

A Background On Point-E

Point-E [39] is a diffusion-based generative model that produces 3D point clouds from text or images. The Point-E pipeline consists of three stages: first, it generates a single synthetic view using a text-to-image diffusion model; second, it produces a coarse, low-resolution 3D point cloud (1024 points) using a second diffusion model which is conditioned on the generated image; third, it upsamples/“densifies” the coarse point cloud to a high-resolution one (4096 points) with a third diffusion model. The two diffusion models operating on point clouds use a permutation invariant transformer architecture with different model sizes. The entire model is trained on Point-E’s curated dataset of several million 3D models and associated metadata which captures a generic distribution of common 3D shapes, providing a suitable and sufficiently diverse prior for robot geometry. The diffused data is a set of points, each point possessing 6 feature dimensions: 3 for spatial coordinates and 3 for colors. We ignore the color channels in this work. The conditioning for the synthesized image in the first stage relies on embeddings computed from a pre-trained ViT-L/14 CLIP model; in the embedding optimization of DiffuseBot, the variables to be optimized is exactly the same embedding. Diffusion as co-design is only performed in the second stage (coarse point cloud generation) since the third stage is merely an upsampling which produces only minor modifications to robot designs. We refer the reader to the original paper [39] for more details.

B Theoretical Motivation

Online learning in embedding optimization. In Section 2.4, we discuss how to online collect a dataset to optimize the embedding toward improved physical utility. Given a simplified version of (5)

$$\min_{\mathbf{c}} \mathbb{E}_{p_{\theta}(\mathbf{x}_0|\mathbf{c})} [g(\mathbf{x}_0, \mathbf{c})] \quad (8)$$

where, for notation simplicity, we drop $t \sim [1, T]$, $\mathcal{N}(\epsilon; \mathbf{0}, \mathbf{I})$ in the sampling distribution, and summarize $\mathbb{E}[\|\epsilon - \epsilon_{\theta}(\mathbf{x}_t(\mathbf{x}_0, \epsilon, t), t, \mathbf{c})\|^2]$ as $g(\mathbf{x}, \mathbf{c})$. We can rewrite the expectation term as,

$$\int p_{\theta}(\mathbf{x}_0) \frac{p_{\theta}(\mathbf{c}|\mathbf{x}_0)}{p_{\theta}(\mathbf{c})} g(\mathbf{x}_0, \mathbf{c}) dx \quad (9)$$

which allows to sample from $p_{\theta}(\mathbf{x}_0)$ (i.e., generating samples from the diffusion model) and reweight the loss with $\frac{p_{\theta}(\mathbf{c}|\mathbf{x}_0)}{p_{\theta}(\mathbf{c})}$; the latter scaling term is essentially proportional to a normalized task performance. Empirically, we can maintain a buffer for the online dataset and train the embedding with the sampling distribution biased toward higher task performance; we use a list to store samples with top-k performance in our implementation (we also tried reshaping the sampling distribution like prioritized experience replay in reinforcement learning but we found less stability and more hyperparameters required in training compared to our simpler top-k approach).

Connection to MCMC. In diffusion sampling, the simplest way to perform reverse denoising process as in Section 2.2 follows [24],

$$p_{\theta}(\mathbf{x}_{t-1}|\mathbf{x}_t) = \mathcal{N}\left(\frac{1}{\sqrt{\alpha_t}}(\mathbf{x}_t - \frac{\beta_t}{\sqrt{1-\alpha_t}}\epsilon_{\theta}(\mathbf{x}_t, t)), \frac{1-\bar{\alpha}_{t-1}}{1-\bar{\alpha}_t}\beta_t\mathbf{I}\right) \quad (10)$$

Here, the denoising term can either be unconditional $\epsilon_{\theta}(\mathbf{x}_t, t)$ or conditional via classifier-free guidance $\hat{\epsilon}_{\theta, \text{classifier-free}}(\mathbf{x}_t, t, \mathbf{c})$ as in (3). We use the latter to incorporate the optimized embedding. To further leverage physics-based simulation, we aim to introduce physical utility during diffusion sampling process. One possibility is to utilize classifier-based guidance [12],

$$\hat{\epsilon}_{\theta, \text{classifier-based}} := \epsilon_{\theta}(\mathbf{x}_t, t) - s \cdot \sqrt{1-\bar{\alpha}_t} \nabla_{\mathbf{x}_t} \log p(\mathbf{c}|\mathbf{x}_t) \quad (11)$$

where $p(\mathbf{c}|\mathbf{x}_t)$ can be conceptually viewed as improving physical utility and $\nabla_{\mathbf{x}_t} \log p(\mathbf{c}|\mathbf{x}_t)$ can be obtained using differentiable physics and the unconditional score. Note that we slightly abuse the notation here by overloading \mathbf{c} with conditioning from differentiable simulation during sampling other than the classifier-free guidance using the optimized embedding. However, combining (10) and (11) much less resembles any gradient-based optimization techniques, which are shown to be effective in soft robot co-design with differentiable simulation [26, 2]. Fortunately, drawing a connection to energy-based models [54, 15, 14], yet another alternative to incorporate conditioning in diffusion

Table 4: Configuration of embedding optimization.

| | Balancing | Landing | Crawling | Hurdling | Gripping | Moving a Box |
|----------------------|-----------|---------|----------|----------|----------|--------------|
| Buffer Size | 600 | 600 | 60 | 600 | 60 | 60 |
| Min. Buffer Size | 60 | 60 | 60 | 60 | 60 | 60 |
| Num. Samples / Epoch | 60 | 60 | 60 | 60 | 60 | 60 |
| Train Iter. / Epoch | 1 | 1 | 1 | 1 | 1 | 1 |
| Buffer Top-K | 12 | 12 | 6 | 6 | 6 | 6 |
| Batch Size | 6 | 6 | 6 | 6 | 6 | 6 |

Table 5: Configuration of diffusion as co-design.

| | Balancing | Landing | Crawling | Hurdling | Gripping | Moving a Box |
|--------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|
| t_{\max} | 400 | 150 | 400 | 400 | 400 | 400 |
| t_{\min} | 0 | 0 | 0 | 0 | 0 | 0 |
| Δt | 50 | 25 | 50 | 50 | 50 | 50 |
| K | 3 | 3 | 5 | 5 | 5 | 5 |
| σ | $10^{-4} \cdot \beta$ |
| κ | 10^4 | 10^4 | 10^4 | 10^4 | 10^4 | 10^4 |
| γ | - | - | 0.01 | 0.001 | 0.001 | 0.001 |
| Renorm Scale | 10 | 10 | 10 | 10 | 10 | 10 |

models is Markov Chain Monte Carlo (MCMC) sampling [14], where we use Unadjusted Langevin Dynamics,

$$\mathbf{x}_t = \mathbf{x}_t^{(K)}, \text{ where } \mathbf{x}_t^{(k)} \sim \mathcal{N}\left(\mathbf{x}_t^{(k)}; \mathbf{x}_t^{(k-1)} + \frac{\sigma^2}{2} \nabla_{\mathbf{x}} \log p(\mathbf{c} | \mathbf{x}_t^{(k-1)}), \sigma^2 \mathbf{I}\right) \quad (12)$$

where K is the number of samples in the current MCMC with k as indexing, σ^2 is a pre-defined variance, and $\mathbf{x}_t^{(0)} = \mathbf{x}_{t-1}$. In the context of diffusion models, this procedure is commonly performed within a single diffusion step to drive the sample toward higher-density regime under the intermediate distribution $p(\mathbf{x}_t) = \int q(\mathbf{x}_t | \mathbf{x}_0) p(\mathbf{x}_0) d\mathbf{x}_0$ at diffusion time t . Inspired by its resemblance to gradient ascent with stochasticity from the added Gaussian noise of variance σ^2 , we establish a connection to design optimization, reformulating diffusion process as co-design optimization as in (6) and (7). Specifically, we can apply Bayes rule to decompose the score of $p(\mathbf{c} | \mathbf{x}_t^{(k-1)})$,

$$\nabla_{\mathbf{x}} \log p(\mathbf{c} | \mathbf{x}_t) = \nabla_{\mathbf{x}} \log p(\mathbf{x}_t | \mathbf{c}) - \nabla_{\mathbf{x}} \log p(\mathbf{x}_t) \quad (13)$$

where $\nabla_{\mathbf{x}} \log p(\mathbf{x}_t)$ is simply the denoiser output $\epsilon_{\theta}(\mathbf{x}_t, t)$ and $\nabla_{\mathbf{x}} \log p(\mathbf{x}_t | \mathbf{c})$ is the gradient of task performance with respect to the intermediate sample of the diffusion model from differentiable physical simulation and robotizing process. Overall, this leads to (6).

C Implementation Details In Algorithm

In this section, we provide more implementation details and experimental configurations of Diffuse-Bot and other baselines. In Table 4, we list the configurations of the embedding optimization. With respect to Algorithm 1, “buffer” refers to the online dataset \mathcal{D} .

- *Buffer Size* is the capacity of the dataset.
- *Min. Buffer Size* is the minimum of data filled in the buffer before training starts.
- *Num. Samples / Epoch* is number of new samples collected in each epoch, where epoch here refers to a new round of data collection in online learning.
- *Train Iter. / Epoch* is number of training iterations per epoch.
- *Buffer Top-K* is the number of datapoints with top-k performance being retained in the *Filter* step atop the most up-to-date data.
- *Batch Size* is the batch size in the embedding optimization.

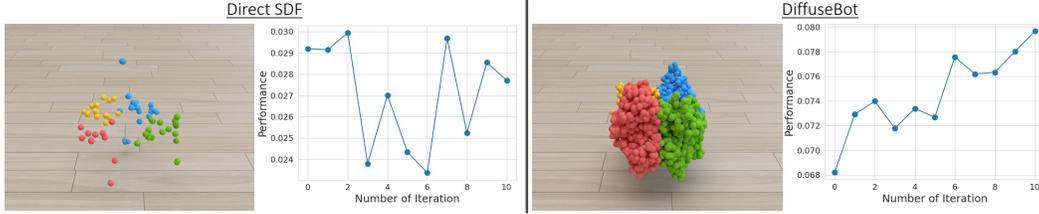


Figure 8: Comparison between *Direct SDF* and *DiffuseBot* on conversion to solid geometry for robotizing.

In Table 5, we list the configurations of diffusion as co-design. We follow (6)(7) for:

- K is number of MCMC sampling steps at the current diffusion time.
- σ is the standard deviation related to the MCMC step size.
- κ is the ratio between two types of design gradients.
- γ is the weight for trading off design and control optimization.
- t_{\max} and t_{\min} are the maximal and minimal diffusion time to perform diffusion as co-design, respectively.
- Δt is the diffusion time interval to perform diffusion as co-design.

For baselines, we use learning rates for control optimization following γ in Table 5; for particle-based and voxel-based approaches, we use learning rate 0.01 for design optimization; for implicit function and diff-CPPN, we use learning rate 0.001 for design optimization. For the inputs of implicit function and diff-CPPN, we use x , y , z coordinates in the local workspace, the distance to the workspace center on the xy , xz , yz planes, and the radius from the center. For the network architecture of implicit function, we use a 2-layer multilayer perceptron with hidden size 32 and Tanh activation. For the network architecture of diff-CPPN, we use Sin, Sigmoid, Tanh, Gaussian, SELU, Softplus, Clamped activations with 5 hidden layers and 28 graph nodes in each layer.

Hyperparameters are chosen mostly based on intuition and balancing numerical scale with very little tuning. In the following, we briefly discuss the design choices of all hyperparameters listed in Table 5 and Table 4. For min buffer size, samples per epoch, training iteration per epoch, and batch size, we roughly make sufficiently diverse the data used in the optimization and use the same setting for all tasks. For buffer size, we start with 60 and if we observe instability in optimization, we increase to 10 times, 600 (similar to online on-policy reinforcement learning); note that buffer size refers to the maximal size and increasing this won't affect runtime. For buffer Top-K, we start with 6 and if we observe limited diversity of generation throughout the optimization (or lack of exploration), we double it. For t_{\max} , t_{\min} , and Δt , we roughly inspect how structured the generation in terms of achieving the desired robotic task to determine t_{\max} and modify Δt accordingly to match the similar number of performing MCMC sampling (e.g., $t_{\max}/\Delta t$: $400/50 \approx 150/25$). For the number of MCMC steps K , we simply set 3 for passive tasks and 5 for active tasks by intuition. For σ , we simply follow one of the settings in [14]. For the guidance scale κ and renorm scale, we check the numerical values between ϵ and gradient from differentiable simulation and try to make them roughly in the similar magnitude, and set the same scale for all tasks for simplicity. For γ , we set 0.001 for trajectory optimization and 0.01 for parameterized controllers based on our experience of working with differentiable physics. Overall, from our empirical findings, the only hyperparameters that may be sensitive include buffer size and buffer Top-K for optimization stability and generation diversity, and guidance scales, which need to be tuned to match the numerical magnitude of other terms so as to take proper effect.

D Details In Task Setup

In this section, we provide more details of all tasks setup including environment configuration and prescribed actuator and controller if any (as mentioned in Section 2.4; it is fixed during embedding optimization and will be co-optimize during diffusion as co-design). We select tasks that

1. can cover a wide spectrum of existing robotics tasks: we briefly categorize tasks into passive dynamics, locomotion, and manipulation. Note that passive dynamics tasks are explicitly considered

here since there is no active control of robot bodies, making optimization on robot design a direct factor toward physical utility.

2. only involve lower-level control/motion without the complications of long-term or higher-level task planning: we select tasks that mostly involve few motor skills, e.g., in manipulation, instead of pick and place, we simply aim at picking up/gripping an object.
3. are commonly considered in other soft robot co-design literature: all proposed active tasks are widely used in the soft robot community, including crawling [8, 10, 47, 61], hurdling/jumping [26, 59, 3], and manipulating objects [5, 11, 38].
4. may induce more visible difference in robot designs between the performing and the non-performing ones to facilitate evaluation and algorithmic development: we select tasks more based on heuristics and intuition, e.g., in crawling, we expect leg-like structures may outperform other random designs.

We build our environments on top of SoftZoo [61] and employ the Material Point Method for simulation. Each environment is composed of boundary conditions that include impenetrable ground, and, in the case of fixed objects or body parts, glued particles. All tasks use units without direct real-world physical correspondence and last for 100 steps with each step consisting of 17 simulation substeps. In the following, when describing the size of a 3D shape without explicitly mentioning the axis, we follow the order: length (x, in the direction when we talk about left and right), height (y, parallel with the gravity direction), and width (z). All tasks are demonstrated in Figure 3 in the main paper.

Balancing. The robot is initialized atop a stick-like platform of shape 0.02-unit \times 0.05-unit \times 0.02-unit. The robot is given an initial upward velocity of a 0.5-unit/second and allowed to free fall under gravity. The goal is for the robot to passively balance itself after dropping again on the platform; the performance is measured as the intersection over union (IoU) between the space occupied by the robot during the first simulation step and the space occupied by the robot during the last simulation step. The robot geometry is confined to a 0.08-unit \times 0.08-unit \times 0.08-unit workspace. There is no prescribed actuator placement or controller (passive dynamics).

Landing. The robot is initialized to be 0.08-unit to the right and 0.045-unit above the landing target with size of 0.02-unit \times 0.05-unit \times 0.02-unit. The robot is given an initial velocity of 0.5-unit/second to the right. The goal of the robot is to land at the target; the performance is measured as the exponential to the power of the negative distance between the target and the robot in the last frame $e^{-\|p_H^{\text{object}} - p_H^{\text{robot}}\|}$, where p_H is the position of the robot or object at the last frame with horizon H . The robot geometry is confined to a 0.08-unit \times 0.08-unit \times 0.08-unit workspace. There is no prescribed actuator placement or controller (passive dynamics).

Crawling. The robot is initialized at rest on the ground. The goal of the robot is to actuate its body to move as far away as possible from the starting position; the performance is measured as the distance traveled $\|p_H^{x,\text{robot}} - p_0^{x,\text{robot}}\|$, where $p^{x,\text{robot}}$ is the position of the robot in the x axis at a certain frame with horizon H . The robot geometry is confined to a 0.08-unit \times 0.08-unit \times 0.08-unit workspace. The actuator placement is computed by clustering the local coordinates of the robot centered at its average position in the xz (non-vertical) plane into 4 groups. Each group contains an actuator with its direction parallel to gravity. The prescribed controller is a composition of four sine waves with frequency as 30hz, amplitude as 0.3, and phases as 0.5π , 1.5π , 0 , and π for the four actuators.

Hurdling. An obstacle of shape 0.01-unit \times 0.03-unit \times 0.12-unit is placed in 0.07-unit front of the robot. The goal of the robot is to jump as far as possible, with high distances achieved only by bounding over the obstacle. The performance is measured as the distance traveled. The robot geometry is confined to a 0.08-unit \times 0.08-unit \times 0.08-unit workspace. The actuator placement is computed by clustering the local coordinates of the robot centered at its average position in the length direction into 2 groups. Each group contains an actuator aligned parallel to gravity. The prescribed controller takes the form of open-loop, per-step actuation sequences, set to linearly-increasing values from (0.0, 0.0) to (1.0, 0.3) between the first and the thirtieth frames and zeros afterward for the two actuators (the first value of the aforementioned actuation corresponds to the actuator closer to the obstacle) respectively.

Gripping. An object of shape 0.03-unit \times 0.03-unit \times 0.03-unit is placed 0.08-unit underneath the robot. The goal of the robot is to vertically lift the object; the performance is measured as the vertical distance of the object being lifted $\|p_H^{y,\text{object}} - p_0^{y,\text{object}}\|$, where $p^{y,\text{object}}$ is the position of the

object in the y axis at a certain frame with horizon H . Within a $0.06\text{-unit} \times 0.08\text{-unit} \times 0.12\text{-unit}$ workspace, we decompose the robot into a base and two attached submodules, and we set the output of DiffuseBot or other robot design algorithms as one of the submodules and make constrain the other submodule to be a mirror copy; conceptually we design “the finger of a parallel gripper.” The base is clamped/glued to the upper boundary in z and given a large polyhedral volume to serve as the attachment point of the gripper finger submodules. The actuator placement is computed by clustering the local coordinates of the robot centered at its average position in the length direction into 2 groups; each submodule has one pair of actuators. Each actuator is aligned parallel to gravity. Overall, one pair of actuators comprises the "inner" part of the gripper and the other comprises the "outer" part. Suppose the actuation of the two pairs of actuators is denoted in the format of (actuation of the outer part, actuation of the inner part), the prescribed controller is per-frame actuation, set to (i) linearly-increasing values from (0.0, 0.0) to (0.0, 1.0) between the first and the fiftieth frames, and (ii) then linearly-decreasing values from (1.0, 0.0) to (0.0, 0.0) between the fiftieth and the last frames.

Moving a box. A $0.03\text{-unit} \times 0.03\text{-unit} \times 0.03\text{-unit}$ cube is placed on the right end of the robot (half a body length of the robot to the right from the robot center). The goal of the robot is to move the box to the left; the performance is measured as the distance that the box is moved to the right with respect to its initial position $\|p_H^{x,\text{object}} - p_0^{x,\text{object}}\|$, where $p^{x,\text{object}}$ is the position of the object in the x axis at a certain frame with horizon H . The robot geometry is confined to a $0.16\text{-unit} \times 0.06\text{-unit} \times 0.06\text{-unit}$ workspace. The actuator placement is computed by clustering the local coordinates of the robot centered at its average position in the height direction into 2 groups. Each group contains an actuator aligned parallel to the ground. The prescribed controller is per-frame actuation, initialized to linearly-increasing values from (0.0, 0.0) to (1.0, 0.0) between the first and last frame for the lower and the upper actuator respectively.

E Analysis On Robotizing

As mentioned in Section 2.3, Material Point Method simulation requires solid geometry for simulation; thus, we need to convert the surface point cloud from Point-E [39] to a volume. The most direct means of converting the point cloud to a solid geometry is to compute the signed distance function (SDF), and populate the interior of the SDF using rejection sampling. We refer to this baseline as *Direct SDF*. Here, we use a pretrained transformer-based model provided by Point-E as the SDF. In Figure 8, we compare Direct SDF with our approach described in Section 2.3 Solid Geometry. We perform robotizing on intermediate samples at $t=300$. We observe that Direct SDF fails to produce well-structured solid geometry since it is trained with watertight on geometry, and thus the conversion cannot gracefully handle generated point clouds that do not exhibit such watertight structure. This is specifically common in the intermediate diffusion sample \mathbf{x}_t as \mathbf{x}_t is essentially a Gaussian-noise-corrupted version of the clean surface point cloud. In contrast, the robotizing of DiffuseBot produces a much well-structured solid geometry since it explicitly handles the noisy interior 3D points by introducing a tailored loss as in Shape As Points optimization [40] (see Section 2.3). In addition, a better-structured robot geometry is critical to obtain not only a more accurate evaluation of a robot design at the forward pass of the simulation but also the gradients from differentiable physics at the backward pass. In Figure 8, we further perform co-optimization on the two robots obtained by Direct SDF and DiffuseBot; we observe a more stably increasing trend of task performance in our approach, demonstrating that the gradient is usually more informative with a better-structured robot. Empirically, while Direct SDF sometimes still provides improving trend with unstructured robots in co-design optimization, the performance gain is not as stable as DiffuseBot.

F Additional Visualizations Of Experiments

In this section, we show more visualization for a diverse set of the experimental analysis. Please visit our project page (<https://diffusebot.github.io/>) for animated videos.

Generated robots performing various tasks. In Figure 9, we display a montage of a generated robot’s motion for each task; this is the unabridged demonstration of Figure 5 in the main paper.

Physics-guided robot generation In Figure 10, we show how DiffuseBot evolves robots throughout the embedding optimization process; this is the full demonstration of Figure 4 in the main paper. Note



Figure 9: Examples of robots bred by DiffuseBot to achieve the all presented tasks.

that DiffuseBot is a generative algorithm and thus results presented are for demonstration purposes and reflect only a single fixed random seed.

Diffusion process of robot generation. In Figure 11, we show how robots evolve throughout the diffusion process at different diffusion time via \mathbf{x}_t with t from T to 0; not to be confused by Figure 10 where all generations are obtained *via* a full diffusion sampling \mathbf{x}_0 and different conditioning embeddings c . This figure demonstrates the robotized samples across diffusion times, gradually converging to the final functioning robots as t approaches 0.

Comparison with baselines. In Figure 12, we present robots generated from different baselines. Note that Point-E is essentially DiffuseBot without embedding optimization and diffusion as co-design.

Incorporating human feedback. In Figure 13, we showcase more examples of incorporating human textual feedback to the generation process of DiffuseBot. As briefly mentioned in Section 3.4 in the main paper, we leverage the compositionality of diffusion-based generative models [34, 14]. Specifically, this is done by combining two types of classifier-free guidance as in (3); one source

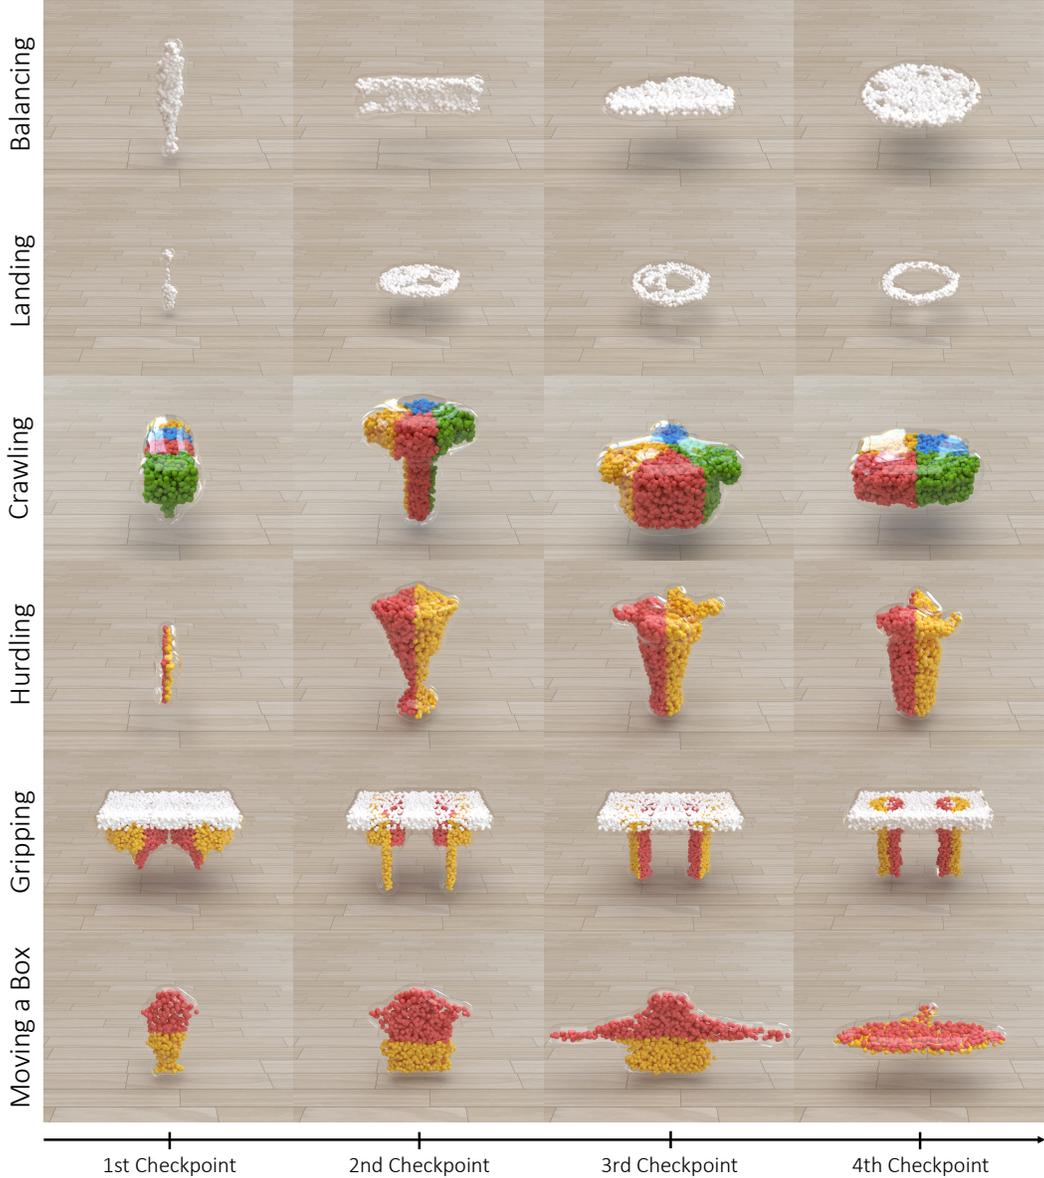


Figure 10: Examples of DiffuseBot evolving robots to solve all presented tasks.

of guidance is derived from the task-driven embedding optimization while the other is derived from embeddings from human-provided textual descriptions. These two conditional scores, namely $\epsilon_{\theta}(\mathbf{x}_t, t, \mathbf{c}_{\text{physics}})$ and $\epsilon_{\theta}(\mathbf{x}_t, t, \mathbf{c}_{\text{human}})$, can then be integrated into the diffusion as co-design framework as in (6). We refer the reader to the original papers for more details and theoretical justification. We demonstrate examples of: the balancing robot with additional text prompt *"a ball"*, the crawling robot with additional text prompt *"a star"*, the hurdling robot with additional text prompt *"a pair of wings"*, and the moving-a-box robot with additional text prompt *"thick"*. Interestingly, human feedback is introduced as an augmented "trait" to the original robot. For example, while the hurdling robot keeps the tall, lengthy geometry that is beneficial for storing energy for jumping, a wing-like structure appears at the top of the robot body instead of overwriting the entire geometry. This allows users to design specific parts for composition while also embedding fabrication and aesthetic priorities and constraints through text. We leave future explorations of these ideas for future work.



Figure 11: Demonstrations of generations changing from noises to robots with physical utility throughout the diffusion sampling process. We present snapshots at diffusion times $t = 600, 400, 200, 0$ for all presented tasks.

G Hardware Design and Fabrication

To create a real-world counterpart of the DiffuseBot, we fabricated a proof-of-concept physical robot gripper, as illustrated in Figure 14(a). The gripper was designed to match the shape of the digital gripper and successfully demonstrated its ability to grasp objects, as shown in Figure 15. A video demonstrating the grasping process can be found on our project page (<https://sites.google.com/view/diffusebot>).

During the translation from the simulated design to a physical robot, we aimed to minimize differences between the simulated and physical designs. However, due to hardware and fabrication limitations, certain modifications were necessary.

One such challenge involved replicating the arbitrary contraction force at the particles from the simulation in the physical world. To address this, we employed tendon transmission as the actuation method for the real-world gripper, which was found to be suitable for emulating interaction forces

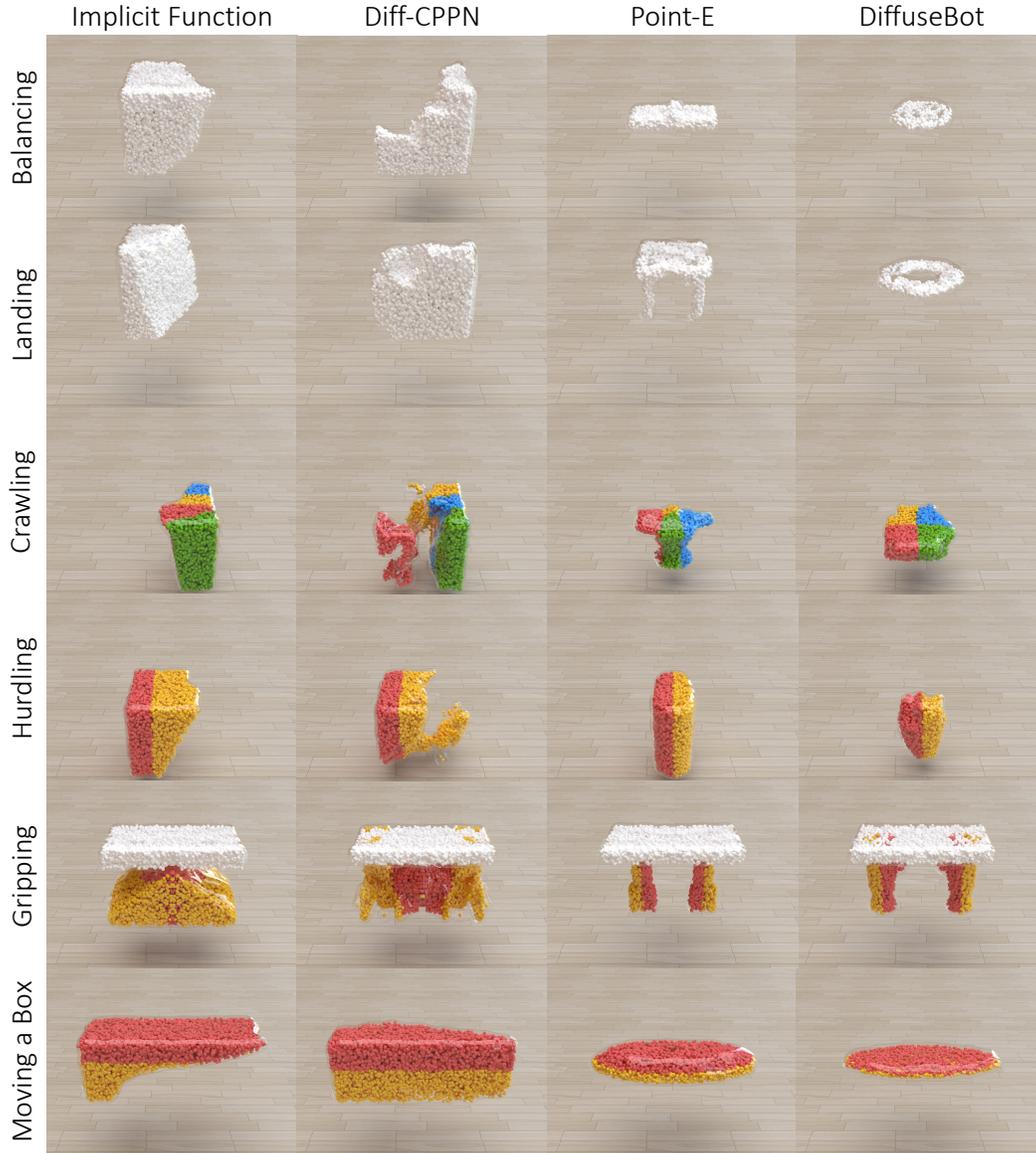


Figure 12: Comparison of robots generated by different baselines.

from the digital world without introducing significant discrepancies. To enhance the stability of the tendon-driven actuation, we incorporated a rigid base ("Base" in Figure 14(a)) above the gripper itself ("Gripper" in Figure 14(a)). This base was designed to withstand the reaction force generated by the tendon-driven actuation. Furthermore, we added four tendon routing points (represented by small dotted circles in Figure 14(b)) on each finger to securely fix the tendon path. By utilizing four Teflon tubes (shown in Figure 14(b)), we were able to position complex components such as motors, batteries, and controllers away from the gripper, reducing its complexity.

The actuation strategy of the simulated gripper applies equal contraction force to both fingers simultaneously. To replicate this strategy, we employed underactuated tendon routing in the development of the gripper. This approach eliminates the need for four separate actuators, thereby reducing the complexity of the robot. We used tendon-driven actuators specifically designed for underactuated tendon-driven robots as they solve practical issues such as size and friction issues that commonly arise in such systems [30].

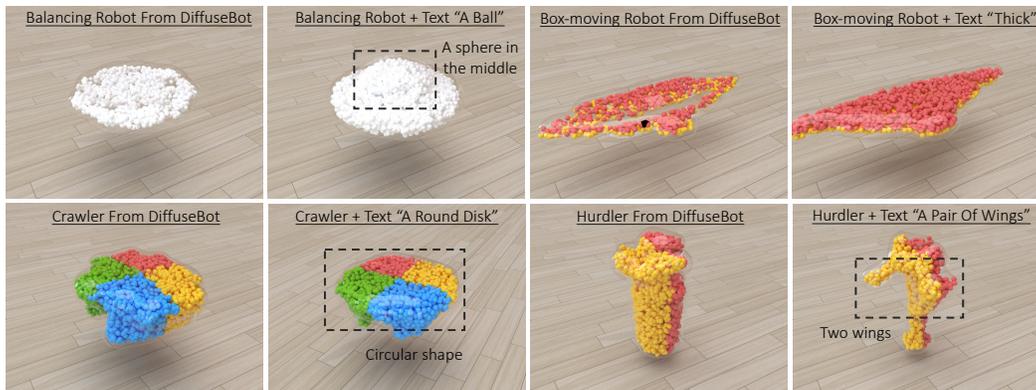


Figure 13: More examples of incorporating human textual feedback into robot generation by DiffuseBot.

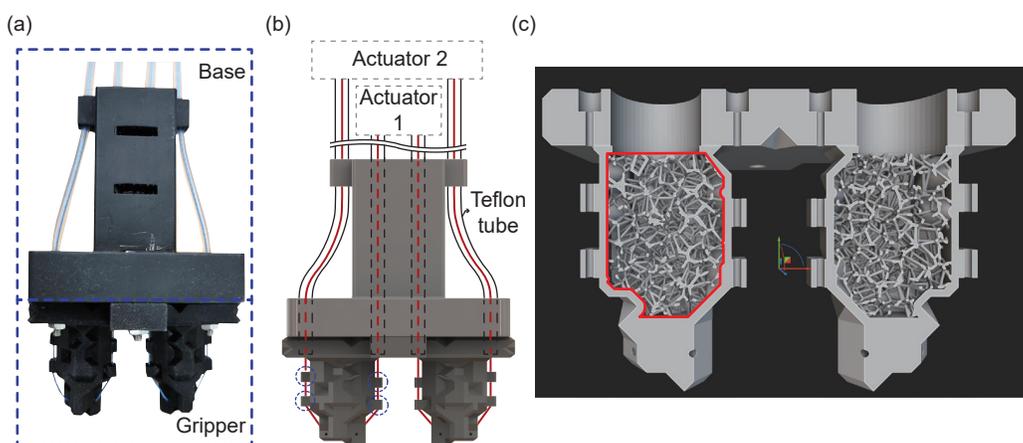


Figure 14: (a) shows a fabricated proof-of-concept physical robot gripper; (b) describes detailed robot configuration; (c) represents interior structure design used to soften the robot body.

The soft gripper was 3D-printed using a digital light projection (DLP) type 3D printer (Carbon M1 printer, Carbon Inc.) and commercially available elastomeric polyurethane (EPU 40, Carbon Inc.). To enhance the softness of the robot body beyond the inherent softness of the material itself, we infilled the finger body with a Voronoi lattice structure, a metamaterial useful for creating a soft structure with tunable effective isotropic stiffness [20, 36]. We generated a Voronoi lattice foam with point spacing of 2.5mm, and a beam thickness of 0.4 mm as shown in the red boundary of Fig. 14(c). Finally, we tested the soft gripper, designed and fabricated as described above, to verify its ability to grasp an object, as shown in Figure 15.

More discussion on fabrication. Although, at present, the compilation of the virtual robot to a physical, digitally fabricated counterpart involves manual post-processing of algorithm’s output, most, if not all of these steps could be automated. Our method outputs a point cloud (defining geometry), actuator placements, and an open-loop controller, along with a prescribed stiffness. Since we can easily convert the point cloud into a 3D triangle mesh, the geometry can be created by almost any 3D printing method. In order to realize an effective stiffness and material behavior, stochastic lattices, specifically Voronoi foams, have been used [36, 20] in the past and employed here in order to match target material properties. Given the actuator placement, tendons [27, 30] can be aligned with the prescribed (contiguous) regions. Since a lattice is used, threading tendons through the robot body is simple, and we note that even more complex routings have been studied in detail in the literature [4]. Creating attachment points for the tendons is a relatively simple geometry processing problem [6]. Thus, converting a virtual robot to a specification that incorporates geometry, material, and actuation can be automated in a straightforward way.

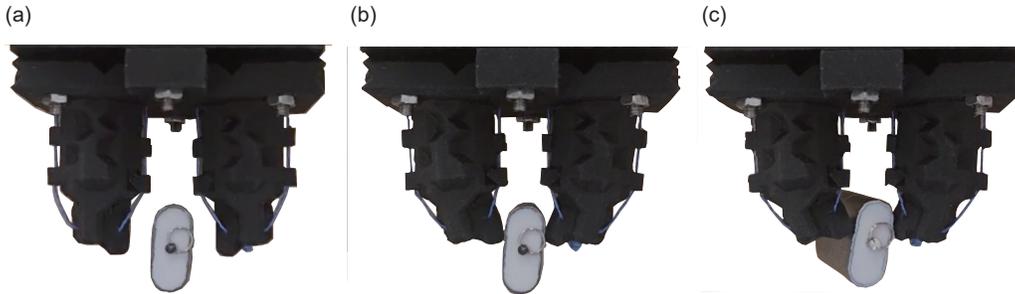


Figure 15: Grasping motion of the real-world soft gripper from DiffuseBot. Time order from (a) to (c). The demo video can be seen at our project page.

We note that when controlled, the physical robot may not always match the virtual robot’s motion. This is the sim-to-real gap, and is significantly harder to overcome in our case than translating the virtual robot to physical hardware. Significant literature has been invested in specifically tackling the sim-to-real gap, and in our case would require its own dedicated project; however, we note that often hardware can be adapted to work by modifying only the control policies using feedback from real-world experiments, often even with little human intervention [21].

Quantitative analysis. In Figure 16, we provide a detailed analysis on the gap between simulation and physical robot. In order to explore the quantitative gap between the behavior of the physical robot and the simulated robot, we conducted an experiment with the following conditions, where similar setups are commonly adopted in soft robot literature [17]. The objective was to measure the change in distance between two tips when we pull/release two tendons - one for closing the gripper (flexion) and the other for opening it (extension). The tendons were pulled or released in increments and decrements of 2mm.

When contracting the tendon to flex or extend the fingers, both simulation and real robot results show log-shaped graphs. The pattern in the physical robot plot is a commonly observed phenomenon called hysteresis. However, the main difference between the simulation and real-world cases can be seen when releasing the tendon from a fully contracted state. In the real robot experiment, the tip distance changes rapidly, while in the simulation, the opposite effect is observed.

One plausible explanation for this disparity could be attributed to the friction direction and elongation of the tendons. During the transition from tendon contraction to tendon release, the tension of the tendon at the end-effector may change suddenly due to the change of the friction direction. Also, since we only control the motor position (not the tendon position) to pull/release the tendon with 2mm step, the exact tendon length may not be exactly the same when we consider the tendon elongation.

Potential solutions. Given that the gap between simulation and real robot performance seems to originate from the actuation/transmission method, our future work will focus on developing a tendon-driven actuation simulation framework. This framework aims to address the differences and improve the accuracy of our simulations. We are exploring other possible explanations for the sim-to-real gap and will investigate any additional factors that may contribute to the observed discrepancies. Overall, as for a high-level general solution, we believe (1) adjusting parameters based on observed sim to real gap and repeat the design process or (2) building a more accurate physics-based simulation (which can be straightforwardly plug-and-played in DiffuseBot) can largely bridge the sim-to-real gap of fabricating physical robots; or more interestingly, connecting generative models to commercial-level design and simulation softwares.

H More Discussions

H.1 Limitation

The major limitation of DiffuseBot is that we make a simplification in the parameterization of actuators and stiffness; we make dependencies of the two design specifications on robot geometry (check more technical details in Section 2.3 paragraph *Actuators and Stiffness*). This works well

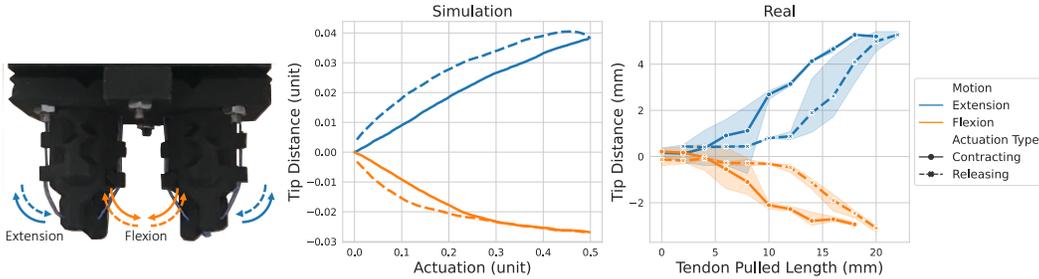


Figure 16: Quantitative analysis of behavior between simulation and physical robots.

with properly-crafted mapping from geometry to the others yet limits the potential by human prior with little use of the generative power. While this may be reasonable as properly-set actuators and stiffness based on geometry (hand-tuned empirically in this work) roughly reflects task performance, a more flexible parameterization can definitely lead to improved performance. Potential remedy can be using part-based 3D generative models for actuators and direct optimization for stiffness. Another limitation is the gap between simulated results and real robots. While the hardware experiment has shown as a proof-of-concept that minimally demonstrates the potential, physical robot fabrication and real-world transfer have countless non-trivial challenges including stiffness and actuator design, sim-to-real gap, etc. This may require studies on more high-fidelity physics-based simulation, which can be straightforwardly plugged into DiffuseBot.

H.2 Conditioning Beyond Text

The use of textual inputs additional to the embeddings optimized toward physical utility is achieved by both being able to be consumed by the diffusion model to produce guidance for the diffusion process. More concretely speaking, in DiffuseBot, we use the CLIP feature extractor as in Point-E and it allows to extract embedding for both text and image modalities, which can then be used as a condition in the diffusion model. Thus, we can also incorporate images as inputs and perform the exact same procedure as that of the textual inputs. Theoretically, the textual inputs are incorporated via following the intuition in end of the paragraph *Diffusion as Co-design* in Section 2.4, where the textual inputs additionally provide gradients toward following the textual specification. Similarly, the image inputs can also be processed to provide gradients since CLIP embeddings live in a joint space of images and languages. More interestingly, if we build DiffuseBot on models other than Point-E, which can consume embeddings for other modalities like audio as conditioning, we can then straightforwardly perform robot design generation guided by the other corresponding input formats (and meanwhile, toward physical utility). Note that this critical feature of compositionality across different sources of guidance throughout the reverse diffusion process is one of the biggest advantages of using diffusion-based models as opposed to other types of generative models.

H.3 Connection To Text Inversion

There is some synergy between text inversion in [19] and embedding optimization in DiffuseBot. Both of them aim at tuning the embedding toward reflecting certain properties of the output generation, i.e., describing the output generated images in [19] and toward improved physical utility in DiffuseBot. The major difference lies in the nuance of the data/samples used to carry out the optimization. Text inversion performs a direct optimization using latent diffusion model loss (Eq. (2) in [19]), which computes losses on noisy samples/latents corrupted from the real dataset. On the other hand, it is tricky to think about real dataset in robot design (as discussed in the second paragraph of Section 1 and the paragraph *Embedding Optimization* in Section 2.4), embedding optimization in DiffuseBot computes losses on noisy samples corrupted from self-generated data filtered by robot performance (as in Algorithm 1 and Section 2.4). Conceptually, it is more like a mixture of diffusion model training and online imitation learning like DAGGER [45].

H.4 Connection To Diffusion Models With Physical Plausibility

A potential and interesting way to adapt DiffuseBot to other applications like motion planning or control [28, 1] is to view a generated robot as one snapshot/frame of a motion/state sequence and the physics prior can be the dynamics constraint across timesteps (e.g., robot dynamics or contact dynamics that enforce non-penetration). The physics prior can be injected similarly to diffusion as co-design that propagates the enforcement of physical plausibility of generated states from differentiable physics-based simulation to diffusion samples. For example, considering states in two consecutive timesteps, we can compute loss in the differentiable simulation to measure the violation of physical constraints regarding robot dynamics or interaction with the environment. Then, we can compute gradients with respect to either control or design variables; for gradients in control, this will essentially augment works like [28, 1] with classifier-based guidance to achieve physical plausibility; for gradients in design, this will much resemble optimizing toward the motion sequence of a shape-shifting robot.

H.5 Parameterization Of Actuator And Stiffness

The goal of DiffuseBot is to demonstrate the potential of using diffusion models to generate soft robot design and to leverage the knowledge of the pre-trained generative models learned from a large-scale 3D dataset. Under this setting, the generated output of the diffusion model can only provide the geometry information of robot designs, leading to our design choice of having full dependency of actuator and stiffness on the geometry. This may be a reasonable simplification as prior works [48] have shown geometry along with properly-set actuator and stiffness (we take manual efforts to design proper mapping from geometry to actuator and stiffness in this work) roughly reflect the performance of a soft robot design. For better generality, one potential remedy is to optimize actuator and stiffness independently from the geometry generated by the diffusion model, i.e., apply DiffuseBot and do direct optimization on actuator and stiffness afterward or at the same time. Another interesting direction may be, for actuators, to leverage part-based models [29] to decompose a holistic geometry into parts (or different actuator regions in soft robots).

References

- [1] Anurag Ajay, Yilun Du, Abhi Gupta, Joshua Tenenbaum, Tommi Jaakkola, and Pulkit Agrawal. Is conditional generative modeling all you need for decision-making? *arXiv preprint arXiv:2211.15657*, 2022.
- [2] Moritz Bächer, Espen Knoop, and Christian Schumacher. Design and control of soft robots using differentiable simulation. *Current Robotics Reports*, 2(2):211–221, 2021.
- [3] Nicholas W Bartlett, Michael T Tolley, Johannes TB Overvelde, James C Weaver, Bobak Mosadegh, Katia Bertoldi, George M Whitesides, and Robert J Wood. A 3d-printed, functionally graded soft robot powered by combustion. *Science*, 349(6244):161–165, 2015.
- [4] James M Bern, Kai-Hung Chang, and Stelian Coros. Interactive design of animated plushies. *ACM Transactions on Graphics (TOG)*, 36(4):1–11, 2017.
- [5] Jagdeep Bhatia, Holly Jackson, Yunsheng Tian, Jie Xu, and Wojciech Matusik. Evolution gym: A large-scale benchmark for evolving soft robots. *Advances in Neural Information Processing Systems*, 34:2201–2214, 2021.
- [6] Xiang’Anthony’ Chen, Stelian Coros, Jennifer Mankoff, and Scott E Hudson. Encore: 3d printed augmentation of everyday objects with printed-over, affixed and interlocked attachments. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology*, pages 73–82, 2015.
- [7] Nicholas Cheney, Jeff Clune, and Hod Lipson. Evolved electrophysiological soft robots. In *Artificial life conference proceedings*, pages 222–229. MIT Press One Rogers Street, Cambridge, MA 02142-1209, USA journals-info . . . , 2014.
- [8] Nick Cheney, Robert MacCurdy, Jeff Clune, and Hod Lipson. Unshackling evolution: evolving soft robots with multiple materials and a powerful generative encoding. *ACM SIGEVOLUTION*, 7(1):11–23, 2014.
- [9] Cheng Chi, Siyuan Feng, Yilun Du, Zhenjia Xu, Eric Cousineau, Benjamin Burchfiel, and Shuran Song. Diffusion Policy: Visuomotor Policy Learning via Action Diffusion. *arXiv preprint arXiv:2303.04137*, 2023.
- [10] Francesco Corucci, Nick Cheney, Francesco Giorgio-Serchi, Josh Bongard, and Cecilia Laschi. Evolving soft locomotion in aquatic and terrestrial environments: effects of material properties and environmental transitions. *Soft robotics*, 5(4):475–495, 2018.
- [11] Raphael Deimel, Patrick Irmisch, Vincent Wall, and Oliver Brock. Automated co-design of soft hand morphology and control strategy for grasping. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1213–1218. IEEE, 2017.
- [12] Prafulla Dhariwal and Alexander Quinn Nichol. Diffusion models beat GANs on image synthesis. In *Advances in Neural Information Processing Systems*, 2021.
- [13] Tao Du, Kui Wu, Pingchuan Ma, Sebastien Wah, Andrew Spielberg, Daniela Rus, and Wojciech Matusik. Diffpd: Differentiable projective dynamics. *ACM Transactions on Graphics (TOG)*, 41(2):1–21, 2021.
- [14] Yilun Du, Conor Durkan, Robin Strudel, Joshua B Tenenbaum, Sander Dieleman, Rob Fergus, Jascha Sohl-Dickstein, Arnaud Doucet, and Will Grathwohl. Reduce, reuse, recycle: Compositional generation with energy-based diffusion models and mcmc. *arXiv preprint arXiv:2302.11552*, 2023.
- [15] Yilun Du and Igor Mordatch. Implicit generation and generalization in energy-based models. *arXiv preprint arXiv:1903.08689*, 2019.
- [16] Yilun Du, Mengjiao Yang, Bo Dai, Hanjun Dai, Ofir Nachum, Joshua B Tenenbaum, Dale Schuurmans, and Pieter Abbeel. Learning universal policies via text-guided video generation. *arXiv e-prints*, pages arXiv–2302, 2023.

- [17] Bin Fang, Fuchun Sun, Linyuan Wu, Fukang Liu, Xiangxiang Wang, Haiming Huang, Wenbing Huang, Huaping Liu, and Li Wen. Multimode grasping soft gripper achieved by layer jamming structure and tendon-driven mechanism. *Soft Robotics*, 9(2):233–249, 2022.
- [18] Chrisantha Fernando, Dylan Banarse, Malcolm Reynolds, Frederic Besse, David Pfau, Max Jaderberg, Marc Lanctot, and Daan Wierstra. Convolution by evolution: Differentiable pattern producing networks. In *Proceedings of the Genetic and Evolutionary Computation Conference 2016*, pages 109–116, 2016.
- [19] Rinon Gal, Yuval Alaluf, Yuval Atzmon, Or Patashnik, Amit H Bermano, Gal Chechik, and Daniel Cohen-Or. An image is worth one word: Personalizing text-to-image generation using textual inversion. *arXiv preprint arXiv:2208.01618*, 2022.
- [20] Debkalpa Goswami, Shuai Liu, Aniket Pal, Lucas G Silva, and Ramses V Martinez. 3d-architected soft machines with topologically encoded motion. *Advanced Functional Materials*, 29(24):1808713, 2019.
- [21] Sehoon Ha, Peng Xu, Zhenyu Tan, Sergey Levine, and Jie Tan. Learning to walk in the real world with minimal human effort. *arXiv preprint arXiv:2002.08550*, 2020.
- [22] Jonathan Hiller and Hod Lipson. Automatic design and manufacture of soft robots. *IEEE Transactions on Robotics*, 28(2):457–466, 2011.
- [23] 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. *arXiv preprint arXiv:2210.02303*, 2022.
- [24] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020.
- [25] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022.
- [26] Yuanming Hu, Jiancheng Liu, Andrew Spielberg, Joshua B Tenenbaum, William T Freeman, Jiajun Wu, Daniela Rus, and Wojciech Matusik. Chainqueen: A real-time differentiable physical simulator for soft robotics. In *2019 International conference on robotics and automation (ICRA)*, pages 6265–6271. IEEE, 2019.
- [27] Hyunki In, Useok Jeong, Haemin Lee, and Kyu-Jin Cho. A novel slack-enabling tendon drive that improves efficiency, size, and safety in soft wearable robots. *IEEE/ASME Transactions on Mechatronics*, 22(1):59–70, 2016.
- [28] Michael Janner, Yilun Du, Joshua Tenenbaum, and Sergey Levine. Planning with diffusion for flexible behavior synthesis. In *International Conference on Machine Learning*, 2022.
- [29] Adrien Kaiser, Jose Alonso Ybanez Zepeda, and Tamy Boubekeur. A survey of simple geometric primitives detection methods for captured 3d data. In *Computer Graphics Forum*, volume 38, pages 167–196. Wiley Online Library, 2019.
- [30] Byungchul Kim, Useok Jeong, Brian Byunghyun Kang, and Kyu-Jin Cho. Slider-tendon linear actuator with under-actuation and fast-connection for soft wearable robots. *IEEE/ASME Transactions on Mechatronics*, 26(6):2932–2943, 2021.
- [31] Sam Kriegman, Douglas Blackiston, Michael Levin, and Josh Bongard. A scalable pipeline for designing reconfigurable organisms. *Proceedings of the National Academy of Sciences*, 117(4):1853–1859, 2020.
- [32] Sam Kriegman, Douglas Blackiston, Michael Levin, and Josh Bongard. Kinematic self-replication in reconfigurable organisms. *Proceedings of the National Academy of Sciences*, 118(49):e2112672118, 2021.
- [33] Yifei Li, Tao Du, Kui Wu, Jie Xu, and Wojciech Matusik. Diffcloth: Differentiable cloth simulation with dry frictional contact. *ACM Transactions on Graphics (TOG)*, 42(1):1–20, 2022.

- [34] Nan Liu, Shuang Li, Yilun Du, Antonio Torralba, and Joshua B Tenenbaum. Compositional visual generation with composable diffusion models. *arXiv preprint arXiv:2206.01714*, 2022.
- [35] Pingchuan Ma, Tao Du, John Z Zhang, Kui Wu, Andrew Spielberg, Robert K Katzschmann, and Wojciech Matusik. Diffaqua: A differentiable computational design pipeline for soft underwater swimmers with shape interpolation. *ACM Transactions on Graphics (TOG)*, 40(4):1–14, 2021.
- [36] Jonàs Martínez, Jérémie Dumas, and Sylvain Lefebvre. Procedural voronoi foams for additive manufacturing. *ACM Transactions on Graphics (TOG)*, 35(4):1–12, 2016.
- [37] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4460–4470, 2019.
- [38] Thomas Morzadec, Damien Marcha, and Christian Duriez. Toward shape optimization of soft robots. In *2019 2nd IEEE International Conference on Soft Robotics (RoboSoft)*, pages 521–526. IEEE, 2019.
- [39] Alex Nichol, Heewoo Jun, Prafulla Dhariwal, Pamela Mishkin, and Mark Chen. Point-e: A system for generating 3d point clouds from complex prompts. *arXiv preprint arXiv:2212.08751*, 2022.
- [40] Songyou Peng, Chiyu "Max" Jiang, Yiyi Liao, Michael Niemeyer, Marc Pollefeys, and Andreas Geiger. Shape as points: A differentiable poisson solver. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- [41] Yi-Ling Qiao, Junbang Liang, Vladlen Koltun, and Ming C Lin. Scalable differentiable physics for learning and control. *arXiv preprint arXiv:2007.02168*, 2020.
- [42] Yiling Qiao, Junbang Liang, Vladlen Koltun, and Ming Lin. Differentiable simulation of soft multi-body systems. *Advances in Neural Information Processing Systems*, 34:17123–17135, 2021.
- [43] 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.
- [44] Aditya Ramesh, Mikhail Pavlov, Scott Gray Gabriel Goh, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. *arXiv preprint arXiv:2102.12092*, 2021.
- [45] Stéphane Ross, Geoffrey Gordon, and Drew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *International Conference on Artificial Intelligence and Statistics*, 2011.
- [46] 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. *arXiv preprint arXiv:2205.11487*, 2022.
- [47] Charles Schaff, Audrey Sedal, and Matthew R Walter. Soft robots learn to crawl: Jointly optimizing design and control with sim-to-real transfer. *arXiv preprint arXiv:2202.04575*, 2022.
- [48] Arne Schneuing, Yuanqi Du, Arian Jamasb Charles Harris, Ilia Igashov, Weitao Du, Tom Blundell, Pietro Lió, Carla Gomes, Michael Bronstein Max Welling, and Bruno Correia. Structure-based drug design with equivariant diffusion models. *arXiv preprint arXiv:2210.02303*, 2022.
- [49] Uriel Singer, Adam Polyak, Thomas Hayes, Xi Yin, Jie An, Songyang Zhang, Qiyuan Hu, Harry Yang, Oron Ashual, Oran Gafni, et al. Make-a-video: Text-to-video generation without text-video data. *arXiv preprint arXiv:2209.14792*, 2022.

- [50] Lawrence Smith, Travis Hainsworth, Jacob Haimes, and Robert MacCurdy. Automated synthesis of bending pneumatic soft actuators. In *2022 IEEE 5th International Conference on Soft Robotics (RoboSoft)*, pages 358–363. IEEE, 2022.
- [51] Lawrence Smith and Robert MacCurdy. Soroforge: End-to-end soft actuator design. *IEEE Transactions on Automation Science and Engineering*, 2023.
- [52] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*, pages 2256–2265. PMLR, 2015.
- [53] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020.
- [54] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 32, 2019.
- [55] Andrew Spielberg, Alexander Amini, Lillian Chin, Wojciech Matusik, and Daniela Rus. Co-learning of task and sensor placement for soft robotics. *IEEE Robotics and Automation Letters*, 6(2):1208–1215, 2021.
- [56] Andrew Spielberg, Tao Du, Yuanming Hu, Daniela Rus, and Wojciech Matusik. Advanced soft robot modeling in chainqueen. *Robotica*, 41(1):74–104, 2023.
- [57] Andrew Spielberg, Allan Zhao, Yuanming Hu, Tao Du, Wojciech Matusik, and Daniela Rus. Learning-in-the-loop optimization: End-to-end control and co-design of soft robots through learned deep latent representations. *Advances in Neural Information Processing Systems*, 32, 2019.
- [58] Kenneth O Stanley. Compositional pattern producing networks: A novel abstraction of development. *Genetic programming and evolvable machines*, 8:131–162, 2007.
- [59] Michael T Tolley, Robert F Shepherd, Michael Karpelson, Nicholas W Bartlett, Kevin C Galloway, Michael Wehner, Rui Nunes, George M Whitesides, and Robert J Wood. An untethered jumping soft robot. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 561–566. IEEE, 2014.
- [60] Merel Van Diepen and Kristina Shea. A spatial grammar method for the computational design synthesis of virtual soft locomotion robots. *Journal of Mechanical Design*, 141(10), 2019.
- [61] Tsun-Hsuan Wang, Pingchuan Ma, Andrew Everett Spielberg, Zhou Xian, Hao Zhang, Joshua B. Tenenbaum, Daniela Rus, and Chuang Gan. Softzoo: A soft robot co-design benchmark for locomotion in diverse environments. In *The Eleventh International Conference on Learning Representations*, 2023.
- [62] Tian Xie, Xiang Fu, Octavian-Eugen Ganea, Regina Barzilay, and Tommi S Jaakkola. Crystal diffusion variational autoencoder for periodic material generation. In *International Conference on Learning Representations*, 2021.
- [63] Minkai Xu, Lantao Yu, Yang Song, Chence Shi, Stefano Ermon, , and Jian Tang. GeoDiff: A geometric diffusion model for molecular conformation generation. In *International Conference on Learning Representations*, 2021.
- [64] Ye Yuan, Jiaming Song, Umar Iqbal, Arash Vahdat, and Jan Kautz. Physdiff: Physics-guided human motion diffusion model. *arXiv preprint arXiv:2212.02500*, 2022.
- [65] Xiaohui Zeng, Arash Vahdat, Francis Williams, Zan Gojcic, Or Litany, Sanja Fidler, and Karsten Kreis. Lion: Latent point diffusion models for 3d shape generation. *arXiv preprint arXiv:2210.06978*, 2022.
- [66] Linqi Zhou, Yilun Du, and Jiajun Wu. 3d shape generation and completion through point-voxel diffusion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5826–5835, 2021.