

Scaling up Robot Skill Learning with Generative Simulation

Pushkal Katara

CMU-RI-TR-24-35

July 1, 2024



The Robotics Institute
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA

Thesis Committee:

Professor Katerina Fragkiadaki, *chair*
Professor David Held
Zhou Xian

*Submitted in partial fulfillment of the requirements
for the degree of Master of Science in Robotics.*

Copyright © 2024 Pushkal Katara. All rights reserved.

To my parents, Geetanjali and Ravindra Katara

Abstract

Generalist robots need to learn a wide variety of skills to perform diverse tasks across multiple environments. Current robot training pipelines rely on humans to either provide kinesthetic demonstrations or program simulation environments with manually-designed reward functions for reinforcement learning. Such human involvement is an important bottleneck towards scaling up robot learning across diverse tasks and environments.

In this thesis, we present Generation to Simulation (Gen2Sim), a method for scaling up robot skill learning in simulation by automating generation of 3D assets, task descriptions, task decompositions and reward functions using large pre-trained generative models of language and vision. We generate 3D assets for simulation by lifting open-world 2D object-centric images to 3D using image diffusion models and querying LLMs to determine plausible physics parameters. Given URDF files of generated and human-developed assets, we chain-of-thought prompt LLMs to map these to relevant task descriptions, temporal decompositions, and corresponding python reward functions for reinforcement learning. We show Gen2Sim succeeds in learning policies for diverse long horizon tasks, where reinforcement learning with non temporally decomposed reward functions fails.

Gen2Sim provides a viable path for scaling up robot skill learning in simulation, both by diversifying and expanding task and environment development, and by facilitating the discovery of reinforcement-learned behaviors through temporal task decomposition in RL. Our work contributes hundreds of simulated assets, tasks and demonstrations, taking a step towards fully autonomous robotic manipulation skill acquisition in simulation.

Acknowledgments

During my two years at CMU, I have experienced tremendous growth, both professionally and personally. This journey would not have been possible without the extraordinary support of my advisor, labmates, friends, and the broader CMU community, to whom I am deeply grateful.

First and foremost, I wish to express my profound gratitude to my advisor, Prof. Katerina Fragkiadaki, for her unwavering support and mentorship. As a fledgling master’s student, she took me under her wing and guided me toward becoming the researcher I am today. Prof. Fragkiadaki instilled in me the importance of pursuing challenging and impactful questions. Her mantra, “Research is all about finding the right problem to work on and relentlessly work towards it,” has become my guiding principle. Her positive energy, enthusiasm, and cheerful attitude have been a constant source of motivation. I am immensely grateful to Prof. Fragkiadaki for believing in me and allowing me to gain invaluable experience under her guidance.

I also wish to sincerely thank Prof. David Held for serving on my thesis committee, offering insightful feedback, and posing thought-provoking questions during my defense.

I am deeply grateful to Xian for his unwavering support throughout my time in the lab. Xian’s wisdom, calm demeanor, and friendly nature made our collaboration both productive and enriching. His meticulous attention to detail, commitment to research, and rigorous coding practices have significantly enhanced my skills as a researcher. I will always cherish our sessions debugging robot simulation and hardware together.

To Ayush, my heartfelt thanks for your steadfast support, friendship, and mentorship throughout my academic journey. Collaborating with you has been both productive and immensely enjoyable, and I truly value the moments we’ve spent working together. Your encouragement, especially during challenging times, has been a constant source of strength.

Engaging with my lab members during group meetings, reading sessions, and spontaneous discussions has been instrumental in my development as a researcher. I have gained valuable insights from the meticulous attention to detail, strong work ethic, and impressive presentation skills

of the senior lab members. I would like to extend my heartfelt gratitude to Nikos, Tsung-Wei, and Gabe. Our countless hours of collaboration, shared frustrations over the relentless pace of research, and moments of lightheartedness have made the journey more manageable.

I also want to express my sincere gratitude to Theophile Gervet, Adam Harley, Wen-Hsuan Chu, Lei Ke, Yufei Wang, Brian Yang, Mayank Singh, and Mihir Prabhudesai. Your collaboration and support have been instrumental in making this thesis possible.

My deepest thanks go to Utkarsha, whose unwavering support has been the foundation of my journey through this master's program. Your encouragement has driven me to reach new heights, unlocking my potential and fostering both personal and professional growth. I am truly grateful for your unwavering belief in me and the continuous support you've provided throughout our journey together.

To say that my friends helped me reach the finish line would be an understatement. I am incredibly grateful for the amazing friends I made at CMU—Pranay Gupta, Sarthak Bhagat, Shagun Uppal, Dvij Kalaria, Shreya Sharma, Anish Madan, Poorvi Hebbar, Bharath Raj, and Sriram Narayanan. You all made the intense grind far less stressful.

Lastly, I extend my deepest gratitude to my parents, Geetanjali and Ravindra Katara, for their unconditional love and support. They shielded me from life's challenges, encouraged me to pursue my dreams, and inspired me to persevere even in times of uncertainty. I am also deeply thankful to my brother, Pranjali Katara, for being my constant pillar of strength. Your encouragement during moments of doubt kept me from settling for less. My heartfelt thanks go to my grandfather, Ramesh Chandra Katara, whose support was instrumental in my pursuit of higher studies. Though living far from my family has been difficult, their spirit and brilliance continue to guide me. I dedicate this thesis to them with great honor and love, as a tribute to their unwavering belief in me and their lasting influence on my life.

Contents

1	Introduction	1
2	Related Work	5
2.1	Large Language Models for task and motion planning in robotics . . .	5
2.2	Automating 3D asset creation with generative models	6
2.3	Procedural demonstration generation using symbolic planners	6
2.4	Simulation environments for robotic skill learning	7
3	Approach	9
3.1	3D Asset Generation	9
3.1.1	Image diffusion models	10
3.1.2	Image-to-3D Mesh using Score Distillation Sampling	11
3.1.3	Texture generation	12
3.1.4	Generating plausible physical properties	12
3.2	Task Generation, Temporal Decomposition and Reward Function Prediction	12
3.3	Sequential Reinforcement Learning for Long Horizon Tasks	15
4	Experiments	17
4.1	Experiments	17
4.1.1	Asset Generation	17
4.1.2	Automated Skill Learning	19
4.1.3	Twin environment construction and sim-to-real world transfer	20
5	Discussion	21
5.1	Limitations	21
5.2	Conclusion	22
	Bibliography	23

List of Figures

1.1	Gen2Sim is an automated generative pipeline of assets, tasks, task decompositions, and rewards functions for autonomous robotic skill learning in simulation. Here we show 32 generated tasks, concerning affordances of diverse types of object assets and their combinations.	2
3.1	Gen2Sim components. Gen2Sim generates 3D assets by lifting object-centric 2D images to 3D. It then uses both generated assets and assets obtained from other publicly available datasets to populate scene environments. Afterwards, it queries LLMs to generate meaningful task descriptions for the assembled scenes, as well as decompose the generated task descriptions to sub-tasks and their reward functions.	10
4.1	3D asset generation from Gen2Sim, RealFusion [50] and Make-It-3D [67]. Gen2Sim uses a view and camera pose conditioned image generative model during score distillation, which helps generate more accurate 3D geometry in comparison to the baselines.	18
4.2	Twin environments constructed and generated tasks for sim-to-real transfer. Left: real-world. Right: simulated.	19

List of Tables

4.1 **Size and physics parameter generated by LLMs** for a number of
generated assets. 18

Chapter 1

Introduction

Scaling up training data has been a driving force behind the recent revolutions in language modeling [4], image understanding [57], speech recognition [58], image generation [60], to name a few. This begs the question: can we scale up robot data to enable a similar revolution in robotic skill learning? One way to scale robot data is in the real world, by having multiple robots explore [38] or by having humans provide kinesthetic demonstrations [2, 3, 65]. This is a promising direction; however, safety concerns and wear and tear of the robots hinder robot exploration in the real-world, and collecting kinesthetic demonstrations scales poorly as it is time-consuming and labor-intensive [2]. Another way to scale robot data is in simulation, by developing simulated environments, defining tasks and their reward functions, and training robot policies with reinforcement learning, augmenting visuals and physics parameters to facilitate transfer of policies to the real world [27]. Such sim2real paradigm has seen recent successes in robot locomotion [16, 32, 36, 37], object re-orientation [9, 55], and drone flight [34]. These examples, though very important and exciting, are still fairly isolated.

A central bottleneck towards scaling up simulation environments and tasks is the laborious manual effort needed for developing the visuals and physics of assets, their spatial arrangement and configurations, the development of task definition and reward functions, or the collection of programmatic demonstrations. Tremendous resources have been invested in developing simulators for autonomous vehicles [14], warehouse robots, articulated objects [81], home environments [19, 62, 66], etc.,

1. Introduction



Figure 1.1: **Gen2Sim** is an automated generative pipeline of assets, tasks, task decompositions, and rewards functions for autonomous robotic skill learning in simulation. Here we show 32 generated tasks, concerning affordances of diverse types of object assets and their combinations.

many of which are proprietary and not open-sourced. Given these considerations, an important question naturally arises: How can we minimize manual effort in simulation development for diverse robotic skill learning?

In this thesis, we explore automating the development of simulation environments, tasks and rewards for robot skill learning, by building upon latest advances in large pre-trained generative models of images and language. Our system strives to automate all stages of robot learning: from generating 3D assets, textures, and physics parameters, to generating task descriptions and reward functions, leading to automated skill learning in diverse scenarios, as shown in Figure 1.1. This generative pipeline was first proposed in a recent position paper [78], described as a promising pathway towards generating diverse data for generalist robot learning. In this thesis, we present Gen2Sim, the first attempt and realization of such a generative paradigm. We automate 3D object asset generation by combining image diffusion models for 3D mesh and texture generation, and LLMs for querying physical parameters information. We showcase how LLMs and image generative models can diversify the

appearances and behaviors of assets by producing plausible ranges of textures, sizes and physical parameters, achieving “intelligent” domain diversification. We automate task description, task decomposition and reward function generation by few-shot prompting of LLMs to generate language descriptions for semantically meaningful tasks, concerning affordances of existing and generated 3D assets, articulated or not, alongside their reward functions. Gen2Sim is able to generate numerous object assets and task variations without any human involvement beyond few LLM prompt designs. We successfully train RL policies using our auto-generated tasks and reward functions. We also demonstrate the usefulness of our simulation-trained policies, by constructing digital-twin environments from given real scenes, allowing a robot to practice skills in the twin simulator and deploying it back to the real world to execute the task.

In summary, we make the following contributions:

- We show how pre-trained generative models of images and language can help automate 3D asset generation and diversification, task description generation, task decomposition and reward function generation that supports reinforcement learning of long horizon tasks in simulation with minimal human involvement.
- We deploy our method to generate hundreds of assets, and hundreds of manipulation tasks, their decompositions and their reward functions, for both human-developed and automatically generated object assets.

For code, videos and qualitative video results, please visit our project website: <https://gen2sim.github.io/>.

1. Introduction

Chapter 2

Related Work

2.1 Large Language Models for task and motion planning in robotics

Large language models (LLMs) map instructions to language subgoals [28, 29, 82, 86] or action programs [40] with appropriate plan-like or program-like prompts. LLMs trained from Internet-scale text have shown impressive zero-shot reasoning capabilities for a variety of downstream language tasks [4] when prompted appropriately, without any weight fine-tuning [5, 7, 44, 76]. LLMs were used to generate task curricula and predict skills to execute in Minecraft worlds [41, 72, 85]. Following the seminal work of Code as Policies, many works map language to programs over given skills [24] or hand-designed motion planners [30]. Our work instead maps task descriptions into task decompositions and reward functions, to guide reinforcement learning in simulation, to discover behaviours that achieve the generated tasks. Work of [83] also uses language for predicting reward functions for robot locomotion, but does not consider task generation and decomposition or interaction with objects. Our work is the first to use LLMs for task decomposition and reward generation, as well as asset generation.

2.2 Automating 3D asset creation with generative models

The traditional process of creating 3D assets typically involves multiple labor-intensive stages, including geometry modeling, shape baking, UV mapping, material creation, texturing and physics parameter estimation, where different software tools and the expertise of skilled artists are often required. It is thus desirable to automate 3D asset generation to automatically generate high-quality assets that support realistic rendering under arbitrary views and have plausible physical behaviours during force application and contacts. The lack of available 3D data and the abundance of 2D image data have stimulated interest in learning 3D models from 2D image generators [6, 53]. The availability of strong 2D image generative models based on diffusion led to high-quality 3D models from text descriptions [8, 42, 56] or single 2D images using the diffusion model as a 2D prior [50, 64, 67]. In this work, instead of a text-conditioned model, we use a view and relative pose conditioned image generative model, which we found to provide better prior for score distillation. Some methods attempt to use videos of assets and differentiable simulations to estimate their physics parameters and/or adapt the simulation environment, in an attempt to close the simulation to reality gap [25, 26, 74]. Our effort is complementary to these works.

2.3 Procedural demonstration generation using symbolic planners

Many recent works procedurally generate scenes and demonstration trajectories using planners that have access to privileged information to solve the task, and distill the demonstration solutions into learning-based policies that operate directly from pixel or point-cloud input [13, 15, 49]. Task and motion planners [33, 46, 51, 70] use predefined symbolic rules and known dynamics models, and infer discrete task plans given instruction with lookahead logic search [20, 33, 33, 46, 51, 70]. These methods predominantly rely on manually-specified symbolic transition rules, planning domains, and grounding, which limits their applicability. Indeed, works of [13, 49] demonstrate their results on relatively simple multi-object box stacking tasks. Scene procedural

generation in the aforementioned works [13, 49, 52] entails randomizing locations and number of given 3D models under weak supervision from a human that defines the task and the possible location candidates. In contrast, we unleash the common sense knowledge and reasoning capabilities provided by LLMs and use them to suggest task descriptions, task decompositions, and reward functions. We then use reinforcement learning to discover solution trajectories instead of TAMP-based search.

2.4 Simulation environments for robotic skill learning

In recent years, improving simulators for robot manipulation has attracted increasingly more attention. Many robotic manipulation environments and benchmarks [1, 35, 81] are built on top of either PyBullet [12] or MuJoCo [69] as their underlying physics engines, which mainly support rigid-body manipulation [22, 23, 71, 77, 79]. Recently, environments supporting soft-body manipulation ([18, 43, 47, 75, 80, 81]) provide capabilities for simulating deformable robots, objects and fluids. Our automated asset and task generation are not tied to any specific simulation platforms and can be used with any of them.

2. Related Work

Chapter 3

Approach

Gen2Sim generates 3D assets from object-centric images using image diffusion models and predicts physical parameters for them using LLMs (Section 3.1). It then prompts LLMs to generate language task descriptions and corresponding reward functions for each generated or human-developed asset, suitable to their affordances (Section 3.2). Finally, we train RL policies in the generated environments using the generated reward functions. We additionally show the applicability of the simulation-trained policy by constructing digital twin environment in simulation, and deploy the trained trajectory in the real world (Section 3.3). See Figure 3.1 for our method overview.

3.1 3D Asset Generation

Gen2Sim automates 3D asset generation by mapping 2D images of objects to textured 3D meshes with plausible physics parameters. The images can be 1) real images taken in the robot’s environment, 2) real images provided by Google search under relevant category names, e.g., “*avocado*”, or 3) images generated by pre-trained text-conditioned diffusion models, such as stable diffusion [61], prompted appropriately to generate uncluttered images of the relevant objects, e.g., “*an image of an individual avocado*”. We query GPT-4 [54] for a list of object categories relevant for manipulation tasks to search online for or to generate, instead of manually designing it. Please, visit our project site for a detailed list of the objects we generated. Given a real or generated 2D image of an object, we lift it to a 3D model by minimizing re-projection

3. Approach

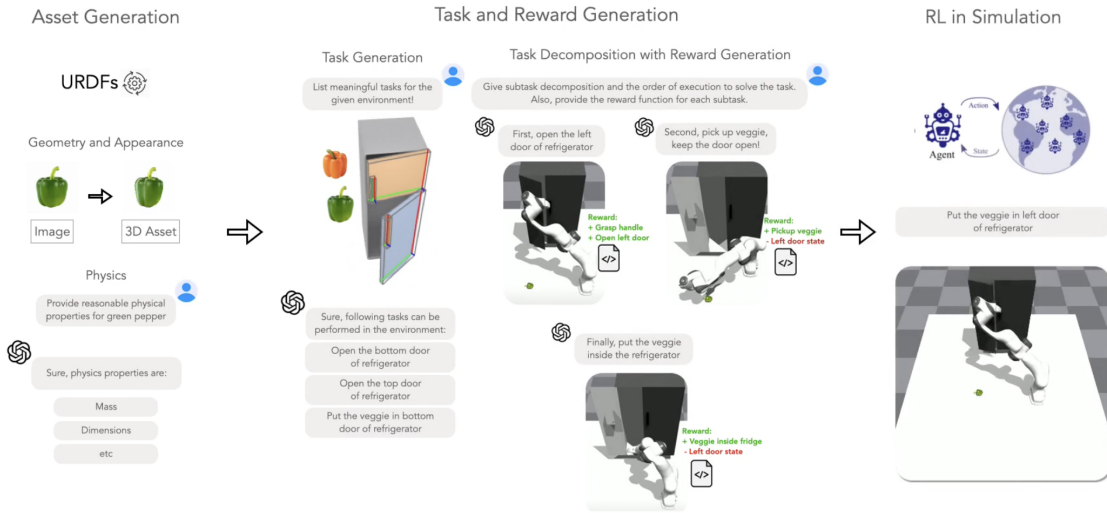


Figure 3.1: **Gen2Sim components.** Gen2Sim generates 3D assets by lifting object-centric 2D images to 3D. It then uses both generated assets and assets obtained from other publicly available datasets to populate scene environments. Afterwards, it queries LLMs to generate meaningful task descriptions for the assembled scenes, as well as decompose the generated task descriptions to sub-tasks and their reward functions.

error and maximizing likelihood of its image renderings using a diffusion model [8, 56]. We provide background on image diffusion models below, before we describe our 3D model fitting approach.

3.1.1 Image diffusion models

A diffusion model learns to model a probability distribution $p(x)$ by inverting a process that gradually adds noise to the image x . The diffusion process is associated with a variance schedule $\{\beta_t \in (0, 1)\}_{t=1}^T$, which defines how much noise is added at each time step. The noisy version of sample x at time t can then be written $x_t = \sqrt{\bar{\alpha}_t}x + \sqrt{1 - \bar{\alpha}_t}\epsilon$ where $\epsilon \sim \mathcal{N}(\cdot, \cdot)$, is a sample from a Gaussian distribution (with the same dimensionality as x), $\alpha_t = 1 - \beta_t$, and $\bar{\alpha}_t = \prod_{i=1}^t \alpha_i$. One then learns a denoising neural network $\hat{\epsilon} = \epsilon_\phi(x_t; t)$ that takes as input the noisy image x_t and the noise level t and tries to predict the noise component ϵ . Diffusion models can be easily extended to draw samples from a distribution $p(x|)$ conditioned on a prompt

, where ϕ can be a text description, a camera pose, and image semantic map, *etc* [39, 60, 84]. Conditioning on the prompt can be done by adding ϕ as an additional input of the network ϵ_ϕ . For 3D lifting, we build on Zero-1-to-3 [45], a diffusion model for novel object view synthesis that conditions on an image view of an object and a relative camera rotation around the object to generate plausible images for the target object viewpoint, $\pi = [I_1, \pi]$. It is trained on a large collection $\mathcal{D}' = \{(x^i, \pi^i)\}_{i=1}^N$ of images paired with views and relative camera orientations as conditioning prompt by minimizing the loss:

$$\mathcal{L}_{\text{diff}}(\phi; \mathcal{D}') = \frac{1}{|\mathcal{D}'|} \sum_{x^i, \pi^i \in \mathcal{D}'} \|\epsilon_\phi(\sqrt{\bar{\alpha}_t}x^i + \sqrt{1 - \bar{\alpha}_t}\epsilon, \pi^i, t) - \epsilon\|^2.$$

3.1.2 Image-to-3D Mesh using Score Distillation Sampling

Given an image and relative camera pose 2D diffusion model $p(I|[I_0, \pi])$, we extract from it a 3D rendition of the input image I_0 , represented by a differential 3D representation using Score Distillation Sampling (SDS) [56, 73]. We do so by randomly sampling a camera pose π , rendering a corresponding view I_π , assessing the likelihood of the view based on a diffusion model $p(I_\pi|[I_0, \pi])$, and updating the differentiable 3D representation to increase the likelihood of the generated view based on the model. Specifically, the diffusion model is frozen and the 3D model is updated as:

$$\nabla(\theta)\mathcal{L}_{SDS}(\theta; \pi, I_0, t) = \mathbb{E}_{t, \epsilon}[w(t) (\epsilon_\phi(a_t I + \sigma_t \epsilon; t, \pi) - \epsilon) \cdot \nabla_\theta I],$$

where $I = R(\theta, \pi)$ is the image rendered from a given viewpoint π . The loss we use to backpropagate to the 3D model parameters θ includes an image re-projection loss for the camera viewpoint of the input image, and score distillation for the other views, using a pre-trained view and pose conditioned image diffusion model of [45] to measure 2D image likelihood. We use a two-stage fitting, wherein the first stage an instantNGP NeRF representation is used, similar to RealFusion [50], and in the second stage a mesh-based representation is initialized from the NeRF and finetuned differentially, similar to Fantasia3D [8].

3.1.3 Texture generation

We augment the textures of our generated assets using the method of TEXTure [59] which iteratively edits a mesh’s texture by rendering the mesh from different viewpoints and updating the rendered 2D images. While domain randomization [68] randomly re-textures simulated assets, TEXTure produces diverse yet plausible texture augmentations.

3.1.4 Generating plausible physical properties

The visual and collision parameters of an asset are generated from the Image-to-Mesh pipeline discussed above. To define 3D sizes and physics parameters for the generated 3D meshes, we query GPT-4 regarding the range of plausible width, height, and depth for each object, and the range of its mass given its category. We then scale the generated 3D mesh based on the generated size range. We feed the mass and 3D mesh information to MeshLab [11] to get the inertia matrix for the asset. Our prompts for querying GPT for mass and 3D object size can be found on our website. We wrap the generated mesh information, its semantic name, as well as the physical parameters into URDF files to be loaded into our simulator.

3.2 Task Generation, Temporal Decomposition and Reward Function Prediction

Given either generated assets or assets obtained from publically available datasets, we prompt LLMs to generate meaningful manipulation tasks considering their affordances, to decompose these tasks into subtasks when possible, and to generate reward functions for each subtask. We train reinforcement learning policies for each (sub)task using the generated reward functions, and then chain them together to solve long horizon tasks. Our LLM prompts contain the following sections:

1. Asset descriptions. We use combinations of assets we generate using the method of Section 3.1, as well as articulated assets from PartNet Mobility [81] and GPartNet dataset [21]. We populate our simulation environment with randomly sampled assets. Then, we extract information from the URDF files including link names, joint types

and limits using automated scripts. For example, an asset `microwave` has parts [door, handle, and body], and joint [door-joint] of type `revolute` with a joint position range [0, 1]. We then describe the extracted configurations of the assets to the LLM, as shown below:

```
The environment contains the following assets:
1. asset_name: "microwave"
   part_configuration:
     Part 1: "body"
     Part 2: "door"
       - link_name: "link_0"
       - joint_name: "joint_0"
       - joint_type: "revolute"
       - limit: [0, 1]
     Part 3: "handle"
       - link_name: "handle_0"
       - joint_name: "handlejoint_0"
       - joint_type: "fixed"
2. asset_name: "cup"
   part_configuration:
     Part 1: cup
       - link_name: "base"
       - joint_name: "base_joint"
       - joint_type: "fixed"
```

2. Instructions. These include function APIs that can be used by the LLM to query the pose of the robot end-effector, as well as different assets in the given environment:

```
Available APIs from the simulator are:
# returns the pose of the link
get_pose_by_link_name(asset_name, link_name)
# returns the pose of the robot gripper
get_robot_gripper_pose(asset_name, link_name)
# returns the state of the joint
get_state_by_joint_name(asset_name, joint_name)
# returns the limit of the joint
get_limits_by_joint_name(asset_name, joint_name)
```

Note:

1. Only use the available APIs from the simulator.
2. Generate the reward function code snippets in Python.

3. Examples of task descriptions and decompositions. These are question-answer pairs that demonstrate task descriptions and their temporal decompositions.

```
List meaningful manipulation tasks that can be performed
in this environment. Give subtask decomposition and the
order of execution to solve the task. Also, provide the
reward function for each subtask.
```

3. Approach

The following tasks can be performed **in** this environment:

1. Open the Microwave Door
2. Close the Microwave Door
3. Pick Cup
4. Place Cup
5. Put the Cup **in** the Microwave

This task needs to be decomposed into sub-tasks:

- Open the Microwave
- Pick Cup
- Place the Cup **in** the Microwave

4. Examples of reward functions. These are task to reward function pairs that present demonstrations of how tasks can be translated to reward functions, as shown below:

```
Task: OpenMicrowaveDoor
Task Description: open the door of the microwave
'''
def compute_reward(env):
    # reward function
    door_handle_pose = env.get_pose_by_link_name("microwave", "handle_0")
    gripper_pose = env.get_robot_gripper_pose()
    distance_gripper_to_handle = torch.norm(door_handle_pose - gripper_pose, dim=-1)
    door_state = env.get_state_by_joint_name("microwave", "joint_0")
    cost = distance_gripper_to_handle - door_state
    reward = - cost

    # success condition
    target_door_state = env.get_limits_by_joint_name("microwave", "joint_0")["upper"]
    success = torch.abs(door_state - target_door_state) < 0.1

    return reward, success
'''
```

For the example above, the reward function is comprised of 1) distance between the end-effector and the target part, and 2) distance between the current and the target pose of an articulated asset, link, or joint.

We show in Section 4.1 that our method can generalize across assets, suggest diverse and plausible tasks, decomposition and reward functions automatically, using a single in-context example in the prompt, without any additional human involvement.

3.3 Sequential Reinforcement Learning for Long Horizon Tasks

We train policies using Proximal Policy Optimization (PPO) [63] maximizing the generated reward functions for each subtask. We train RL for each generated subtask in temporal order. Once policy training for a subtask converges, we proceed to the next subtask, by sampling the initial state of the end-effector and the environment close to the terminal states of the previous subtask. This ensures policies can be temporally chained upon training. Our policies are trained per environment using privileged information of the simulation state to accelerate exploration. Such learned policies can be used as demonstration data and distilled into vision-language transformer policies, similar to [2, 10, 31]; we leave this for future work.

3. Approach

Chapter 4

Experiments

4.1 Experiments

Our experiments aim to answer the following questions:

1. Can Gen2Sim generate plausible geometry, appearance, and physics for diverse types of objects and parts, without human expertise and with minimal human involvement?
2. Can Gen2Sim generate task language goals and reward functions for novel object categories, novel assets with different part configurations, and a combination of multiple assets in an environment?
3. Can the generated environments and reward function lead to successful learning of RL policies?

4.1.1 Asset Generation

We compare our image-to-3D lifting with two baselines:

1. *RealFusion* [50], which uses textual inversion of [17] to learn a word embedding for the depicted object concept in an image, and uses text-conditioned diffusion with this text embedding during score distillation.
2. *Make-It-3D* [67], which uses the same NeRF and textured mesh two-stage fitting as Gen2Sim, but does not use a view and pose conditioned generative model, rather a text-based image diffusion model, similar to [56].

4. Experiments

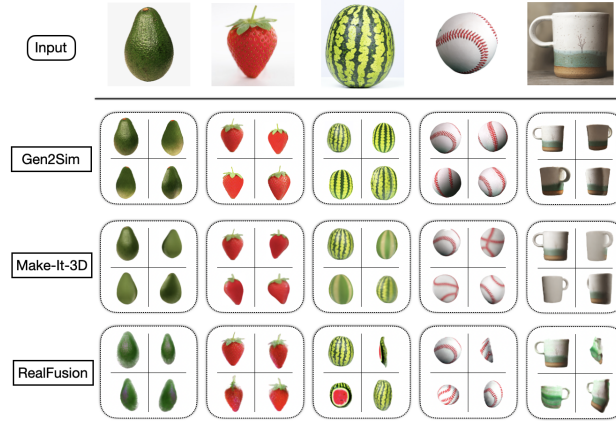


Figure 4.1: **3D asset generation** from Gen2Sim, RealFusion [50] and Make-It-3D [67]. Gen2Sim uses a view and camera pose conditioned image generative model during score distillation, which helps generate more accurate 3D geometry in comparison to the baselines.

	Mass (gram)	Length (cm)	Width	Height
Papaya	500-1000	15-20	10-15	10-15
Cucumber	200-300	15-20	5-7	5-7
Watermelon	5000-7000	30-40	20-30	20-30
Raspberry	3-5	2-3	2-3	2-3
Coconut	600-800	10-15	8-12	8-12
Corn	50-100	10-15	8-12	8-12
Pumpkin	2000-5000	20-40	20-40	20-40
Avocado	150-250	10-12	6-8	4-5

Table 4.1: **Size and physics parameter generated by LLMs** for a number of generated assets.

We show comparisons in Figure 4.1, with images rendered from 4 different views. Our model generates more plausible 3D model as our image diffusion prior comes from an image and pose-conditioned model in comparison to approaches like Fantasia3D or RealFusion which uses text conditioning.

We show generated values for 3D sizes and mass for a number of example objects in Table 4.1. We see that the common sense knowledge encoded in LLMs can produce reasonable physical parameters.

4.1.2 Automated Skill Learning

Gen2Sim generates diverse task descriptions, task decompositions and reward functions automatically for hundreds of assets, with different category labels and number of joints, **given only a single in-context prompt example** regarding the task decomposition and reward function of the task “putting a cup in a Microwave” . Then, the model can generalize to different scenes, asset articulated structures and task temporal lengths. We show some examples of such generated task descriptions in Figure 1.1 and more on our website. We show examples of task decompositions in Figure 3.1. We provide our prompts in our project website, alongside examples of the LLM’s responses.

We learn policies that optimize LLM generated rewards with PPO, an off-the-shelf model-free RL algorithm [63]. We make use of GPU-parallel data sampling in IsaacGym [48] for reinforcement learning. Our robotic setup uses a Franka Panda arm with a mobile base. It is equipped with a parallel-jaw gripper. Our state representation for PPO includes the robot’s joint position $q \in \mathbb{R}^{11}$, velocity $\dot{q} \in \mathbb{R}^{11}$ (7-DoF arm, x and y for the mobile base and 2 extra DoFs from the gripper jaws), orientation of the gripper $r \in SO(3)$, and poses and joint configurations of the assets present in the scene. We use position control and at each timestep t our policy produces target gripper pose and configurations which is converted to target robot configurations through inverse kinematics. A low-level PID torque controller provided by IsaacGym is used to produce low-level joint torque commands. We can successfully learn useful manipulation policies, and the policies are able to solve the tasks upon convergence. We show videos of such policies on our website.




Figure 4.2: **Twin environments** constructed and generated tasks for sim-to-real transfer. Left: real-world. Right: simulated.

4.1.3 Twin environment construction and sim-to-real world transfer

In order to validate the usefulness of the policies trained in simulation, we construct a twin simulated environment of our lab’s real-robot setup (Figure 4.2). We detect, segment, and estimate the poses of the objects in the scene. For non-articulated assets, we use our model to lift the detected object image to corresponding 3D models; for articulated objects, we select the most similar asset from the [81], and populated the simulated environment. We train RL policies in simulation and transfer the joint space trajectory back to our real-world setup. Our method allows successful execution of the generated tasks. For more videos of the trained policies and their task executions in simulation, as well as the sim2real transfer, please refer to our website.

Chapter 5

Discussion

5.1 Limitations

Gen2Sim has currently the following two important points to address towards materializing into a platform for large-scale robot skill learning that are deployable in real-world:

1. Sim2real transfer of closed-loop policies: Our current real-world experiments transfer open loop trajectories optimized in the constructed twin environment. For closed-loop policies to transfer to the real world and consume realistic sensory input, we would need to generate large-scale augmentations for both visual appearances and dynamics for each task and sub-task, and then distil the state-based RL policies to a foundational vision-language policy network. This is a direct avenue for our future work.

2. Beyond rigid asset generation: The assets we can currently generate are rigid or mostly rigid objects, which do not deform significantly under external forces. For articulated assets, we are using existing manually designed and labelled datasets ([21, 81]). To generate articulated objects, deformable objects and liquids, accurate fine-grained video perception is required in combination with generative priors to model the temporal dynamics of their geometry and appearance. This is an exciting and challenging direction for future work.

5.2 Conclusion

We have presented Gen2Sim, a method for automating the development of simulation environments, tasks and reward functions with pre-trained generative models of vision and language. We presented methods that create and augment geometry, textures and physics of object assets from single images, parse URDF files of assets, generate task descriptions, decompositions and reward python functions, and train reinforcement learning policies to solve the generated long horizon tasks. Addressing the limitations including generating diverse assets with more complex physical properties, and transferring trained policies to real world are direct avenues for our future work. We believe generative models of images and language will play an important role in automating and supersizing robot training data in simulation, and in crossing the sim2real gap, necessary for delivering robot generalists in the real world. Gen2Sim takes one first step in that direction.

Bibliography

- [1] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016. 2.4
- [2] Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Joseph Dabis, Chelsea Finn, Keerthana Gopalakrishnan, Karol Hausman, Alex Herzog, Jasmine Hsu, et al. Rt-1: Robotics transformer for real-world control at scale. *arXiv preprint arXiv:2212.06817*, 2022. 1, 3.3
- [3] Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Xi Chen, Krzysztof Choromanski, Tianli Ding, Danny Driess, Avinava Dubey, Chelsea Finn, et al. Rt-2: Vision-language-action models transfer web knowledge to robotic control. *arXiv preprint arXiv:2307.15818*, 2023. 1
- [4] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners, 2020. URL <https://arxiv.org/abs/2005.14165>. 1, 2.1
- [5] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners, 2020. 2.1
- [6] Eric R Chan, Marco Monteiro, Petr Kellnhofer, Jiajun Wu, and Gordon Wetzstein. pi-gan: Periodic implicit generative adversarial networks for 3d-aware image synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5799–5809, 2021. 2.2

- [7] Stephanie C. Y. Chan, Adam Santoro, Andrew K. Lampinen, Jane X. Wang, Aaditya Singh, Pierre H. Richemond, Jay McClelland, and Felix Hill. Data distributional properties drive emergent in-context learning in transformers, 2022. [2.1](#)
- [8] Rui Chen, Yongwei Chen, Ningxin Jiao, and Kui Jia. Fantasia3d: Disentangling geometry and appearance for high-quality text-to-3d content creation. *arXiv preprint arXiv:2303.13873*, 2023. [2.2](#), [3.1](#), [3.1.2](#)
- [9] Tao Chen, Jie Xu, and Pulkit Agrawal. A system for general in-hand object re-orientation. In *Conference on Robot Learning*, pages 297–307. PMLR, 2022. [1](#)
- [10] Sammy Christen, Wei Yang, Claudia Pérez-D’Arpino, Otmar Hilliges, Dieter Fox, and Yu-Wei Chao. Learning human-to-robot handovers from point clouds, 2023. [3.3](#)
- [11] Paolo Cignoni, Marco Callieri, Massimiliano Corsini, Matteo Dellepiane, Fabio Ganovelli, and Guido Ranzuglia. MeshLab: an Open-Source Mesh Processing Tool. In Vittorio Scarano, Rosario De Chiara, and Ugo Erra, editors, *Eurographics Italian Chapter Conference*. The Eurographics Association, 2008. ISBN 978-3-905673-68-5. doi: 10.2312/LocalChapterEvents/ItalChap/ItalianChapConf2008/129-136. [3.1.4](#)
- [12] Erwin Coumans and Yunfei Bai. Pybullet, a python module for physics simulation for games, robotics and machine learning. <http://pybullet.org>, 2016. [2.4](#)
- [13] Murtaza Dalal, Ajay Mandlekar, Caelan Garrett, Ankur Handa, Ruslan Salakhutdinov, and Dieter Fox. Imitating task and motion planning with visuomotor transformers, 2023. [2.3](#)
- [14] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. Carla: An open urban driving simulator, 2017. [1](#)
- [15] Adam Fishman, Adithyavairan Murali, Clemens Eppner, Bryan Peele, Byron Boots, and Dieter Fox. Motion policy networks, 2022. [2.3](#)
- [16] Zipeng Fu, Ashish Kumar, Ananye Agarwal, Haozhi Qi, Jitendra Malik, and Deepak Pathak. Coupling vision and proprioception for navigation of legged robots. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17273–17283, 2022. [1](#)
- [17] 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, 2022. [4.1.1](#)
- [18] Chuang Gan, Jeremy Schwartz, Seth Alter, Martin Schrimpf, James Traer, Julian De Freitas, Jonas Kubilius, Abhishek Bhandwaldar, Nick Haber, Megumi Sano, et al. Threedworld: A platform for interactive multi-modal physical simulation.

- arXiv preprint arXiv:2007.04954*, 2020. 2.4
- [19] Chuang Gan, Jeremy Schwartz, Seth Alter, Damian Mrowca, Martin Schrimpf, James Traer, Julian De Freitas, Jonas Kubilius, Abhishek Bhandwaldar, Nick Haber, Megumi Sano, Kuno Kim, Elias Wang, Michael Lingelbach, Aidan Curtis, Kevin Feigelis, Daniel M. Bear, Dan Gutfreund, David Cox, Antonio Torralba, James J. DiCarlo, Joshua B. Tenenbaum, Josh H. McDermott, and Daniel L. K. Yamins. Threedworld: A platform for interactive multi-modal physical simulation, 2021. 1
- [20] Caelan Reed Garrett, Tomás Lozano-Pérez, and Leslie Pack Kaelbling. Stripstream: Integrating symbolic planners and blackbox samplers. *CoRR*, abs/1802.08705, 2018. URL <http://arxiv.org/abs/1802.08705>. 2.3
- [21] Haoran Geng, Helin Xu, Chengyang Zhao, Chao Xu, Li Yi, Siyuan Huang, and He Wang. Gapartnet: Cross-category domain-generalizable object perception and manipulation via generalizable and actionable parts, 2023. 3.2, 5.1
- [22] Theophile Gervet, Zhou Xian, Nikolaos Gkanatsios, and Katerina Fragkiadaki. Act3d: Infinite resolution action detection transformer for robotic manipulation. *arXiv preprint arXiv:2306.17817*, 2023. 2.4
- [23] Nikolaos Gkanatsios, Ayush Jain, Zhou Xian, Yunchu Zhang, Christopher Atkeson, and Katerina Fragkiadaki. Energy-based models as zero-shot planners for compositional scene rearrangement. *arXiv preprint arXiv:2304.14391*, 2023. 2.4
- [24] Huy Ha, Pete Florence, and Shuran Song. Scaling up and distilling down: Language-guided robot skill acquisition. *arXiv preprint arXiv:2307.14535*, 2023. 2.1
- [25] Eric Heiden, Christopher E. Denniston, David Millard, Fabio Ramos, and Gaurav S. Sukhatme. Probabilistic inference of simulation parameters via parallel differentiable simulation, 2022. 2.2
- [26] Eric Heiden, Miles Macklin, Yashraj Narang, Dieter Fox, Animesh Garg, and Fabio Ramos. Disect: A differentiable simulator for parameter inference and control in robotic cutting, 2022. 2.2
- [27] Sebastian Höfer, Kostas E. Bekris, Ankur Handa, Juan Camilo Gamboa, Melissa Mozifian, Florian Golemo, Christopher G. Atkeson, Dieter Fox, Ken Goldberg, John Leonard, C. Karen Liu, Jan Peters, Shuran Song, Peter Welinder, and Martha White. Sim2real in robotics and automation: Applications and challenges. *IEEE Trans Autom. Sci. Eng.*, 18(2):398–400, 2021. doi: 10.1109/TASE.2021.3064065. URL <https://doi.org/10.1109/TASE.2021.3064065>. 1
- [28] Wenlong Huang, Pieter Abbeel, Deepak Pathak, and Igor Mordatch. Language models as zero-shot planners: Extracting actionable knowledge for embodied agents. *arXiv preprint arXiv:2201.07207*, 2022. 2.1

- [29] Wenlong Huang, Fei Xia, Ted Xiao, Harris Chan, Jacky Liang, Pete Florence, Andy Zeng, Jonathan Tompson, Igor Mordatch, Yevgen Chebotar, Pierre Sermanet, Noah Brown, Tomas Jackson, Linda Luu, Sergey Levine, Karol Hausman, and Brian Ichter. Inner monologue: Embodied reasoning through planning with language models, 2022. URL <https://arxiv.org/abs/2207.05608>. 2.1
- [30] Wenlong Huang, Chen Wang, Ruohan Zhang, Yunzhu Li, Jiajun Wu, and Li Fei-Fei. Voxposer: Composable 3d value maps for robotic manipulation with language models. *arXiv preprint arXiv:2307.05973*, 2023. 2.1
- [31] Andrew Jaegle, Sebastian Borgeaud, Jean-Baptiste Alayrac, Carl Doersch, Catalin Ionescu, David Ding, Skanda Koppula, Daniel Zoran, Andrew Brock, Evan Shelhamer, et al. Perceiver io: A general architecture for structured inputs & outputs. *arXiv preprint arXiv:2107.14795*, 2021. 3.3
- [32] Yandong Ji, Gabriel B. Margolis, and Pulkit Agrawal. Dribblebot: Dynamic legged manipulation in the wild, 2023. 1
- [33] Leslie Pack Kaelbling and Tomás Lozano-Pérez. Hierarchical task and motion planning in the now. In *2011 IEEE International Conference on Robotics and Automation*, pages 1470–1477, 2011. doi: 10.1109/ICRA.2011.5980391. 2.3
- [34] Elia Kaufmann, Leonard Bauersfeld, Antonio Loquercio, Matthias Mueller, Vladlen Koltun, and Davide Scaramuzza. Champion-level drone racing using deep reinforcement learning. *Nature*, 620:982–987, 08 2023. doi: 10.1038/s41586-023-06419-4. 1
- [35] Eric Kolve, Roozbeh Mottaghi, Winson Han, Eli VanderBilt, Luca Weihs, Alvaro Herrasti, Daniel Gordon, Yuke Zhu, Abhinav Gupta, and Ali Farhadi. Ai2-thor: An interactive 3d environment for visual ai. *arXiv preprint arXiv:1712.05474*, 2017. 2.4
- [36] Ashish Kumar, Zipeng Fu, Deepak Pathak, and Jitendra Malik. Rma: Rapid motor adaptation for legged robots, 2021. 1
- [37] Joonho Lee, Jemin Hwangbo, Lorenz Wellhausen, Vladlen Koltun, and Marco Hutter. Learning quadrupedal locomotion over challenging terrain. *CoRR*, abs/2010.11251, 2020. URL <https://arxiv.org/abs/2010.11251>. 1
- [38] Sergey Levine, Peter Pastor, Alex Krizhevsky, and Deirdre Quillen. Learning hand-eye coordination for robotic grasping with deep learning and large-scale data collection, 2016. 1
- [39] Yuheng Li, Haotian Liu, Qingyang Wu, Fangzhou Mu, Jianwei Yang, Jianfeng Gao, Chunyuan Li, and Yong Jae Lee. Gligen: Open-set grounded text-to-image generation. *arXiv preprint arXiv:2301.07093*, 2023. 3.1.1
- [40] Jacky Liang, Wenlong Huang, Fei Xia, Peng Xu, Karol Hausman, Brian Ichter,

- Pete Florence, and Andy Zeng. Code as policies: Language model programs for embodied control, 2022. URL <https://arxiv.org/abs/2209.07753>. 2.1
- [41] Shalev Lifshitz, Keiran Paster, Harris Chan, Jimmy Ba, and Sheila McIlraith. Steve-1: A generative model for text-to-behavior in minecraft, 2023. 2.1
- [42] Chen-Hsuan Lin, Jun Gao, Luming Tang, Towaki Takikawa, Xiaohui Zeng, Xun Huang, Karsten Kreis, Sanja Fidler, Ming-Yu Liu, and Tsung-Yi Lin. Magic3d: High-resolution text-to-3d content creation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 2.2
- [43] Xingyu Lin, Yufei Wang, Jake Olkin, and David Held. Softgym: Benchmarking deep reinforcement learning for deformable object manipulation. *arXiv preprint arXiv:2011.07215*, 2020. 2.4
- [44] Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *CoRR*, abs/2107.13586, 2021. URL <https://arxiv.org/abs/2107.13586>. 2.1
- [45] Ruoshi Liu, Rundi Wu, Basile Van Hoorick, Pavel Tokmakov, Sergey Zakharov, and Carl Vondrick. Zero-1-to-3: Zero-shot one image to 3d object, 2023. 3.1.1, 3.1.2
- [46] Daoming Lyu, Fangkai Yang, Bo Liu, and Steven Gustafson. SDRL: interpretable and data-efficient deep reinforcement learning leveraging symbolic planning. *CoRR*, abs/1811.00090, 2018. URL <http://arxiv.org/abs/1811.00090>. 2.3
- [47] Miles Macklin, Matthias Müller, Nuttapong Chentanez, and Tae-Yong Kim. Unified particle physics for real-time applications. *ACM Transactions on Graphics (TOG)*, 33(4):1–12, 2014. 2.4
- [48] Viktor Makoviychuk, Lukasz Wawrzyniak, Yunrong Guo, Michelle Lu, Kier Storey, Miles Macklin, David Hoeller, Nikita Rudin, Arthur Allshire, Ankur Handa, and Gavriel State. Isaac gym: High performance gpu-based physics simulation for robot learning, 2021. 4.1.2
- [49] Michael James McDonald and Dylan Hadfield-Menell. Guided imitation of task and motion planning, 2021. 2.3
- [50] Luke Melas-Kyriazi, Christian Rupprecht, Iro Laina, and Andrea Vedaldi. Real-fusion: 360 reconstruction of any object from a single image. In *CVPR*, 2023. URL <https://arxiv.org/abs/2302.10663>. (document), 2.2, 3.1.2, 4.1.1, 4.1
- [51] Toki Migimatsu and Jeannette Bohg. Object-centric task and motion planning in dynamic environments. *CoRR*, abs/1911.04679, 2019. URL <http://arxiv.org/abs/1911.04679>. 2.3
- [52] Adithyavairavan Murali, Arsalan Mousavian, Clemens Eppner, Adam Fishman,

- and Dieter Fox. Cabinet: Scaling neural collision detection for object rearrangement with procedural scene generation, 2023. 2.3
- [53] Thu Nguyen-Phuoc, Chuan Li, Lucas Theis, Christian Richardt, and Yong-Liang Yang. Hologan: Unsupervised learning of 3d representations from natural images, 2019. 2.2
- [54] OpenAI. Gpt-4 technical report, 2023. 3.1
- [55] OpenAI, Ilge Akkaya, Marcin Andrychowicz, Maciek Chociej, Mateusz Litwin, Bob McGrew, Arthur Petron, Alex Paino, Matthias Plappert, Glenn Powell, Raphael Ribas, Jonas Schneider, Nikolas Tezak, Jerry Tworek, Peter Welinder, Lilian Weng, Qiming Yuan, Wojciech Zaremba, and Lei Zhang. Solving rubik’s cube with a robot hand, 2019. 1
- [56] Ben Poole, Ajay Jain, Jonathan T. Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. *arXiv*, 2022. 2.2, 3.1, 3.1.2, 4.1.1
- [57] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. *CoRR*, abs/2103.00020, 2021. URL <https://arxiv.org/abs/2103.00020>. 1
- [58] Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. Robust speech recognition via large-scale weak supervision, 2022. 1
- [59] Elad Richardson, Gal Metzer, Yuval Alaluf, Raja Giryes, and Daniel Cohen-Or. Texture: Text-guided texturing of 3d shapes, 2023. 3.1.3
- [60] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022. 1, 3.1.1
- [61] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models, 2022. 3.1
- [62] Manolis Savva, Abhishek Kadian, Oleksandr Maksymets, Yili Zhao, Erik Wijmans, Bhavana Jain, Julian Straub, Jia Liu, Vladlen Koltun, Jitendra Malik, Devi Parikh, and Dhruv Batra. Habitat: A platform for embodied ai research, 2019. 1
- [63] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms, 2017. 3.3, 4.1.2
- [64] Bokui Shen, Xinchen Yan, Charles R. Qi, Mahyar Najibi, Boyang Deng, Leonidas Guibas, Yin Zhou, and Dragomir Anguelov. Gina-3d: Learning to generate

- implicit neural assets in the wild, 2023. [2.2](#)
- [65] Maximilian Sieb, Zhou Xian, Audrey Huang, Oliver Kroemer, and Katerina Fragkiadaki. Graph-structured visual imitation. In *Conference on Robot Learning*, pages 979–989. PMLR, 2020. [1](#)
- [66] Sanjana Srivastava, Chengshu Li, Michael Lingelbach, Roberto Martín-Martín, Fei Xia, Kent Vainio, Zheng Lian, Cem Gokmen, Shyamal Buch, C. Karen Liu, Silvio Savarese, Hyowon Gweon, Jiajun Wu, and Li Fei-Fei. Behavior: Benchmark for everyday household activities in virtual, interactive, and ecological environments, 2021. [1](#)
- [67] Junshu Tang, Tengfei Wang, Bo Zhang, Ting Zhang, Ran Yi, Lizhuang Ma, and Dong Chen. Make-it-3d: High-fidelity 3d creation from a single image with diffusion prior, 2023. ([document](#)), [2.2](#), [4.1.1](#), [4.1](#)
- [68] Josh Tobin, Rachel Fong, Alex Ray, Jonas Schneider, Wojciech Zaremba, and Pieter Abbeel. Domain randomization for transferring deep neural networks from simulation to the real world, 2017. [3.1.3](#)
- [69] Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ international conference on intelligent robots and systems*, pages 5026–5033. IEEE, 2012. [2.4](#)
- [70] Marc Toussaint. Logic-geometric programming: An optimization-based approach to combined task and motion planning. In *Proceedings of the 24th International Conference on Artificial Intelligence, IJCAI’15*, page 1930–1936. AAAI Press, 2015. ISBN 9781577357384. [2.3](#)
- [71] Hsiao-Yu Fish Tung, Zhou Xian, Mihir Prabhudesai, Shamit Lal, and Katerina Fragkiadaki. 3d-oes: Viewpoint-invariant object-factorized environment simulators. *arXiv preprint arXiv:2011.06464*, 2020. [2.4](#)
- [72] Guanzhi Wang, Yuqi Xie, Yunfan Jiang, Ajay Mandlekar, Chaowei Xiao, Yuke Zhu, Linxi Fan, and Anima Anandkumar. Voyager: An open-ended embodied agent with large language models, 2023. [2.1](#)
- [73] Haochen Wang, Xiaodan Du, Jiahao Li, Raymond A. Yeh, and Greg Shakhnarovich. Score jacobian chaining: Lifting pretrained 2d diffusion models for 3d generation, 2022. [3.1.2](#)
- [74] Kun Wang, William R. Johnson III au2, Shiyang Lu, Xiaonan Huang, Joran Booth, Rebecca Kramer-Bottiglio, Mridul Aanjaneya, and Kostas Bekris. Real2sim2real transfer for control of cable-driven robots via a differentiable physics engine, 2023. [2.2](#)
- [75] Tsun-Hsuan Wang, Andrew Everett Spielberg, Pingchuan Ma, Zhou Xian, Hao Zhang, Joshua B. Tenenbaum, and Chuang Gan. Softzoo: A soft robot co-design

- benchmark for locomotion in diverse environments. In *International Conference on Learning Representations*, 2023. 2.4
- [76] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Ed H. Chi, Quoc Le, and Denny Zhou. Chain of thought prompting elicits reasoning in large language models. *CoRR*, abs/2201.11903, 2022. URL <https://arxiv.org/abs/2201.11903>. 2.1
- [77] Zhou Xian, Shamit Lal, Hsiao-Yu Tung, Emmanouil Antonios Platanios, and Katerina Fragkiadaki. Hyperdynamics: Meta-learning object and agent dynamics with hypernetworks. *arXiv preprint arXiv:2103.09439*, 2021. 2.4
- [78] Zhou Xian, Theophile Gervet, Zhenjia Xu, Yi-Ling Qiao, and Tsun-Hsuan Wang. Towards a foundation model for generalist robots: Diverse skill learning at scale via automated task and scene generation. *arXiv preprint arXiv:2305.10455*, 2023. 1
- [79] Zhou Xian, Nikolaos Gkanatsios, Theophile Gervet, Tsung-wei Ke, and Katerina Fragkiadaki. Chaineddiffuser: Unifying trajectory diffusion and keypose prediction for robotic manipulation. *Conference on Robot Learning*, 2023. 2.4
- [80] Zhou Xian, Bo Zhu, Zhenjia Xu, Hsiao-Yu Tung, Antonio Torralba, Katerina Fragkiadaki, and Chuang Gan. Fluidlab: A differentiable environment for benchmarking complex fluid manipulation. In *International Conference on Learning Representations*, 2023. 2.4
- [81] Fanbo Xiang, Yuzhe Qin, Kaichun Mo, Yikuan Xia, Hao Zhu, Fangchen Liu, Minghua Liu, Hanxiao Jiang, Yifu Yuan, He Wang, Li Yi, Angel X. Chang, Leonidas J. Guibas, and Hao Su. Sapien: A simulated part-based interactive environment, 2020. 1, 2.4, 3.2, 4.1.3, 5.1
- [82] Danfei Xu, Roberto Martín-Martín, De-An Huang, Yuke Zhu, Silvio Savarese, and Li Fei-Fei. Regression planning networks. *CoRR*, abs/1909.13072, 2019. URL <http://arxiv.org/abs/1909.13072>. 2.1
- [83] Wenhao Yu, Nimrod Gileadi, Chuyuan Fu, Sean Kirmani, Kuang-Huei Lee, Montse Gonzalez Arenas, Hao-Tien Lewis Chiang, Tom Erez, Leonard Hasenclever, Jan Humplik, et al. Language to rewards for robotic skill synthesis. *arXiv preprint arXiv:2306.08647*, 2023. 2.1
- [84] Lvmin Zhang and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. *arXiv preprint arXiv:2302.05543*, 2023. 3.1.1
- [85] Xizhou Zhu, Yuntao Chen, Hao Tian, Chenxin Tao, Weijie Su, Chenyu Yang, Gao Huang, Bin Li, Lewei Lu, Xiaogang Wang, Yu Qiao, Zhaoxiang Zhang, and Jifeng Dai. Ghost in the minecraft: Generally capable agents for open-world environments via large language models with text-based knowledge and memory, 2023. 2.1

- [86] Yifeng Zhu, Jonathan Tremblay, Stan Birchfield, and Yuke Zhu. Hierarchical planning for long-horizon manipulation with geometric and symbolic scene graphs. *CoRR*, abs/2012.07277, 2020. URL <https://arxiv.org/abs/2012.07277>. 2.1