

# Optimizing Diffusion Noise Can Serve As Universal Motion Priors

Korrawe Karunratanakul<sup>1</sup> Konpat Preechakul<sup>2</sup> Emre Aksan<sup>4</sup>  
 Thabo Beeler<sup>4</sup> Supasorn Suwajanakorn<sup>3</sup> Siyu Tang<sup>1</sup>

<sup>1</sup>ETH Zürich, Switzerland <sup>2</sup>UC Berkeley  
<sup>3</sup>VISTEC, Thailand <sup>4</sup>Google

<https://korrawe.github.io/dno-project/>




Figure 1. Our proposed Diffusion Noise Optimization (DNO) can leverage the existing human motion diffusion models as universal motion priors. We demonstrate its capability in the motion editing tasks where DNO can preserve the content of the original model and accommodates a diverse range of editing modes, including changing trajectory, pose, joint location, and avoiding newly added obstacles.

## Abstract

We propose *Diffusion Noise Optimization (DNO)*, a new method that effectively leverages existing motion diffusion models as motion priors for a wide range of motion-related tasks. Instead of training a task-specific diffusion model for each new task, DNO operates by optimizing the diffusion latent noise of an existing pre-trained text-to-motion model. Given the corresponding latent noise of a human motion, it propagates the gradient from the target criteria defined on the motion space through the whole denoising process to update the diffusion latent noise. As a result, DNO supports any use cases where criteria can be defined as a function of motion. In particular, we show that, for mo-

*tion editing and control, DNO outperforms existing methods in both achieving the objective and preserving the motion content. DNO accommodates a diverse range of editing modes, including changing trajectory, pose, joint locations, or avoiding newly added obstacles. In addition, DNO is effective in motion denoising and completion, producing smooth and realistic motion from noisy and partial inputs. DNO achieves these results at inference time without the need for model retraining, offering great versatility for any defined reward or loss function on the motion representation.*

## 1. Introduction

Many applications of great interest to the motion modeling community can be framed as “finding a plausible motion that fulfills a set of criteria”, typically formulated as minimizing a cost function addressing the given constraints. These include but are not limited to generating motions that follow a trajectory, target locations for joints, keyframes, or avoiding obstacles; denoising, completing missing parts of a motion or editing an existing one. Though they may seem diverse, a unified framework should be able to address these diverse yet relevant tasks, which typically follow a task-agnostic motion prior that ensures plausible motions and task-specific cost functions. Such high-quality motion priors are highly sought after, often with efforts that focus on improving architectures. Yet, a truly versatile framework utilizing the underlying motion prior expressively is still lacking.

Among many human motion priors, the diffusion-based models [14, 32, 40, 45, 65, 79] have become the most prominent ones by achieving impressive performance on certain motion modeling tasks. Avatar Grow Legs [14] and EgoEgo [40] can generate full-body motions from limited joint specifications, such as head and hand poses, but their specific scope precludes their use as general motion priors. PriorMDM [65] introduces a more flexible setting; however, the underlying root-relative motion representation requires dense condition signals for individual joints or keyframes, limiting its use to a small set of applications. OmniControl [79] addresses this limitation by augmenting the motion diffusion model with another network to encode input conditions. While this improves the fidelity and quantity of control signals, it is not always straightforward to support many other constraints. For instance, generating locomotion in a scene with obstacles requires a perception of the physical scene, which is unclear how to cover arbitrary scenes effectively. It is only exacerbated when the scene dynamically changes over time, such as a scene populated by other human agents. GMD [32] presents a more versatile approach by limiting the conditioning to task-based objectives, which theoretically supports arbitrary use cases. Yet there is a trade-off between this flexibility and the fidelity of controls and quality of motions.

To this end, we propose a simple yet effective approach to utilize a motion diffusion model as a motion prior. By treating the denoising process as a black box, we can frame motion-related tasks, such as motion editing and refinement, as optimization problems on the latent manifold of the diffusion model, similar to other classes of generative models such as GANs [49] and VAEs [58] where optimization is seen as iteratively updating the solution on the learned motion manifold (Fig. 2a). In this work, we demonstrate that it is possible and feasible to back-propagate gradients through the full-chain diffusion process, and then optimize the noise

vector based on user-provided criteria in the motion space. This simple approach, dubbed **Diffusion Noise Optimization (DNO)**, which we show to be effective at editing and preserving contents, is a unified and versatile framework that supports a wide variety of applications without the need for fine-tuning the underlying motion model for each specific application. By changing the optimization objective, defined as any differentiable loss function computed on the output motion, DNO enables the diffusion model to effectively serve as a motion prior.

Our experiments show that our unified framework DNO, without any model fine-tuning, produces high-quality motions for motion editing, outperforming existing methods in both preserving the motion content and fulfilling a wide range of objectives including changing trajectory, pose, joint location, and avoiding obstacles. In addition, we demonstrate that it can produce smooth and realistic motion from noisy and partial inputs. Lastly, we provide extensive studies to validate our design choices, which can serve as a basis to effectively extend our framework to other motion-related tasks.

## 2. Related Works

**Motion synthesis, editing, and completion.** The human motion generation task aims to generate motions either conditionally or unconditionally [50, 61, 80, 88, 89]. Various conditioning signals have been explored such as partial poses [15, 20], trajectories [33, 65, 76, 87], images [8, 58], music [37, 38, 41], objects [78], action labels [16, 51], scene [27], or text [1, 18, 19, 34, 52]. Recently, the focus has been on generating motion based on natural-language descriptions using diffusion-based models [6, 72]. These models, utilizing the CLIP model [56], have shown significant improvements in text-to-motion generation [7, 83, 86] and support conditioning on partial motions or music [2, 73]. Alternatively, the motion can be treated as a new language and embedded into the language model framework [29, 84]. Nevertheless, they lack the ability to handle spatial conditioning signals, such as keyframe locations or trajectories. GMD [32] handles this problem with classifier-based guidance at inference time to steer the denoising process toward the target conditions. OmniControl [79] combined GMD with ControlNet [85] to improve realism but limits the conditioning signals to text and partial motion observations. However, motion editing under the motion synthesis framework remains unexplored due to the lack of an explicit mechanism to retain the original motion content.

In parallel, the generation of full-body poses from sparse tracking signals of body joints has gained considerable interest within the community. Previous work such as EgoEgo [40], AGRoL [14], and AvatarPoser [30] while showing impressive results, tend to be specialized models trained ex-

PLICITLY for the motion completion or denoising task. Therefore, it is unclear how we can leverage the motion priors in these trained models for solving other tasks. In this work, we focus on using the trained motion diffusion model to tackle various motion-related tasks under a unified framework including editing, completion, and refinement.

**Diffusion models and guidance.** Diffusion-based probabilistic generative models (DPMs), a class of generative models based on learning to progressively denoising the input data, [24, 66, 69, 70] have gained significant attention across multiple fields of research. They have been successfully applied to tasks such as image generation [13], image super-resolution [39, 62], speech synthesis [36, 36, 54], video generation [25, 26], 3D shape generation [53, 77], and reinforcement learning [28]. The growing interest in the diffusion-based model stems from their superior results and impressive generation controllability, for example, in text-conditioned generation [57, 60, 63] and image-conditioned editing [3, 4, 9, 22, 47]. In terms of conditioning, there are various methods for the diffusion-based models such as imputation and inpainting [9, 10, 47], classifier-based guidance [10, 13], and classifier-free guidance [23, 57, 60, 63]. For refinement, SDEdit [47] enables the repetition of the denoising process to gradually improve output quality but lacks the ability to provide editing guidance. Recently, DOODL [74] demonstrates a direct latent optimization approach for image editing with the help of an invertible ODE [75]. Being inspired by that work, we propose a related method with an improved optimization algorithm that speeds up optimization. DNO can effectively be used in various motion-related tasks making it suitable as a versatile human motion priors. Additionally, we discover that an invertible ODE is not required which makes DNO simpler and convenient to use with minimal effort.

### 3. Background

#### 3.1. Motion generation with diffusion model

A diffusion probabilistic model is a denoising model that learns to invert a diffusion process. A diffusion process is defined as  $q(\mathbf{x}_t|\mathbf{x}_0) = \mathcal{N}(\sqrt{\alpha_t}\mathbf{x}_0, (1 - \alpha_t)\mathbf{I})$  where  $\mathbf{x}_0$  is a clean motion and  $\mathbf{x}_t$  is a noisy motion at the level of  $t$  defined by noise schedule  $\alpha_t$ . With the diffusion process, we can infer an inverse denoising process  $q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0)$ . Now, we can train the diffusion model by predicting  $q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0)$  with a learned  $p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t, c)$ , which is parameterized by a function  $d_\theta(\mathbf{x}_t, c)$  where  $c$  is additional conditions, e.g. text prompts. Conditioning is particularly useful for specifying what kind of motion activity we want from the model.

While diffusion models are stochastic, there exist deterministic sampling processes that share the same marginal distribution. These processes include those defined by prob-

ability flow ODE [70] or by reformulating the diffusion process to be non-Markovian as in DDIM [67]. To obtain a sample based on deterministic sampling, we can solve the associated ODE using an ODE solver starting from  $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ .

#### 3.2. Diffusion model inversion

Diffusion inversion is a process that retrieves the corresponding noise map  $\mathbf{x}_T$  given an input  $\mathbf{x}_0$ . Under deterministic sampling, the associated ODE not only describes the denoising process from  $\mathbf{x}_T$  to  $\mathbf{x}_0$  but also its reverse from  $\mathbf{x}_0$  to  $\mathbf{x}_T$ , which is achievable by solving the ODE backward. However, the fidelity of this inversion relies on the smoothness assumption  $d_\theta(\mathbf{x}_t) \approx d_\theta(\mathbf{x}_{t-1})$ , which is unlikely to hold true when solving the ODEs with just a few discretization steps. For tasks requiring multiple back-and-forth evaluations between  $\mathbf{x}_T$  and  $\mathbf{x}_0$ , an alternative inversion method is available [75].

#### 3.3. Motion representation

The relative-root representation [17] has been widely adopted for text-to-motion diffusion models [72]. The representation is a matrix of human joint features over the motion frames with shape  $D \times M$ , where  $D = 263$  and  $M$  are the representation size and the number of motion frames, respectively. Each motion frame represents root relative rotation and velocity, root height, joint locations, velocities, rotations, and foot contact labels. As the relative representation abstracts away the absolute root location, it can improve the generalization of motion models but makes the controllable generation more challenging [32, 72].


### 4. Diffusion Noise Optimization

A straightforward way to obtain a motion that fulfills a criterion  $\mathcal{L}(\mathbf{x})$  is by optimizing  $\mathbf{x}^* = \arg \min_{\mathbf{x}} \mathcal{L}(\mathbf{x})$  via iterative optimization such as gradient descent. However, optimization in this representation space often yields implausible results as most motion samples  $\mathbf{x} \in \mathbb{R}^{D \times M}$  do not encode plausible motions. This motivates performing optimization on an expressive *latent* space  $\mathbf{z}$  which provides valid motion samples when decoded. The new optimization task becomes

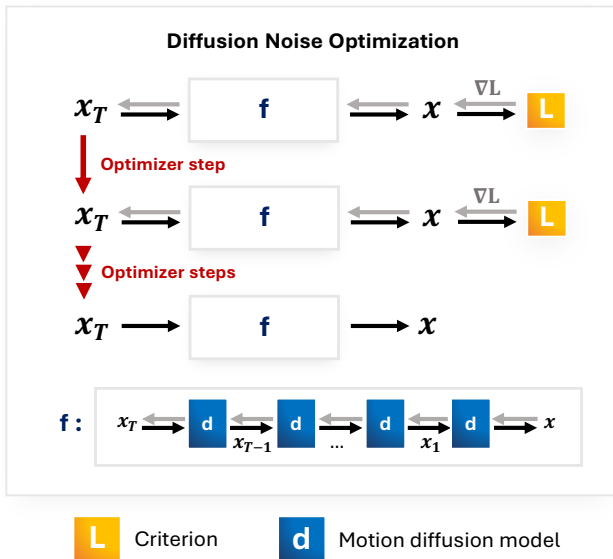
$$\mathbf{z}^* = \arg \min_{\mathbf{z}} \mathcal{L}(f(\mathbf{z})) \quad (1)$$

Such optimization has led to the success in image editing of GANs [21, 31], enabled by its smooth latent space  $\mathbf{z}$  whose mapping is parameterized by a powerful generative model  $f(\mathbf{z})$  and learned from a large dataset. In this work, we show that the latent optimization extends well to  $f$ , parameterized as pretrained diffusion models in the motion domain.

For diffusion models, our choice of the latent variable is the diffusion’s noise at time  $T$ ,  $\mathbf{z} = \mathbf{x}_T$ . While  $\mathbf{x}_T$  is not



(a) At each optimization step, DNO maintains the output motion equality by making a step in the latent space  $\mathbf{x}_T$ , which is decodable to a realistic motion almost everywhere.



(b) DNO’s step direction is obtained from the gradient of a task-specific criterion  $\mathcal{L}$  via backpropagation through an ODE solver,  $f(\mathbf{x}_T) = \text{ODESolver}(d(\cdot), \mathbf{x}_T)$ . At convergence, the optimized  $x_T$  is ultimately decoded via the ODE solver to get the prediction  $x_0$ .

Figure 2. Diffusion Noise Optimization (DNO).

perfectly smooth [55, 67],  $\mathbf{x}_T$  offers the highest-level abstraction readily available. Now, getting the output  $\mathbf{x}$  from a diffusion model  $d(\cdot)$  requires solving an ordinary differential equation (ODE) using an ODE solver [70]. We now see Equation 1 as

$$\mathbf{x}_T^* = \arg \min_{\mathbf{x}_T} \mathcal{L}(\text{ODESolver}(d(\cdot), \mathbf{x}_T)) \quad (2)$$

This allows us to approach many motion tasks simply by adjusting the task-specific criterion  $\mathcal{L}$  while keeping the motion model,  $d$ , intact. In this work, we use the DDIM-ODE [67] and its solver.

This optimization is iteratively solved using gradient descent. Starting from a certain noise  $\mathbf{x}_T$ , we solve the ODE, arrive at a prediction  $\mathbf{x}$ , and evaluate the criterion function  $\mathcal{L}(\mathbf{x})$ . Then, we obtain the gradient  $\nabla_{\mathbf{x}_T} \mathcal{L}(\mathbf{x})$  by backpropagating through the ODE solver. An optimizer updates

$\mathbf{x}_T$  based on the gradient, possibly with a small random perturbation to encourage exploration [31, 74]. We repeat this until convergence. The final output is the motion obtained by solving the ODE one last time with the optimized  $\mathbf{x}_T$ . We call the above algorithm Diffusion Noise Optimization (DNO) and summarize it in Algo. 1.

Maintaining the intermediate activations for solving the ODE during backpropagation can be memory-intensive. This issue can be addressed with gradient checkpointing [11] or an invertible ODE [74, 75], at the cost of more computation or model complexity. In the motion domain, impressive results can be achieved from an ODE solver with a minimal number of function evaluations and feasible memory overhead. Given the ongoing efforts to reduce diffusion sampling steps, as evidenced in distillation [42, 64, 71] and high-order solvers studies [43, 44, 90], our simple design, despite requiring back-propagation through all steps, can become increasingly relevant and applicable.

**Optimization through ODE solver.** Empirically, gradients via backpropagation through the ODE solver have norms spanning multiple orders of magnitude, making the optimization unstable and slow. We instead propose to normalize the gradient to have a unit norm, which optimizes faster in practice. Normalized gradient methods are also found to help escape saddle-point in the loss landscape faster [12, 48]. When combined with momentum, the update becomes a moving average of directions, which is robust to gradient norm outliers.

### Algorithm 1 Diffusion Noise Optimization

**Require:**  $\mathbf{x}_T$ , motion model  $d$ , ODESolver, Optimizer, criterion  $\mathcal{L}$ , learning rate  $\eta$ , perturbation  $\gamma$  (default 0)

- 1: **while** not converged **do**
- 2:    $\mathbf{x} \leftarrow \text{ODESolver}(d(\cdot), \mathbf{x}_T)$
- 3:    $\nabla \leftarrow \nabla_{\mathbf{x}_T} \mathcal{L}(\mathbf{x})$  ▷ Task-specific
- 4:    $\mathbf{x}_T \leftarrow \text{Optimizer}(\mathbf{x}_T, \nabla / \|\nabla\|, \eta) + \gamma \mathcal{N}(\mathbf{0}, \mathbf{1})$
- 5: **end while**
- 6: **return** ODESolver( $\mathbf{x}_T, f$ )

**DNO fundamentally differs from a guided motion diffusion method [32]** in how the criterion  $\mathcal{L}(\cdot)$  is computed at each iteration. For guided diffusion methods, the criterion is computed on an expected  $\hat{x} = \mathbb{E}[\mathbf{x}_0 | \mathbf{x}_t]$  at a denoising step  $t$  [68, 82] that is  $\mathcal{L}(\mathbf{x}_t) = \mathbb{E}_{p(\mathbf{x}_0 | \mathbf{x}_t)} \mathcal{L}(\mathbf{x}_0) \approx \mathcal{L}(\hat{x})$ , resulting in severe error when  $\text{Var}[\mathbf{x}_0 | \mathbf{x}_t]$  is large. For DNO, the criterion is exactly computed on  $\mathbf{x}$  after the full-chain denoising which eliminates the approximation error. We further discuss these differences in the supp. mat.

## 5. Applications

We demonstrate the versatility of DNO on a wide range of conditional motion synthesis tasks. These tasks can be

solved by designing a task-specific criterion  $\mathcal{L}$ , along with any additional auxiliary loss functions. The method for initializing the noise  $\mathbf{x}_T$  can vary depending on the task. In the following sections,  $\mathbf{x}$  refers to an output from an ODE solver which is a function of  $\mathbf{x}_T$ .

### 5.1. Motion editing and control

The goal of motion editing tasks is to modify a given reference motion  $\mathbf{x}_{\text{ref}}$  to satisfy certain objectives. These objectives may include following a specific trajectory, conforming to specified poses in all or some keyframes, and avoiding static or moving obstacles, all while preserving the key characteristics of the input motion.

**Editing to follow trajectories and poses.** Editing a motion to follow a specific trajectory or to match specific poses can be easily achieved by minimizing the average distance between each generated joint location and its corresponding target location. Target locations can be specified for any subset of joints in any subset of motion frames.

Specifically, let  $\mathbf{c}_j^k \in \mathbb{R}^3$  be the target location for joint  $j$  at keyframe  $k$ , and  $\hat{\mathbf{c}}_j^k$  be its generated location from the current motion  $\mathbf{x}$ . The loss function can be defined as:

$$\mathcal{L}_{\text{pose}}(\mathbf{x}, O) := \frac{1}{|O|} \sum_{(j,k) \in O} \|\hat{\mathbf{c}}_j^k(\mathbf{x}) - \mathbf{c}_j^k\|_1, \quad (3)$$

where  $O$  is the *observed* set containing  $(j, k)$  pairs denoting the target joints and keyframes determined by the task.

**Editing to avoid obstacles.** Obstacles can be represented as a signed distance function (SDF), whose gradient field defines the repelling direction away from the obstacles. The loss function can incorporate a safe distance threshold  $\tau$ , beyond which the gradient becomes zero such that

$$\mathcal{L}_{\text{obs}}(\mathbf{x}) := \sum_{j,k} -\min \left[ \text{SDF}^k(\hat{\mathbf{c}}_j^k(\mathbf{x})), \tau \right], \quad (4)$$

where  $\text{SDF}^k$  is the signed distance function for obstacles in frame  $k$ , which may vary across frames in the case of moving obstacles.

**Preserving original characteristics.** When modifying, for example, a jumping motion to align with a certain trajectory, it is crucial to retain key aspects of the jump like its rhythm, height, and overall body coordination. This can be achieved using two means. First, we invert the reference motion  $\mathbf{x}_{\text{ref}}$  using diffusion inversion to obtain the corresponding noise sample  $\mathbf{x}_{T_{\text{ref}}} = \text{ODESolver}^{-1}(d(\cdot), \mathbf{x}_{\text{ref}})$  and use  $\mathbf{x}_{T_{\text{ref}}}$  as the initial value for  $\mathbf{x}_T$  we are optimizing. Second, we penalize the distance between  $\mathbf{x}_{T_{\text{ref}}}$  and  $\mathbf{x}_T$  to ensure that it remains close to the reference motion during optimization:

$$\mathcal{L}_{\text{cont}}(\mathbf{x}_T) := \|\mathbf{x}_{T_{\text{ref}}} - \mathbf{x}_T\|_2. \quad (5)$$

Many motion editing tasks can be solved using a combination of these loss functions with a set of balancing weights

$$\mathcal{L}(\cdot) = \mathcal{L}_{\text{pose}}(\mathbf{x}) + \lambda_{\text{obs}} \mathcal{L}_{\text{obs}}(\mathbf{x}, O) + \lambda_{\text{cont}} \mathcal{L}_{\text{cont}}(\mathbf{x}_T) \quad (6)$$

### 5.2. Motion refinement and completion

The tasks in this category is to reconstruct a motion from noisy and/or incomplete input. This class of tasks includes completing a motion with missing frames or joints, seamlessly blending motions together, denoising a noisy motion, or any combination of these tasks.

**Motion refinement.** Given a noisy input motion, motion refinement seeks to enhance the input so that it becomes more realistic. We can solve this problem simply by starting the optimization from a random  $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ , setting the observed set  $O$  to include all joints from the input noisy motion, and optimizing  $\mathcal{L}_{\text{pose}}(\mathbf{x}, O)$ . While using  $\mathcal{L}_{\text{pose}}$  here may seem counterintuitive, as it seeks to match the predicted motion to the original noisy input, DNO is able to generate a plausible motion by eliminating the noise components from the motion.

Since this optimization begins with a random noise  $\mathbf{x}_T$ , the initial prediction can be far from the desired predicted motion, requiring significant changes on  $\mathbf{x}_T$  with several optimization steps, which tend to correlate neighboring noises and reduce the representation capacity. We empirically find that regularizing the noise to decorrelate each latent dimension alleviates the foot skating problem. Inspired by StyleGAN2 [31], we introduce a latent decorrelation loss across motion frames:

$$\mathcal{L}_{\text{decorr}}^m(\mathbf{x}_T) = \frac{1}{mD} \sum_{i=1}^m \mathbf{x}_{T,m}(i)^\top \mathbf{x}_{T,m}(i+1). \quad (7)$$

We apply this loss at several scales of temporal resolutions  $m \in \{M, M/2, M/4, \dots, 2\}$ . Specifically, starting with the original length  $M$ , we downsample the sequence’s temporal resolution by half via average pooling two consecutive frames successively.

This loss, summed over resolutions, encourages more plausible motions and can be used together with  $\mathcal{L}_{\text{pose}}$  for motion refinement:

$$\mathcal{L}(\cdot) = \mathcal{L}_{\text{pose}}(\mathbf{x}, O) + \lambda_{\text{decorr}} \sum_m \mathcal{L}_{\text{decorr}}^m(\mathbf{x}_T) \quad (8)$$

**Motion completion.** Unlike motion refinement, where its input contains complete joint information for all frames, this task involves taking an incomplete motion, that may be noisy, as input and seek to fill in the missing information. We begin the optimization with  $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$  and apply the same loss in Equation 8, but with  $O$  containing the existing joints in the input motion for  $\mathcal{L}_{\text{pose}}(\mathbf{x}, O)$  term.

## 6. Experiments

**Datasets.** When applicable, we evaluate generated motions on the HumanML3D [18] dataset, which contains 44,970 motion annotations of 14,646 motion sequences from AMASS [46] and HumanAct12 [16] datasets.

**Motion diffusion model.** As DNO is plug-and-play and model-agnostic, it works with any trained motion diffusion model. When paired with MDM [72], which is trained on the HumanML3D dataset, we name the combination DNO-MDM. We retrained MDM with Exponential Moving Averaging [24] with no further modifications to help stabilize the model, which was previously found to produce inconsistent output between checkpoints.

**DNO implementation details.** We use Adam optimizer [35] with a learning rate of 0.05, a linear warm-up for the first 50 steps, and a cosine learning rate schedule to zero for the entire optimization. We use a unit-normalized gradient. We set the coefficients  $\lambda_{\text{cont}} = 0.01$ ,  $\lambda_{\text{decorr}} = 10^3$ ,  $\lambda_{\text{obs}} = 1.0$ , and the perturbation amount  $\gamma = 0$ . This design choice is explored in Section 6.3.

For editing tasks, we obtain the initial  $\mathbf{x}_T$  from DDIM-100 inversion on MDM [72] without text conditions. The subsequent optimization is run for 300 steps, which takes approximately 3 minutes on an Nvidia 3090 GPU. For refinement tasks, the optimization is run for 500 steps. DNO optimization with DDIM-10 and a batch size of 16 requires 18 GB of GPU memory. Optimization in all tasks is also done without text conditions to the diffusion model. Gradient checkpointing [5, 11] can be used with DNO to reduce memory usage further for a low-memory GPU at the cost of more computation.

### 6.1. Motion Editing

Given an input motion, we generate 16 edited motions using DNO, each with a single randomly chosen keyframe between frames 60 and 90 (i.e., 2-3 seconds after the start). The objective is to change the pelvis position to a random target location within an area ranging from -2m to 2m with respect to the initial position. We conducted this experiment on 6 input motions, resulting in a total of 96 edited motions. To generate the input motions, we use the prompt “a person is [action]” with a predetermined [action]. During editing, DNO does not leverage the text prompt or the action class.

**Evaluation metrics.** We focus on three main aspects: (1) is the output motion realistic?, (2) can it fulfill the editing objective?, and (3) how much are the original motion characteristics preserved after editing?

To evaluate the realism of the output, we measure *Jitter* and *Foot skating ratio*, following prior work [18, 32, 81]. *Jitter* is a proxy for motion smoothness, measuring the mean changes in acceleration of all joints over time in

Table 1. Motion editing evaluation on specific actions generated from MDM given the text prompts. We focus on actions that can be distinctly classified per frame basis. The Content Preservation scores are computed against the inputs.

Action	Content ↑ Preserve	Objective ↓ Error (m)	Foot ↓ skating ratio	Jitter ↓
“jumping”				
Input	1.00	1.62	0.01	1.42
GMD [32]	0.64	0.22	0.12	3.07
DNO-MDM Edit	<b>0.95</b>	<b>0.00</b>	<b>0.05</b>	<b>1.31</b>
“doing a long jump”				
Input	1.00	3.03	0.01	1.20
GMD	0.59	0.22	0.15	4.10
DNO-MDM Edit	<b>0.92</b>	<b>0.00</b>	<b>0.07</b>	<b>1.34</b>
“walking with raised hand”				
Input	1.00	2.76	0.01	0.24
GMD	0.65	0.16	0.11	0.51
DNO-MDM Edit	<b>0.92</b>	<b>0.00</b>	<b>0.05</b>	<b>0.32</b>
“crawling”				
Input	1.00	1.83	0.06	0.47
GMD	0.79	0.15	<b>0.04</b>	2.76
DNO-MDM Edit	<b>0.94</b>	<b>0.00</b>	0.08	<b>0.48</b>

$10^2 m/s^3$ . *Foot skating ratio* is a proxy for the incoherence between the trajectory and human motion, which measures the proportion of frames in which a foot skates for more than a certain distance (2.5 cm) while maintaining the contact with the ground (foot height < 5 cm).

To evaluate the success in the editing task, we use *Objective error*, which measures the distance between the target location of the edited joint at a specific frame and the corresponding joint position in the output motion. When the editing is performed on the ground plane, this metric measures the 2D distance.

While metrics leveraging action classifiers can provide a coarse evaluation of whether the action class is preserved or not after the editing, we require a more precise approach to quantify the content similarity. We report the ratio of frames in which both motions, before editing and after, perform the same action throughout the entire sequence, which we call *Content preservation ratio*. The criteria to determine if two actions are the same are action-dependent. The ratio will be 1 if both motions perform the same action in all frames. Since there are no established methods for per-frame content similarity, in this work, we consider only distinct actions for which the boundaries of the actions can be clearly defined: For “jumping”, the action is defined as having both feet more than 5cm above the ground. For “raised hands”, both hands need to be above the head. For the “crawling” motion, the head needs to be below 150 cm.

**Results.** As shown in Table 1, our proposed DNO outperforms GMD [32], the state of the art in spatially-conditioned motion generation, in all actions and metrics, except for *Foot skating ratio* on “crawling.” Overall, our method




Figure 3. Qualitative results from motion editing task. Each line indicates the starting and target location of the selected joint at a specific keyframe.

achieves significantly lower *Jitter*. Furthermore, our higher *Content preservation ratio* suggests that our approach retains the original content if editing is not necessary while successfully achieving the editing objectives, as indicated by the low *Objective errors*.

We present qualitative results of the edited motions across various tasks in the supplementary video. To follow a desired trajectory, we can use a set of target locations for the pelvis. Similarly, our DNO can edit motions to avoid obstacles without requiring explicit target keyframes. Furthermore, our method enables fine-grained editing by targeting individual joints such as the hand joint or an entire pose at a specific keyframe. Please refer to the supp. mat. for additional qualitative results.

## 6.2. Motion Refinement

We evaluate motion refinement performance on a random subset of size 300 from the test set of the HumanML3D dataset in two setups. First, we assess if DNO can recover a clean and plausible motion from a noisy input motion by adding Gaussian noise with a standard deviation of 5 cm to all axes of all joints. We then evaluate the motion refinement performance for an input motion with incomplete and noisy joint information in three scenarios: (1) *six joints*: head + two hands + two legs + pelvis, (2) *eight joints*: six joints + two shoulders, (3) *ten joints*: eight joints + two knees.

**Evaluation metrics.** As in the motion editing tasks, we report *Foot skating ratio* and *Jitter* as well as the established *Mean per-joint pose error* (MPJPE) between the ground-truth and predicted motions, and *FID* which measures the distance between the ground-truth and synthetic data distributions using a pretrained motion encoder [16].

**Results.** As shown in Table 2, we observe that, in every experiment, DNO successfully improves the signal by reducing MPJPE beyond the input level and producing smooth and realistic motions. We demonstrate that DNO’s performance scales with the base model by pairing it with

Table 2. Noisy motion refinement results (noise std. 5 cm) on a subset of HumanML3D [18] dataset. All experiments were run with  $N = 300$ . FIDs are computed against *Real* except *Real*’s FIDs which are computed against a holdout set from the dataset. HuMoR\* means we exclude the sequence when its optimization fails. DNO-MLD\* runs with 1,000 optimization steps.

	MPJPE ↓ observed (cm)	FID ↓	Foot ↓ skating ratio	Jitter ↓
Real	0.0	0.50	0.08	0.50
<b>All joints</b>				
Noisy	11.1	58.76	0.66	28.60
HuMoR* [58]	7.2	0.87	0.13	<b>0.33</b>
GMD	25.7	6.91	0.08	0.81
DNO-MDM	9.1	0.69	<b>0.07</b>	0.36
DNO-MLD*	10.41	0.27	0.11	1.37
DNO-GMD	<b>7.0</b>	<b>0.10</b>	0.08	0.89
<b>Six joints</b>				
Noisy	11.6	58.73	0.66	28.61
HuMoR*	8.8	1.40	0.10	<b>0.20</b>
GMD	31.2	7.07	0.08	0.80
DNO-MDM	8.8	1.15	0.07	0.38
DNO-MLD*	11.4	0.56	0.10	1.35
DNO-GMD	<b>7.5</b>	<b>0.36</b>	0.07	0.97
<b>Eight joints</b>				
Noisy	11.4	58.73	0.66	28.61
HuMoR*	8.5	1.19	0.14	<b>0.20</b>
GMD	29.9	6.99	0.08	0.80
DNO-MDM	9.3	0.85	<b>0.07</b>	0.36
DNO-MLD*	11.6	0.42	0.11	1.36
DNO-GMD	<b>7.2</b>	<b>0.13</b>	0.08	0.96
<b>Ten joints</b>				
Noisy	11.3	58.73	0.66	28.61
HuMoR*	9.3	1.06	0.13	<b>0.21</b>
GMD	28.4	6.88	0.09	0.80
DNO-MDM	9.0	0.66	0.08	0.39
DNO-MLD*	11.0	0.43	0.11	1.35
DNO-GMD	<b>7.1</b>	<b>0.12</b>	<b>0.07</b>	0.96

three motion diffusion models (ordered by FID from high-to-low on motion generation): MDM [72], MLD [7], and GMD [32]. The DNO-GMD combination outperforms the weaker combinations, DNO-MLD and DNO-MDM, and the optimization-based motion prior HuMoR [59]. And any DNO combination outperforms the guided diffusion method, GMD, in the motion refinement task. For the GMD baseline, we use the spatially-conditioned generation as in the original paper. While it also successfully denoises the motion (low Jitter 0.8), it struggles to satisfy fine pose signals and can only follow trajectory guidance leading to high MPJPE > 25 cm. We also tried SDEdit [47] on motion refinement, however, it cannot deal with such a level of noise while being able to retain the original motion. Additional details and tasks such as motion completion, blending, and in-betweening, are discussed in the supplementary. Note that, unlike the task-specific methods such as [14, 30], our method is never trained specifically for these tasks.

Table 3. Ablation study on the noisy motion refinement task (std. 1 cm) results on a subset of HumanML3D [18]. All experiments were run with  $N = 300$ . FIDs are computed against *Real* except *Real*'s FIDs which are computed against a holdout set from the dataset.

	MPJPE ↓ all (cm)	FID ↓	Jitter ↓	Foot ↓ skating ratio
Real	0.0	0.50	0.50	0.08
Noisy	6.4	9.91	5.87	0.15
DNO	8.7	0.66	0.33	0.07
– Normalized grad.	30.2	4.13	0.34	0.06
– Norm grad. & $\mathcal{L}_{\text{decorr}}$	25.5	4.02	0.34	0.06
– $\mathcal{L}_{\text{decorr}}$	6.8	0.65	0.40	0.10
– LR scheduler & warmup	8.4	0.58	0.37	0.07
Perturb $\gamma = 0$ (DNO)	8.7	0.66	0.33	0.07
Perturb $\gamma = 2 \times 10^{-4}$	8.5	0.75	0.34	0.07
Perturb $\gamma = 5 \times 10^{-4}$	8.4	0.73	0.36	0.07
Perturb $\gamma = 10^{-3}$	9.1	0.91	0.34	0.07
Optimize for 300 steps	12.1	1.48	0.31	0.07
Optimize for 500 steps (DNO)	8.7	0.66	0.33	0.07
Optimize for 700 steps	7.2	0.54	0.36	0.07
DDIM 5 steps	9.8	0.90	0.36	0.09
DDIM 10 steps (DNO)	8.7	0.66	0.33	0.07
DDIM 20 steps	7.9	0.92	0.34	0.07

### 6.3. Ablation studies

To motivate and justify DNO’s design choices, we conducted experiments on a surrogate motion refinement task, where we add noise with a standard deviation of 1 cm to an input motion and seek to recover the original motion.

**Normalized gradients.** As  $\text{ODESolver}(\cdot)$  involves iterative calls to the diffusion model, the gradients can become highly unstable. In our experiments, we observe that the gradient norms can span multiple orders of magnitude, leading to slow or poor convergence. Table 3 demonstrates that our choice of using normalized gradients has improved the solutions by reducing MPJPE (which corresponds to the lower  $\mathcal{L}_{\text{pose}}$  value) from 30.2 to 8.7 cm.

**Decorrelating noise.** DNO’s main motivation is to identify a latent sample that is capable of generating a plausible motion while fulfilling the task-driven constraints. However, in practice, not all latent samples yield a plausible motion sample. One of the factors accounting for this is the uncorrelated random noise samples that the diffusion models are trained with. Correlated latent samples often results in poor motion samples. For example, the latent sample with all zeros has the highest likelihood, yet often generates low-quality motions. Our  $\mathcal{L}_{\text{decorr}}$ , motivated by this observation, discourages correlations across the time axis in  $\mathbf{x}_T$ . Table 3 demonstrates that this loss directly contributes to the motion quality. It significantly improves the *Foot skating ratio* from 0.10 to 0.07 and the *Jitter* from 0.40 to 0.33, though at the cost of an increased optimization difficulty as indicated

by the rise in MPJPE from 6.8 to 8.7. However, while it is rather straightforward to improve the MPJPE by increasing optimization steps, it is more challenging to fix the artifacts in the generated motion.

**Random perturbation.** As suggested in multiple studies [31, 74], random perturbation is theoretically motivated as a facilitator of exploration, which helps optimization escape from local minima. To ensure that  $\gamma$  converges to zero, we tie  $\gamma$  with the learning rate scheduler, which has a warm-up period and a cosine decay. We found that small  $\gamma$ ’s may decrease the optimization errors only marginally, while a larger  $\gamma = 10^{-3}$  has a negative impact on the optimization, yet they all produced motions with worse FIDs. While the results are not exhaustive, they serve as a piece of evidence to support our choice of  $\gamma = 0$ .

**ODE solver steps.** We experimented with varying the number of DDIM sampling steps from 5, 10, 20, which produce MPJPE scores of 9.8, 8.7, and 7.9, respectively. This suggests that higher DDIM steps may be able to capture finer motion detail. We chose 10-step sampling as a balanced choice between result quality and resource usage.

## 7. Discussion and Limitations

In this work, we proposed DNO a simple and versatile method to use a pretrained motion diffusion model as universal motion priors. We show that DNO can leverage the motion priors to achieve precise and fine-grained control for motion editing. Apart from editing, we demonstrate that our formulation can be extended to a wide range of motion tasks. Our method is plug-and-play and does not require training a new model for every new task. DNO is, however, not a perfect motion priors. It works better when the observation has good coverage of the human body. While the decorrelation loss does help, there are situations where DNO cannot easily project to a realistic motion. Ultimately, the effectiveness of our method is limited by the performance of the underlying diffusion model; nevertheless, we expect the performance of the base model to increase as the community is searching for better model designs and collecting more training data. As the inference speed of the diffusion model keeps increasing, we hope to overcome the speed limitation of our optimization framework to achieve interactive motion editing. In summary, our extensive studies on diffusion noise optimization effectiveness can serve as a useful basis for leveraging existing diffusion models to solve a wider range of tasks.

**Acknowledgement.** This work was supported by the SNSF project grant 200021 204840. Konpat Preechakul is funded by DARPA MCS. We sincerely thank Bram Wallace for an insightful discussion on EDICT and DOODL which inspired this project.



## References

- [1] Chaitanya Ahuja and Louis-Philippe Morency. Language2pose: Natural language grounded pose forecasting. In *2019 International Conference on 3D Vision (3DV)*, pages 719–728. IEEE, 2019. 2
- [2] Simon Alexanderson, Rajmund Nagy, Jonas Beskow, and Gustav Eje Henter. Listen, denoise, action! audio-driven motion synthesis with diffusion models. *arXiv preprint arXiv:2211.09707*, 2022. 2
- [3] Yogesh Balaji, Seungjun Nah, Xun Huang, Arash Vahdat, Jiaming Song, Karsten Kreis, Miika Aittala, Timo Aila, Samuli Laine, Bryan Catanzaro, Tero Karras, and Ming-Yu Liu. eDiff-I: Text-to-Image diffusion models with an ensemble of expert denoisers. 2022. 3
- [4] T Brooks, A Holynski, and A A Efros. InstructPix2Pix: Learning to follow image editing instructions. *arXiv preprint arXiv:2211.09800*, 2022. 3
- [5] Tianqi Chen, Bing Xu, Chiyuan Zhang, and Carlos Guestrin. Training deep nets with sublinear memory cost. 2016. 6
- [6] Xin Chen, Biao Jiang, Wen Liu, Zilong Huang, Bin Fu, Tao Chen, Jingyi Yu, and Gang Yu. Executing your commands via motion diffusion in latent space. 2022. 2
- [7] Xin Chen, Biao Jiang, Wen Liu, Zilong Huang, Bin Fu, Tao Chen, Jingyi Yu, and Gang Yu. Executing your commands via motion diffusion in latent space. *arXiv preprint arXiv:2212.04048*, 2022. 2, 7
- [8] Xin Chen, Zhuo Su, Lingbo Yang, Pei Cheng, Lan Xu, Bin Fu, and Gang Yu. Learning variational motion prior for video-based motion capture. *arXiv preprint arXiv:2210.15134*, 2022. 2
- [9] Jooyoung Choi, Sungwon Kim, Yonghyun Jeong, Youngjune Gwon, and Sungroh Yoon. ILVR: Conditioning method for denoising diffusion probabilistic models. 2021. 3
- [10] Hyungjin Chung, Byeongsu Sim, Dohoon Ryu, and Jong Chul Ye. Improving diffusion models for inverse problems using manifold constraints. 2022. 3
- [11] Kevin Clark, Paul Vicol, Kevin Swersky, and David J Fleet. Directly Fine-Tuning diffusion models on differentiable rewards. 2023. 4, 6
- [12] Jorge Cortés. Finite-time convergent gradient flows with applications to network consensus. *Automatica*, 42(11):1993–2000, 2006. 4
- [13] Prafulla Dhariwal and Alex Nichol. Diffusion models beat GANs on image synthesis. 2021. 3
- [14] Yuming Du, Robin Kips, Albert Pumarola, Sebastian Starke, Ali Thabet, and Artsiom Sanakoyeu. Avatars grow legs: Generating smooth human motion from sparse tracking inputs with diffusion model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 481–490, 2023. 2, 7
- [15] Yinglin Duan, Tianyang Shi, Zhengxia Zou, Yenan Lin, Zhehui Qian, Bohan Zhang, and Yi Yuan. Single-shot motion completion with transformer. *arXiv preprint arXiv:2103.00776*, 2021. 2
- [16] Chuan Guo, Xinxin Zuo, Sen Wang, Shihao Zou, Qingyao Sun, Annan Deng, Minglun Gong, and Li Cheng. Action2motion: Conditioned generation of 3d human motions. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 2021–2029, 2020. 2, 6, 7
- [17] Chuan Guo, Shihao Zou, Xinxin Zuo, Sen Wang, Wei Ji, Xingyu Li, and Li Cheng. Generating diverse and natural 3D human motions from text. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5152–5161. IEEE, 2022. 3
- [18] Chuan Guo, Shihao Zou, Xinxin Zuo, Sen Wang, Wei Ji, Xingyu Li, and Li Cheng. Generating diverse and natural 3d human motions from text. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5152–5161, 2022. 2, 6, 7, 8
- [19] Chuan Guo, Xinxin Zuo, Sen Wang, and Li Cheng. Tm2t: Stochastic and tokenized modeling for the reciprocal generation of 3d human motions and texts. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXV*, pages 580–597. Springer, 2022. 2
- [20] Félix G Harvey, Mike Yurick, Derek Nowrouzezahrai, and Christopher Pal. Robust motion in-betweening. *ACM Transactions on Graphics (TOG)*, 39(4):60–1, 2020. 2
- [21] Zhenliang He, Wangmeng Zuo, Meina Kan, Shiguang Shan, and Xilin Chen. AttGAN: Facial attribute editing by only changing what you want. *IEEE Trans. Image Process.*, 28(11):5464–5478, 2019. 3
- [22] Amir Hertz, Ron Mokady, Jay Tenenbaum, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Prompt-to-Prompt image editing with cross attention control. 2022. 3
- [23] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. In *NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications*. 3
- [24] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. 2020. 3, 6
- [25] Jonathan Ho, William Chan, Chitwan Saharia, Jay Whang, Ruiqi Gao, Alexey Gritsenko, Diederik P Kingma, Ben Poole, Mohammad Norouzi, David J Fleet, and Tim Salimans. Imagen video: High definition video generation with diffusion models. 2022. 3
- [26] Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J Fleet. Video diffusion models. 2022. 3
- [27] Siyuan Huang, Zan Wang, Puhao Li, Baoxiong Jia, Tengyu Liu, Yixin Zhu, Wei Liang, and Song-Chun Zhu. Diffusion-based generation, optimization, and planning in 3d scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16750–16761, 2023. 2
- [28] Michael Janner, Yilun Du, Joshua B Tenenbaum, and Sergey Levine. Planning with diffusion for flexible behavior synthesis. 2022. 3
- [29] Biao Jiang, Xin Chen, Wen Liu, Jingyi Yu, Gang Yu, and Tao Chen. Motiongpt: Human motion as a foreign language. *arXiv preprint arXiv:2306.14795*, 2023. 2
- [30] Jiayi Jiang, Paul Strel, Huajian Qiu, Andreas Fender, Larissa Laich, Patrick Snape, and Christian Holz. Avatarposer: Articulated full-body pose tracking from sparse motion sensing. In *European Conference on Computer Vision*, pages 443–460. Springer, 2022. 2, 7

- [31] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of StyleGAN. 2019. [3](#), [4](#), [5](#), [8](#)
- [32] Korrawe Karunratanakul, Konpat Preechakul, Supasorn Suwajanakorn, and Siyu Tang. Guided motion diffusion for controllable human motion synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2151–2162, 2023. [2](#), [3](#), [4](#), [6](#), [7](#)
- [33] Manuel Kaufmann, Emre Aksan, Jie Song, Fabrizio Pece, Remo Ziegler, and Otmar Hilliges. Convolutional autoencoders for human motion infilling. In *2020 International Conference on 3D Vision (3DV)*, pages 918–927. IEEE, 2020. [2](#)
- [34] Jihoon Kim, Jiseob Kim, and Sungjoon Choi. Flame: Free-form language-based motion synthesis & editing. *arXiv preprint arXiv:2209.00349*, 2022. [2](#)
- [35] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. [6](#)
- [36] Zhifeng Kong, Wei Ping, Jiaji Huang, Kexin Zhao, and Bryan Catanzaro. DiffWave: A versatile diffusion model for audio synthesis. 2020. [3](#)
- [37] Hsin-Ying Lee, Xiaodong Yang, Ming-Yu Liu, Ting-Chun Wang, Yu-Ding Lu, Ming-Hsuan Yang, and Jan Kautz. Dancing to music. *Advances in neural information processing systems*, 32, 2019. [2](#)
- [38] Buyu Li, Yongchi Zhao, Shi Zhelun, and Lu Sheng. Danceformer: Music conditioned 3d dance generation with parametric motion transformer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 1272–1279, 2022. [2](#)
- [39] Haoying Li, Yifan Yang, Meng Chang, Huajun Feng, Zhihai Xu, Qi Li, and Yueting Chen. SRDiff: Single image Super-Resolution with diffusion probabilistic models. 2021. [3](#)
- [40] Jiaman Li, Karen Liu, and Jiajun Wu. Ego-Body pose estimation via Ego-Head pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17142–17151, 2023. [2](#)
- [41] Ruilong Li, Shan Yang, David A Ross, and Angjoo Kanazawa. Ai choreographer: Music conditioned 3d dance generation with aist++. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13401–13412, 2021. [2](#)
- [42] Xingchao Liu, Xiwen Zhang, Jianzhu Ma, Jian Peng, and Qiang Liu. InstaFlow: One step is enough for High-Quality Diffusion-Based Text-to-Image generation. 2023. [4](#)
- [43] Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. DPM-Solver++: Fast solver for guided sampling of diffusion probabilistic models. 2022. [4](#)
- [44] Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. DPM-Solver: A fast ODE solver for diffusion probabilistic model sampling in around 10 steps. In *Advances in Neural Information Processing Systems*, pages 5775–5787. Curran Associates, Inc., 2022. [4](#)
- [45] Jianxin Ma, Shuai Bai, and Chang Zhou. Pretrained diffusion models for unified human motion synthesis. *arXiv preprint arXiv:2212.02837*, 2022. [2](#)
- [46] Naureen Mahmood, Nima Ghorbani, Nikolaus F Troje, Gerard Pons-Moll, and Michael J Black. Amass: Archive of motion capture as surface shapes. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 5442–5451, 2019. [6](#)
- [47] Chenlin Meng, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. SDEdit: Image synthesis and editing with stochastic differential equations. 2021. [3](#), [7](#)
- [48] Ryan Murray, Brian Swenson, and Soumya Kar. Revisiting normalized gradient descent: Fast evasion of saddle points. 2017. [4](#)
- [49] Xingang Pan, Ayush Tewari, Thomas Leimkühler, Lingjie Liu, Abhimitra Meka, and Christian Theobalt. Drag your gan: Interactive point-based manipulation on the generative image manifold. In *ACM SIGGRAPH 2023 Conference Proceedings*, pages 1–11, 2023. [2](#)
- [50] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed AA Osman, Dimitrios Tzionas, and Michael J Black. Expressive body capture: 3d hands, face, and body from a single image. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10975–10985, 2019. [2](#)
- [51] Mathis Petrovich, Michael J Black, and Gül Varol. Action-conditioned 3d human motion synthesis with transformer vae. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10985–10995, 2021. [2](#)
- [52] Mathis Petrovich, Michael J Black, and Gül Varol. Temos: Generating diverse human motions from textual descriptions. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXII*, pages 480–497. Springer, 2022. [2](#)
- [53] Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. DreamFusion: Text-to-3D using 2D diffusion. 2022. [3](#)
- [54] Vadim Popov, Ivan Vovk, Vladimir Gogoryan, Tasnima Sadekova, and Mikhail Kudinov. Grad-TTS: A diffusion probabilistic model for Text-to-Speech. 2021. [3](#)
- [55] Konpat Preechakul, Nattanat Chatthee, Suttisak Widadwongsa, and Supasorn Suwajanakorn. Diffusion autoencoders: Toward a meaningful and decodable representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10619–10629, 2022. [4](#)
- [56] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. [2](#)
- [57] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with CLIP latents. *ArXiv*, 2022. [3](#)
- [58] Davis Remppe, Tolga Birdal, Aaron Hertzmann, Jimei Yang, Srinath Sridhar, and Leonidas J Guibas. HuMoR: 3D human motion model for robust pose estimation. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*. IEEE, 2021. [2](#), [7](#)

- [59] Davis Rempe, Tolga Birdal, Aaron Hertzmann, Jimei Yang, Srinath Sridhar, and Leonidas J Guibas. Humor: 3d human motion model for robust pose estimation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 11488–11499, 2021. 7
- [60] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-Resolution image synthesis with latent diffusion models. 2021. 3
- [61] Alla Safonova and Jessica K. Hodgins. Construction and optimal search of interpolated motion graphs. *ACM Trans. Graph.*, 26(3):106–es, 2007. 2
- [62] Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J Fleet, and Mohammad Norouzi. Image Super-Resolution via iterative refinement. 2021. 3
- [63] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S Sara Mahdavi, Rapha Gontijo Lopes, Tim Salimans, Jonathan Ho, David J Fleet, and Mohammad Norouzi. Photorealistic Text-to-Image diffusion models with deep language understanding. 2022. 3
- [64] Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models. 2022. 4
- [65] Yonatan Shafir, Guy Tevet, Roy Kapon, and Amit H Bermano. Human motion diffusion as a generative prior. 2023. 2
- [66] Jascha Sohl-Dickstein, Eric A Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. 2015. 3
- [67] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. 2020. 3, 4
- [68] Jiaming Song, Qinsheng Zhang, Hongxu Yin, Morteza Mardani, Ming-Yu Liu, Jan Kautz, Yongxin Chen, and Arash Vahdat. Loss-guided diffusion models for plug-and-play controllable generation. In *International Conference on Machine Learning*, pages 32483–32498. PMLR, 2023. 4
- [69] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. 2019. 3
- [70] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-Based generative modeling through stochastic differential equations. 2020. 3, 4
- [71] Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever. Consistency models. 2023. 4
- [72] Guy Tevet, Sigal Raab, Brian Gordon, Yonatan Shafir, Daniel Cohen-Or, and Amit H Bermano. Human motion diffusion model. In *The Eleventh International Conference on Learning Representations*, 2023. 2, 3, 6, 7
- [73] Jonathan Tseng, Rodrigo Castellon, and C Karen Liu. Edge: Editable dance generation from music. *arXiv preprint arXiv:2211.10658*, 2022. 2
- [74] Bram Wallace, Akash Gokul, Stefano Ermon, and Nikhil Naik. End-to-End diffusion latent optimization improves classifier guidance. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 7280–7290, 2023. 3, 4, 8
- [75] Bram Wallace, Akash Gokul, and Nikhil Naik. EDICT: Exact diffusion inversion via coupled transformations. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2023. 3, 4
- [76] Jiashun Wang, Huazhe Xu, Jingwei Xu, Sifei Liu, and Xiaolong Wang. Synthesizing long-term 3d human motion and interaction in 3d scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9401–9411, 2021. 2
- [77] Daniel Watson, William Chan, Ricardo Martin-Brualla, Jonathan Ho, Andrea Tagliasacchi, and Mohammad Norouzi. Novel view synthesis with diffusion models. 2022. 3
- [78] Yan Wu, Jiahao Wang, Yan Zhang, Siwei Zhang, Otmar Hilliges, Fisher Yu, and Siyu Tang. Saga: Stochastic whole-body grasping with contact. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part VI*, pages 257–274. Springer, 2022. 2
- [79] Yiming Xie, Varun Jampani, Lei Zhong, Deqing Sun, and Huaizu Jiang. OmniControl: Control any joint at any time for human motion generation. 2023. 2
- [80] Sijie Yan, Zhizhong Li, Yuanjun Xiong, Huahan Yan, and Dahua Lin. Convolutional sequence generation for skeleton-based action synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4394–4402, 2019. 2
- [81] Xinyu Yi, Yuxiao Zhou, Marc Habermann, Soshi Shimada, Vladislav Golyanik, Christian Theobalt, and Feng Xu. Physical inertial poser (pip): Physics-aware real-time human motion tracking from sparse inertial sensors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13167–13178, 2022. 6
- [82] Jiwen Yu, Yinhuai Wang, Chen Zhao, Bernard Ghanem, and Jian Zhang. Freedom: Training-free energy-guided conditional diffusion model. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 23174–23184, 2023. 4
- [83] Ye Yuan, Jiaming Song, Umar Iqbal, Arash Vahdat, and Jan Kautz. Physdiff: Physics-guided human motion diffusion model. *arXiv preprint arXiv:2212.02500*, 2022. 2
- [84] Jianrong Zhang, Yangsong Zhang, Xiaodong Cun, Shaoli Huang, Yong Zhang, Hongwei Zhao, Hongtao Lu, and Xi Shen. T2m-gpt: Generating human motion from textual descriptions with discrete representations. *arXiv preprint arXiv:2301.06052*, 2023. 2
- [85] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models, 2023. 2
- [86] Mingyuan Zhang, Zhongang Cai, Liang Pan, Fangzhou Hong, Xinying Guo, Lei Yang, and Ziwei Liu. Motiondiffuse: Text-driven human motion generation with diffusion model. *arXiv preprint arXiv:2208.15001*, 2022. 2
- [87] Siwei Zhang, Yan Zhang, Federica Bogo, Marc Pollefeys, and Siyu Tang. Learning motion priors for 4d human body capture in 3d scenes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11343–11353, 2021. 2

- [88] Yan Zhang, Michael J Black, and Siyu Tang. Perpetual motion: Generating unbounded human motion. *arXiv preprint arXiv:2007.13886*, 2020. [2](#)
- [89] Rui Zhao, Hui Su, and Qiang Ji. Bayesian adversarial human motion synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6225–6234, 2020. [2](#)
- [90] Kaiwen Zheng, Cheng Lu, Jianfei Chen, and Jun Zhu. DPM-Solver-v3: Improved diffusion ODE solver with empirical model statistics. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. [4](#)