

# DragAPart: Learning a Part-Level Motion Prior for Articulated Objects

Ruining Li, Chuanxia Zheng, Christian Rupprecht, and Andrea Vedaldi

Visual Geometry Group, University of Oxford  
 {ruining, cxzheng, chrisr, vedaldi}@robots.ox.ac.uk  
[dragapart.github.io](https://github.com/vedaldi/dragapart)



**Fig. 1: Examples of our DragAPart.** Each drag in DragAPart represents a *part-level* interaction, resulting in a physically plausible deformation of the object shape. DragAPart is trained on a new synthetic dataset, Drag-a-Move, for this task, and generalizes well to *real* data and even *unseen categories*. The trained model also can be used for segmenting movable parts and analyzing motion prompted by a drag.

**Abstract.** We introduce DragAPart, a method that, given an image and a set of drags as input, generates a new image of the same object that responds to the action of the drags. Differently from prior works that focused on repositioning objects, DragAPart predicts part-level interactions, such as opening and closing a drawer. We study this problem as a proxy for learning a generalist motion model, not restricted to a specific kinematic structure or object category. We start from a pre-trained image generator and fine-tune it on a new synthetic dataset, Drag-a-Move, which we introduce. Combined with a new encoding for the drags and dataset randomization, the model generalizes well to real images and different categories. Compared to prior motion-controlled generators, we demonstrate much better part-level motion understanding.

## 1 Introduction

We consider the problem of learning an *interactive* image generator that allows moving the parts of an object by dragging them [44, 62] instead of repositioning

the object as a whole. For example, dragging on the door of a cabinet should result in the image of the same cabinet but with the door open (Fig. 1).

Besides applications to controlled image generation, we explore dragging as a way of learning and probing generalist models of motion. Modelling deformable objects often uses ad-hoc models specific to each object type. For example, SMPL [39] only represents humans and SMAL [79] only certain mammals. Some works [25, 34, 71, 72] have suggested training similar models automatically from raw Internet images but still assuming that objects share the same kinematic structure (a skeleton) which limits generality. In stark contrast, foundation models like CLIP [49], GPT-4 [42], DALL-E [50] and Stable Diffusion (SD) [52] take a generalist approach. They do not pose any restriction on the type of content in the images or text. Likewise, a foundation model of motion should be able to understand the motion of any object, whether it is a human, a jellyfish, a towel, or a piece of furniture. However, there is no universal deformable template that can express the pose and deformation of all objects in a homogeneous manner — for example, the pose of a human is hardly comparable to that of a cabinet.

We posit that a model of motion does *not* require to refer to a template; it is enough that the model understands the possible physical configurations of an object and their transitions. Dragging provides a way to probe such a model without using a template. A drag specifies how a single physical point of the object moves, leaving the model to ‘fill in’ the details by predicting a plausible motion for the object’s parts. For example, dragging a single drawer in a cabinet should not result in all other drawers moving together because their motion is independent, but dragging the handle of a microwave oven’s door should open the door itself, as the motion is, in this case, perfectly correlated (Fig. 1).

The dragging approach also lends itself to large-scale training, potentially from millions of Internet images and videos, by exploiting the strengths of recent generative models [19, 50, 52, 59]. In fact, it is natural to ask whether existing off-the-shelf image generators [50, 52] might already understand dragging and thus motion. Like others [17, 41, 58, 62], we found that these models do respond to dragging, but only in the sense of shifting an entire object, but *fail to capture more nuanced effects*. Hence, the challenge is to encourage generative models to produce fine-grained and, thus informative dragging outcomes.

Our first contribution is to develop Drag-a-Move, a dataset to align image generators with part-level dragging. Previous works [32, 38, 57, 78] have proposed to use synthetic 3D objects [13, 14] to fine-tune pre-trained generative models for new view synthesis (NVS). Because the base model is trained on billions of images, a relatively small number of such objects, in the order of millions, is sufficient to train generalizable NVS. We thus propose to take a similar approach and contribute data rendered from an existing 3D dataset with rich part-level annotations [18]. Augmented with drag annotations, this data can be used to fine-tune the generator and learn non-trivial drag responses. Even so, our synthetic data is far less abundant than that available for NVS, which includes several million objects [13, 14]. Instead, our data is in the hundreds of different synthetic 3D objects due to the cost of annotating parts and kinematics. While the number

of objects is small, we show that *texture randomization* can significantly improve generalization to out-of-domain cases, including real data and unseen categories at test time.

We then use our dataset Drag-a-Move to train a new model, DragAPart, for the task of part-oriented drag control. Key to the design of DragAPart is how the drags are encoded. Compared to prior works that considered similar tasks, we propose a new way of encoding drags with benefits that are transversal to the underlying model architecture. We explore different network architectures for this task: DiT [45] with a transformer [66] backbone trained from scratch and Latent Diffusion architecture [52] with a U-Net [53] backbone pre-trained on billions of images, and study their performance on synthetic and real data.

Finally, we also explore some downstream applications of DragAPart. First, we show that it can be used to optimize the motion parameters of a given articulated 3D object to predict *how* its movable part is likely to move, subject to a drag. Second, we show that the model’s implicit part-level understanding can be leveraged to segment moving parts in an image prompted by a drag.

## 2 Related Work

*Generative Modeling.* Generative models, such as diffusion models [22, 60, 61], can generate high-quality images [50, 52, 55], videos [4, 21] and 3D assets [31, 36, 63]. Conditioning plays a key role in controlling the generated content and its quality. For instance, DALL-E 3 [3] emphasizes the importance of quality textual prompts used for training. ControlNet [76] shows how to inject new kinds of controls in a pre-trained diffusion model and is at the basis of a new model zoo [47], widely used in applications such as text-to-image/video/3D [20, 48, 54] and image-to-image/video/3D [23, 25, 69, 78].

*Drag-Conditioned Image & Video Synthesis.* Authors have considered controlling image and video generation via dragging. *Training-free* methods iteratively update the source image to match user-specified drags. Among those, DragGAN [44] optimizes a latent representation of the image in StyleGAN [29] to match the user-specified drags. DragDiffusion [58] ports this idea to SD [52]. DragonDiffusion [41] and Motion Guidance [17] combine SD with a guidance term [15] that captures feature correspondences and a flow loss. *Training-based* methods, on the other hand, learn drag-based control using ad-hoc training data for this task. For instance, iPoke [5] trains a variational autoencoder (VAE) to synthesize videos with objects in motion. MCDiff [10] and YODA [12] train a diffusion model using DDPM and flow matching, respectively. Li *et al.* [33] use a Fourier-based representation of motion suitable for natural, oscillatory dynamics characteristic of objects like trees and candles and generates motion with a diffusion model. DragNUWA [75] and MotionCtrl [68] extend text-to-video generators with drag-based control.

Many of these methods show impressive results in terms of repositioning objects in a scene, but *do not* allow to control motion at the level of object

parts. Here, we address the latter task by introducing a new synthetic dataset and a new encoding technique for the drags.

### 3 Method

We develop *DragAPart*, an *interactive* generative model that, given as input a single object-centric RGB image  $y$  and one or more drags  $\mathcal{D}$ , synthesizes a second image  $x \sim \mathbb{P}(x|y, \mathcal{D})$  that reflects the effect of the drags. We first provide a formal definition of the problem and discuss the challenges of learning such a motion prior for articulated objects in Sec. 3.1. To overcome these challenges, we introduce a novel mechanism to encode the drags  $\mathcal{D}$  (*i.e.*, the sparse motion condition), which enables more efficient information propagation in Sec. 3.2. Since we train our model only on synthetic renderings, in Sec. 3.3, we propose a simple yet effective way to mitigate the sim-to-real gap. In Sec. 3.4, we suggest downstream applications of DragAPart.

#### 3.1 Preliminaries

We consider images  $x, y \in \mathbb{R}^{3 \times \Omega}$  as tensors defined on the spatial grid  $\Omega = \{1, \dots, H\} \times \{1, \dots, W\}$ . A *drag*  $d$  is a pair  $(u, v) \in \Omega \times \mathbb{Z}^2$ . Given the image  $y$  and one or more drags  $\mathcal{D} \subset \Omega \times \mathbb{Z}^2$ , the goal is to draw samples from the conditional distribution  $\mathbb{P}(x|y, \mathcal{D})$ . The origin  $u$  of the drag is always contained in the domain  $\Omega$  of image  $y$ , but the termination  $v$  may not be since the selected point may move outside the camera’s field of view.

The distribution  $\mathbb{P}(x|y, \mathcal{D})$  is defined by the expected effect of the drags. Each image contains a certain object, such as a cabinet or a microwave oven, and the effect of each drag is to interact with the corresponding object part. Formally, let  $M \subset \mathbb{R}^3$  be the (unknown) surface of the 3D object contained in the image and consider a drag  $(u, v) \in \Omega \times \mathbb{Z}^2$ . Let  $p \in M$  be the physical object point that corresponds to the drag origin  $u$ . Then, the effect of the drag  $(u, v)$  is to cause the object to deform such that point  $p$  moves to a new location  $q \in \mathbb{R}^3$  that re-projects as close as possible to the drag termination  $v$ , while obeying the physical rules that govern the object’s deformation. The effect of multiple drags is to specify multiple such constraints, which must be satisfied simultaneously. Because there is no explicit model of the 3D shape or kinematic properties of the object in the image, the model must learn these properties implicitly.

Note that not all drags are physically plausible, particularly if the drags are given by a user, so they may not be perfectly physically consistent with respect to the underlying 3D object’s geometry. Hence, the model must simulate a motion of the object that corresponds to the drags as well as possible, but *without breaking physical plausibility*. On the other hand, drags often under-specify the motion, which justifies using a stochastic rather than deterministic prediction model. In particular, a single drag can often be satisfied in a physically plausible manner by translating the object as a whole. This solution is not very informative about

the object’s dynamics. Instead, we are interested in *modelling nuanced physical interactions at the level of the individual object parts*.

Such nuanced interactions are typical in natural data — for example, dragging on the handle of a drawer normally results in opening the drawer instead of translating the entire cabinet — so, in principle, should be learned automatically from, *e.g.*, raw Internet data. Perhaps surprisingly, however, we found that this is *not* necessarily the case. Following the lead of prior motion-conditioned image generators [17, 41, 44, 58], the obvious approach is to start from a pre-trained foundation model like Stable Diffusion [52], and then modify its internal representation of the image  $y$  to respond to drags. The advantage is that foundation models are trained on billions of Internet images and, as such, should have an understanding of almost any object and its configurations. However, we found that such methods tend to *move the object as a whole* at best, which is uninteresting for us. We hypothesize that the reason is that the training data of foundation models, while abundant in quantity, inevitably entangles multiple factors such as camera viewpoint and multi-object interaction, making it difficult to understand different *part-level* object configurations.

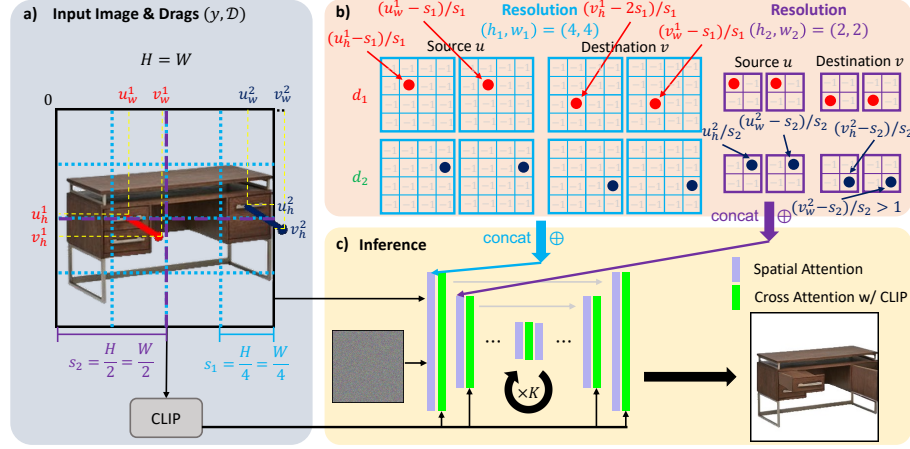
### 3.2 DragAPart: Architecture and Drag Encoding

Given the challenges identified in Sec. 3.1, we now describe DragAPart, which learns a motion prior for articulated objects from a pre-trained image generator. Key to our model is to fine-tune the motion generator, which models the distribution  $\mathbb{P}(x|y, \mathcal{D})$ , on a synthetic dataset of triplets  $(x, y, \mathcal{D})$ . This dataset will be described in Sec. 4. In addition to this data, which is not very large compared to datasets used to train image generators, we must also address the challenge of generalization. The rest of this section describes our architectural choices which improve the model’s generalizability, both *across domains* to real-world images and *across categories* to those that are not present in the dataset.

*Latent Diffusion Architecture.* We build our model on large pre-trained diffusion models for image generation, such as Stable Diffusion (SD) [52]. These models are trained on a very large number of natural images (*i.e.* LAION [56]) and are open-ended. However, they are controlled by textual prompts, whereas our conditioning input consists of a reference image  $y$  and a set of drags  $\mathcal{D}$ . Hence, our goal is to modify the model to accept the pair  $(y, \mathcal{D})$  as conditions.

In more detail, DragAPart uses a Latent Diffusion Model (LDM) architecture and is equipped with an encoder  $E$ , a decoder  $D$  and a denoiser  $\Phi$ , all implemented as neural networks. For training, the image  $x$  is first mapped to a latent code  $z = E(x)$  to which we add Gaussian noise  $\epsilon \sim \mathcal{N}(0, \mathbf{I})$  to obtain  $z_t = \sqrt{1 - \sigma_t^2}z + \sigma_t\epsilon$ , where  $\sigma_0, \dots, \sigma_T$  is a noise schedule that increases monotonically from 0 to 1. For efficiency, the code  $z \in \mathbb{R}^{4 \times h \times w}$  is designed to have reduced resolution  $(h, w)$  compared to the image resolution  $(H, W)$ . The denoising model  $\Phi$  is trained by optimizing the objective function

$$\min_{\Phi} \mathbb{E}_{(x, y, \mathcal{D}), t, \epsilon \sim \mathcal{N}(0, \mathbf{I})} [\|\epsilon - \Phi(z_t, t, y, \mathcal{D})\|_2^2]. \quad (1)$$



**Fig. 2: The Overall Pipeline of DragAPart.** (a) Our model takes as input a single RGB image  $y$  and one or more drags  $\mathcal{D}$ , and generates a second image  $x$  that reflects the effect of the drags (Sec. 3.1). (b) We propose a novel flow encoder (Sec. 3.2), which enables us to inject the motion control into the latent diffusion model at different resolutions more efficiently (The resolutions 4 and 2 are for illustrative purposes. Our model generates  $256 \times 256$  images, and the first two latent blocks have resolutions 32 and 16.). (c) At inference time, our model generalizes to real data, synthesizing physically-plausible part-level dynamics.

For the network  $\Phi$ , we experiment with two diffusion architectures, SD and DiT. The SD model is a UNet with several transformer blocks  $l = 1, 2, \dots, L$  operating at different spatial resolutions  $(h_l, w_l)$ . It is conditioned on a textual prompt encoded using CLIP [49]. At each UNet block  $l$ , a self-attention layer generates spatial queries, keys and values  $Q_l, K_l, V_l \in \mathbb{R}^{C_l \times h_l \times w_l}$  at resolution  $(h_l, w_l)$ . In order to inject the source image  $y$  as a condition, we follow [8, 70] and pass  $y$  (without adding noise to it) through the same network  $\Phi$ . Then, the keys and values  $K_l^y, V_l^y$  from image  $y$  are used to replace the keys and values  $K_l, V_l$  for the generated image  $x$ , thus turning self-attention into cross attention between the noised image and the clean reference image  $y$ . In order to encode the drags  $\mathcal{D}$ , we consider a drag encoder  $F_l(\mathcal{D}) \in \mathbb{R}^{C_l \times h_l \times w_l}$  outputting an encoding of the drags at the resolution compatible with transformer block  $l$ . These features are concatenated along the channel dimension to the input of each transformer block. The transformer layers are modified to handle the additional input channels, zero-initializing the corresponding model weights.

The DiT model has several attention blocks at the same latent resolution. For each block, we follow [45] and inject the (tokenized) drag encoding via a token-wise adaLN-zero block [46]. We defer the details to the sup. mat.

*Drag Encoding.* We now define the encoding  $F_l(\mathcal{D})$  of the drags. Prior works [68, 75] represent  $\mathcal{D}$  as a sparse optical flow image  $F(\mathcal{D}) \in \mathbb{R}^{2 \times \Omega}$  which is all zero except that, for each drag  $(u, v) \in \mathcal{D}$ , location  $u$  contains the displacement  $v - u$ .

This mechanism has two shortcomings. First, the transformer can sense well the origin  $u$  of the drag, which is encoded spatially, but cannot do so for the termination  $v$ . Second, because all drags are encoded by a single optical flow image and the latter must be reduced to resolution  $(h_l, w_l)$  before it can be injected into the LDM transformer block  $l$ , distinct drags may overlap. For instance, [68, 75] first blur the sparse flow image  $F$  using a Gaussian filter and then apply further convolutional layers to the result before injecting the information into the LDM. As demonstrated in Sec. 5, this design can confuse different drags and diffuse a drag so that it affects more than the intended part (Fig. 5 left).

In order to mitigate these downsides, we introduce a new encoding  $F_l$  for the drags, illustrated in Fig. 2. First, we assume a maximum number  $|\mathcal{D}| \leq N$  of drags and assign different channels to different ones, so that *they do not interfere with each other in the encoding*. Second, we use *separate channels* to encode the drag source  $u$  and the drag termination  $v$ . Specifically, given a single drag  $d = (u, v)$  where  $u = (u_h, u_w)$ , let  $F_l(u, v) \in \mathbb{R}^{2 \times h_l \times w_l}$  be the tensor of all negative ones<sup>1</sup> except for  $F_l(\lfloor \frac{u_h \cdot h_l}{H} \rfloor, \lfloor \frac{u_w \cdot w_l}{W} \rfloor) = (\frac{u_h \cdot h_l}{H} - \lfloor \frac{u_h \cdot h_l}{H} \rfloor, \frac{u_w \cdot w_l}{W} - \lfloor \frac{u_w \cdot w_l}{W} \rfloor) \in [0, 1]^2$ . We then define  $F_l(\mathcal{D}) = \bigoplus_{(u,v) \in \mathcal{D}} F_l(u, v) \oplus F_l(v, u) \in \mathcal{R}^{4N \times h \times w}$  to be the concatenation of the encoding of all drags in  $\mathcal{D}$  along the channel dimension. Note that we encode each drag twice, swapping source and termination.<sup>2</sup> If there are less than  $N$  drags, we simply pad the tensor with  $-1$ . We show in Sec. 5 that our proposed *multi-resolution drag encoding* technique boosts the generation quality in both U-Net and attention-based LDMs.

### 3.3 Generalization to Real Data via Domain Randomization

Most objects in GPartNet [18], which we use to construct the dataset used to train our model (see Sec. 4), have synthetic textures not representative of real-world objects. Since these textures usually lack variability in color, models trained on these renderings can potentially “cheat” by generating the same pixel values everywhere without an implicit understanding of where the pixel comes from in the original image  $y$ . Empirically, we find training the model jointly on a version of the renderings in which each object part is assigned a random color as shown in Fig. 1 and Fig. 3 (right), sometimes called *domain randomization* [65], significantly improves the model’s performance on real-world images. The effectiveness of this training strategy is validated in Sec. 5.

### 3.4 Downstream Applications

While our primary goal is to demonstrate that an image generator can learn a motion prior for articulated objects at *part-level*, we also explore the potential

<sup>1</sup> So we distinguish between no drag and drags precisely between two latent pixels.

<sup>2</sup> There is a caveat when the drag termination  $v$  is outside the image, in which case  $(\lfloor \frac{u_h \cdot h}{H} \rfloor, \lfloor \frac{u_w \cdot w}{W} \rfloor)$  is outside resolution  $(h, w)$ , so we define the termination  $v$  at the nearest spatial location *inside*  $(h, w)$ . The corresponding values will also be relative to this new coordinate, so they might not lie in  $[0, 1]$  (see Fig. 2 for an example).



of our DragAPart in several downstream applications. By operating at the level of object parts, our approach is complementary to prior works that focus on the task of repositioning [68, 75] and can address different applications than them.

*Segmenting Moving Parts.* First, DragAPart can be used to segment the moving parts of an object. Several prior works [2, 28, 64, 74] have used pre-trained generators for 2D segmentation, but focus on part semantic rather than part mobility. We suggest that DragAPart, as a motion prior, can be used to segment the part that would move in response to a drag (useful, *e.g.*, for predicting affordances).

In order to segment mobile parts, we follow [2] and extract internal features from the forward pass of the denoiser  $\Phi$ . However, [2] looks at semantics and only extracts a single set of features from the input image. Instead, because our goal is to obtain the segmentation of the part that moves prompted by a drag, we run the forward pass twice, once with the drag conditioning and once without, and compute the difference between the resulting features. This difference represents the effect of the drag in the diffusion model’s internal representation.

*Motion Analysis for Articulated Objects.* Next, we consider motion analysis, *i.e.*, the problem of understanding how the parts of an object move [1, 16, 18, 26, 35, 40]. Here, we suggest a method to use DragAPart for motion analysis. Given a 3D mesh  $\mathcal{M} = \mathcal{M}_{\text{static}} \cup \mathcal{M}_{\text{moving}}$  with moving parts pre-segmented and a 3D drag on  $\mathcal{M}$ , the goal is to output the type of motion  $t$  and the parameters  $p_{\text{motion}}$ <sup>3</sup>.

In order to estimate the parts’ motion types and parameters using DragAPart, we start by rendering  $K$  images  $y_k$  of the object  $\mathcal{M}$  from random cameras  $C_k$ . We also project the 3D drag to pixel space  $d_{\text{pixel}}^k \in \Omega \times \mathbb{Z}^2$ . Then, we use DragAPart to generate images  $x_k \sim \mathbb{P}(x_k | y_k, \{d_{\text{pixel}}^k\})$  of the dragged object under  $C_k$ . Following [35], using the generated images, we estimate the motion parameters  $(t, p_{\text{motion}})$  by minimizing the objective function

$$\underset{t, p_{\text{motion}}}{\operatorname{argmin}} \frac{1}{K} \sum_{k=1}^K \|R(\mathcal{M}_{\text{static}} \cup T(\mathcal{M}_{\text{moving}}; t, p_{\text{motion}}), C_k) - x_k\|_2^2 \quad (2)$$

where  $R$  is the rendering function and  $T$  transforms the moving parts based on the motion estimation. While [35] solves Eq. (2) using back-propagation, we find gradient-based optimization prone to local optima, leading to sub-optimal estimation. Since the search space has relatively low dimensionality<sup>4</sup>, we optimize Eq. (2) directly using grid search.


## 4 The Drag-a-Move Dataset

Key to training our model is a suitable dataset with drag annotations. It must thus contain a collection of triplets  $(x, y, \mathcal{D})$ , where image  $y$  represents the initial

<sup>3</sup> Following [35] for each part. The motion type is one of {revolute, prismatic}. For a revolute joint, the motion parameter  $p_{\text{motion}}$  is represented by a pivot point  $\mathbf{p} \in \mathbb{R}^3$  and a rotation in the form of a quaternion  $\mathbf{q} \in \mathbb{R}^4$ ; for a prismatic joint, the parameter specifies the joint axis  $\mathbf{a} \in \mathbb{R}^3$  and translation distance  $s \in \mathbb{R}$  along this axis.

<sup>4</sup> 7 for revolute type and 4 for prismatic type.





**Fig. 3: Animations from the Drag-a-Move dataset.** We visualize two objects with diverse articulation states: the left is rendered with the original texture and the right with each part in a single random color.

state of the object,  $\mathcal{D}$  is the collection of drags applied to  $y$  and image  $x$  denotes the same object in a new state compatible with the action of the drags. Moreover, we wish the state change to be minimal, in the sense that each drag only affects the smallest subset of object parts that still result in a physically plausible state.

Creating such a dataset using real-world images would require carefully capturing photos of an articulated object in two articulation states while ensuring consistency in camera position, lighting condition and background across captures, and then manually annotating or reconstructing the drags. Fortunately, recent works on novel-view synthesis have indicated that generative models can be effectively fine-tuned using synthetic data only. We thus choose to create a synthetic dataset, which allows us to control the articulation and lighting conditions, and to obtain the ground-truth drags directly from the 3D model.

We build our dataset on an existing 3D synthetic dataset, GAPartNet [18], which adds rich, part-level annotations to the objects of [37, 73]. By rendering the assets under different articulations, we can obtain synthetic data with desired properties, as described next. Among all the objects in GAPartNet, we select those that have at least one part which is annotated with one of the labels “Hinge Handle”, “Slider Drawer”, “Hinge Lid” and “Hinge Door”, which results in 763 candidate 3D models spanning 16 categories. See the sup. mat. for details.

*Animating the Assets.* The annotation of each object contains its kinematic tree structure, where nodes are rigid parts and edges are joints (fixed or not). Our goal is to sample a diversity of starting and ending states (before and after drag) of the different parts, including cases in which only a subset of parts move. We generate with equal probability two kinds of animations. In the first, some joints are fixed and locked to their “closed” or “rest” state, whereas the other joints transition continuously between two extreme states (*e.g.*, a drawer fully closed or fully open, Fig. 3 right). The second kind is similar, but the stationary joints are set to a random state, and the moving joints transition between two randomly selected states (Fig. 3 left). This is further detailed in the sup. mat.

*Generating the Drags.* Since we have ground-truth 3D data with articulation annotation, we can project the motion  $(p, q) \in \mathbb{R}^3 \times \mathbb{R}^3$  of any visible 3D point on the moving part back to pixel space to obtain a drag  $(u, v)$ . However, this would result in a dense set of drags, akin to optical flow, whereas we are interested in only selecting one drag per part. In order to generate a sparse subset of drags, for each moving joint we first sample a part in the subtree (of the entire kinematic

structure) rooted at that joint uniformly at random. Then, we sample a point on the visible surface of that part with probabilities proportional to its displacement due to the parent joint’s motion. We use the trajectory of that point as the drag for the moving parts. This process is fast, enabling us to generate different drag combinations *on the fly* during training to augment the dataset further.

*Rendering.* For each 3D asset, we render 48 articulation animations, with 36 frames each, under random lighting conditions and camera viewpoints by means of the SAPIEN [73] renderer. This amounts to 40 million unique image pairs  $(x, y)$  for training purposes. We also render a version of the dataset where each part of the 3D model is given a random, monochromatic color texture (Fig. 3 right), which we find useful for the model’s generalizability (Sec. 3.3).

## 5 Experiments

We show that our method outperforms prior works both quantitatively and qualitatively, and conduct ablation studies to validate our design choices.

### 5.1 Datasets and Implementation Details

*Articulated Objects.* We assess DragAPart on the test split of our Drag-a-Move dataset (Sec. 4), which contains 150 image pairs from 23 objects not seen during training spanning 10 categories. Furthermore, we also conduct evaluations on real images sourced from the Common Object in 3D (CO3D) dataset [51] and the Amazon Berkeley Objects (ABO) dataset [11]. For these, we use SAM [30] for object segmentation and manually define plausible drags.

As there is no dataset of real images for this task to test generalizability, we generate a second version of the test split using Blender and extract the drags using an off-the-shelf point tracker [27]. This is closer to real-world data because Blender generates more realistic images and the tracker’s inherent inaccuracies produce trajectories more representative of the user-defined drags compared to using ground-truth 3D information. We refer to this version of the test set as Out-of-Distribution (O.O.D.) and the original version as In-Distribution (I.D.).

*Humans.* We train a separate model for humans on the Human3.6M dataset [24]. Following previous works [5, 10], we use actors S1, S5, S6, S7 and S8 for training and S9 and S11 for testing. We randomly sample two frames, as the conditioning image and target, within a video sequence. The corresponding drags are derived using the annotated paths of keypoints that show sufficient displacement.

*Implementation Details.* We evaluate SD [52] and DiT [45] as LDM architectures. In our implementation, we minimally alter the original architecture to enable a more accurate comparison of the capacities of both architectures. Detailed modifications and our training schedule are provided in the sup. mat.

A notable difference in training between the two architectures is that we fine-tune the U-Net LDM from a pre-trained checkpoint SD v1.5 [52], while we train from scratch for DiT [45], due to the lack of publicly available DiT

Method	I.D. Drag-a-Move			Human3.6M [24]		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
iPoke [5]	16.79	0.883	0.150	21.43	0.856	0.258
DragDiffusion [58]	15.30	0.773	0.226	18.06	0.707	0.280
DragonDiffusion [41]	17.63	0.852	0.162	19.45	0.765	0.258
DragNUWA [75]	13.58	0.765	0.277	15.16	0.668	0.292
DragAPart (Ours)	<b>21.38</b>	<b>0.925</b>	<b>0.066</b>	<b>23.82</b>	<b>0.870</b>	<b>0.091</b>

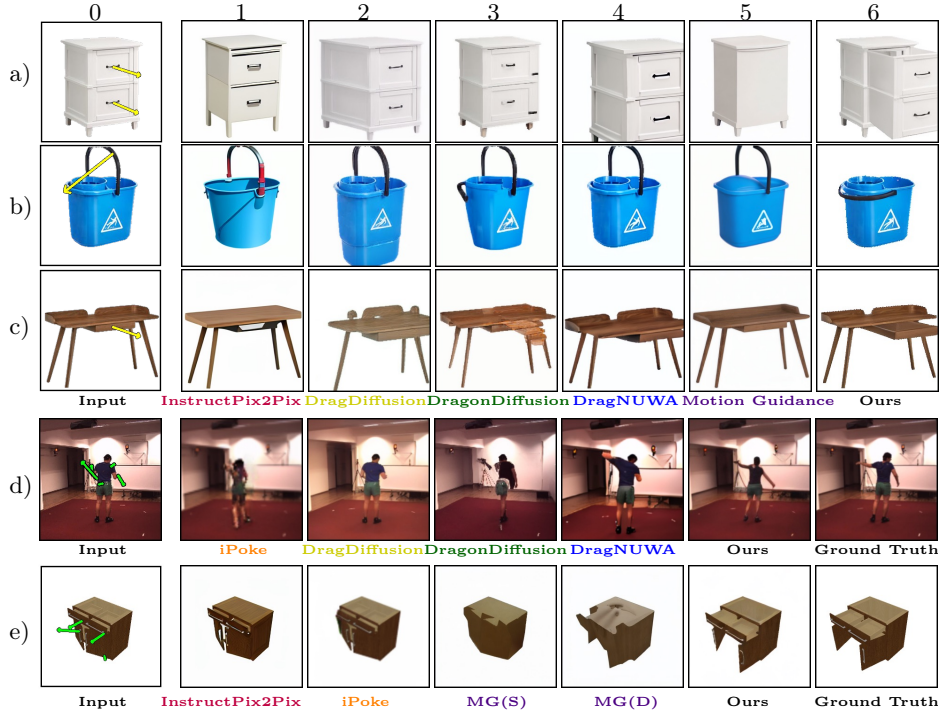
**Table 1: Quantitative Comparisons** on our Drag-a-Move and Human3.6M [24].

checkpoints trained on Internet-scale data. We compare the performance of the two architectures on the I.D. Drag-a-Move in Tab. 2. It should be noted that the conditional DiT lacks pre-training, thereby limiting its generalizability to real-world scenarios. In the rest of this section, unless stated otherwise, our results are obtained with the fine-tuned U-Net LDM checkpoints.

## 5.2 Comparisons with Prior Work

*Baselines.* We perform a thorough comparison of DragAPart to the state-of-the-art generative models which are applicable to this task: iPoke [5], DragDiffusion [58], DragonDiffusion [41], DragNUWA [75], and Motion Guidance [17]. For the training-free methods (DragDiffusion, DragonDiffusion and Motion Guidance), we use their official implementations. Since Motion Guidance is designed for dense motion manipulation, we include two versions in our comparison, one given a dense motion field as input and the other adapted to take only sparse drags. We train iPoke using the publicly available code on our Drag-a-Move dataset and use its released checkpoint trained on Human3.6M [24]. For DragNUWA, we use the released checkpoint. As it is a large-scale open-domain video generator, we do not fine-tune it with our data. iPoke and DragNUWA generate videos, from which we take the last frame for comparison. Additionally, we compare with another naive baseline, InstructPix2Pix [6], which fine-tunes SD [52] on a dataset of paired images and corresponding text edits generated with a large language model [7]. As it is conditioned on textual instructions, we manually craft prompts to convey the manipulation of the drags.

*Quantitative Comparisons* are conducted on the test split of I.D. Drag-a-Move and Human3.6M in Tab. 1. DragAPart significantly outperforms the state-of-the-art methods on all standard metrics, including pixel-level PSNR, patch-level SSIM [67], and feature-level LPIPS [77]. This includes iPoke [5], which has been trained with the identical experimental setting as DragAPart, and DragNUWA, which is a large pre-trained generative model with motion control, but mainly focuses on moving objects as a whole and does not yield motion resulting from part-level interactions. InstructPix2Pix and Motion Guidance are excluded in quantitative comparisons because curating a textual instruction for every image pair in the test set is impractical, and we could not run Motion Guidance on the entire test set as it requires one hour per instance.



**Fig. 4: Qualitative Comparisons** on real images from the ABO [11] dataset with manually defined drags (a-c) and the Human3.6M [24] dataset (d) and a rendered image from our Drag-a-Move test split (e). The images generated by our model are more realistic and capture nuanced part-level dynamics.

*Qualitative Comparisons* are shown in Fig. 4. Despite our best efforts, **InstructPix2Pix** is never able to deform the object significantly, which motivates using dragging for conditioning instead of language. Methods that build on Stable Diffusion [52] latent features (**DragDiffusion**, **DragonDiffusion**, **Motion Guidance**) tend to copy the appearance of the handle points to the target points without moving the entire part (a-d3, e4). For **DragDiffusion** and **DragNUWA**, the object is re-scaled (a2, b2, c4) or repositioned (a4, c2), rather than performing a *part-level* motion. For **Motion Guidance**, even when the dense motion field is given (e4), it fails to preserve the object identity very well (a-b5, e4). By contrast, by using a data-driven approach, our model learns to generate images that preserve faithfully the identity and that exhibit physically plausible motion. Applied to humans (d), our method produces a result that is closest to the ground truth, without changing the identity of the actor, as opposed to **iPoke** (d1).

*Qualitative Results on Real Data.* Fig. 6 shows real images generated by our model, which can (a) preserve fine-grained texture details, (b) generate reasonable shades, (c) handle thin structures, (d) compose multi-part motion, (e) “dream up” the internal structures of the object, and (f) generalize to categories not seen during training. More results are presented in the sup. mat.

Flow Encoding	DiT [45]			SD [52, 53]		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
Conv. Input Only	19.56	0.910	0.095	19.97	0.914	0.077
Conv. Every Block	20.58	0.922	0.078	21.10	<b>0.925</b>	0.067
Multi-Res. Enc. Every Block	<b>21.11</b>	<b>0.925</b>	<b>0.074</b>	<b>21.38</b>	<b>0.925</b>	<b>0.066</b>

**Table 2: Ablations on Architecture and Flow Encoding.** The flow encoding we proposed in Sec. 3.2, when injected into every LDM block at different resolutions, enhances the model’s performance, regardless of the underlying LDM architecture.

Training Data	I.D. Drag-a-Move			O.O.D. Drag-a-Move		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
w/o Domain Random.	21.38	0.925	<b>0.066</b>	18.03	0.897	0.104
w/ Domain Random.	<b>21.82</b>	<b>0.928</b>	<b>0.066</b>	<b>19.74</b>	<b>0.920</b>	<b>0.083</b>

**Table 3: Ablations on Domain Randomization.** Training jointly on the random texture renderings (Sec. 3.3) bridges the gap between in-domain performance and out-of-distribution robustness.




**Fig. 5: Qualitative Ablations.** The visual results are consistent with the numerical ones in Tabs. 2 and 3 and validate our design choices.


### 5.3 Ablation Analysis and Limitations

We present ablation studies on the two key components of DragAPart: the multi-resolution flow encoding (Sec. 3.2) and training with random texture (Sec. 3.3). As shown in Tab. 2, it is not sufficient to inject the drag conditioning through the input layer only (“Input Only” in Tab. 2). As previous works [68, 75] have discovered, injecting this motion conditioning at every LDM block (“Every Block” in Tab. 2) is essential. Our multi-resolution drag encoding better enforces disentanglement among different part motion than using a convolution network (Fig. 5 left). It outperforms the previous methods and the improvement generalizes to both LDM architectures (Tab. 2). Note also the importance of our domain randomization strategy (Sec. 3.3). Without training on the random texture renderings, the model’s performance on out-of-distribution data drops significantly (Tab. 3), and generalizes poorly to unseen categories (Fig. 5 right).

*Limitations.* Currently, we do not explicitly enforce consistency of the generated images of the same object across different viewpoints and drag conditions. We also train separate models for everyday objects and humans. Extending a single model to all moving entities could help us obtain a universal motion prior.



**Fig. 6: Qualitative Results on Real Data.** Trained only on synthetic renderings, DragAPart shows excellent generalizability to both *real images* and *novel categories*.



**Fig. 7: DragAPart provides a useful motion model for applications**, including motion analysis for articulated objects (left) and moving part segmentation (right).

#### 5.4 Application Results

Fig. 7 (left) shows two motion analysis results, where we accurately recover the motion type and parameters specified by a drag. This also indicates that our DragAPart model is relatively consistent across different viewpoints. In Fig. 7 (right) we show that DragAPart’s implicit part-level understanding can be leveraged to produce a coarse segmentation of moving parts. Notably, the segmentation faithfully reflects the input drag. In the sup. mat., we provide a quantitative comparison of the segmentation quality on I.D. Drag-a-Move (with available ground-truth masks) against baselines using CLIP [49], DINO [9, 43] and SD [52].

## 6 Conclusion

We presented DragAPart, an image generator that uses drags as an interface for *part-level* dynamics. By using a small amount of synthetic data, which we contributed, and domain randomization, DragAPart can respond to drags by interpreting them as fine-grained part-level interactions with the underlying object. Partly thanks to a new drag encoder, we have obtained better results than other methods on this task. The model generalizes well to real-world images and novel categories. We have also demonstrated some applications of the motion prior that DragAPart provides.



*Acknowledgments.* We thank Minghao Chen, Junyu Xie and Laurynas Karazija for insightful discussions and our sponsors Toshiba, ERC-CoG UNION 101001212 and EPSRC SYN3D EP/Z001811/1.

## References

1. Abbatematteo, B., Tellex, S., Konidaris, G.D.: Learning to generalize kinematic models to novel objects. In: CoRL (2019)
2. Baranchuk, D., Voynov, A., Rubachev, I., Khrulkov, V., Babenko, A.: Label-efficient semantic segmentation with diffusion models. In: ICLR (2021)
3. Betker, J., Goh, G., Jing, L., Brooks, T., Wang, J., Li, L., Ouyang, L., Zhuang, J., Lee, J., Guo, Y., et al.: Improving image generation with better captions. Tech. rep., OpenAI (2023), <https://cdn.openai.com/papers/dall-e-3.pdf>
4. Blattmann, A., Dockhorn, T., Kulal, S., Mendelevitch, D., Kilian, M., Lorenz, D., Levi, Y., English, Z., Voleti, V., Letts, A., et al.: Stable video diffusion: Scaling latent video diffusion models to large datasets. arXiv preprint arXiv:2311.15127 (2023)
5. Blattmann, A., Milbich, T., Dorkenwald, M., Ommer, B.: iPOKE: Poking a still image for controlled stochastic video synthesis. In: ICCV (2021)
6. Brooks, T., Holynski, A., Efros, A.A.: Instructpix2pix: Learning to follow image editing instructions. In: CVPR (2023)
7. Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J.D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al.: Language models are few-shot learners. In: NeurIPS (2020)
8. Cao, M., Wang, X., Qi, Z., Shan, Y., Qie, X., Zheng, Y.: Masactrl: Tuning-free mutual self-attention control for consistent image synthesis and editing. In: ICCV (2023)
9. Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers. In: ICCV (2021)
10. Chen, T.S., Lin, C.H., Tseng, H.Y., Lin, T.Y., Yang, M.H.: Motion-conditioned diffusion model for controllable video synthesis. arXiv preprint arXiv:2304.14404 (2023)
11. Collins, J., Goel, S., Deng, K., Luthra, A., Xu, L., Gundogdu, E., Zhang, X., Yago Vicente, T.F., Dideriksen, T., Arora, H., Guillaumin, M., Malik, J.: Abo: Dataset and benchmarks for real-world 3d object understanding. In: CVPR (2022)
12. Davtyan, A., Favaro, P.: Learn the force we can: Enabling sparse motion control in multi-object video generation. In: AAAI (2024)
13. Deitke, M., Liu, R., Wallingford, M., Ngo, H., Michel, O., Kusupati, A., Fan, A., Laforte, C., Voleti, V., Gadre, S.Y., Vanderbilt, E., Kembhavi, A., Vondrick, C., Gkioxari, G., Ehsani, K., Schmidt, L., Farhadi, A.: Objaverse-XL: A universe of 10M+ 3D objects. In: NeurIPS (2023)
14. Deitke, M., Schwenk, D., Salvador, J., Weihs, L., Michel, O., Vanderbilt, E., Schmidt, L., Ehsani, K., Kembhavi, A., Farhadi, A.: Objaverse: A universe of annotated 3d objects. In: CVPR (2023)
15. Dhariwal, P., Nichol, A.Q.: Diffusion models beat GANs on image synthesis. In: Proc. NeurIPS (2021)
16. Gadre, S.Y., Ehsani, K., Song, S.: Act the part: Learning interaction strategies for articulated object part discovery. In: ICCV (2021)



17. Geng, D., Owens, A.: Motion guidance: Diffusion-based image editing with differentiable motion estimators. In: ICLR (2024)
18. Geng, H., Xu, H., Zhao, C., Xu, C., Yi, L., Huang, S., Wang, H.: Gapartnet: Cross-category domain-generalizable object perception and manipulation via generalizable and actionable parts. In: CVPR (2023)
19. Girdhar, R., Singh, M., Brown, A., Duval, Q., Azadi, S., Rambhatla, S.S., Shah, A., Yin, X., Parikh, D., Misra, I.: Emu video: Factorizing text-to-video generation by explicit image conditioning. arXiv preprint arXiv:2311.10709 (2023)
20. Guo, Y., Yang, C., Rao, A., Liang, Z., Wang, Y., Qiao, Y., Agrawala, M., Lin, D., Dai, B.: Animatediff: Animate your personalized text-to-image diffusion models without specific tuning. In: ICLR (2024)
21. Ho, J., Chan, W., Saharia, C., Whang, J., Gao, R., Gritsenko, A., Kingma, D.P., Poole, B., Norouzi, M., Fleet, D.J., et al.: Imagen video: High definition video generation with diffusion models. arXiv preprint arXiv:2210.02303 (2022)
22. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. In: NeurIPS (2020)
23. Hu, M., Zheng, J., Liu, D., Zheng, C., Wang, C., Tao, D., Cham, T.J.: Cocktail: Mixing multi-modality control for text-conditional image generation. In: NeurIPS (2023)
24. Ionescu, C., Papava, D., Olaru, V., Sminchisescu, C.: Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. PAMI (2014)
25. Jakab, T., Li, R., Wu, S., Rupprecht, C., Vedaldi, A.: Farm3D: Learning articulated 3d animals by distilling 2d diffusion. In: 3DV (2024)
26. Jiang, Z., Hsu, C.C., Zhu, Y.: Ditto: Building digital twins of articulated objects from interaction. In: CVPR (2022)
27. Karaev, N., Rocco, I., Graham, B., Neverova, N., Vedaldi, A., Rupprecht, C.: Cotracker: It is better to track together. arXiv preprint arXiv:2307.07635 (2023)
28. Karazija, L., Laina, I., Vedaldi, A., Rupprecht, C.: Diffusion models for zero-shot open-vocabulary segmentation. arXiv preprint arXiv:2306.09316 (2023)
29. Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: CVPR (2019)
30. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., Dollar, P., Girshick, R.: Segment anything. In: ICCV (2023)
31. Lei, J., Deng, C., Shen, B., Guibas, L., Daniilidis, K.: Nap: Neural 3d articulation prior. In: NeurIPS (2023)
32. 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. In: ICLR (2024)
33. Li, Z., Tucker, R., Snavely, N., Holynski, A.: Generative image dynamics. arXiv preprint arXiv:2309.07906 (2023)
34. Li, Z., Litvak, D., Li, R., Zhang, Y., Jakab, T., Rupprecht, C., Wu, S., Vedaldi, A., Wu, J.: Learning the 3d fauna of the web. In: CVPR (2024)
35. Liu, J., Mahdavi-Amiri, A., Savva, M.: Paris: Part-level reconstruction and motion analysis for articulated objects. In: ICCV (2023)
36. Liu, J., Tam, H.I.I., Mahdavi-Amiri, A., Savva, M.: CAGE: controllable articulation generation. In: CVPR (2024)
37. Liu, L., Xu, W., Fu, H., Qian, S., Yu, Q., Han, Y., Lu, C.: AKB-48: A real-world articulated object knowledge base. In: CVPR (2022)

38. Liu, R., Wu, R., Hoorick, B.V., Tokmakov, P., Zakharov, S., Vondrick, C.: Zero-1-to-3: Zero-shot one image to 3d object. In: ICCV (2023)
39. Loper, M., Mahmood, N., Romero, J., Pons-Moll, G., Black, M.J.: SMPL: a skinned multi-person linear model. In: ACM TOG (2015)
40. Mo, K., Guibas, L.J., Mukadam, M., Gupta, A.K., Tulsiani, S.: Where2act: From pixels to actions for articulated 3d objects. In: ICCV (2021)
41. Mou, C., Wang, X., Song, J., Shan, Y., Zhang, J.: Dragondiffusion: Enabling drag-style manipulation on diffusion models. In: ICLR (2024)
42. OpenAI: Gpt-4 technical report. arXiv preprint arXiv:2303.08774 (2023)
43. Oquab, M., Darcet, T., Moutakanni, T., Vo, H., Szafraniec, M., Khalidov, V., Fernandez, P., Haziza, D., Massa, F., El-Nouby, A., et al.: Dinov2: Learning robust visual features without supervision. arXiv preprint arXiv:2304.07193 (2023)
44. Pan, X., Tewari, A., Leimkühler, T., Liu, L., Meka, A., Theobalt, C.: Drag your gan: Interactive point-based manipulation on the generative image manifold. In: ACM SIGGRAPH (2023)
45. Peebles, W., Xie, S.: Scalable diffusion models with transformers. In: ICCV (2023)
46. Perez, E., Strub, F., de Vries, H., Dumoulin, V., Courville, A.C.: FiLM: Visual reasoning with a general conditioning layer. In: AAAI (2018)
47. von Platen, P., Patil, S., Lozhkov, A., Cuenca, P., Lambert, N., Rasul, K., Davaadorj, M., Wolf, T.: Diffusers: State-of-the-art diffusion models. <https://github.com/huggingface/diffusers> (2022)
48. Poole, B., Jain, A., Barron, J.T., Mildenhall, B.: DreamFusion: Text-to-3d using 2d diffusion. In: ICLR (2023)
49. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: ICML (2021)
50. Ramesh, A., Pavlov, M., Goh, G., Gray, S., Voss, C., Radford, A., Chen, M., Sutskever, I.: Zero-shot text-to-image generation. In: Proc. ICML (2021)
51. Reizenstein, J., Shapovalov, R., Henzler, P., Sbordon, L., Labatut, P., Novotny, D.: Common objects in 3d: Large-scale learning and evaluation of real-life 3d category reconstruction. In: ICCV (2021)
52. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: CVPR (2022)
53. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: MICCAI (2015)
54. Ruiz, N., Li, Y., Jampani, V., Pritch, Y., Rubinstein, M., Aberman, K.: Dream-booth: Fine tuning text-to-image diffusion models for subject-driven generation. In: CVPR (2023)
55. Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E.L., Ghasemipour, K., Gontijo Lopes, R., Karagol Ayan, B., Salimans, T., et al.: Photorealistic text-to-image diffusion models with deep language understanding. In: NeurIPS (2022)
56. Schuhmann, C., Beaumont, R., Vencu, R., Gordon, C., Wightman, R., Cherti, M., Coombes, T., Katta, A., Mullis, C., Wortsman, M., et al.: Laion-5b: An open large-scale dataset for training next generation image-text models. In: NeurIPS (2022)
57. Shi, Y., Wang, P., Ye, J., Long, M., Li, K., Yang, X.: MVDream: Multi-view diffusion for 3D generation. In: ICLR (2024)
58. Shi, Y., Xue, C., Pan, J., Zhang, W., Tan, V.Y., Bai, S.: Dragdiffusion: Harnessing diffusion models for interactive point-based image editing. In: CVPR (2024)

59. Singer, U., Polyak, A., Hayes, T., Yin, X., An, J., Zhang, S., Hu, Q., Yang, H., Ashual, O., Gafni, O., Parikh, D., Gupta, S., Taigman, Y.: Make-a-video: Text-to-video generation without text-video data. In: ICLR (2022)
60. Song, Y., Ermon, S.: Generative modeling by estimating gradients of the data distribution. In: NeurIPS (2019)
61. Song, Y., Sohl-Dickstein, J., Kingma, D.P., Kumar, A., Ermon, S., Poole, B.: Score-based generative modeling through stochastic differential equations. In: ICLR (2021)
62. Teng, Y., Xie, E., Wu, Y., Han, H., Li, Z., Liu, X.: Drag-a-video: Non-rigid video editing with point-based interaction. arXiv.cs **abs/2312.02936** (2023)
63. Tevet, G., Raab, S., Gordon, B., Shafir, Y., Cohen-or, D., Bermano, A.H.: Human motion diffusion model. In: ICLR (2022)
64. Tian, J., Aggarwal, L., Colaco, A., Kira, Z., Gonzalez-Franco, M.: Diffuse, attend, and segment: Unsupervised zero-shot segmentation using stable diffusion. arXiv preprint arXiv:2308.12469 (2023)
65. Tobin, J., Fong, R., Ray, A., Schneider, J., Zaremba, W., Abbeel, P.: Domain randomization for transferring deep neural networks from simulation to the real world. In: Proc.IROS (2017)
66. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: NIPS (2017)
67. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. IEEE Trans. on Image Processing **13**(4) (2004)
68. Wang, Z., Yuan, Z., Wang, X., Chen, T., Xia, M., Luo, P., Shan, Y.: Motionctrl: A unified and flexible motion controller for video generation. arXiv preprint arXiv:2312.03641 (2023)
69. Watson, D., Chan, W., Brualla, R.M., Ho, J., Tagliasacchi, A., Norouzi, M.: Novel view synthesis with diffusion models. In: ICLR (2023)
70. Weng, H., Yang, T., Wang, J., Li, Y., Zhang, T., Chen, C.L.P., Zhang, L.: Consistent123: Improve consistency for one image to 3D object synthesis. arXiv (2023)
71. Wu, S., Jakab, T., Rupperecht, C., Vedaldi, A.: DOVE: Learning deformable 3d objects by watching videos. IJCV (2023)
72. Wu, S., Li, R., Jakab, T., Rupperecht, C., Vedaldi, A.: Magicpony: Learning articulated 3d animals in the wild. In: CVPR (2023)
73. Xiang, F., Qin, Y., Mo, K., Xia, Y., Zhu, H., Liu, F., Liu, M., Jiang, H., Yuan, Y., Wang, H., Yi, L., Chang, A.X., Guibas, L.J., Su, H.: SAPIEN: A simulated part-based interactive environment. In: CVPR (2020)
74. Xu, J., Liu, S., Vahdat, A., Byeon, W., Wang, X., De Mello, S.: Open-vocabulary panoptic segmentation with text-to-image diffusion models. In: CVPR (2023)
75. Yin, S., Wu, C., Liang, J., Shi, J., Li, H., Ming, G., Duan, N.: Dragnuwa: Fine-grained control in video generation by integrating text, image, and trajectory. arXiv preprint arXiv:2308.08089 (2023)
76. Zhang, L., Rao, A., Agrawala, M.: Adding conditional control to text-to-image diffusion models. In: Proc. CVPR (2023)
77. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: Proc. CVPR. pp. 586–595 (2018)
78. Zheng, C., Vedaldi, A.: Free3d: Consistent novel view synthesis without 3d representation. In: CVPR (2024)
79. Zuffi, S., Kanazawa, A., Jacobs, D.W., Black, M.J.: 3D menagerie: Modeling the 3D shape and pose of animals. In: CVPR (2017)