuous user control over several 3D-aware attributes, including time-of-day illumination, bird wing orientation, dollyzoom effect, and object poses. Our method is capable of conditioning image creation with multiple Continuous 3D Words and text descriptions simultaneously while adding no overhead to the generative process. Project Page:* https://ttchengab.github.io/continuous_3d_words

---

*Work was done during internship at Adobe Research.

1. Introduction

Photography is fascinating because it enables very detailed control over the composition and aesthetics of the final image. On the one hand, this is simply the result of application of physical laws to achieve image acquisition. On the other, the slightest changes in the moment captured, illumination, object orientation, or camera parameters bring a completely different feeling to the viewer. While the giant leap of modern text-to-image diffusion can bring generated 2-D images to close proximity with real photos, text prompts are inherently limited to high-level descriptions, far removed from the detailed controls one has over actual photography. This is mainly due to the scarcity of such descriptions in the training dataset — very few would describe a photo based on exact object movements and camera parameters like the wing pose of a bird or the rotation of a person's head in degrees. On the other hand, 3D rendering engines allow us to mimic many of these 3D controls that photographers enjoy. We can render images of objects with predefined camera, illumination and pose changes at a very fine-grained scale. However, creating detailed 3D worlds is incredibly laborious, which limits the diversity of the scenes that can be generated by non-specialized practitioners. In that regard, using text-to-image diffusion to create images is a much more accessible technology, whereas precise 3D scene control remains firmly in the domain of experts.

In this work, we aim to bring together the best of two worlds by expanding the vocabulary of text-to-image diffusion models with very few samples generated from rendering engines. Specifically, we render meshes based on the attribute we aim to control, creating images with color and other useful information to generate a small set of data samples. The goal is to disentangle these abstract attributes from the original object and encode them into the textual space in a controllable manner – we term these attributes *Continuous 3D Words*. They allow users to create custom sliders that enable fine-grained control during image generation and can be seamlessly used along text prompts.

At the heart of our approach is an algorithm to learn a continuous vocabulary. The benefits of continuity are two-fold: *i)* the association between different values of the same attribute makes it much easier to learn, rather than having to learn hundreds of discrete tokens as an approximation and *ii)* we learn an MLP that would allow interpolation during inference to generate an actual continuous control. On top of this, we also propose two training strategies to prevent degenerate solutions and enable generalization on new objects beyond the training category. First, we apply a two-stage training strategy: we first apply the Dreambooth [25] approach to learn the object identity of the underlying mesh used for rendering, then sequentially learn the various attribute values disentangled from the object identity. This prevents the model from falling into a degenerate solution of encoding

each value of an attribute as a new object, which would prevent us from generalizing the attribute to new objects. Second, we apply ControlNet [30] with various conditioned images to generate a set of additional images with varying backgrounds and object textures. This prevents the model from overfitting to the artificial backgrounds of rendered images. The entire training was done in a light-weight Lower Rank Adaptation (LoRA) [14] manner, making it fast and accessible with single GPUs.

We implement our continuous vocabulary and training method, across various sets of single (e.g., dollyzoom extracted from chairs) and multiple (e.g., object pose and illumination extracted from a dog mesh) attributes, and show through quantitative user studies and qualitative comparisons that our method can properly reflect various attributes while maintaining the aesthetics of the image — significantly outperforming competitive baselines.

In summary, we present 1) Continuous 3D Words, a new method of gaining 3D-aware, continuous attribute controls over text-to-image generation that can be easily tailored to a plethora of new conditions, 2) a series of training strategies to disentangle the attributes from object identity to enhance the improvements in image generation and 3) extensive qualitative and quantitative studies to showcase our approach in various interesting applications.

2. Related Work

Conditional Diffusion Models. Ever since diffusion models [13, 27] pushed the quality and generalizability of image generation beyond GANs [10], the vision community has introduced a diverse range of modalities that can be used as conditions to control the image generation process. The most common condition is currently text. Works such as DALLE [2, 22, 23] and Imagen [26] used large scale text-image datasets and strong language understandings from pretrained LLMs [4, 8, 21] to guide the generation process. Stable Diffusion further popularized this class of methods by employing memory efficient models through latent-space diffusion [24]. Other works built on top of these models by adding other forms of conditioning [19, 30]. Highly relevant to our work is ControlNet [30], which proposes a general pipeline with zero-convolutions for conditioning on text and image data (*e.g.*, depth maps, canny maps, sketches). Despite their impressive image quality, it is not clear how to use these models to control other attributes of images like illumination or object orientation.

Other set of works explored how to perform image edits using textual instructions. Given a text-generated image, they demonstrate how the user can edit the image by amending the prompt, yet still preserve some aspects of the original image [3, 12, 20]. While convenient, these approaches do not allow *fine-grained* control over image elements since they are ultimately restricted by the user's ability to describe
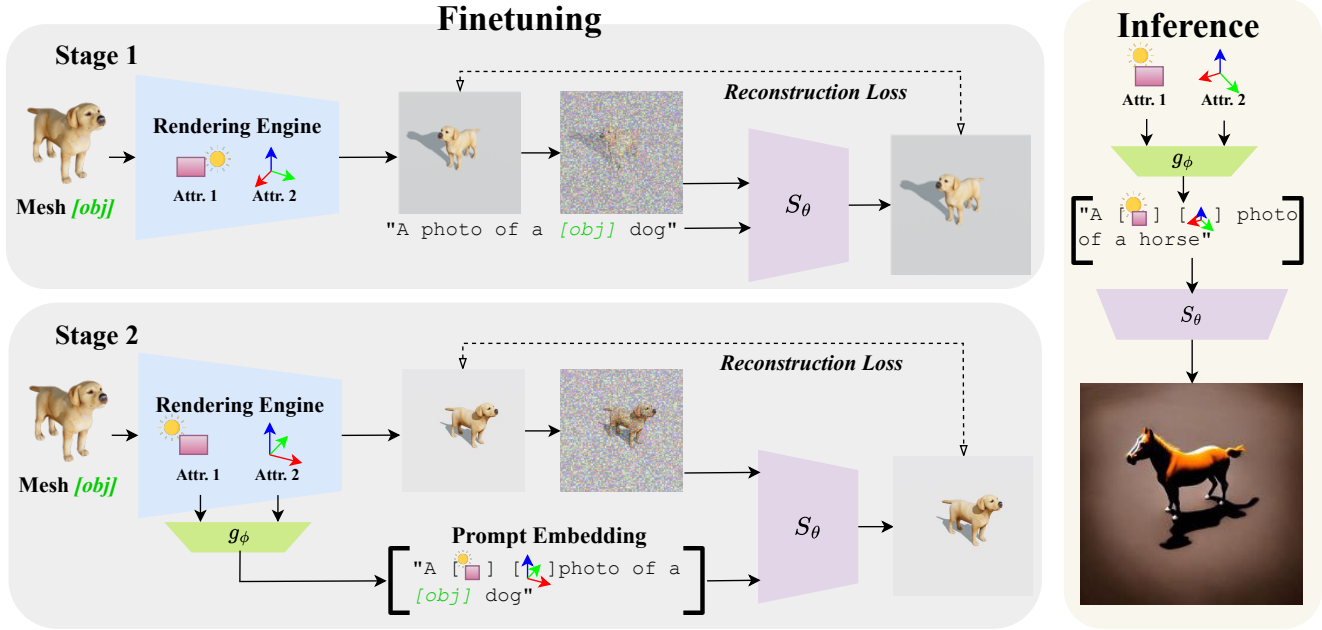
Figure 2. **Method Overview.** *Finetuning:* Our finetuning is divided into two stages. In the first stage, we render a series of images using different attribute values (e.g., illumination and pose). We feed them into the text-to-image diffusion model to learn token embedding *[Obj]* representing the *single mesh* used for training. In the second stage, we add the tokens representing individual attributes into the prompt embedding. The two stage training allows us to better disentangle the individual attributes against *[Obj]*. *Inference:* Attributes can be applied to different objects for text-to-image generation.
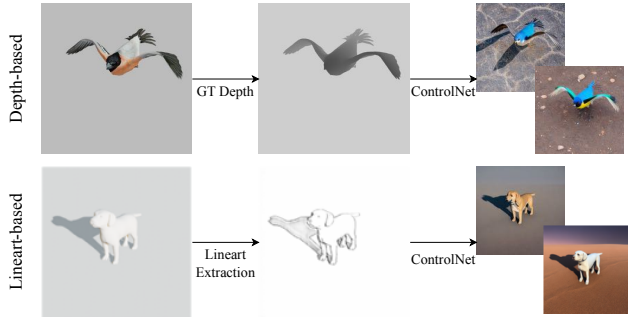

Figure 3. **ControlNet Augmentations.** Depth ControlNet is used for attributes creating direct shape changes. Lineart ControlNet is applied for more subtle changes that cannot be reflected by depths (e.g., illumination).

visual content through text – e.g., it would be very difficult to change the illumination direction by a precise angle such as 11°.

Recently, as the amount of 3D data available significantly increased, Liu et al. [16] introduced Zero-1-to-3, a diffusion model trained on various viewpoints of 3D rendering that enables viewpoint editing given an image of a single object. Similarly, works like DreamSparse [29] also employ diffusion models to synthesize novel views on open-set categories. Differently from our approach, these techniques are

focused only on object orientation and rely on vast 3D shape datasets. On the other hand, we investigate how to learn several continuous concepts (*e.g.* illumination [🔆] , wing pose [⌄] , dolly zoom [▣]) that can be directly used in text-to-image scenarios; *i.e.* we don't generate an image to then change the orientation or illumination later, but instead we use *Continuous Words* directly on text prompts.

Learning new concepts on diffusion models. With diffusion models being trained on unforeseen quantities in images and texts, a stream of work focused on adding specific concepts with very few data samples. For example, given a small set of images representing one particular object instance, textual inversion learns a new word embedding to describe the object, such that the word can be applied with new text prompts for image generation [9]. NETI [1] extended the word embedding to a time-space conditioned neural mapper for better generation while preserving quality. Similarly, Dreambooth [25] aims to achieve the same goal, but by using a repurposed token rarely used in text and finetuning the entire diffusion model with an additional constraint to prevent generative loss. There are numerous subsequent works showing improvements on finetuning different layers/weights and by improving the training strategy [11, 15, 17].

Despite the advances in adding new personalized entities

to existing models, few works focus on learning general concepts that can be applied to a variety of scenarios. A concurrent work, ViewNETI [5], is the first to learn viewpoints as a concept, but we hypothesise that the 3D awareness of large text-to-image diffusion models goes far beyond merely viewpoints, allowing us to associate and even create interactions with multiple 3D-aware concepts like illumination, pose and camera parameters *at the same time*. Our method, despite being trained only using a single mesh, shows superior generalization properties – while trained to learn illumination and orientation from renderings of a single dog, we are capable of employing the learned concepts to generate cars, horses (Figure 4), polar bears (Figure 8), lions (Figure 1) and so on.

3. Method

3.1. Preliminaries

We define an image I that captures object O from category C as a function of several attributes $I = f(a_1, a_2, a_3, ..., a_n)$, where a_i belongs to a vast set of image attributes $\mathcal{A}$: shape, material reflectivity, rotation/translation, camera parameters, shape deformations, etc. Some of these components can be translated to other categories while others cannot, so for simplicity we assume they only work for a single category. In the experimental section of this study, we will demonstrate that the definition of category for some attributes is rather loose and the user is capable of generating images with *continuous words* depicting objects very different from the ones seen during training. Notice that images annotated with the attributes we are interested are very rare, so the models do not have a very precise knowledge of them, except for what is already described in text-image pairs.

Given an image set I^O with images capturing an object O, previous methods like Dreambooth [25], Custom Diffusion [15], or Textual-Inversion [9] aim to minimize the following objective:

$$\mathbb{E}_{\hat{I}_{\epsilon,i}, T_O, I_i \in I^O}\left[\left\| S_\theta(\hat{I}_{\epsilon,i}, P(T_O)) - I_i \right\|_2^2 \right], \quad (1)$$

where S_θ is a Text-to-Image diffusion model [24], $\hat{I}_{\epsilon,i}$ is a noised image $\alpha_t I_i + \sigma_t \epsilon$ with noise ϵ and noise schedulers α_t, σ_t. $P(T_O)$ is the prompt condition which contains a token embedding T_O used as an identifier of object O. In practise $P(\cdot)$ is the text encoder from CLIP. The finetuned network S_θ can then generate new images containing O when given a new prompt condition $P'(T_O)$ and some Gaussian noise. Unlike previous methods, our goal is not to add concepts representing specific objects, but rather have them describe some attributes $a_i \in \mathcal{A}$, by learning to disentangle them using as few objects as possible within C – most of the time just one object suffices. Our model can also be easily extended to controlling multiple attributes *at the same time*.

3.2. Continuous Control

A naive way to control an attribute a is to use some realistic rendering engine to generate images of the available objects that have the same value $a = x$, and then apply similar approaches to previous works by assigning a token T_x to identify images with this particular value. This is not ideal, however, since a is often continuous and have infinitely many values – we would require an unfeasible number of tokens to gain fine-grained control over these attributes.

Therefore, we propose to instead learn a continuous function $g_\phi(\mathbf{a}) : \mathcal{D} \to \mathcal{T}$ that maps a set of attributes from some continuous domain $\mathcal{D}$ to the token embedding domain $\mathcal{T}$. We use positional encoding to first cast each attribute $a \in \mathbf{a}$ to a higher frequency space before feeding into the function, which is represented by a very simple 2-layer MLP. The output of this network is named *Continuous 3D Word* and will allow users to easily control continuous attributes from text prompts augmented by these tokens. Finally, our training objective can then be formulated as:

$$\arg\min_{\theta, \phi} \mathbb{E}_{\hat{I}_{\epsilon,\mathbf{a}}, \mathbf{a}}\left[\left\| S_\theta(\hat{I}_{\epsilon,\mathbf{a}}, P(g_\phi(\mathbf{a}))) - I_\mathbf{a} \right\|_2^2 \right]. \quad (2)$$

3.3. Disentangling Object Identity and Attributes

In practice, when we only utilize a single object to learn attributes, a degenerate solution occurs when directly optimizing (2), where S_θ treats the same object with different values for a as different objects. This hinders the generalization capability of changing the attribute when given an image of a new object. To this end, we propose two strategies, one during training and one during inference, to disentangle the object identity and individual attributes.

Training: Learning identity and attributes. We provide a simple regularization by explicitly forcing the model to use the same identifier for images representing the same object. The objective can thus be formulated as:

$$\arg\min_{\theta, \phi} \mathbb{E}_{\hat{I}_{\epsilon,\mathbf{a}}, \mathbf{a}}\left[\left\| S_\theta(\hat{I}_{\epsilon,\mathbf{a}}, P(T_O, g_\phi(\mathbf{a}))) - I_\mathbf{a} \right\|_2^2 \right]. \quad (3)$$

Optimizing both T_O and $g_\phi(\mathbf{a})$ concurrently proved to be difficult from our experiments (see Section 5). We propose a two-stage training strategy as depicted in Figure 2. For every available image in I^O with varying $\mathbf{a}$, we first use the same prompt condition $P(T_O)$ to associate them to the same object, then learn the diffusion model parameters θ and our Continuous 3D Words MLP g_ϕ by using the prompt condition $P(T_O, g_\phi(\mathbf{a}))$.

Inference: Negatively prompting object identifier. To further the disentanglement of attributes against object identities. We propose a simple trick during inference by adding the object identity as a negative prompt. Specifically, for each sampling step, we swap the null-text embedding used

Training Renderings

Results

ControlNet (0.5) ControlNet (1.0) Ours

a)

"A [🕊] [↗]
photo of an
eagle in the
forest"

"A [🕊] [↗]
photo of a
parrot in the
forest"

b)

"A [☀🟪] [↗]
photo of a horse"

"A [☀🟪] [↗]
photo of a taxi"

c)
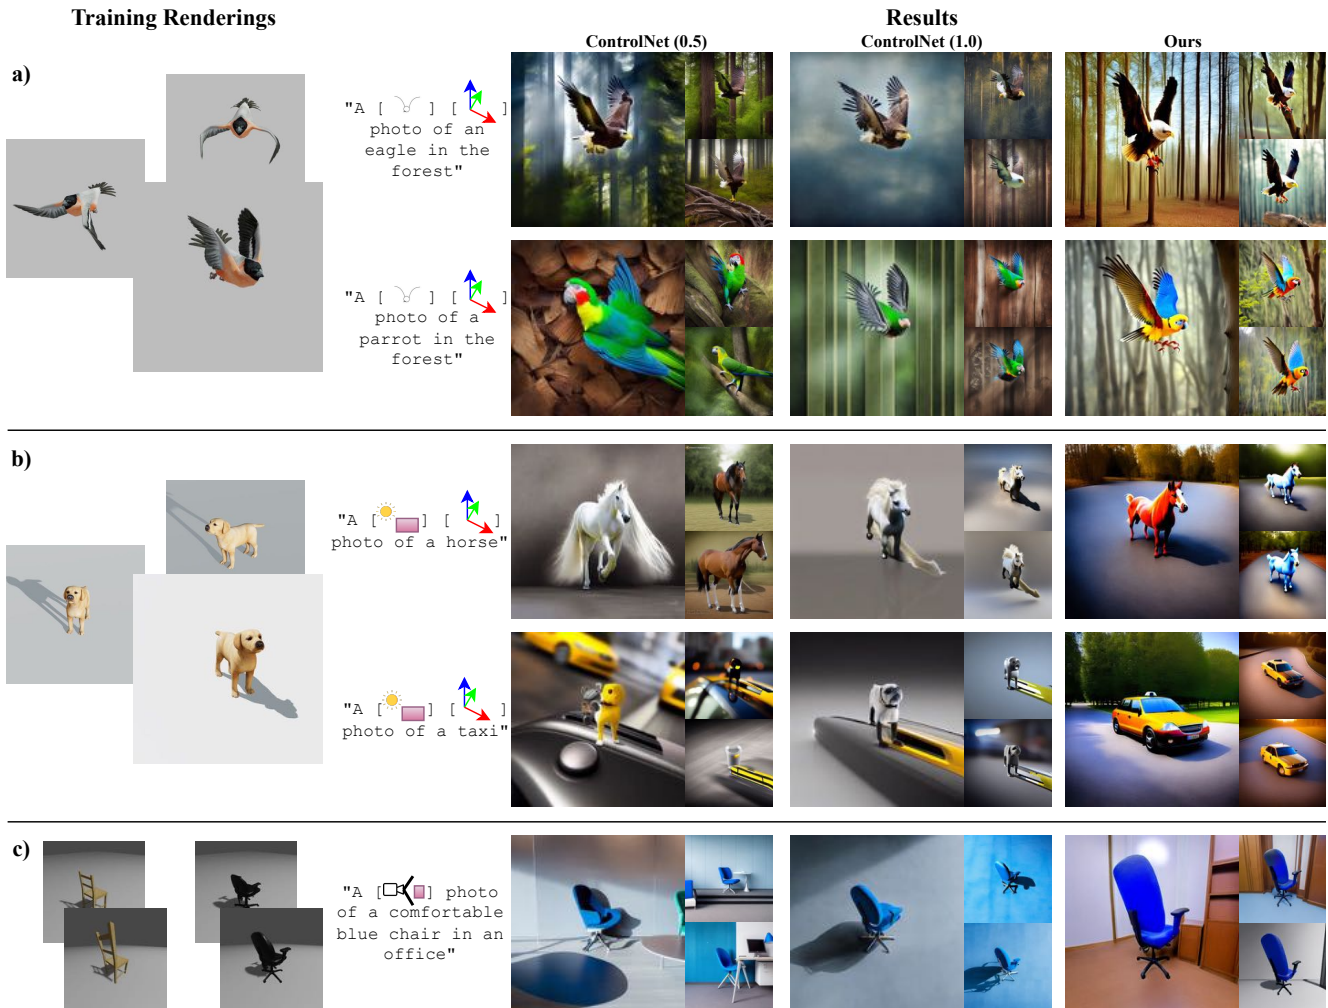
"A [📷🟪] photo
of a comfortable
blue chair in an
office"

Figure 4. **Qualitative Comparisons.** We compare our Continuous 3D Words trained under three settings against ControlNet of various strengths. Note that the dollyzoom setup was trained with trained with multiple chair meshes, so we give additional ControlNet by manually picking the chair rendering that best follows the prompt (i.e., "comfortable" and in "the office").

"Photorealistic, in a forest park
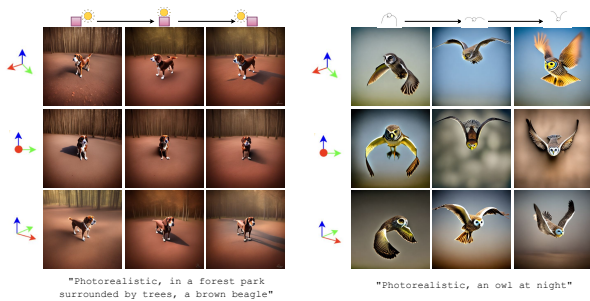surrounded by trees, a brown beagle"

"Photorealistic, an owl at night"

Figure 5. **Disentangling Multiple Attributes.** We show four examples of controlling multiple Continuous 3D words in addition to text descriptions. Left: trained with a single golden retriever mesh. Right: trained with a single animated dove.

by classifier-free guidance with T_O. Intuitively, since our training happened mostly using O, we want to disincentivize the model to generate images containing such object.

3.4. ControlNet Augmentation

To prevent the fine-tuning process from overfitting to a simple white backgrounds and pre-defined object textures, we augment the backgrounds and textures in the rendering process. However, directly doing this in simulation engines is time consuming specially if one is targeting realistic scenes. Thus, we propose an automated solution by utilizing pre-trained ControlNets.

Figure 3 shows two types of ControlNet augmentations we used in our framework. For attributes that can be directly reflected on shape changes (e.g., wing pose), we directly render the ground-truth depth maps to use as the condition for ControlNet. On the other hand, for attributes that cannot be reflected directly from depths (e.g., illumination), we first render the images without textures, then use a lineart extractor to obtain a "sketch" of the image. This captures

subtle changes such as shades and shadows in the pixel space, which can then be used as the condition for Lineart ControlNet.

We add additional prompts describing the object appearance and background during the ControlNet generation. It is important to note that prompts deviating away too much from the original mesh category would lead to degenerate images, so our prompts are in general very simple and straightforward (details in supplemental material). We include the ControlNet generated images as a small set of data augmentation, and we use the same prompt which we use to guide ControlNet generation to guide our Stage 2 fine-tuning.

4. Experiments

We use off-the-shelf Stable Diffusion v2.1 [24] as the backbone of our method. We resort to the recent Low-Rank Adaptation (LoRA) [14] for the fine-tuning of the denoising U-Net and text encoder, allowing us to train on a single A10 GPU occupying roughly 16GB of memory. Thanks to the low-rank optimization, our models have a very small size (approximately 6MB). Training time varies by the complexity of the single/multiple attributes to learn, but falls within 15k to 20k steps, which generally takes around 3-4 hours in a single GPU. For ControlNet augmentation, we use the official implementation of ControlNet v1.1 [30].

We implement our Continuous 3D words under five different attribute settings. For single attributes we implement 1) illumination [🔆] using a single dog mesh, 2) wing pose [🕊] using a single animated dove mesh, and 3) Dolly zoom [🎥] with five Pix3D chairs [28]. For multi-concepts, we train 4) illumination and object orientation [🔆] + [↗] using a single dog mesh and 5) wing pose and orientation [🕊] + [↗] using a single animated dove mesh. Settings 1) and 4) use lineart images [6] while the others use depth map to compute the ControlNet background augmentation (see Section 3.4).

4.1. Comparison with Baselines

Baseline Design. We design a very competitive baseline that enables fine-grained attribute control in image generation by combining the mesh training data we used in our experiments, a rendering engine and ControlNet [30]. Specifically, we take a *novel* text prompt for a *training* object, and grab a corresponding condition map in the training set rendered with the intended attributes. For example, if the prompt is `a [ 🕊 ] eagle flying in a forest`, we select the frame of the dove mesh that corresponds to the user-prescribed wing pose [🕊], render its depth map using a rendering engine and pass it through ControlNet with the same prompt but removing the *continuous 3D word*. The strength of the ControlNet guidance is a critical hyperparameter — increase in strength could increase the accuracy of reflecting the attribute but decrease the robustness to general-

	User Preference (%) ↑				
Method	[🕊]	[🔆]	[🕊]/[↗]	[🔆]/[↗]	Avg.
ControlNet (1.0)	28.3%	16.2%	35.0%	15.0%	23.6%
ControlNet (0.5)	10.0%	28.8%	12.5%	32.5%	21.0%
Ours	61.7%	55.0%	52.5%	52.5%	55.4%

	Average User Ranking ↑				
Method	[🕊]	[🔆]	[🕊]/[↗]	[🔆]/[↗]	Avg.
ControlNet (1.0)	2.07 ± 0.70	1.49 ± 0.76	2.20 ± 0.68	1.60 ± 0.74	1.84 ± 0.72
ControlNet (0.5)	1.38 ± 0.66	2.11 ± 0.67	1.38 ± 0.70	2.06 ± 0.76	1.73 ± 0.70
Ours	2.55 ± 0.62	2.40 ± 0.73	2.43 ± 0.67	2.33 ± 0.77	2.43 ± 0.70

Table 1. **User study results.** We asked users to rank three images according to their preference and how well they followed the given conditions – text prompt and one or two continuous controls. The controls were described by representative images (*i.e.* arrows for orientation, shaded sphere for illumination, *etc.*). *See supplemental material for more details*. We evaluate four different control types: wing pose [🕊], illumination [🔆], wing pose [🕊] + orientation [↗], and illumination [🔆] + orientation [↗]. Cells colored as red and yellow represent the best and second best method, respectively. Our method was selected as the favorite for the majority of users in all evaluated setups.

ize to the text-prompt intended object. Therefore, we present the ControlNet baseline with both full and half strength in terms of guidance. We also explored an interpolation of null-text embedding [18] baseline, but the results failed to capture even the simplest attributes, so they were omitted from our analysis.

Quantitative Results. Due to the complexity and abstract nature of the attributes we are analyzing, automatically measuring whether a generated image reflects a set of values is a considerable challenge. Thus, following previous papers [12, 25, 30], we create a user study to evaluate the quality of the generated images while controlling several attributes. For each setting 1, 2, 4, and 5, we randomly sample a set of 3D conditions and prompts, generating over 60 questions per setting. Similarly to the user studies above, we then invite 20 participants, showing them all the given constraints we want the image to pertain (prompts and attributes), and ask them to rank each image from best to worst.

Table 1 shows the results in both the percentage of user preference (image chose as the best one) and the overall ranking. Best results are highlighted as red. Our *Continuous 3D Words* won the majority of votes (over 50%) in all analyzed scenarios. Interestingly, the second best (as highlighted in yellow) alternates between the two guidance strengths of ControlNet. For wing-pose based controls, having a weaker control diminishes the strong prior offered by the depth map, resulting in inaccurate poses. Conversely, strong guidance for illumination forces ControlNet to generate objects around the shadow outline (further shown in Figure 4). Differently from the baselines, our training strategy is one-size-fits-all without any hyperparameter required.

Qualitative Analyses We present a detailed comparison of Continuous 3D Words ... ance stre... ... training s... ... nation and orientation, and c) dollyzoom. We observe that the dollyz... ... to be don... ... ControlN... ... followed ... ... portantly, ... when the ... ... the training data. For example, even when trained on a single dog mesh ... ... attributes ... ... b), Figure ...

4.2. Multi-Concep...

Just like sentences wh... ... but each disentangled from one another when controlling the image generation, our Continuous 3D Words can do the same. We show four examples in Figure 5, two from [🔶] and [✈] , and two [↻] and [✈] , where we can keep one attribute fixed but change the other without sabotaging the quality of image generations. They can be jointly used with complex prompts describing the background and object texture. Moreover, while all these words are learned from a single mesh, we can easily transfer the attribute to objects with fairly close semantics (e.g., a labrador mesh to a polar bear, or a dove mesh to a parrot).

4.3. Real World Image Editing

Our Continuous 3D Words can be directly applied to real world images to perform editing. To do so, we simply have to encode a real world image to a rare token via Dreambooth [25]. Then, we only have to use that token in conjunction with our Continuous 3D Words to generate the edited image. We show 4 examples in Figure 6, changing [🔶] , [↻] , [✈] . As we can see, the Dreambooth token preserves most of the image appearance, while our Continuous Words understands and brings edits to the main subject.

Comparing with Zero-1-to-3 While our approach is focused on enhancing text-to-image, given the capabilities of real-world image editing, we can use the same setup to provide a comparison with *only* with object orientation changes (Figure 6). Notice that Zero1-to-3 [16] operates in foreground-only images so we also had to segment the object [7], inpaint the background [24], and place the novel orientation back to to the image. Each step contain errors that, when compounded, hurt the quality of the final result. More importantly, our method allows controls beyond orientation changes without relying on massive 3D datasets.
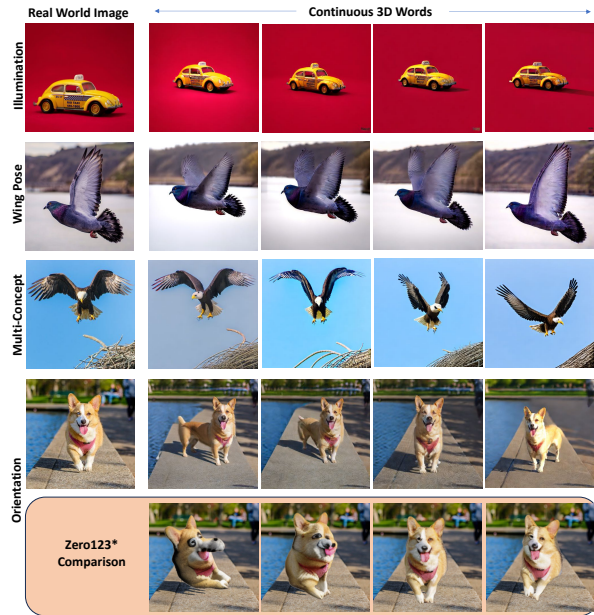

Figure 6. **Real World results & Comparison w/ Zero1-to-3.** We learn tokens representing single images and use along with *Continuous 3D Words* for image editing. The learned image token encodes most of the image appearance, while the *Continuous 3D Word* modifies its relevant aspects. Zero123, on the other hand, tends to yield images where the modified object is deformed or not properly harmonized.
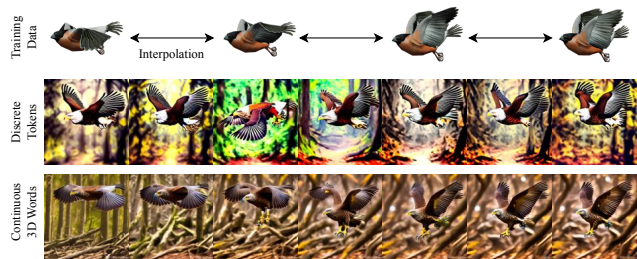

Figure 7. **Interpolating continuous 3D words.** We present two results of interpolation, one with discrete tokens and one with our Continuous 3D Words. Our method preserves the attributes better and enables interpolation between two values.

5. Discussion and Limitations

Why Not Discrete Tokens? The benefits of fitting a single MLP instead of multiple tokens on an attribute with different values are twofold. First, the MLP learns a continuous function, allowing us to interpolate between two training data samples whereas fitting multiple tokens leads to two datapoints as two discrete mappings. Second, finetuning a model to learn multiple custom tokens simultaneously is very hard. These benefits are illustrated in Figure 7. We show a comparison of learning a single *Continuous 3D Word* [↻]
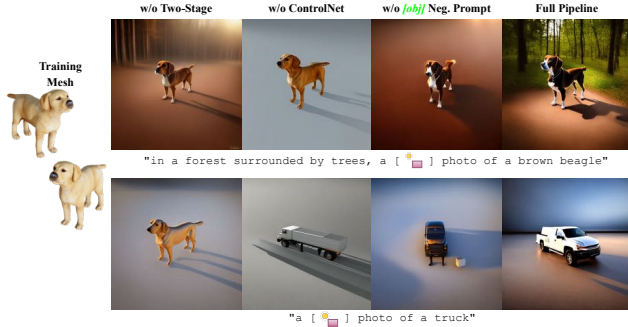
Figure 8. **Ablation Study.** We present our ablations for: w/o two-stage training, w/o ControlNet augmentation, w/o *[obj]* as negative prompt, and our full pipeline. Notice that, when the prompt deviates a lot from the training data (truck in prompt, mesh of a dog for training), the ablated version fails to follow the prompt. Without two-stage training, it ignores the prompt and creates a dog; without the other parts it yields deformed shadows and dog-shaped trucks.

and 18 discrete tokens for different values of a wing pose. Both were learned with the same training method (2-stage with ControlNet augmentation). We present the generated results of three attribute values that are present in the training set, as well as the results when you interpolate the value in between. Interpolation is straightforward for our case where we just input the intermediate value into our Continuous 3D Words MLP. For discrete tokens, we interpolate between two nearest discrete bins. Notice that the discrete tokens have difficulty not only in interpolating results but also in learning the concept simultaneously. The eagle wings in the second row of Figure 7 do not follow the training data and sometimes generate a completely different pose (third column, second row). On the other hand, our images closely follow the pose prescribed by the user (top row) while yielding appealing images even when the values were not seen during training (columns 2, 4 and 6; second row).

Ablation study. Figure 8 shows an ablation of our training strategy. We remove each component of our training strategy (w/o two-stage training, w/o ControlNet augmentation, w/o *[obj]* as negative prompt) and compare it with our full pipeline. Two examples are presented: one where the prompt is similar to the training and another one where it is significantly different. Without the two-stage training, the model fails to disentangle object identity with our attributes, hindering the generalization capability to new objects. This is particularly noticeable for the bottom row when the text prompt is a `truck` but a dog similar to the training mesh is generated when the two-stage training is removed. Without ControlNet, the finetuning process often overfits to the background training renderings, resulting in an inability to generate realistic backgrounds. Finally, adding *[obj]* as the negative prompt serves as a minor improvement in further disentangling both the backgrounds and object shape seen during training, resulting in a more aesthetic image.
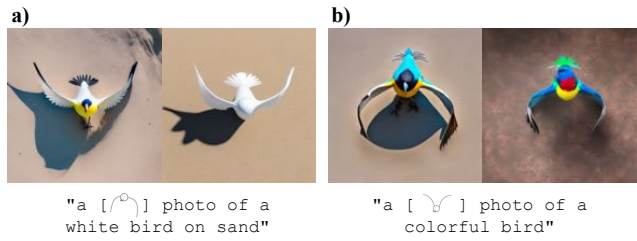

Figure 9. **Condition v.s. Accuracy in User Study.** a) shows two images generated by the prompt "a [🕊] photo of a white bird on sand". b) shows two images generated by the prompt "a [🕊] photo of a colorful bird". Left is ours, right is ControlNet (1.0). Users preferred right over left for both.
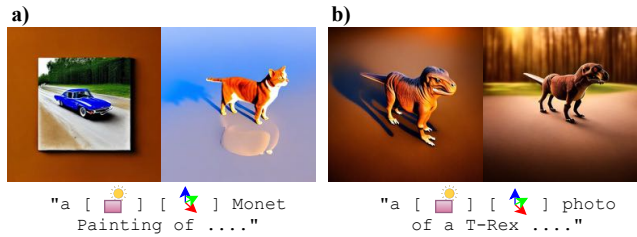

Figure 10. **Failure cases.** a) shows images generated with prompt "a [🖼] [🪁] Monet Painting of .....". b) shows images generated with "a [🖼] [🪁] photo of a T-Rex".

Condition v.s. Generated Accuracy in User Study. During our user study, we realize that occasionally users select the image which more strictly follows the exact conditions over the more physically probable image. For example, in Figure 9 a) and b), the head of the bird is generated poorly, but users preferred them as they are "whiter" and "more colorful".

Failure Cases. We present in Figure 10 two examples of typical failure cases in our results. First, our model currently fails on more difficult scenarios where the style is given by the text prompt. In a), the image cannot fully reflect the "Monet painting" style imposed by our prompt. Second, the generated object may sometimes still overfit to the training set. In b), the T-Rex had four feet on the ground instead of standing with two claws in air – an attribute that is similar to the training dog mesh used to learn illumination.

6. Conclusion

We presented *Continuous 3D Words*, a framework that allows us to learn 3D-aware attributes reflected in renderings of meshes as special words, which can then be injected into the text prompts for fine-grained text-to-image generation. Extensive study on learning both single and multiple continuous words shows that we can control challenging attributes. With the lightweight design and promising results, we hope that this work opens up interesting applications in the vision community to create their own 3D words with a single mesh and an accessible rendering engine.

References

[1] Yuval Alaluf, Elad Richardson, Gal Metzer, and Daniel Cohen-Or. A neural space-time representation for text-to-image personalization. *arXiv preprint arXiv:2305.15391*, 2023. 3

[2] James Betker, Gabriel Goh, Li Jing, † TimBrooks, Jianfeng Wang, Linjie Li, † LongOuyang, † JuntangZhuang, † JoyceLee, † YufeiGuo, † WesamManassra, † PrafullaDhariwal, † CaseyChu, † YunxinJiao, and Aditya Ramesh. Improving image generation with better captions, 2023. 2

[3] Tim Brooks, Aleksander Holynski, and Alexei A. Efros. Instructpix2pix: Learning to follow image editing instructions. In *CVPR*, 2023. 2

[4] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020. 2

[5] James Burgess, Kuan-Chieh Wang, and Serena Yeung. Viewpoint textual inversion: Unleashing novel view synthesis with pretrained 2d diffusion models. *arXiv preprint arXiv:2309.07986*, 2023. 4

[6] Caroline Chan, Fredo Durand, and Phillip Isola. Learning to generate line drawings that convey geometry and semantics. In *CVPR*, 2022. 6

[7] Danielgatis. Danielgatis/rembg: Rembg is a tool to remove images background. 7

[8] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. 2

[9] Rinon Gal, Yuval Alaluf, Yuval Atzmon, Or Patashnik, Amit Haim Bermano, Gal Chechik, and Daniel Cohen-or. An image is worth one word: Personalizing text-to-image generation using textual inversion. In *The Eleventh International Conference on Learning Representations*, 2022. 3, 4

[10] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014. 2

[11] Ligong Han, Yinxiao Li, Han Zhang, Peyman Milanfar, Dimitris Metaxas, and Feng Yang. Svdiff: Compact parameter space for diffusion fine-tuning. *ICCV*, 2023. 3

[12] Amir Hertz, Ron Mokady, Jay Tenenbaum, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Prompt-to-prompt image editing with cross attention control. In *arXiv preprint arXiv:2208.01626*, 2022. 2, 6

[13] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. 2

[14] Edward J Hu, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2021. 2, 6

[15] Nupur Kumari, Bingliang Zhang, Richard Zhang, Eli Shechtman, and Jun-Yan Zhu. Multi-concept customization of text-to-image diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 3, 4

[16] 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, 2023. 3, 7

[17] Zhiheng Liu, Yifei Zhang, Yujun Shen, Kecheng Zheng, Kai Zhu, Ruili Feng, Yu Liu, Deli Zhao, Jingren Zhou, and Yang Cao. Cones 2: Customizable image synthesis with multiple subjects. *arXiv preprint arXiv:2305.19327*, 2023. 3

[18] Ron Mokady, Amir Hertz, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Null-text inversion for editing real images using guided diffusion models. *arXiv preprint arXiv:2211.09794*, 2022. 6

[19] Chong Mou, Xintao Wang, Liangbin Xie, Jian Zhang, Zhongang Qi, Ying Shan, and Xiaohu Qie. T2i-adapter: Learning adapters to dig out more controllable ability for text-to-image diffusion models. *arXiv preprint arXiv:2302.08453*, 2023. 2

[20] Gaurav Parmar, Krishna Kumar Singh, Richard Zhang, Yijun Li, Jingwan Lu, and Jun-Yan Zhu. Zero-shot image-to-image translation. In *ACM SIGGRAPH 2023 Conference Proceedings*, pages 1–11, 2023. 2

[21] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research*, 21(1):5485–5551, 2020. 2

[22] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *International Conference on Machine Learning*, pages 8821–8831. PMLR, 2021. 2

[23] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. In *arXiv preprint arXiv:2204.06125*, page 3, 2022. 2

[24] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models, 2021. 2, 4, 6, 7

[25] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22500–22510, 2023. 2, 3, 4, 6, 7

[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] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations*, 2020. 2

[28] Xingyuan Sun, Jiajun Wu, Xiuming Zhang, Zhoutong Zhang, Chengkai Zhang, Tianfan Xue, Joshua B Tenenbaum, and William T Freeman. Pix3d: Dataset and methods for single-

image 3d shape modeling. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2974–2983, 2018. 6

[29] Paul Yoo, Jiaxian Guo, Yutaka Matsuo, and Shixiang Shane Gu. Dreamsparse: Escaping from plato's cave with 2d diffusion model given sparse views. *arXiv preprint arXiv:2306.03414*, 2023. 3

[30] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *IEEE International Conference on Computer Vision (ICCV)*, 2023. 2, 6