

# URDFormer: A Pipeline for Constructing Articulated Simulation Environments from Real-World Images

Zoey Chen<sup>1</sup>, Aaron Walsman<sup>1</sup>, Marius Memmel<sup>1</sup>, Kaichun Mo<sup>2</sup>, Alex Fang<sup>1</sup>,  
 Dieter Fox<sup>\*1,2</sup>, Abhishek Gupta<sup>\*1,2</sup>

<sup>1</sup>University of Washington    <sup>2</sup> Nvidia

[urdfomer.github.io](https://urdfomer.github.io)

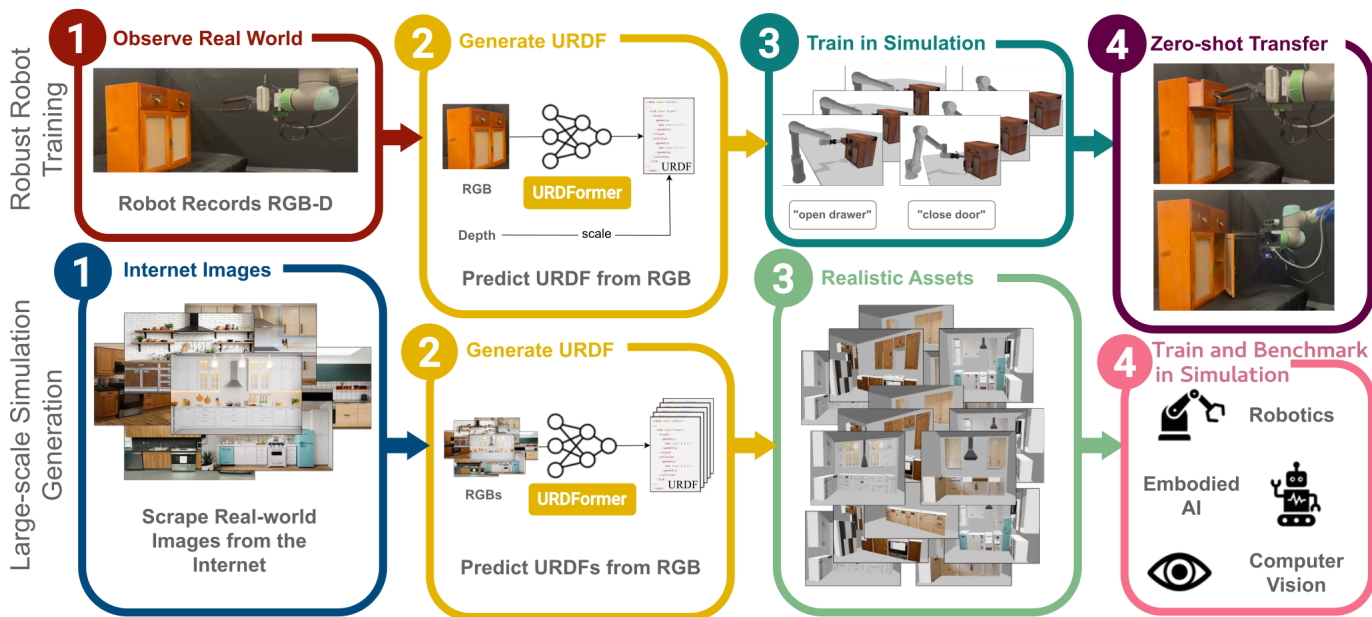


Fig. 1: The *URDFormer* predicts realistic, kinematic scenes from images that 1) allow for zero-shot real-to-sim-to-real transfer through targeted randomization, and 2) generate internet-scale simulation assets valuable to a multitude of applications.

**Abstract**—Constructing accurate and targeted simulation scenes that are both visually and physically realistic is a problem of significant practical interest in domains ranging from robotics to computer vision. This problem has become even more relevant as researchers wielding large data-hungry learning methods seek new sources of training data for physical decision-making systems. However, building simulation models is often still done by hand - a graphic designer and a simulation engineer work with predefined assets to construct rich scenes with realistic dynamic and kinematic properties. While this may scale to small numbers of scenes, to achieve the generalization properties that are required for data-driven robotic control, we require a pipeline that is able to synthesize large numbers of realistic scenes, complete with “natural” kinematic and dynamic structure. To attack this problem, we develop models for inferring structure and generating simulation scenes from natural images, allowing for scalable scene generation from web-scale datasets. To train these image-to-simulation models, we show how controllable text-to-image generative models can be used in generating paired training data that allows for modeling of the *inverse* problem, mapping from realistic images back to complete scene models. We show how this paradigm allows us to build large datasets of scenes

in simulation with semantic and physical realism. We present an integrated end-to-end pipeline that generates simulation scenes complete with articulated kinematic and dynamic structures from real-world images and use these for training robotic control policies. We then robustly deploy in the real world for tasks like articulated object manipulation. In doing so, our work provides both a pipeline for large-scale generation of simulation environments and an integrated system for training robust robotic control policies in the resulting environments.

## I. INTRODUCTION

Simulation has become a cornerstone of a plethora of applied machine learning problems - from the natural sciences such as physics, chemistry, and biology [1, 2] to robotics [3, 4] and computer vision [5, 6]. Simulation allows for scalable and cheap data collection while providing an easy way to encode domain-specific prior knowledge into end-to-end machine learning problems. This is particularly important for data-scarce problems such as robotics, where collecting real data can lead to costly and unsafe failures or may require

expensive human supervision. Critical to each of these endeavors is a rich and accurate simulation environment, complete with assets depicting complex scene layouts and kinematic structure. For instance, advances in robotic mobile manipulation in the Habitat simulator [7], are critically dependent on the Matterport dataset for realistic scenes [8]. The creation and curation of these simulation scenes and assets is an important but often overlooked part of the process.

The most common approaches for generating simulation content are either manual [9, 10], procedural [11] or, more recently, purely generative [12, 13, 14, 15, 16]. The manual process for creating simulation scenes requires a designer to characterize, identify, and model a particular real-world scene, a painstaking and impractical process. While this approach can produce high quality results, it often leads to content that lacks diversity due to the amount of human effort required. On the other hand, rule-based procedural generation methods [11, 17] have been applied in robotics applications such as navigation, but often struggle to capture the natural complexity of the real world for problems such as manipulation. Moreover, the procedural generation process is not *controllable*, making it hard to generate simulation content corresponding to a *particular* real-world environment, which is often important in real-world robotic learning pipelines. Recently introduced generative methods for content creation [12, 15, 16, 13] can generate visually appealing 3-D geometries for particular objects, but often result in undesired physical simulation behavior and lack kinematic structure like articulation.

What are the desiderata for a content creation method for simulation? To enable a variety of downstream use cases such as robotic learning, scalable content creation in simulation must be (1) realistic enough such that machine learning models trained in the constructed simulation environments transfer back to the real world, (2) diverse in a way that captures the statistics of natural environments so as to enable learning generalizable models and policies (3) controllable in a way that allows for targeted generation of particular scenes of interest.

To generate content of this nature, we develop a pipeline to train a transformer-based network, *URDFormer*, that maps directly from individual real-world images to corresponding simulation content (expressed as a Unified Robot Description File (URDF)) that could plausibly represent the semantics, kinematics, and structure of the scene (Fig1). We leverage controllable text-to-image generative models [18] to generate a large-scale paired dataset of structured simulation scenes and closely corresponding, realistic images. This paired dataset is then *inverted* to train *URDFormer*, which maps from RGB images directly to plausible simulation environments with semantic and kinematic structure.

*URDFormer* can then naturally be used in several use cases — (1) diverse content generation for simulation: generating a large and diverse set of realistic simulation environments that correspond directly to uncured, real-world RGB images (e.g scraped off the web), or (2) targeted generation: generating a simulation environment (or narrow distribution of environments) corresponding to a particular set of desired images. We show that the generated simulation environments are a useful

tool for robotic learning in a real-to-sim-to-real pipeline for training robotic control policies.

## II. URDFORMER - A PIPELINE FOR SCALABLE CONTENT CREATION FOR SIMULATION FROM REAL-WORLD IMAGES

Generating simulated scenes with a high degree of visual realism that supports rich kinematic and dynamic structure, while reflecting the natural statistics of the real world is a challenging problem. Downstream applications in robotics and computer vision typically require data that is simultaneously *realistic*, *diverse*, and *controllable*.

To accomplish these requirements, we attempt to generate rich kinematic scenes directly from RGB images, so that we can use either a large web-scale dataset of images harvested from the internet to generate a large number of realistic scenes, or a small number of domain-specific images to train robot models for more specific targeted environments. Unfortunately, no large dataset containing images along with corresponding articulated scene descriptions is publicly available to train such a model. We address this by building a new dataset using generative image models to enhance the visual quality of synthetic renders of procedurally sampled scenes. This process requires us to reason about two distinct pipelines. The *forward* pipeline takes a procedurally sampled environment, renders it, and augments the resulting images with a generative model to produce a dataset of scene/image pairs. The second *inverse* pipeline then trains a model which takes in the images produced in the forward pipeline and produces the kinematic scene description that was used to create it.

Once the inverse model is trained, it allows for scene generation that is *realistic* since it can be used to build scenes that correspond to real images. The generated scenes are *diverse* since large web-scale image datasets with diverse content can be used to seed this generation process. Lastly, the generation is *controllable* since curated images of particular target environments can be used to generate corresponding simulation assets. We first define the inverse problem of synthetic scene generation from real-world images, then describe how to learn inverse models to solve this problem with supervised learning on a paired dataset generated using controllable text-to-image generative models [18]. Finally, we show how the learned inverse model can be used with real-world image datasets for scalable content creation. This tool can then be used to instantiate a real-to-simulation-real pipeline for robot learning, as described in Section III.

### A. Problem Formulation

To formalize the problem of simulation scene generation from real-world images, let us consider a kinematic scene description  $z$  drawn from a target scene distribution  $P(z)$  in the real world. For our purposes, the scene can be described as a list of objects  $z = \{o_1 \dots o_n\}$ , where each object  $o_i$  contains a base class label  $c_i$ , a 3D bounding box  $b_i \in \mathbb{R}^6$ , a 3D transform  $T_i \in SE(3)$ , a kinematic parent  $p_i \in [1 \dots i-1]$  and a joint type  $j_i$  that specifies how that object can move relative to its parent  $o_i = (c_i, b_i, T_i, p_i, j_i)$ . This resembles the typical representation of scenes and robots using the unified robot description format (URDF). Let’s consider a kitchen scenario

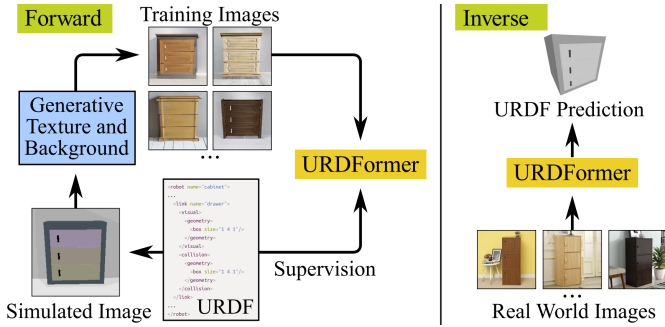


Fig. 2: The *URDFormer* is trained on a large paired dataset of simulation assets and realistic renderings (forward). During inference, this process is inverted and it predicts the URDF from a real image (inverse).

that contains a row of cabinets next to a stove. The cabinets and the stove will be the kinematic children of the wall, and the kinematic parents of their respective doors. Similarly, these doors will be the kinematic parents of their handles. As the example shows, the kinematic structure  $z$  for a particular real-world scenario is unknown without extensive human labeling effort, and instead, we only have access to the result  $x$  of an indirect “forward” function  $f$ ,  $x = f(z)$ . For example,  $x$  could be a photograph of the real environment, or a point cloud captured with a LIDAR scanner. The goal of this work is to recover the entire kinematic and semantic structure of the scene from just having access to the forward evaluation  $x$ , requiring complete inference of a rich scene representation  $z$ .

Unfortunately, since the scene structure  $z$  is unknown for most complex real-world scenes and difficult to generate manually, it is challenging to solve the “inverse” generation problem to infer the scene description  $z$  from the forward rendered images (or alternative sensor readings)  $x$ ,  $z = f^{-1}(x)$ . Had there been a sizeable dataset  $\mathcal{D} = \{(z_i, x_i)\}_{i=1}^N$  of scene descriptors  $z_i$  in simulation and their corresponding real-world counterparts  $x_i$ , the inverse problem could be solved using supervised learning (minimizing a loss  $\mathcal{L}$  like the cross entropy loss or a MSE loss) to learn an  $f_{\theta}^{-1}$  that approximates the scene structure  $\hat{z}$  given an input forward-rendered image  $x$ .

However, such a paired dataset does not readily exist, making direct application of supervised learning methods challenging. In this work we take an inversion through synthesis approach—leveraging pre-trained generative models to convert procedurally generated scenes in simulation into a large paired dataset of scene content  $z$  and corresponding realistic RGB images  $x$ . This process can generate a large and diverse dataset of image and scene-description  $(x, z)$  pairs that we can use to train an approximate inverse model  $f_{\theta}^{-1}(x)$  that generates scene descriptions  $\hat{z}$  from real RGB images  $x$ . Since most scenes that we consider are object-centric, we decompose the inverse problem into two parts: (1) object-level prediction that focuses on the kinematic structure of individual objects, *e.g.*, cabinets, and (2) global-scene prediction, *e.g.*, kitchens, that focuses on the structure of an overall scene. We next discuss the process of generating a large paired dataset for these two components (Section II-B) and then show the training process for the inverse model in detail (Section II-C).

## B. Controlled Generation of Paired Datasets with Pretrained Generative Models

Given a simulated scene  $z$  (drawn from a dataset such as PartNet [19], or procedurally generated), we use the fact that controllable generative models [18] are both diverse and realistic enough to take an unrealistic simulation rendering of a scene and generate a distribution of corresponding *realistic* images. This allows the scene in simulation with unrealistic appearance and texture to be translated into a diverse set of visually realistic images that plausibly match the same underlying environment, as shown on the left side of Figure 2. To ensure piecewise consistency and realism of the generated images, we use two different dataset generation techniques for the global scene structure and local object structure respectively. These share the same conceptual ideas but differ to account for consistency properties in each case.

**Scene-Level Dataset Generation:** To generate training data for the scene model, we feed a poorly rendered image from simulation along with a templated text prompt to an image-and-text guided diffusion model [18], as shown in Fig 3. This generates a new image that attempts to simultaneously match the content described in the text prompt while retaining the global scene layout from the provided image. We found that this model is able to reliably maintain the scene layout, but it may change some individual components of the scene, *e.g.*, replacing objects with a different but plausible category, or changing the number of subcomponents within an object such as the drawers or handles. Despite these failures, the large-scale structural consistency still provides a useful source of training data. After running our simulated image through the generative model, we have realistic images that contain known high-level object positions and spatial relationships, but unknown category and low-level part structures (*e.g.* the parts for articulated objects such as cabinets), since these may have been modified by the generative model in its forward pass. This means that the scene model dataset contains complete images, but incomplete labels. Rather than complete  $(x, z)$  pairs, we have a dataset  $\mathcal{D}_{\text{scene}} = \{(x, \tilde{z})\}$  of  $(x, \tilde{z})$  pairs where  $\tilde{z}$  only contains the bounding boxes, transforms and parents of the high-level (non-part) objects  $\tilde{z} = \{(b_1, T_1, p_1) \dots (b_n, T_n, p_n)\}$  but lacks accurate low-level information.

**Object-Level Dataset Generation:** The process for generating object-level training data is similar but requires more care due to the tendency of text-to-image generative models to modify low-level details. For objects with complex kinematic structures, such as cabinets, we procedurally generate a large number of examples of these objects and render them in isolation from different angles. Rather than using a generative model to construct entirely new images, we use it to produce diverse texture images, which are overlaid in the appropriate locations on the image using perspective warping. We then change the background of the image using the generative model with appropriate masking derived from the original rendering. This part-by-part texture-based rendering process ensures diversity while maintaining consistency in low-level details. Unlike the scene dataset which contains complete



Fig. 3: **Controlled Generation:** Rendering URDF models in simulation and generating paired images with a guided diffusion model.

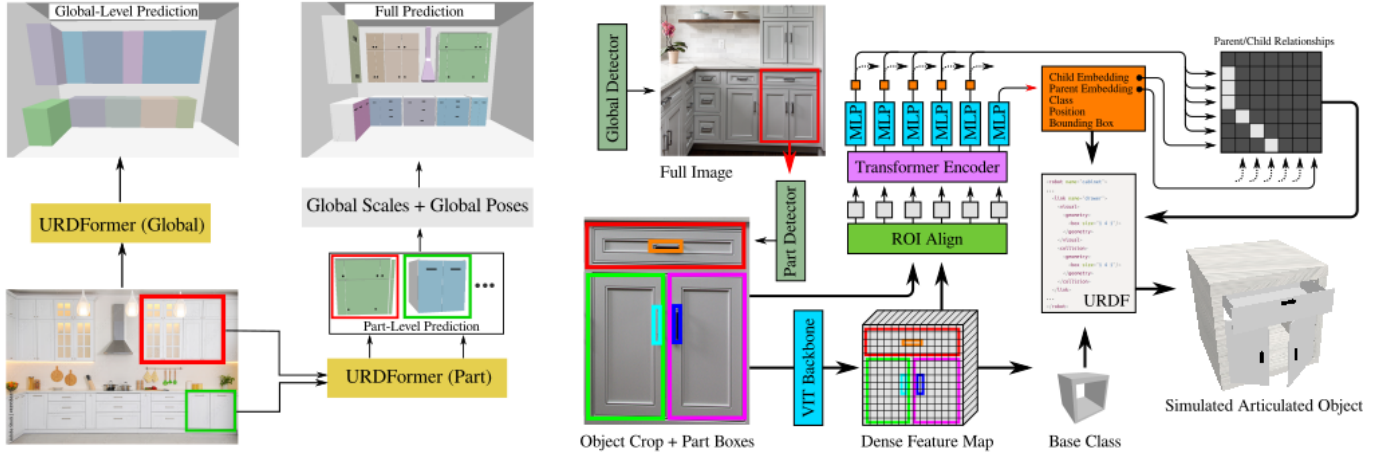


Fig. 4: Depiction of the *URDFormer* Training Procedure and Architecture. **(Left)** Given an RGB image of the scene, *i.e.* a kitchen, we train two separate networks: *URDFormer* (Global) focuses on predicting parent and spatial info of how to place the object. *URDFormer* (Part) takes the cropped image containing each object and predicts detailed structure. The results of the two predictions are combined and create the full scene prediction. **(Right)** The *URDFormer* architecture takes as input a cropped RGB image and object part boxes and predicts a hierarchy consisting of a base class and parent-child relations that make up the final URDF file.

images but partial labels, the object dataset contains partial images (in the sense that they contain only a single object), but complete labels for the object and its kinematic parts. We can say that this dataset  $\mathcal{D}_{\text{object}}$  contains  $(\tilde{x}, z)$  pairs where  $\tilde{x}$  is an image of a single object rather than a full scene (hence the partial  $x$ ), and  $z$  is complete for the single object and its parts.

The result of these two data generation processes is a high-level scene structure dataset  $\mathcal{D}_{\text{scene}}$ , and a low-level object dataset  $\mathcal{D}_{\text{object}}$ , that can subsequently be used to train an object-level and a scene-level inverse model, as shown on right side in Figure 2.

### C. Learning Inverse Generative Models for Scene Synthesis

Given the datasets  $\mathcal{D}_{\text{object}} = (\tilde{x}, z)$  and  $\mathcal{D}_{\text{scene}} = (x, \tilde{z})$  constructed as described above, we can use supervised learning

methods to learn an *inverse model* that maps images of a complex object or scene to the corresponding simulation asset. To take advantage of these partially complete datasets, we must add some structure to our prediction model. We do this by splitting our learned inverse model in correspondence with the split in our forward data generation process: we train one network  $f_{\theta}^{-1}$  to predict the high-level scene structure using dataset  $\mathcal{D}_{\text{scene}}$  and another network  $g_{\phi}^{-1}$  to predict the low-level part structure of objects using  $\mathcal{D}_{\text{object}}$ .

To model both the scene-level prediction model ( $f_{\theta}^{-1}$ ) and the low-level part prediction model ( $g_{\phi}^{-1}$ ), we propose a novel network architecture — *URDFormer*, that takes an RGB image and predicts URDF primitives as shown in Fig4. Note that both the scene-level prediction and the low-level part prediction use the same network architecture, the scene-level simply operates

on full images with object bounding boxes extracted, while the part-level operates on crops of particular objects with object parts extracted. In the *URDFormer* architecture, the image is first fed into a vision transformer [20] (ViT) visual backbone to extract global features. We then obtain bounding boxes of the objects in the image using the masks rendered from the original procedurally generated scene in simulation (these are known at training time, and can be extracted using detection models at test time). We then use ROI alignment [21] to extract features for each of these bounding boxes. These feature maps are combined with an embedding of the bounding box coordinates and then fed through a Transformer [22] to produce a feature for each object in the scene. An MLP then decodes these features into an optional base class label (used only when training the object-level model), and a discretized 3D position and bounding box. In addition, it also produces a child embedding and a parent embedding that are used to predict the hierarchical relationships in the scene (object to its parent and so on). To construct these relationships, the network uses a technique from scene graph generation [23] that produces an  $n \times n$  relationship score matrix by computing the dot product of every possible parent with every possible child. The scene-level model also specially has a set of learned embeddings for six different root objects corresponding to the four walls, the floor, and the ceiling so that large objects like countertops and sinks can be attached to the room.

#### Test-time scene generation from real-world RGB images:

Due to the unpredictable nature of the text-to-image generative transforms that are used to perform global dataset generation, which may change the base class identities, *e.g.*, the diffusion model might turn a fridge into a cabinet, only the position, bounding box, and relationship information is used when computing the high-level scene structure. To generate a full approximation of the scene structure  $\hat{z}$  from a natural image at test time, the image and a list of high-level bounding boxes (from a detection model) are first fed into the scene prediction model  $f_{\theta}^{-1}$ , which predicts the global structure, *i.e.* the location and parent for each object. The image regions corresponding to these boxes are then extracted and a second detection model is used to produce part-level bounding boxes. Each of these image regions that correspond to a particular object (*e.g.* a cabinet or a fridge) and the corresponding part-level boxes (*e.g.* individual drawers or doors) are then fed into the part prediction model  $g_{\phi}^{-1}$  to compute the kinematic structure of the low-level objects and parts. This nested prediction structure can be used to generate entire scenes from web-scraped RGB images drawn from any image dataset to generate novel simulation content both at the scene level and at the object level. We visualize this process in Fig 4

### III. USING *URDFormer* FOR ROBOTIC CONTROL VIA A REAL-TO-SIMULATION-TO-REAL PIPELINE

As described above, *URDFormer* provides the ability to generate realistic, diverse and controllable scenes in simulation through inverse modeling. This suggests that on deployment, *URDFormer* allows a system to take a single picture of a deployment scene and then construct a fully articulated scene in simulation with *minimal* human effort. In this section, we

describe how this type of controllable inverse modeling can serve as a useful tool for robotic learning through a real-to-simulation-to-real pipeline.

The most straightforward way of using *URDFormer* for robotic control is a model-based one - first synthesize a “digital twin” for a deployment time scene, and directly use this for planning and control in the real world with a model-based approach. This is challenging for several reasons - (1) the constructed simulations may not be perfectly accurate, (2) there is no access to Lagrangian environment state in the real world. Instead, we take a learning-based approach to the problem; we use *URDFormer* to generate not precisely the test time scene in simulation, but rather a narrow, representative distribution of simulation environments via a targeted randomization procedure. This distribution of environments in simulation can be used to learn generalizable robotic policies that operate from raw perceptual input, directly transferring back from simulation to the real-world. Doing so closes the real-to-simulation-to-real loop, obtaining real-world robotic policies with minimal human effort in the process. This pipeline has three major components:

**Scene Generation:** Given a robot’s RGB pointcloud observation of an unseen environment, we use *URDFormer* to generate a URDF file from RGB that captures the kinematic and dynamic structure of the real-world scene. We further resize the URDF to fit the pointcloud’s scale. Importing the URDF into simulation then provides a playground for data collection and policy training. Most importantly, this simulation is not just an arbitrary model but an approximate representation of the real world scene of interest on deployment.

**Targeted Randomization:** In simulation, we have access to ground truth information which we can exploit to inexpensively collect trajectory data for policy learning. We use an efficient motion planner [24] to quickly collect approximately optimal trajectories solving multiple tasks in simulation. To increase the simulation diversity and decrease the sim2real gap, we additionally randomize over minor details, *e.g.*, texture variants, shapes, and sizes of parts. This “targeted” randomization is in contrast to procedural generation which covers many different environment variations but is not informed by the real-world environment.

**Policy Synthesis:** To synthesize a policy from the collected data, we can train a language-conditioned behavior cloning policy [25] operating from RGB point-clouds in simulation, applying image augmentations during training to enable policy transfer back to the real world. We stress that the proposed pipeline is not limited to data collection with motion planning and behavior cloning but can also be used to train policies with other policy search methods [26, 27, 28]. The proposed pipeline results in a robust policy that can successfully solve tasks in the real-world without the manual burden of constructing environments in simulation [11], expensive human data collection [30] or real-world reinforcement learning [31].

## IV. EXPERIMENTS

In this section, we aim to answer the following questions:

A. Does integrating *URDFormer* in a real2sim2real pipeline

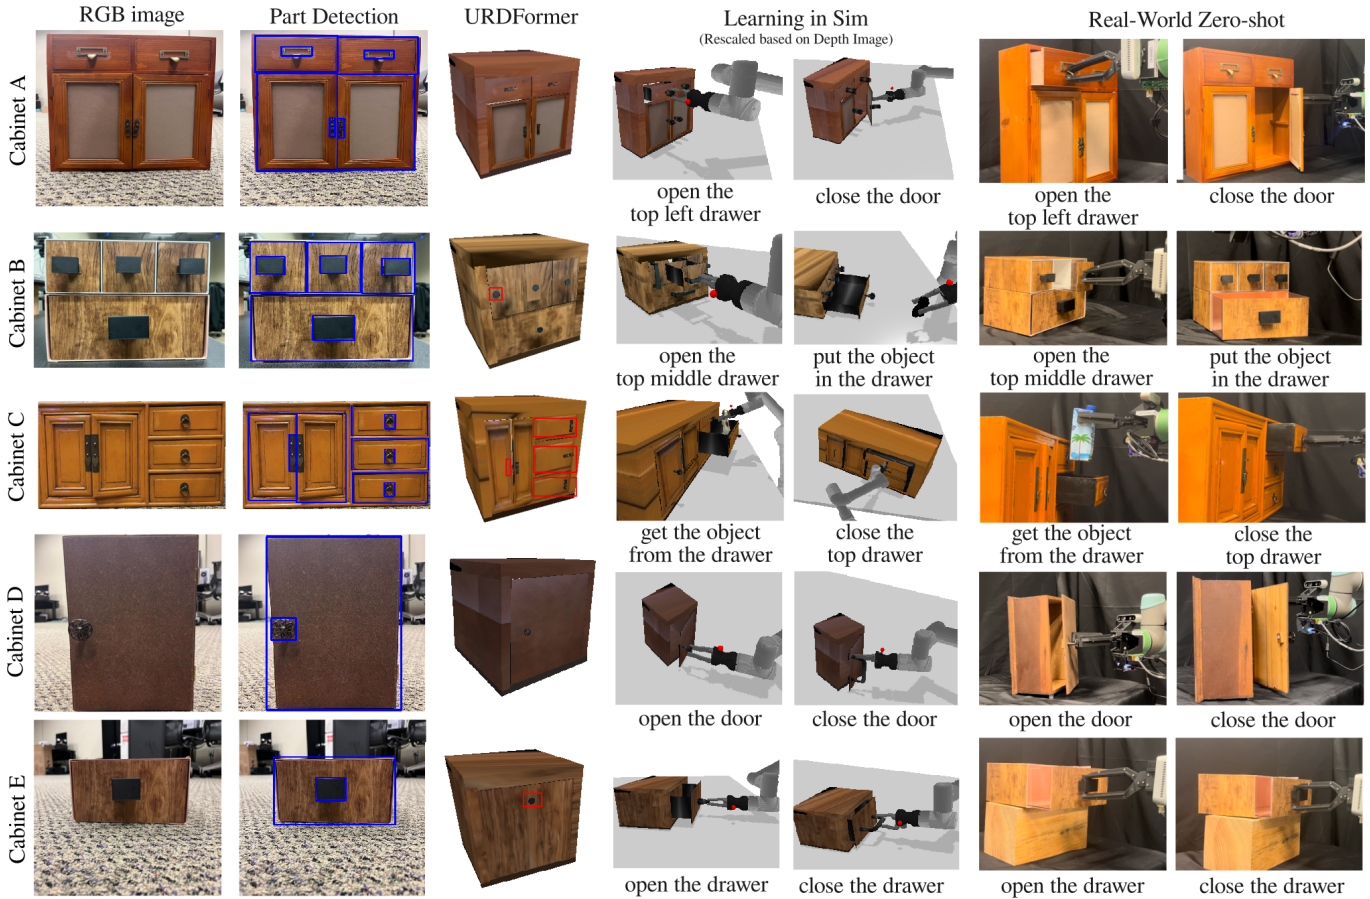


Fig. 5: **Qualitative Results for Real-world Robot Experiments:** An RGB image is pre-processed by detecting bounding boxes of relevant parts. The *URDFormer* then predicts the corresponding URDF of the cabinet. When importing the cabinet into the simulation, it is re-scaled using depth measurements. Furthermore, the real-world texture is cropped using the bounding boxes and projected onto the cabinet. This realistic simulation can then be used to generate massive data with the help of motion planning, ground truth information, and targeted domain randomization. Finally, we show that training a language-conditioned multi-task policy can be zero-shot transferred to the real world to solve several opening and closing tasks.

improve policies that can transfer zero-shot to the real world?

- B. Can *URDFormer* generate plausible and accurate simulation content from internet images?
- C. Can *URDFormer* generalize to diverse objects and scenes?
- D. Can *URDFormer* support different robots and tasks?

A. Does integrating *URDFormer* in a *real2sim2real* pipeline improve policies that can transfer zero-shot to the real world?

As described in Section III, *URDFormer* can be used to instantiate a real-to-simulation-to-real pipeline for robot learning. In this section, we describe a concrete instantiation of one such pipeline and provide a detailed evaluation of the resulting robotic behavior in the real world.

**Real-to-sim-to-real pipeline:** We implement our real-to-sim-to-real approach on a UR5 robot equipped with a custom-made 3D printed 2-fingered gripper and an Intel RealSense D435i mounted on the endeffector. We evaluate our pipeline on five different cabinets with varying sizes, shapes, textures, joint types (revolute and prismatic), and handles, with two tasks per cabinet. The details of each pipeline element is as

follows:

**1. Real-to-sim:** First, our system takes an image of the scene and uses a finetuned Grounding DINO [32] with model soup approach[33] (Detailed in IV-B2) to automatically detect the parts of the cabinet, *i.e.*, door, drawer, and handles. *URDFormer* then generates the corresponding URDF from the RGB image and the predicted bounding boxes using *URDFormer*, as described in Section II-C.

**2. Policy Learning in Sim:** The URDF is then imported into a physics simulator (PyBullet [34]) and scaled appropriately using depth measurement. Next, a motion planner, in this case cuRobo [24], generates trajectories that solve a variety of tasks, *e.g.*, closing/opening a drawer/door by utilizing privileged information in simulation. To account for possible prediction inaccuracies of *URDFormer* (red highlights in Fig 5) and to robustify the trained policy, we apply targeted randomization (TR) that randomizes the scene while maintaining its semantic configurations. In particular, doors, drawers, and handles are randomly replaced with their PartNet [19] equivalents while the size and base of the object are kept fixed. For each door or drawer that was generated in this process, the geometry was randomly replaced with alternate geometry from the

Task	Cabinet A		Cabinet B		Cabinet C	
	Open the top left drawer	close the door	put object in bottom drawer	open top middle drawer	get object from middle drawer	close the top drawer
OWL-ViT [29]	0/5	0/5	0/5	0/5	0/5	0/5
DR	0/5	0/5	0/5	0/5	0/5	0/5
URDFormer-ICP	2/5	1/5	—	3/5	1/5	3/5
URDFormer-TR	4/5	5/5	2/5	4/5	3/5	3/5

Task	Cabinet D		Cabinet E		Average
	Open the door	close the door	open the drawer	close the drawer	
OWL-ViT [29]	0/5	0/5	0/5	0/5	0/50
DR	0/5	2/5	2/5	5/5	9/50
URDFormer-ICP	0/5	3/5	3/5	2/5	24/45
URDFormer-TR	5/5	4/5	4/5	5/5	39/50

TABLE I: **Quantitative Results for our Real-world Robot Experiments:** Our Real2Sim2Real pipeline with the *URDFormer* and targeted (domain) randomization in simulation results in a 78% success rate across all tasks. Results are reported in success / number of trials.

same PartNet class, but rescaled to be the appropriate size. Additionally each handle or knob that was generated was similarly replaced with alternate geometry, but also randomly translated in the plane of the door or drawer that it was attached to. Textures are randomized by cropping out the real texture using the bounding boxes, generating variations by prompting Stable Diffusion [18] and fitting it back onto the shape. Finally, the RGB input is augmented by adding standard augmentations such as Gaussian noise and color jitter. After automatically collecting a dataset of successful simulation trajectories using the motion planner, we train a behavior cloning policy network that predicts end-effector poses from point clouds. The network architecture follows M2T2 [25] and predicts 6D end-effector poses from RGB pointclouds given language instructions specifying the task. Appendix D-C provides the full training procedure and architectural details.

**3. Sim-to-Real:** In order to transfer back to the real world, the policy takes an RGB pointcloud, current end-effector pose and a natural language instruction, and predicts the next end-effector pose, using a PD controller to execute these predictions.

We term this specific instantiation of our real2sim2real pipeline *URDFormer-TR* and stress that the data generation process, the policy network, and the transfer procedure are highly flexible and can be replaced depending on the exact problem setting.

**Baselines:** We compare *URDFormer-TR* with multiple variations of our real-to-simulation-real pipeline with varying degrees of access to real-world information. The results are shown in Table I.

(1) **OWL-ViT [29]:** First, we evaluate against a zero-shot vision-language model baseline. Inspired by VoxPoser [35], given a language instruction (Appendix D-B), we used the same open-vocabulary detector (OWL-ViT [29]) to predict bounding boxes for the prompted parts and handles that are important to the task. After mapping the detection to the observed pointcloud, a motion planner can then generate plans to solve the task directly in the real world.

(2) **Domain Randomization (DR) [36]:** Another approach is Domain Randomization (DR). We follow the augmentations from *URDFormer-TR* but on randomly generated cabinets with different configurations. Similarly to *URDFormer-TR*,

since we assume depth observations during inference, we also scale the generated cabinets to the real-world size. We follow the same approach for trajectory generation, policy training, and real-world transfer procedure of *URDFormer-TR*.

(3) ***URDFormer-ICP:*** This presents a learning-free, digital twin-style approach based on the Iterative Closest Point (ICP) [37] algorithm. First, we construct a simulation with the URDF created by the *URDFormer* and scale it according to the depth observation. We then use the ground truth simulation to compute end effector poses which we execute directly in the real world. When the cabinet’s pose changes, we use ICP to transform the computed end effector poses to the new cabinet pose.

**Real-world Results of Real-to-Sim-Real Training:** Running the zero-shot OWL-ViT, we find that the model fails to predict the fine-grained details required to solve the tasks, *i.e.*, ”top middle drawer”, ”right door”, and ”handles”. We visualize qualitative examples of results from OWL-ViT in Appendix D. Without localizing these regions of interest, the motion planning cannot solve the tasks leading to 0% success on all tasks.

While DR works surprisingly well on simple cabinets, *e.g.*, Cabinet E which only has a single drawer, it fails to solve any of the more complicated configurations. We observe *URDFormer-TR* outperforms DR by 40% on average, showing the benefits of targeting the domain randomization procedure to the real-world configuration.

*URDFormer-ICP* shows average performance across all opening and closing tasks. When tasking it with ”put object in bottom drawer”, we observe the limitations of the approach. Since the object is a lot smaller than the cabinet, ICP matches the pointcloud of the cabinet instead of the object. This results in an inability to transform the endeffector pose with respect to the graspable object and results in failure to solve the task. In general, neither baseline can reliably solve putting and getting objects in/from the drawer. While *URDFormer-TR* succeeds 50% of the time, it showcases the difficulty of the task and leaves space for future improvements.

Overall, the proposed real2sim2real pipeline using *URDFormer* and targeted randomization shows a **78%** success rate across all cabinet variations and tasks and an **85%** success rate on opening/closing tasks. Even though the *URDFormer*’s




Fig. 6: **Reality Gym**: A simulation environment with a variety of assets originated from internet images (black box) using *URDFormer*. We predict URDFs of internet images which can be loaded in any simulator. These URDFs are randomized with meshes from the Partner dataset. We introduce 4 main tasks: (1) Open any articulated parts (2) close any articulated parts (3) fetch objects and (4) collect objects



Fig. 7: **Generated Kitchen Scenes**: Examples of kitchen scenes predicted by *URDFormer* from internet images given labeled bounding boxes. Examples of failure parts are highlighted in red boxes.

predictions are not always accurate, as indicated by the red boxes in Fig5, the targeted randomization robustifies the policy network to allow for zero-shot real-world transfer more effectively than un-targeted simulation generation.

### B. Can *URDFormer* generate plausible and accurate simulation content from internet images?

1) *Paired Training Dataset Generation*: To synthesize our paired training dataset, we first procedurally generate a set of URDF representations of scenes in simulation both for global scenes like kitchens and for single objects like ovens, cabinets, and fridges (Fig 3). We then follow the procedure in II-B to control the data generation of paired images, generating a large dataset of simulation scenes and paired realistic RGB images (Fig 3). For objects with diverse parts, we find that depth-guided Stable Diffusion [18] often ignores the semantic details of local parts, leading to inconsistencies as visualized in the Appendix A-A). As described in Section II-B we generated a large and diverse set of texture templates by using images of existing textures downloaded from the internet to guide the depth-guided stable diffusion. During training, we then randomly chose one template texture and warped it back to the original part region using perspective transformation. Finally, we applied a stable diffusion in-painting model [18] to smooth the boundary of the parts and generate background content (Details are described in Appendix A-A). In total,

we generated approximately 118K image-URDF pairs across 7 categories of single articulated objects, and approximately 200K image-URDF pairs of global kitchen scenes.

2) *Object and Part Detection during Inverse Phase*: *URDFormer* takes both the RGB image and bounding boxes of the object parts and predicts URDFs. During the inverse phase, we adopt an off-the-shelf open vocabulary object detector GroundingDINO [32].



Fig. 8: Comparison among pretrained, finetuned and model soup GroundingDINO on cabinet dataset

However, if we directly apply GroundingDINO on detecting parts such as drawers and handles, the detection performance is unsatisfying with an F1 score of 53.4%. Instead, if we use the same generated dataset that was used to train *URDFormer* (Visualized in Appendix A) and finetune groundingDINO, we observe an improvement with an F1 score of 66.2%. However, We also observe that compared to the pretrained



GroundingDINO, the finetuned GroundingDINO often fails to detect parts that have unique shapes or patterns. Inspired by recent work Model Soup[33], we simply average the pretrained weights and the finetuned weights. This leads to surprising improvement, with an F1 score of 79.7%. We additionally apply post-processing to remove duplicated boxes. Figure 8 shows Model Soup’s influence on the bounding box detection.

3) *Real World Evaluation Datasets*: We create two types of test sets for the evaluation of *URDFormer*: (a) *Object-Level* set includes labeled URDFs of 300 internet images of individual objects from 5 categories including 100 cabinets, 50 ovens, 50 dishwashers, 50 fridges and 50 laundry machines. (b) *Kitchen* set includes URDFs of 54 internet images of kitchens, with 5-15 articulated objects per kitchen. For each scene, we manually label the bounding box for each object and its parts, as well as the URDF primitives including mesh types, parent id, positions, and scales relative to its parent. We used the mesh types such as “left door”, and “right door” to infer link axis and joint types. All the position values and scale values are discretized into 12 bins.

4) *Evaluation Metrics*: Evaluating entire scenes is challenging given the mixed structure and subjective nature of human labeling. We therefore measured accuracy of the predicted model using three individual sub-tasks: category accuracy, parent accuracy and spatial error. In order to compute these statistics, we must first align predicted objects in the scene to their ground truth counterparts. This is accomplished by computing the intersection-over-union (IOU) of the 2D detected boxes used to instantiate the scene objects and the ground-truth bounding boxes of all the objects in the scene. Hungarian matching is then used to assign each detected box to a single ground truth box. Note that if the number of predicted boxes is not the same as the number of ground truth boxes, there will be some false positives (predicted objects that do not correspond to any ground truth objects) or false negatives (ground truth objects that do not have any associated predictions). We therefore also record the overall precision and recall of the objects in the scene. Once we have a set of aligned detected and ground-truth boxes, we say that the category of a predicted object is correct if it matches the category of the assigned ground truth object. Similarly the predicted parent is correct if it matches the parent of the assigned ground truth object. Finally the spatial error is the average absolute error of the predicted discretized spatial coordinates. For larger objects, the model predicts four coordinates  $x_1, y_1, x_2, y_2$  representing the bounding box of the object relative to its parent. For small objects such as handles and knobs, we predict only the object center  $x_1, y_1$  and so the spatial error only considers these two values. Finally, we separate out these scores for high-level object predictions (cabinets, dishwashers, etc.) and low-level part predictions (doors, handles, etc.) so that the effects of various ablations on larger and smaller parts are more interpretable.

5) *Qualitative Results*: Figure 7 shows several examples of kitchens generated by *URDFormer* using internet images. While the model makes some mistakes, it is able to reproduce kitchen environments that largely match the structure of

the original images. Figures 9 and 10 show successful and unsuccessful reconstructions of individual objects and scenes respectively. These images also show results when ground-truth boxes are provided. Again, while the model makes some mistakes, it largely captures the overall configuration of the objects and scenes.

6) *Ablation Study*: We compare the prediction of *URDFormer* trained with (1) **Ours**: realistic texture generated by part-consistency stable diffusion (2) **Random**: random texture downloaded from Describable Textures Dataset [38] (3) **Sim**: random 3-channels RGB color and (4) **Selected**: carefully selected texture images that matches the object categories, such as wood texture (cabinet), metal texture (dishwasher) and so on. The results are shown in Table II. We observe that generated realistic texture is particularly helpful in global prediction, including identifying object types (cabinet or oven), parents (which wall the cabinet belongs to), and where to put the object. Surprisingly, we found the texture realism does not affect so much when predicting part structures and sometimes is slightly worse than random texture by 1% to 2%. This is likely due to using bounding box position features is sufficient for predicting simple low-level structures. For example, if a small box A is in the center of another box B, box B is likely a drawer instead of a door or a handle. On average, prediction using finetuned object detector performance is worse than using GT boxes due to detection error. However, one surprising observation is that using detected boxes helps slightly with identifying global object types. We hypothesize this is because boxes labeled by humans sometimes group multiple cabinets into one single box if they are close, making mesh prediction slightly challenging.

7) *Reality Gym*: With the ability to cheaply generate an arbitrary number of realistic simulation assets directly from the internet, we introduce *RealityGym*, a robot learning suite with a collection of realistic simulation assets and scenes created from real world RGB images using *URDFormer*. We provide an initial set of 300 objects (Cabinets, Ovens, Fridges, wahsers and Dishwashers) and 50 kitchen scenes from internet images, with future work looking to expand this into a bigger dataset. In addition, we also provide 84 meshes of cabinet frames, 20 door meshes, 59 drawers, 440 handles and 116 knobs from PartNet[19]. We can randomly incorporate these meshes into the initial URDFs to generate diverse scenes complete with articulated objects. We define 4 main tasks: (1) Open any articulated parts *i.e.* top middle drawer (2) Close any articulated parts (3) Fetch objects (4) Collect objects. We automatically generate tasks and their language descriptions, and use a motion planner (Curobo [24]) to complete the tasks. Fig 6 and Fig 7 show examples of robots performing in RealityGym on a variety of generated simulation environments and assets. Details about RealityGym can be found in Appendix C.

### C. Can *URDFormer* generalize to diverse objects and scenes?

In order to demonstrate the generalization capability of *URDFormer*, we used the same techniques discussed on Section II to create five additional object categories and four additional scene categories for qualitative evaluation. The




Fig. 9: Successful and unsuccessful examples comparing URDFFormer prediction on different articulated objects using (1) manually labeled bounding boxes and (2) bounding boxes provided by fine-tuned GroundingDINO. Image results of fine-tuned GroundingDINO (Red) compared with manually labeled boxes (Blue) are shown in the first row. Note that textures are removed for better visual comparison.




Fig. 10: Successful and unsuccessful examples on kitchens comparing URDFFormer prediction based on boxes generated by (1) manual labeling and (2) fine-tuned GroundingDINO. The first row also shows the boxes detected by fine-tuned GroundingDINO (Red) compared with manually labeled boxes (Blue)

		Kitchen				Object-Level												
		GT boxes				Finetuned Grounding DINO				GT Boxes				Finetuned Grounding DINO				
		Ours	Random	Sim	Selected	Ours	Random	Sim	Selected	Ours	Random	Sim	Selected	Ours	Random	Sim	Selected	
Global	Mesh Acc ( $\uparrow$ )	<b>0.578</b>	0.42	0.407	0.576	<b>0.603</b>	0.533	0.354	0.462	<b>0.740</b>	0.463	0.520	0.677	<b>0.688</b>	0.440	0.520	0.630	
	Parent Acc ( $\uparrow$ )	<b>0.833</b>	0.831	0.813	0.807	0.816	<b>0.819</b>	0.810	<b>0.819</b>	—	—	—	—	—	—	—	—	
	Spatial Err ( $\downarrow$ )	<b>0.809</b>	0.845	0.848	0.885	<b>0.987</b>	1.036	1.038	1.032	—	—	—	—	—	—	—	—	
	Recall ( $\uparrow$ )	1	1	1	1	0.726	0.726	0.726	0.726	—	—	—	—	—	—	—	—	
	Precision ( $\uparrow$ )	1	1	1	1	0.951	0.951	0.951	0.951	—	—	—	—	—	—	—	—	
Parts	Mesh Acc ( $\uparrow$ )	0.704	<b>0.719</b>	0.662	0.675	0.537	<b>0.558</b>	0.505	0.522	0.903	<b>0.927</b>	0.867	0.861	0.851	<b>0.865</b>	0.811	0.803	
	Parent Acc ( $\uparrow$ )	0.765	<b>0.773</b>	0.763	0.75	<b>0.711</b>	0.696	0.707	0.711	0.874	<b>0.878</b>	0.857	0.857	<b>0.826</b>	0.821	0.806	0.805	
	Spatial Err ( $\downarrow$ )	1.799	<b>1.699</b>	1.891	1.981	2.957	2.885	<b>2.798</b>	3.107	0.478	<b>0.420</b>	0.649	0.867	0.791	<b>0.753</b>	0.917	1.129	
	Recall ( $\uparrow$ )	1	1	1	1	0.495	0.495	0.495	0.495	1	1	1	1	0.853	0.832	0.832	0.832	
		Pred Precision ( $\uparrow$ )	1	1	1	1	0.927	0.927	0.927	0.927	1	1	1	1	0.986	0.987	0.987	0.987

TABLE II: **Ablation Study.** We analyze which part training with generated texture benefits the most by comparing URDFFormer trained with other textures, across GT boxes and boxes from the finetuned Grounding DINO detector.

new object categories are toilet, microwave, desk, laptop and chair, while the four scene categories are bedroom, bathroom, laundry room and study room. We trained a single new part model that incorporates these additional categories by adding approximately 6k training examples per object category to the original training dataset. We took the same approach to train a single global scene model and added approximately 10k training examples for the bathroom and bedroom. We found that the laundry room and study could be adequately captured with no new global scene examples, as these categories only contain objects that are present in the other scene categories. Figure 11 shows example training data generated by the forward pipeline, while Figure 12 shows qualitative examples of objects and scenes inferred by a *URDFFormer* trained on this data. More examples are available in Appendix F.

#### D. Can URDFFormer support different robots and tasks?

To demonstrate that *URDFFormer* supports multi-step tasks and different robots, we train an additional policy for a Stretch robot to place an object in a desk drawer as shown in Figure 13. Figure 14 shows the data generation process for training the mobile robot’s policy. First, a single image of the robot’s environment is passed to a pretrained *URDFFormer* model to produce a scene description of the desk. Then, this scene description is used to generate training data in simulation that can instruct the robot how to accomplish its multi-stage objective. When generating this data, we use an inpainting model [18] to reduce the sim2real gap by inpainting the pixels covered by the simulated object overlaid onto the original image. Finally the robot policy, which predicts a per-step affordance map from an initial image, is trained on this new



Fig. 11: The forward phase can be applied to additional objects and create a diverse dataset for training *URDFormer*.

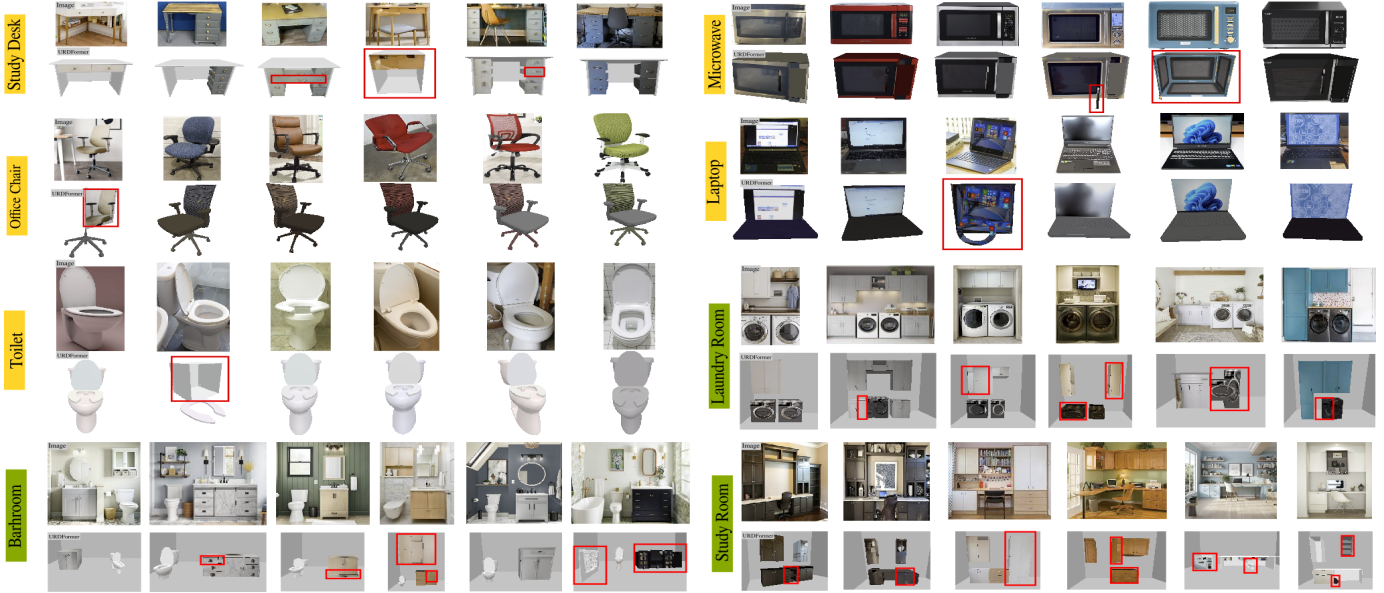


Fig. 12: Qualitative study of *URDFormer* generalization to other scenes and objects with successful and unsuccessful (highlighted in red) examples. In particular, We train *URDFormer* on the new dataset with added categories shown in Fig11 and evaluate on internet images for both object-level and global prediction. Interestingly, we did not train *URDFormer* on the Laundry Room and Study Room but found *URDFormer* can generalize to these unseen categories, which is likely due to the two-stage training of *URDFormer* Global and *URDFormer* Part. Please note that *URDFormer* does not reconstruct accurate meshes or predict properties such as friction or mass, which is detailed in the limitation and future work sections.



Fig. 13: We trained a Stretch robot on a multi-step task "Clean Up the Table Surface" using *URDFormer* prediction

generated data. This policy is implemented as a UNet, which takes the initial observation, as well as the task embedding and predicts multiple object affordance maps. These affordance maps are used with a motion planner to guide the robot at each step. Details on training the policy can be found in Appendix D-C.

## V. RELATED WORK

### A. Asset Creation for Robot Manipulation

Constructing realistic and diverse assets for robot learning and control has been an important goal for many years. Early efforts [39, 40, 41, 42] used geometric approaches to reconstruct static 3D scenes. Recently learned approaches [43, 44, 45] have in some cases improved the accuracy and visual realism of these approaches. While these techniques can offer incredible visual realism, they do not usually produce articulated models that allow for dynamic interaction.

To address this, many prior works have explored utilizing assets that are manually curated [9, 46, 47, 48], procedural-generated [11, 49], or scanned from the real world [50, 51, 52], to train robot manipulation policies. While the learned policies are supposed to apply directly to test scenes, there could be challenging scenes in which the systems fail to faithfully recognize and interact with the scene due to the scene complexity or sim-to-real gaps. In our work, we resort to a targeted system that first reconstructs a URDF given the test scene and then performs scene-specific training to learn a specialized model in the scene. There is also a rich body of literature on building digital twins of articulated objects from the real

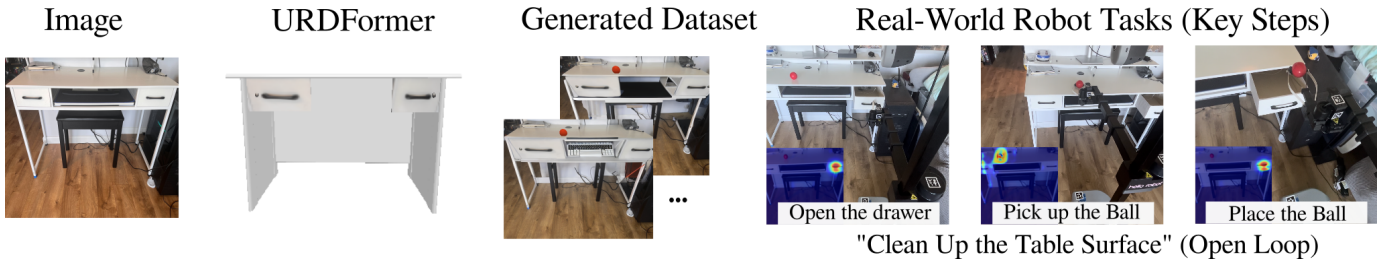


Fig. 14: *URDFormer* can be applied to a different robot such as a Stretch Robot to perform a multi-step task such as "clean up the table surface". We apply *URDFormer* to predict the URDF of a study desk, and render a dataset to train a vision-based policy to predict affordance map at each step.

world, either based on passive visual observations such as 2D multi-view images or videos [53, 54, 55, 56, 57], 3D RGB-D, depth images or point clouds [58, 59, 60, 61, 62], full-scene 3D scans [63, 64, 65], or through interactive perception which requires few-shot interactions with the object [66, 67, 68, 69, 70, 71, 72]. While in some cases, these methods can successfully produce high quality models, they require either video, depth sensing or interaction with the objection which is more difficult to aquire and is rarely available in online internet data. Our work, instead, advocates for a faster and cheaper pipeline where we reconstruct a simple URDF model, which is sufficient for downstream learning, from a single image without the need for laborious scanning or object interaction.

### B. Synthetic Data Augmentation with Generative Models

Everyone can freely generate realistic image content nowadays with the booming development of the latest generative AI technology [73, 18]. Researchers have also explored the use of such powerful generative models in creating large-scale high-quality synthetic data for training perception systems in different applications, such as image classification [74, 75], object segmentation [76, 77, 78], and representation learning [79, 80, 81]. In the field of robotics, several studies have attempted to reduce sim-to-real gaps and broaden the domain coverage using synthetic data augmented with generative models [82, 83, 84, 85, 86, 87]. In our case, we look to augment synthetic asset renderings with realistic textures produced by generative models. While recent studies have proposed controllable image generation algorithms such as ControlNet [88] and 3D mesh texturing methods [89, 90], we find them lacking in preserving part-level structure details and consequently propose a novel controlled texture generation method to produce realistic synthetic data for training.

## VI. LIMITATIONS AND FUTURE WORK

Please see the full list of limitations and future work in Appendix F.

**Part Detection** *URDFormer* relies on the performance of bounding box detection. Although the finetuned Grounding DINO improves performance than the pretrained model, there is still a gap for improvement, especially on global scene detection.

**Texture and Meshes** *URDFormer* focuses on predicting kinematic URDF structures and uses predefined meshes that might not match the real-world scenes. To apply textures, we simply assume all parts are rectangular shapes, and use the bounding box of each object part to crop the image and import it into a uv map template. However, this does not work

for irregular meshes such as a donut-shape door, or when the object in the image is tilted.

**Limited URDF Primitives** *URDFormer* currently only supports articulated objects that have limited joint types such as prismatic and revolute, and cannot predict complex objects such as cars and lamps.

**Link Collisions** *URDFormer* only predicts URDF primitives for each bounding box, which sometimes leads to a collision between two links. Further post-processing is required to resolve this issue.

**Multiple Trained Components** Our pipeline is not trained end-to-end and consists of multiple learning components. While this increases the complexity of the system, it is necessary to ensure consistency when using the generative model, and make the most use of existing pretrained components.

**Inferred Physical Properties** Our system does not presently attempt to infer physical properties such as mass, inertial moments or friction directly from observations in the scene. In theory the visual information present in the images may allow for a rough approximations of these quantities. We expect this to be a productive direction for future work.

## VII. CONCLUSION

We present a scalable pipeline for creating articulated simulation assets from real-world images. In particular, we introduce a forward-inverse framework that generates realistic and consistent images of articulated objects and trains a transformer-based network *URDFormer* to predict their corresponding URDFs. We present RealityGym, a robot learning suite with realistic simulation assets generated from real-world images. Additionally, we show that the predicted URDFs with targeted domain randomization enable better zero-shot performance in the real-world. Our pipeline provides the first step towards cheaper and scalable realistic scene generation for robot learning.

## VIII. ACKNOWLEDGEMENT

We thank all members of UW Robotics Lab for all the support and discussions. In particular, we thank all the lovely people from WEIRD Lab and RSE lab for insightful feedback. We especially thank Mateo Guaman Castro for his detailed feedback on the paper draft, and Marcel Torne for fruitful discussion. We also thank Microsoft for UR5 robots, and Adam Fishman, Rohan Baijal, Tyler Han and Yuquan Deng for helping me move these two heavy robots, and Jacob Berg for letting me borrow his real-sense camera cable. The work was also funded in part by Amazon Science Hub.

## REFERENCES

- [1] Xiaowei Jia, Yiqun Xie, Sheng Li, Shengyu Chen, Jacob Zwart, Jeffrey M. Sadler, Alison P. Appling, Samantha Oliver, and Jordan S. Read. Physics-guided machine learning from simulation data: An application in modeling lake and river systems. In James Bailey, Pauli Miettinen, Yun Sing Koh, Dacheng Tao, and Xindong Wu, editors, *IEEE International Conference on Data Mining, ICDM 2021, Auckland, New Zealand, December 7-10, 2021*, pages 270–279. IEEE, 2021. doi: 10.1109/ICDM51629.2021.00037. URL <https://doi.org/10.1109/ICDM51629.2021.00037>.
- [2] Mark S. Alber, Adrian Buganza Tepole, William R. Cannon, Suvranu De, Salvador Dura-Bernal, Krishna C. Garikipati, George E. Karniadakis, William W. Lytton, Paris Perdikaris, Linda R. Petzold, and Ellen Kuhl. Integrating machine learning and multiscale modeling - perspectives, challenges, and opportunities in the biological, biomedical, and behavioral sciences. *npj Digit. Medicine*, 2, 2019. doi: 10.1038/s41746-019-0193-y. URL <https://doi.org/10.1038/s41746-019-0193-y>.
- [3] Jack Collins, Shelvin Chand, Anthony Vanderkop, and David Howard. A review of physics simulators for robotic applications. *IEEE Access*, 9:51416–51431, 2021. doi: 10.1109/ACCESS.2021.3068769. URL <https://doi.org/10.1109/ACCESS.2021.3068769>.
- [4] Yashraj S. Narang, Kier Storey, Iretoiyo Akinola, Miles Macklin, Philipp Reist, Lukasz Wawrzyniak, Yunrong Guo, Ádám Moravánszky, Gavriel State, Michelle Lu, Ankur Handa, and Dieter Fox. Factory: Fast contact for robotic assembly. In Kris Hauser, Dylan A. Shell, and Shoudong Huang, editors, *Robotics: Science and Systems XVIII, New York City, NY, USA, June 27 - July 1, 2022*, 2022. doi: 10.15607/RSS.2022.XVIII.035. URL <https://doi.org/10.15607/RSS.2022.XVIII.035>.
- [5] Matthias Müller, Vincent Casser, Jean Lahoud, Neil Smith, and Bernard Ghanem. Sim4cv: A photo-realistic simulator for computer vision applications. *Int. J. Comput. Vis.*, 126(9):902–919, 2018. doi: 10.1007/s11263-018-1073-7. URL <https://doi.org/10.1007/s11263-018-1073-7>.
- [6] Marius Memmel, Roman Bachmann, and Amir Zamir. Modality-invariant visual odometry for embodied vision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21549–21559, 2023. URL [https://openaccess.thecvf.com/content/CVPR2023/papers/Memmel\\_Modality-Invariant\\_Visual\\_Odometry\\_for\\_Embodied\\_Vision\\_CVPR\\_2023\\_paper.pdf](https://openaccess.thecvf.com/content/CVPR2023/papers/Memmel_Modality-Invariant_Visual_Odometry_for_Embodied_Vision_CVPR_2023_paper.pdf).
- [7] Andrew Szot, Alexander Clegg, Eric Undersander, Erik Wijmans, Yili Zhao, John Turner, Noah Maestre, Mustafa Mukadam, Devendra Singh Chaplot, Oleksandr Maksymets, Aaron Gokaslan, Vladimir Vondrus, Sameer Dharur, Franziska Meier, Wojciech Galuba, Angel X. Chang, Zsolt Kira, Vladlen Koltun, Jitendra Malik, Manolis Savva, and Dhruv Batra. Habitat 2.0: Training home assistants to rearrange their habitat. In Marc’Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan, editors, *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 251–266, 2021. URL <https://proceedings.neurips.cc/paper/2021/hash/021bbc7ee20b71134d53e20206bd6feb-Abstract.html>.
- [8] Karmesh Yadav, Ram Ramrakhya, Santhosh Kumar Ramakrishnan, Théophile Gervet, John Turner, Aaron Gokaslan, Noah Maestre, Angel Xuan Chang, Dhruv Batra, Manolis Savva, Alexander William Clegg, and Devendra Singh Chaplot. Habitat-matterport 3d semantics dataset. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023*, pages 4927–4936. IEEE, 2023. doi: 10.1109/CVPR52729.2023.00477. URL <https://doi.org/10.1109/CVPR52729.2023.00477>.
- [9] Eric Kolve, Roozbeh Mottaghi, Winson Han, Eli VanderBilt, Luca Weihs, Alvaro Herrasti, Matt Deitke, Kiana Ehsani, Daniel Gordon, Yuke Zhu, et al. Ai2-thor: An interactive 3d environment for visual ai. *arXiv preprint arXiv:1712.05474*, 2017. URL <https://arxiv.org/pdf/1712.05474.pdf>.
- [10] Arjun Gupta, Max Shepherd, and Saurabh Gupta. Predicting motion plans for articulating everyday objects. In *International Conference on Robotics and Automation (ICRA)*, 2023. URL <https://ieeexplore.ieee.org/document/10160752>.
- [11] Matt Deitke, Eli VanderBilt, Alvaro Herrasti, Luca Weihs, Kiana Ehsani, Jordi Salvador, Winson Han, Eric Kolve, Aniruddha Kembhavi, and Roozbeh Mottaghi. Proctor: Large-scale embodied ai using procedural generation. *Advances in Neural Information Processing Systems*, 35:5982–5994, 2022. URL <https://arxiv.org/pdf/2206.06994.pdf>.
- [12] Ben Poole, Ajay Jain, Jonathan T. Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023. URL <https://openreview.net/pdf?id=FjNys5c7VyY>.
- [13] Yufei Wang, Zhou Xian, Feng Chen, Tsun-Hsuan Wang, Yian Wang, Katerina Fragkiadaki, Zackory Erickson, David Held, and Chuang Gan. Robogen: Towards unleashing infinite data for automated robot learning via generative simulation. *arXiv preprint arXiv:2311.01455*, 2023. URL <https://arxiv.org/pdf/2311.01455.pdf>.
- [14] Lirui Wang, Yiyang Ling, Zhecheng Yuan, Mohit Shridhar, Chen Bao, Yuzhe Qin, Bailin Wang, Huazhe Xu, and Xiaolong Wang. Gensim: Generating robotic simulation tasks via large language models. In *Arxiv*, 2023. URL <https://arxiv.org/pdf/2310.01361.pdf>.
- [15] Minghua Liu, Chao Xu, Haian Jin, Linghao Chen, Zexiang Xu, Hao Su, et al. One-2-3-45: Any single image to 3d mesh in 45 seconds without per-shape optimization. *arXiv preprint arXiv:2306.16928*, 2023. URL

- <https://arxiv.org/pdf/2306.16928.pdf>.
- [16] Chuan Fang, Xiaotao Hu, Kunming Luo, and Ping Tan. Ctrl-room: Controllable text-to-3d room meshes generation with layout constraints. *arXiv preprint arXiv:2310.03602*, 2023. URL <https://arxiv.org/pdf/2310.03602.pdf>.
- [17] Alexander Raistrick, Lahav Lipson, Zeyu Ma, Lingjie Mei, Mingzhe Wang, Yiming Zuo, Karhan Kayan, Hongyu Wen, Beining Han, Yihan Wang, Alejandro Newell, Hei Law, Ankit Goyal, Kaiyu Yang, and Jia Deng. Infinite photorealistic worlds using procedural generation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023*, pages 12630–12641. IEEE, 2023. doi: 10.1109/CVPR52729.2023.01215. URL <https://doi.org/10.1109/CVPR52729.2023.01215>.
- [18] 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. URL [https://openaccess.thecvf.com/content/CVPR2022/papers/Rombach\\_High-Resolution\\_Image\\_Synthesis\\_With\\_Latent\\_Diffusion\\_Models\\_CVPR\\_2022\\_paper.pdf](https://openaccess.thecvf.com/content/CVPR2022/papers/Rombach_High-Resolution_Image_Synthesis_With_Latent_Diffusion_Models_CVPR_2022_paper.pdf).
- [19] Kaichun Mo, Shilin Zhu, Angel X. Chang, Li Yi, Subarna Tripathi, Leonidas J. Guibas, and Hao Su. Partnet: A large-scale benchmark for fine-grained and hierarchical part-level 3d object understanding. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 909–918. Computer Vision Foundation / IEEE, 2019. doi: 10.1109/CVPR.2019.00100. URL [http://openaccess.thecvf.com/content\\_CVPR\\_2019/html/Mo\\_PartNet\\_A\\_Large-Scale\\_Benchmark\\_for\\_Fine-Grained\\_and\\_Hierarchical\\_Part-Level\\_3D\\_CVPR\\_2019\\_paper.html](http://openaccess.thecvf.com/content_CVPR_2019/html/Mo_PartNet_A_Large-Scale_Benchmark_for_Fine-Grained_and_Hierarchical_Part-Level_3D_CVPR_2019_paper.html).
- [20] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. URL <https://arxiv.org/pdf/2010.11929.pdf>.
- [21] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. URL <https://arxiv.org/pdf/1703.06870.pdf>.
- [22] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. URL [https://proceedings.neurips.cc/paper\\_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf).
- [23] Jing Kang Yang, Wenxuan Peng, Xiangtai Li, Zujin Guo, Liangyu Chen, Bo Li, Zheng Ma, Kaiyang Zhou, Wayne Zhang, Chen Change Loy, and Ziwei Liu. Panoptic video scene graph generation. In *CVPR, 2023*. URL [https://openaccess.thecvf.com/content/CVPR2023/papers/Yang\\_Panoptic\\_Video\\_Scene\\_Graph\\_Generation\\_CVPR\\_2023\\_paper.pdf](https://openaccess.thecvf.com/content/CVPR2023/papers/Yang_Panoptic_Video_Scene_Graph_Generation_CVPR_2023_paper.pdf).
- [24] Balakumar Sundaralingam, Siva Kumar Sastry Hari, Adam Fishman, Caelan Garrett, Karl Van Wyk, Valts Blukis, Alexander Millane, Helen Oleynikova, Ankur Handa, Fabio Ramos, Nathan Ratliff, and Dieter Fox. curobo: Parallelized collision-free minimum-jerk robot motion generation, 2023. URL <https://arxiv.org/pdf/2310.17274.pdf>.
- [25] Wentao Yuan, Adithyavairavan Murali, Arsalan Mousavian, and Dieter Fox. M2t2: Multi-task masked transformer for object-centric pick and place. In *7th Annual Conference on Robot Learning*, 2023. URL <https://arxiv.org/pdf/2311.00926.pdf>.
- [26] Marius Memmel, Puze Liu, Davide Tateo, and Jan Peters. Dimensionality reduction and prioritized exploration for policy search. In *International Conference on Artificial Intelligence and Statistics*, pages 2134–2157. PMLR, 2022. URL <https://proceedings.mlr.press/v151/memmel22a.html>.
- [27] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017. URL <https://arxiv.org/pdf/1707.06347.pdf>.
- [28] Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In Jennifer G. Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pages 1856–1865. PMLR, 2018. URL <http://proceedings.mlr.press/v80/haarnoja18b.html>.
- [29] Matthias Minderer, Alexey Gritsenko, Austin Stone, Maxim Neumann, Dirk Weissenborn, Alexey Dosovitskiy, Aravindh Mahendran, Anurag Arnab, Mostafa Dehghani, Zhuoran Shen, et al. Simple open-vocabulary object detection. In *European Conference on Computer Vision*, 2022. URL [https://www.ecva.net/papers/eccv\\_2022/papers\\_ECCV/papers/136700714.pdf](https://www.ecva.net/papers/eccv_2022/papers_ECCV/papers/136700714.pdf).
- [30] Open X-Embodiment Collaboration, Abhishek Padalkar, Acorn Pooley, Ajinkya Jain, Alex Bewley, Alex Herzog, Alex Irpan, Alexander Khazatsky, Anant Rai, Anikait Singh, Anthony Brohan, Antonin Raffin, Ayzaan Wahid, Ben Burgess-Limerick, Beomjoon Kim, Bernhard Schölkopf, Brian Ichter, Cewu Lu, Charles Xu, Chelsea Finn, Chenfeng Xu, Cheng Chi, Chenguang Huang, Christine Chan, Chuer Pan, Chuyuan Fu, Coline Devin, Danny Driess, Deepak Pathak, Dhruv Shah, Dieter Büchler, Dmitry Kalashnikov, Dorsa Sadigh, Edward Johns, Federico Ceola, Fei Xia, Freek Stulp, Gaoyue Zhou, Gaurav S. Sukhatme, Gautam Salhotra, Ge Yan, Giulio Schiavi, Hao Su, Hao-Shu Fang, Haochen Shi, Heni Ben Amor, Henrik I Christensen, Hiroki Furuta, Homer Walke, Hongjie Fang, Igor Mordatch, Ilija Radosavovic, Isabel Leal, Jacky Liang, Jaehyung Kim,

- Jan Schneider, Jasmine Hsu, Jeannette Bohg, Jeffrey Bingham, Jiajun Wu, Jialin Wu, Jianlan Luo, Jiayuan Gu, Jie Tan, Jihoon Oh, Jitendra Malik, Jonathan Tompson, Jonathan Yang, Joseph J. Lim, João Silvério, Junhyek Han, Kanishka Rao, Karl Pertsch, Karol Hausman, Keegan Go, Keerthana Gopalakrishnan, Ken Goldberg, Kendra Byrne, Kenneth Oslund, Kento Kawaharazuka, Kevin Zhang, Keyvan Majd, Krishan Rana, Krishnan Srinivasan, Lawrence Yunliang Chen, Lerrel Pinto, Liam Tan, Lionel Ott, Lisa Lee, Masayoshi Tomizuka, Maximilian Du, Michael Ahn, Mingtong Zhang, Mingyu Ding, Mohan Kumar Srirama, Mohit Sharma, Moo Jin Kim, Naoaki Kanazawa, Nicklas Hansen, Nicolas Heess, Nikhil J Joshi, Niko Suenderhauf, Norman Di Palo, Nur Muhammad Mahi Shafullah, Oier Mees, Oliver Kroemer, Pannag R Sanketi, Paul Wohlhart, Peng Xu, Pierre Sermanet, Priya Sundaesan, Quan Vuong, Rafael Rafailov, Ran Tian, Ria Doshi, Roberto Martín-Martín, Russell Mendonca, Rutav Shah, Ryan Hoque, Ryan Julian, Samuel Bustamante, Sean Kirmani, Sergey Levine, Sherry Moore, Shikhar Bahl, Shivin Dass, Shuran Song, Sichun Xu, Siddhant Halder, Simeon Adebola, Simon Guist, Soroush Nasiriany, Stefan Schaal, Stefan Welker, Stephen Tian, Sudeep Dasari, Suneel Belkhale, Takayuki Osa, Tatsuya Harada, Tatsuya Matsushima, Ted Xiao, Tianhe Yu, Tianli Ding, Todor Davchev, Tony Z. Zhao, Travis Armstrong, Trevor Darrell, Vidhi Jain, Vincent Vanhoucke, Wei Zhan, Wenxuan Zhou, Wolfram Burgard, Xi Chen, Xiaolong Wang, Xinghao Zhu, Xuanlin Li, Yao Lu, Yevgen Chebotar, Yifan Zhou, Yifeng Zhu, Ying Xu, Yixuan Wang, Yonatan Bisk, Yoonyoung Cho, Youngwoon Lee, Yuchen Cui, Yueh hua Wu, Yujin Tang, Yuke Zhu, Yunzhu Li, Yusuke Iwasawa, Yutaka Matsuo, Zhuo Xu, and Zichen Jeff Cui. Open X-Embodiment: Robotic learning datasets and RT-X models. <https://arxiv.org/abs/2310.08864>, 2023.
- [31] Abhishek Gupta, Justin Yu, Tony Z. Zhao, Vikash Kumar, Aaron Rovinsky, Kelvin Xu, Thomas Devlin, and Sergey Levine. Reset-free reinforcement learning via multi-task learning: Learning dexterous manipulation behaviors without human intervention. In *IEEE International Conference on Robotics and Automation, ICRA 2021, Xi'an, China, May 30 - June 5, 2021*, pages 6664–6671. IEEE, 2021. doi: 10.1109/ICRA48506.2021.9561384. URL <https://doi.org/10.1109/ICRA48506.2021.9561384>.
- [32] Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Chunyuan Li, Jianwei Yang, Hang Su, Jun Zhu, et al. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. *arXiv preprint arXiv:2303.05499*, 2023. URL <https://arxiv.org/pdf/2303.05499.pdf>.
- [33] Mitchell Wortsman, Gabriel Ilharco, Samir Ya Gadre, Rebecca Roelofs, Raphael Gontijo-Lopes, Ari S Morcos, Hongseok Namkoong, Ali Farhadi, Yair Carmon, Simon Kornblith, et al. Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time. In *International Conference on Machine Learning*, pages 23965–23998. PMLR, 2022.
- [34] Erwin Coumans and Yunfei Bai. Pybullet, a python module for physics simulation for games, robotics and machine learning. <http://pybullet.org>, 2016–2021.
- [35] 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.
- [36] 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. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2017, Vancouver, BC, Canada, September 24-28, 2017*, pages 23–30. IEEE, 2017. doi: 10.1109/IROS.2017.8202133. URL <https://doi.org/10.1109/IROS.2017.8202133>.
- [37] Paul J Besl and Neil D McKay. Method for registration of 3-d shapes. In *Sensor fusion IV: control paradigms and data structures*, volume 1611, pages 586–606. Spie, 1992. URL [https://www.researchgate.net/publication/3191994\\_A\\_method\\_for\\_registration\\_of\\_3-D\\_shapes\\_IEEE\\_Trans\\_Pattern\\_Anal\\_Mach\\_Intell](https://www.researchgate.net/publication/3191994_A_method_for_registration_of_3-D_shapes_IEEE_Trans_Pattern_Anal_Mach_Intell).
- [38] M. Cimpoi, S. Maji, I. Kokkinos, S. Mohamed, , and A. Vedaldi. Describing textures in the wild. In *Proceedings of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2014. URL <https://arxiv.org/pdf/1311.3618.pdf>.
- [39] Shahram Izadi, David Kim, Otmar Hilliges, David Molyneaux, Richard Newcombe, Pushmeet Kohli, Jamie Shotton, Steve Hodges, Dustin Freeman, Andrew Davison, et al. Kinectfusion: real-time 3d reconstruction and interaction using a moving depth camera. In *Proceedings of the 24th annual ACM symposium on User interface software and technology*, pages 559–568, 2011.
- [40] Sameer Agarwal, Yasutaka Furukawa, Noah Snavely, Ian Simon, Brian Curless, Steven M Seitz, and Richard Szeliski. Building rome in a day. *Communications of the ACM*, 54(10):105–112, 2011.
- [41] Peter Henry, Michael Krainin, Evan Herbst, Xiaofeng Ren, and Dieter Fox. Rgb-d mapping: Using kinect-style depth cameras for dense 3d modeling of indoor environments. *The international journal of Robotics Research*, 31(5):647–663, 2012.
- [42] Raul Mur-Artal, Jose Maria Martinez Montiel, and Juan D Tardos. Orb-slam: a versatile and accurate monocular slam system. *IEEE transactions on robotics*, 31(5):1147–1163, 2015.
- [43] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. Deepsdf: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 165–174, 2019.
- [44] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for

- view synthesis. *Communications of the ACM*, 65(1):99–106, 2021.
- [45] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics (ToG)*, 42(4):1–14, 2023.
- [46] Fanbo Xiang, Yuzhe Qin, Kaichun Mo, Yikuan Xia, Hao Zhu, Fangchen Liu, Minghua Liu, Hanxiao Jiang, Yifu Yuan, He Wang, et al. Sapien: A simulated part-based interactive environment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11097–11107, 2020. URL [https://openaccess.thecvf.com/content\\_CVPR\\_2020/papers/Xiang\\_SAPIEN\\_A\\_Simulated\\_Part-Based\\_Interactive\\_Environment\\_CVPR\\_2020\\_paper.pdf](https://openaccess.thecvf.com/content_CVPR_2020/papers/Xiang_SAPIEN_A_Simulated_Part-Based_Interactive_Environment_CVPR_2020_paper.pdf).
- [47] Andrew Szot, Alexander Clegg, Eric Undersander, Erik Wijmans, Yili Zhao, John Turner, Noah Maestre, Mustafa Mukadam, Devendra Singh Chaplot, Oleksandr Maksymets, et al. Habitat 2.0: Training home assistants to rearrange their habitat. *Advances in Neural Information Processing Systems*, 34:251–266, 2021.
- [48] Matt Deitke, Dustin Schwenk, Jordi Salvador, Luca Weihs, Oscar Michel, Eli VanderBilt, Ludwig Schmidt, Kiana Ehsani, Aniruddha Kembhavi, and Ali Farhadi. Objaverse: A universe of annotated 3d objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13142–13153, 2023.
- [49] Yue Yang, Fan-Yun Sun, Luca Weihs, Eli VanderBilt, Alvaro Herrasti, Winson Han, Jiajun Wu, Nick Haber, Ranjay Krishna, Lingjie Liu, et al. Holodeck: Language guided generation of 3d embodied ai environments. *arXiv preprint arXiv:2312.09067*, 2023.
- [50] Roberto Martín-Martín, Clemens Eppner, and Oliver Brock. The rbo dataset of articulated objects and interactions. *The International Journal of Robotics Research*, 38(9):1013–1019, 2019.
- [51] Fei Xia, William B Shen, Chengshu Li, Priya Kasimbeg, Micael Edmond Tchammi, Alexander Toshev, Roberto Martín-Martín, and Silvio Savarese. Interactive gibbon benchmark: A benchmark for interactive navigation in cluttered environments. *IEEE Robotics and Automation Letters*, 5(2):713–720, 2020.
- [52] Liu Liu, Wenqiang Xu, Haoyuan Fu, Sucheng Qian, Qiaojun Yu, Yang Han, and Cewu Lu. Akb-48: A real-world articulated object knowledge base. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14809–14818, 2022.
- [53] Gengshan Yang, Deqing Sun, Varun Jampani, Daniel Vlasic, Forrester Cole, Huiwen Chang, Deva Ramanan, William T Freeman, and Ce Liu. Lasr: Learning articulated shape reconstruction from a monocular video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15980–15989, 2021.
- [54] Shengyi Qian, Linyi Jin, Chris Rockwell, Siyi Chen, and David F. Fouhey. Understanding 3d object articulation in internet videos. In *CVPR*, 2022.
- [55] Fangyin Wei, Rohan Chabra, Lingni Ma, Christoph Lassner, Michael Zollhöfer, Szymon Rusinkiewicz, Chris Sweeney, Richard Newcombe, and Mira Slavcheva. Self-supervised neural articulated shape and appearance models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15816–15826, 2022.
- [56] Nick Heppert, Muhammad Zubair Irshad, Sergey Zakharov, Katherine Liu, Rares Andrei Ambrus, Jeannette Bohg, Abhinav Valada, and Thomas Kollar. Carto: Category and joint agnostic reconstruction of articulated objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21201–21210, 2023.
- [57] Jiayi Liu, Ali Mahdavi-Amiri, and Manolis Savva. Paris: Part-level reconstruction and motion analysis for articulated objects. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 352–363, 2023.
- [58] Xiaogang Wang, Bin Zhou, Yahao Shi, Xiaowu Chen, Qingping Zhao, and Kai Xu. Shape2motion: Joint analysis of motion parts and attributes from 3d shapes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8876–8884, 2019.
- [59] Zihao Yan, Ruizhen Hu, Xingguang Yan, Luanmin Chen, Oliver Van Kaick, Hao Zhang, and Hui Huang. Rpmnet: recurrent prediction of motion and parts from point cloud. *arXiv preprint arXiv:2006.14865*, 2020.
- [60] Yijia Weng, He Wang, Qiang Zhou, Yuzhe Qin, Yueqi Duan, Qingnan Fan, Baoquan Chen, Hao Su, and Leonidas J Guibas. Captra: Category-level pose tracking for rigid and articulated objects from point clouds. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13209–13218, 2021.
- [61] Hameed Abdul-Rashid, Miles Freeman, Ben Abbatemateo, George Konidaris, and Daniel Ritchie. Learning to infer kinematic hierarchies for novel object instances. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 8461–8467. IEEE, 2022.
- [62] Eric Heiden, Ziang Liu, Vibhav Vineet, Erwin Coumans, and Gaurav S Sukhatme. Inferring articulated rigid body dynamics from rgb-d video. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 8383–8390. IEEE, 2022.
- [63] Zhengqin Li, Ting-Wei Yu, Shen Sang, Sarah Wang, Meng Song, Yuhan Liu, Yu-Ying Yeh, Rui Zhu, Nitesh B. Gundavarapu, Jia Shi, Sai Bi, Hong-Xing Yu, Zexiang Xu, Kalyan Sunkavalli, Milos Hasan, Ravi Ramamoorthi, and Manmohan Chandraker. Openrooms: An open framework for photorealistic indoor scene datasets. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, pages 7190–7199. Computer Vision Foundation / IEEE, 2021. doi: 10.1109/CVPR46437.2021.00711. URL [https://openaccess.thecvf.com/content/CVPR2021/html/Li\\_OpenRooms\\_An\\_Open\\_Framework\\_for\\_Photorealistic\\_](https://openaccess.thecvf.com/content/CVPR2021/html/Li_OpenRooms_An_Open_Framework_for_Photorealistic_)



- Indoor\_Scene\_Datasets\_CVPR\_2021\_paper.html.
- [64] Yongsen Mao, Yiming Zhang, Hanxiao Jiang, Angel Chang, and Manolis Savva. Multiscan: Scalable rgbd scanning for 3d environments with articulated objects. *Advances in Neural Information Processing Systems*, 35: 9058–9071, 2022.
- [65] Matt Deitke, Rose Hendrix, Ali Farhadi, Kiana Ehsani, and Aniruddha Kembhavi. Phone2proc: Bringing robust robots into our chaotic world. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023*, pages 9665–9675. IEEE, 2023. doi: 10.1109/CVPR52729.2023.00932. URL <https://doi.org/10.1109/CVPR52729.2023.00932>.
- [66] Dov Katz, Moslem Kazemi, J Andrew Bagnell, and Anthony Stentz. Interactive segmentation, tracking, and kinematic modeling of unknown 3d articulated objects. In *2013 IEEE International Conference on Robotics and Automation*, pages 5003–5010. IEEE, 2013.
- [67] Jürgen Sturm, Vijay Pradeep, Cyrill Stachniss, Christian Plagemann, Kurt Konolige, and Wolfram Burgard. Learning kinematic models for articulated objects.
- [68] Neil Nie, Samir Yitzhak Gadre, Kiana Ehsani, and Shuran Song. Structure from action: Learning interactions for articulated object 3d structure discovery. *arXiv preprint arXiv:2207.08997*, 2022.
- [69] Zhenyu Jiang, Cheng-Chun Hsu, and Yuke Zhu. Ditto: Building digital twins of articulated objects from interaction. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [70] Liqian Ma, Jiaojiao Meng, Shuntao Liu, Weihang Chen, Jing Xu, and Rui Chen. Sim2real2: Actively building explicit physics model for precise articulated object manipulation. In *International Conference on Robotics and Automation (ICRA)*, 2023.
- [71] Cheng-Chun Hsu, Zhenyu Jiang, and Yuke Zhu. Ditto in the house: Building articulation models of indoor scenes through interactive perception. *arXiv preprint arXiv:2302.01295*, 2023. URL <https://arxiv.org/abs/2302.01295>.
- [72] Marius Memmel, Andrew Wagenmaker, Chuning Zhu, Dieter Fox, and Abhishek Gupta. Asid: Active exploration for system identification and reconstruction in robotic manipulation. In *The Twelfth International Conference on Learning Representations*, 2023.
- [73] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. URL <https://arxiv.org/abs/2006.11239>.
- [74] Shekoofeh Azizi, Simon Kornblith, Chitwan Saharia, Mohammad Norouzi, and David J Fleet. Synthetic data from diffusion models improves imagenet classification. *arXiv preprint arXiv:2304.08466*, 2023.
- [75] Brandon Trabucco, Kyle Doherty, Max Gurinas, and Ruslan Salakhutdinov. Effective data augmentation with diffusion models. *arXiv preprint arXiv:2302.07944*, 2023.
- [76] Ziyi Li, Qinye Zhou, Xiaoyun Zhang, Ya Zhang, Yanfeng Wang, and Weidi Xie. Open-vocabulary object segmentation with diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7667–7676, 2023.
- [77] Xinyi Yu, Guanbin Li, Wei Lou, Siqi Liu, Xiang Wan, Yan Chen, and Haofeng Li. Diffusion-based data augmentation for nuclei image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 592–602. Springer, 2023.
- [78] Weijia Wu, Yuzhong Zhao, Hao Chen, Yuchao Gu, Rui Zhao, Yefei He, Hong Zhou, Mike Zheng Shou, and Chunhua Shen. Datasetdm: Synthesizing data with perception annotations using diffusion models. *arXiv preprint arXiv:2308.06160*, 2023.
- [79] Stephanie Fu, Netanel Tamir, Shobhita Sundaram, Lucy Chai, Richard Zhang, Tali Dekel, and Phillip Isola. Dreamsim: Learning new dimensions of human visual similarity using synthetic data. *CoRR*, abs/2306.09344, 2023. doi: 10.48550/arXiv.2306.09344. URL <https://doi.org/10.48550/arXiv.2306.09344>.
- [80] Yonglong Tian, Lijie Fan, Phillip Isola, Huiwen Chang, and Dilip Krishnan. Stablerep: Synthetic images from text-to-image models make strong visual representation learners. *CoRR*, abs/2306.00984, 2023. doi: 10.48550/arXiv.2306.00984. URL <https://doi.org/10.48550/arXiv.2306.00984>.
- [81] Ali Jahanian, Xavier Puig, Yonglong Tian, and Phillip Isola. Generative models as a data source for multiview representation learning. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, 2022. URL <https://openreview.net/forum?id=qhAeZjs7dCL>.
- [82] Daniel Ho, Kanishka Rao, Zhuo Xu, Eric Jang, Mohi Khansari, and Yunfei Bai. Retinagan: An object-aware approach to sim-to-real transfer. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 10920–10926. IEEE, 2021.
- [83] Zoey Qiuyu Chen, Shosuke C. Kiami, Abhishek Gupta, and Vikash Kumar. Genau: Retargeting behaviors to unseen situations via generative augmentation. In Kostas E. Bekris, Kris Hauser, Sylvia L. Herbert, and Jingjin Yu, editors, *Robotics: Science and Systems XIX, Daegu, Republic of Korea, July 10-14, 2023*, 2023. doi: 10.15607/RSS.2023.XIX.010. URL <https://doi.org/10.15607/RSS.2023.XIX.010>.
- [84] Zhao Mandi, Homanga Bharadhwaj, Vincent Moens, Shuran Song, Aravind Rajeswaran, and Vikash Kumar. Cacti: A framework for scalable multi-task multi-scene visual imitation learning. *arXiv preprint arXiv:2212.05711*, 2022.
- [85] Tianhe Yu, Ted Xiao, Jonathan Tompson, Austin Stone, Su Wang, Anthony Brohan, Jaspiar Singh, Clayton Tan, Dee M, Jodilyn Peralta, Karol Hausman, Brian Ichter, and Fei Xia. Scaling robot learning with semantically imagined experience. In Kostas E. Bekris, Kris Hauser,

Sylvia L. Herbert, and Jingjin Yu, editors, *Robotics: Science and Systems XIX, Daegu, Republic of Korea, July 10-14, 2023*, 2023. doi: 10.15607/RSS.2023.XIX.027. URL <https://doi.org/10.15607/RSS.2023.XIX.027>.

- [86] Homanga Bharadhwaj, Jay Vakil, Mohit Sharma, Abhinav Gupta, Shubham Tulsiani, and Vikash Kumar. Roboagent: Generalization and efficiency in robot manipulation via semantic augmentations and action chunking. *arXiv preprint arXiv:2309.01918*, 2023.
- [87] Brandon Trabucco, Kyle Doherty, Max Gurinas, and Ruslan Salakhutdinov. Effective data augmentation with diffusion models. *CoRR*, abs/2302.07944, 2023. doi: 10.48550/arXiv.2302.07944. URL <https://doi.org/10.48550/arXiv.2302.07944>.
- [88] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3836–3847, 2023. URL <https://arxiv.org/abs/2302.05543>.
- [89] Elad Richardson, Gal Metzer, Yuval Alaluf, Raja Giryes, and Daniel Cohen-Or. Texture: Text-guided texturing of 3d shapes. *arXiv preprint arXiv:2302.01721*, 2023. URL <https://arxiv.org/abs/2302.01721>.
- [90] Dave Zhenyu Chen, Haoxuan Li, Hsin-Ying Lee, Sergey Tulyakov, and Matthias Nießner. Scenetex: High-quality texture synthesis for indoor scenes via diffusion priors. *arXiv preprint arXiv:2311.17261*, 2023. URL <https://arxiv.org/abs/2311.17261>.