

Scaling Single Human Demonstrations for Imitation Learning using Generative Foundational Models

Nick Heppert^{*1,2}, Minh Quang Nguyen^{*1}, Abhinav Valada¹

Abstract—Imitation learning is a popular paradigm to teach robots new tasks, but collecting robot demonstrations through teleoperation or kinesthetic teaching is tedious and time-consuming. In contrast, directly demonstrating a task using our human embodiment is much easier and data is available in abundance, yet transfer to the robot can be non-trivial. In this work, we propose Real2Gen to train a manipulation policy from a single human demonstration. Real2Gen extracts required information from the demonstration and transfers it to a simulation environment, where a programmable expert agent can demonstrate the task arbitrarily many times, generating an unlimited amount of data to train a flow matching policy. We evaluate Real2Gen on human demonstrations from three different real-world tasks and compare it to a recent baseline. Real2Gen shows an average increase in the success rate of 26.6% and better generalization of the trained policy due to the abundance and diversity of training data. We further deploy our purely simulation-trained policy zero-shot in the real world. We make the data, code, and trained models publicly available at <https://real2gen.cs.uni-freiburg.de>.

I. INTRODUCTION

In the future, robots should be able to easily acquire new manipulation skills with minimal to no overhead in teaching them. To achieve this, imitation learning has crystallized as one of the primary paradigms for teaching a robot skills [1], [2]. Classically, imitation learning uses a dataset consisting of demonstrations of aligned robot observations as well as actions collected through teleoperation [3] or kinesthetic teaching [4]. While the user operates the robot, the robot actively records its observations and actions to form a training demonstration dataset. We call these types of demonstrations robot demonstrations. One advantage of robot demonstrations is the embodiment alignment of the training data with the robot. In contrast, while collecting the data itself is not only time-consuming, tele-operating also requires a skilled operator. Similarly, for kinesthetic teaching, the operator needs to be physically present in the scene and, for example, can block cameras.

A recent trend aims to mitigate the tedious collection process by directly learning from demonstrations of humans in their own embodiment [5], [6]. We coin these types of demonstrations as human demonstrations. While human




Fig. 1: Overview of Real2Gen. Real2Gen takes a single human demonstration as input and produces simulatable meshes using 3D generative foundational models, which can be used in a generative simulation setup.

demonstrations compared to robot demonstrations are easy to collect or are even widely available on the internet [7], they impose a significant challenge as the embodiment between human and robot differs. This challenge was tackled by researchers, for example, through the use of handcrafted heuristics to match a human hand to the robot gripper [8]–[10] or by constraining the possible human movements in a demonstration to the robot’s capabilities [11]. In contrast, works such as DITTO [12], ORION [13], or OKAMI [14] propose closing the embodiment gap by finding explicit object correspondences between the human demonstration and the robot environment at inference time, allowing for the transfer of motion in an object-centric manner.

In this work, we aim to combine the best of human and robot demonstrations and propose a “Reality to Simulation” (Real2Sim) [15] method called Real2Gen that processes a single human demonstration to set up a simulator that automatically allows us to collect robot demonstrations. Thus, we leverage the ease of collecting demonstrations using our human embodiment, but can train a policy directly on data in the robot embodiment. While previous Real2Sim approaches needed to scan the environment in an offline step [16]–[18], Real2Gen uses 3D generative foundational models such as Point-E [19] or Zero-1-to-3 [20] to directly generate 3D assets of task-relevant objects given their appearance in the human demonstration. We then use a foundational feature matcher [21] to align the generated 3D assets with the human demonstration, thereby retrieving realistic scale and canonical orientation. Similar to other simulation works [22], [23], we first randomly sample poses of the assets and then use a scripted expert agent to generate robot demonstrations. Finally, we train a flow matching policy [24] on the collected robot demonstrations.

We evaluate our policy in a simulation environment consisting of unseen object instances retrieved from a curated large-scale 3D object set such as Objaverse [25]. Compared to the DITTO [12] baseline, we show an average success

^{*}These authors contributed equally and are ordered alphabetically.

¹Computer Science, University of Freiburg, Germany, ²Zuse School ELIZA This work was partially funded by the Carl Zeiss Foundation with the ReScaLe project. Nick Heppert is supported by the Konrad Zuse School of Excellence in Learning and Intelligent Systems (ELIZA) through the DAAD programme Konrad Zuse Schools of Excellence in Artificial Intelligence, sponsored by the Federal Ministry of Education and Research. The authors would like to thank Eugenio Chisari and Max Argus for discussion as well as Martin Büchner for measuring objects.

rate improvement of 26.6%. Lastly, we also demonstrate zero-shot transfer of our purely simulation-trained policy to a real-world setup, successfully performing the task as demonstrated by the human. While in this paper we show the capabilities of Real2Gen in an imitation learning setup, our approach also transfers without any modifications to a reinforcement learning setup.

Our main contributions are as follows:

- A framework to process human demonstrations for generating robot demonstrations.
- Evaluations of mesh retrieval and generation rates as well as of scale and pose estimation methods for meshes.
- Successful transfer from reality to generative simulation to reality from a single human demonstration.
- Publicly available code, trained models, and data at <https://real2gen.cs.uni-freiburg.de>.

II. RELATED WORK

This section reviews previous works that use human demonstrations to learn a robotic manipulation policy as well as highlights progress in the field of scaling robot learning data through generative or procedural simulations.

A. Learning from Human Demonstrations

While learning from human demonstrations, rather than robot demonstrations, a prominent challenge is the gap between human and robot embodiment. Here, at the core, it is first necessary to disentangle task-relevant information from the human embodiment and second, to provide an interface to that information such that a robot can use it to execute the task using its robot embodiment. To store and access this information, a wide range of representations have been proposed, ranging from concrete object trajectories [12], [13] to affordance [8], [26] or pixel-level [27], [28], which we discuss in detail.

Trajectories: An efficient way to represent tasks is through object-centric trajectories [3], [29]. Works such as DITTO [12], ORION [13], and OKAMI [14] propose detecting and tracking objects in human demonstrations using foundational models. DITTO [12] uses an off-the-shelf hand-object detector to extract relevant object masks of the task, followed by tracking correspondences throughout the scene using, e.g., LoFTR [30]. Using RGB-D images, the correspondences are lifted to 3D, and relative object movement is extracted between timesteps. At test-time, the trajectory is transferred using correspondences and executed using off-the-shelf grasping and motion planning libraries. ORION [13] follows a similar approach to extract object trajectories and for test-time inference, but instead of hand-object detectors requires a language instruction to extract relevant objects. Additionally, ORION [13] segments the task demonstration into key frames. For each key frame, ORION [13] stores contact information and keypoints in so-called Open-World Object Graphs. While also leveraging foundational keypoints through KAT [31], R+X [8] neither requires a hand-object detector nor language annotations,

but instead relies on multiple demonstrations provided in a full-length video to identify common task-relevant keypoints.

Affordance: The concept of affordance is a well-studied topic in robotics [32]. Recently, researchers have begun to investigate affordance as a common representation between humans and robots, enabling easy transfer [26], [33]. Different to trajectories, affordance usually consists of a region and a simple short-horizon trajectory [34]. Similar to R+X [8], the affordance-based method RAM [33] retrieves a demonstration from an unstructured source. Another body of work used human affordance as a means of pre-training an encoder for visual policy learning [26], [34], [35].

Pixel-Level: Moving away from structured and semantic interpretable representations, such as trajectories or affordances, pixel-level motion inference and subsequent transfer to the robot is another promising avenue. Here, most notably, Track2Act [28] and Im2Flow2Act [27], follow a similar paradigm of first pre-training an open-loop 2D point inference module either on internet data [28] or mixed human and robot data [27] and second training a complementary closed-loop robot policy on robot data.

Our method Real2Gen builds upon these advances and aims to generate diverse training data using extracted object trajectories and key frames.

B. Learning through Procedural and Generative Simulation

The previously discussed works either require open-loop motion planning [12], [13] or robot demonstrations to fine-tune a closed-loop policy [27], [28]. To combine the advantages of both, the ease and determinism of motion planning and the robustness of closed-loop policy execution, automatic robot demonstration generation in real-world aligned simulations is a promising candidate [18], [22], [23].

The rise of Large Language Models (LLMs) has also fueled this paradigm as LLMs can be used to, e.g., generate task-specific scenarios, success criteria, and demonstration code [23], [36] given a fixed set of curated object CAD models, such as Google scanned objects [37], to choose from. As this limits the system in terms of object diversity, Gen2Sim [38] proposes an automated process that employs differential rendering to generate 3D assets from any 2D image of a singulated object. The images can either stem from the robot’s workspace, the web, or an image generation model, e.g., Stable Diffusion [39]. To ground the generated assets with realistic sizes and masses Gen2Sim [38] uses a Vision Language Model (VLM). Another LLM then generates meaningful tasks for given scenes as well as rewards to train a reinforcement learning agent. Given a live observation GRS [40] uses a VLM to describe and match all objects in their 3D asset set to the shown scene. Unlike Gen2Sim [38], the authors assume that the 3D asset set is already provided, e.g., through Objaverse [25].

RoboGen [41] aims to combine all these previous approaches by allowing scene generation with either assets generated from 3D foundational models or from a large dataset, such as Objaverse [25]. Similarly as Gen2Sim [38]

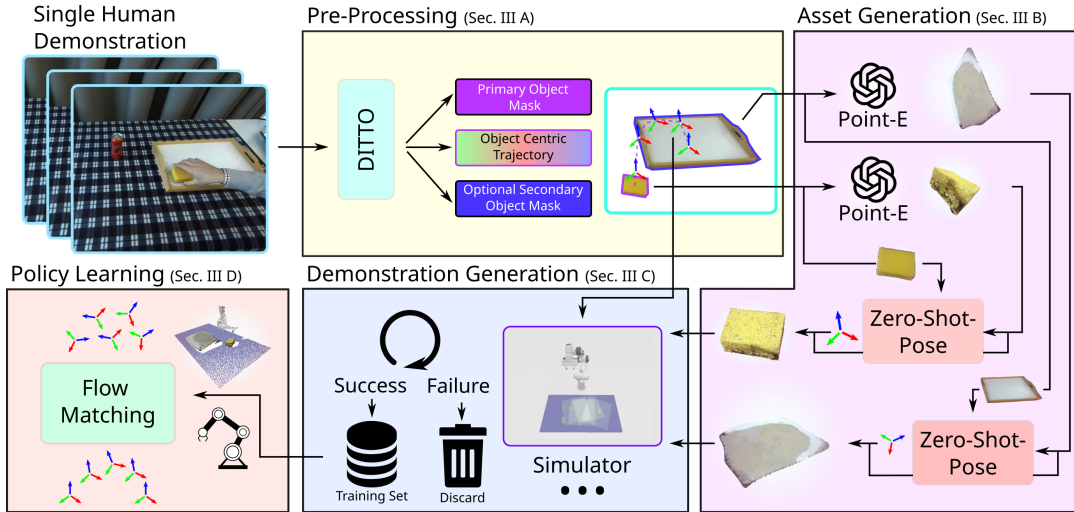


Fig. 2: Technical approach of Real2Gen. Real2Gen uses a single human demonstration as input, consisting of a sequence of RGB-D images. We pre-process (Sec. III-A) these images using DITTO [12] to retrieve a primary and, if applicable, a secondary object mask as well as an object-centric trajectory of the object. In the second step, asset generation (Sec. III-B), we pass object images to Point-E [19] to generate 3D meshes in a canonical space. We then use Zero-Shot-Pose (ZSP) [21] to scale and align the meshes to the human demonstration. We then use the generated meshes combined with object-centric trajectories to set up a simulation (Sec. III-C). Using grasp and motion planning, we use the simulation to generate an expert dataset of policy rollouts. In the last step, policy learning (Sec. III-D), we use the collected dataset to train a conditional flow matching policy [24].

and GRS [40], RoboGen [41] verifies meshes using a VLM. Unlike other works, they also tackle long-horizon tasks by allowing an LLM to decide whether sub-skills should be learned through reinforcement learning, gradient-based trajectory optimization, or motion planning. RialTo [17] goes a slightly different path by not relying on any of these expert ways but rather uses a trained policy trained on real robot demonstrations and only saves successful executions to form a final dataset. RialTo [17] generates their simulation environment through scanning the real scene and annotating it.

Different from these works, Real2Gen directly uses generative 3D foundational models to generate 3D CAD models aligned with the human demonstration, as well as provides an automatic and robust alignment process of the generated CAD models to the demonstration. As a result, Real2Gen needs no additional manual curation of meshes or scenes, nor does Real2Gen require text-input that describes the task or iterative code generation of LLMs.

III. TECHNICAL APPROACH

In this section, we detail our Real2Gen approach. We start by describing the pre-processing procedure in Sec. III-A, how we generate 3D assets in Sec. III-B and robot training data in Sec. III-C, and lastly detail our policy learning approach in Sec. III-D. An overview of Real2Gen approach is presented in Fig. 2.

A. Pre-Processing Human Demonstrations

The input to Real2Gen is a single human demonstration consisting of a sequence of T RGB-D images o_h . We pre-process the sequence o_h using DITTO [12], an object trajectory extractor, but other options such as ORION [13] can also be used. Similarly to these methods, we assume there is a primary object p moving, either in free-form or depending on another secondary goal object s . The goal of

the pre-processing step is to extract segmentation masks of all relevant objects, i.e., m_p and m_s if applicable, in the first RGB image I_0 as well as the object-centric trajectory of the primary object J_p consisting of relative poses. By segmenting the first RGB image I_0 using m_p and m_s masks, we get unobstructed object-centric RGB reference images I_p and I_s of our primary and secondary objects, respectively, before the human interacts with them.

B. Asset Generation

Given one of the reference object RGB images I_f , i.e., either I_p or I_s if applicable, we use an off-the-shelf 3D generative foundational model, specifically Point-E [19] combined with marching cubes, to generate a raw 3D mesh m^c . We decided on this combination because of its ease of use, but other models such as Zero-1-To-3 [20], Stable-Dreamfusion [42], or Shape-E [43] can also be used. As these models return object meshes in a canonical non-metric frame, they are not directly usable in a robot simulation. Although VLMs were previously used to verify the mesh and guess a metric dimension [38], we use our provided human demonstration and match the generated mesh to it. As our generated mesh and the object in the human demonstration share semantic similarities but are not the exact instance, methods like FoundationPose [44] are not applicable here. Instead, we propose to use a foundational feature matcher, Zero-Shot-Pose (ZSP) [21], to first generate correspondences and then subsequently align the correspondences through a 7-DoF affine transformation with a closed form solution [45]. We detail this approach in the following.

Given a generated mesh m^c from our 3D generative foundational model, we render N_v views V of it using spherical Fibonacci sampling [46] to sample the polar and azimuthal angles, similar to what is done by Giammaorino *et al.* [47]. The Fibonacci sampling procedure ensures maximum

diversity in views compared to densely sampling the polar and azimuthal angles. For each view $v_r = \{I_r, D_r\}$, we render an RGB image I_r and a depth image D_r . We collectively give all views V and the initial reference image I_f to ZSP [21]. ZSP [21] first extracts descriptors for each view’s RGB image I_r , matches these descriptors to the reference image I_f to select the view \hat{v} with the top-K similar descriptors. Second, ZSP [21] uses the most similar descriptors as correspondences and projects them in 3D using the selected views \hat{v} depth image \hat{D} , and the reference depth image D_f . Lastly, ZSP [21] solves a closed-form seven degree of freedoms (7-DoF) least squares problem [45] to estimate a full affine transformation ${}^c\mathbf{T}_m$ to transform the canonical mesh m^c to a metric mesh m^m of which the up-orientation matches the human demonstration. Despite using the top-K similar descriptors across all views, there might be outliers present. Thus, the estimate is done multiple times using RANSAC. This process can be repeated arbitrarily many times to continuously generate new and unseen meshes usable in a robotic simulation. In practice, we generate a fixed amount of N_m meshes M before continuing with our demonstration generation. We additionally verify all matches before proceeding, which, with improvements in 3D generative foundational models, will become obsolete. Throughout our experiments, we sampled $N_v = 41$ views in the upper hemisphere and used the top-30 descriptors in ZSP [21].

C. Demonstration Generation

After we generate a set of meshes M_p for the primary object and if applicable, a set M_s for the secondary object, we use them in a simulation environment to generate large-scale robot data. Specifically, we use SAPIEN [48] as our simulator with a Franka Panda Robot mounted on a table. To observe the scene, we set up an external RGB-D camera as well as attach an RGB-D wrist camera to the robot. We assign a default density of 1000 kg m^{-3} to generated meshes.

To generate a robot demonstration, we first sample random meshes, one for the primary object $m_p^m \in M_p$ and if needed, one for the secondary object $m_s^m \in M_s$. We then sample random poses, constraining the meshes to be on the table and randomly rotated around the up-axis. Next, using our privileged simulation information, we analytically calculate grasps on the primary object mesh using antipodality constraints¹ and pick a random grasp. At this point, we explicitly differentiate between tasks only moving the primary object and tasks with an additional secondary object. For tasks that only involve a primary object, we apply the extracted object-centric trajectory J_p to the grasp. For tasks consisting of an additional secondary object, we constrain the primary object to be placed in the center of the secondary object approached through a bottleneck pose above the table. While directly using the object-centric trajectory is also possible, in practice though, we observed higher generation failure rates, prolonging the generation unnecessary. To finally execute the

¹In practice, this process was done in a pre-processing step for each mesh and stored. At demonstration generation time, we only perform collision checks to filter out grasps.

task, we use a motion planner². During execution, we record all observations o_r and the robots proprioceptive states as end-effector poses. After the roll-out is completed, using the proposed pose error metric by GraspNet [49], we compare the distance between the actual pose of the primary object and the expected pose to determine whether the roll-out was a success or failure. Only successful roll-outs d are added to our demonstration dataset D .

D. Policy Learning

The last component of Real2Gen consists of learning a manipulation policy given our previously generated demonstration dataset D . As our policy model, we chose to use PointFlowMatch [24], a flow matching-based imitation learning method. Similar to PointFlowMatch [24], our policy learns to move a Gaussian distribution to a target distribution being our ground truth actions where the process is conditioned on the current robot observations. The robot observations are the last two point clouds encoded using a PointNet [50] encoder as well as the last two robot proprioceptive states, in our case, the end-effector pose in $SE(3)$. Different from PointFlowMatch, we only use a partial point cloud from the wrist and an external camera, and do not aggregate point clouds across multiple cameras observing the scene. Our inferred actions are future end-effector poses with a fixed horizon length H . See Sec. B for our hyperparameter setup.

IV. EXPERIMENTAL EVALUATION

In our experiments, we quantitatively evaluate how well Real2Gen can transfer the human demonstration to a robot compared to a DITTO [12] baseline. We also investigate the difference in quality and effort when using a 3D generative foundational model over a large-scale object dataset, such as Objaverse [25]. Further, we compare our proposed matching procedure to a VLM-based approach. Last, we show zero-shot applicability of Real2Gen on a real-robot system.

A. Quantitative Full Pipeline Evaluation

For our evaluation, we selected three tasks from DITTO [12], *Sponge on Tray*, *Coke on Tray*, and *Paperroll upright* (see Sec. A for further details), and, following DITTO [12], a single human demonstration for each task. We generate five meshes and 800 demonstrations. We use a simulator with the same fixed seeds across all evaluations for a fair and reproducible comparison. We evaluate Real2Gen against DITTO [12] baselines as well as ablate Real2Gen.

Evaluation Setup: To set up and align our evaluation simulator with our tasks, we use unseen realistic 3D meshes from Objaverse [25]. All meshes were manually checked and scaled to have reasonable geometry and realistic sizes. For the *Sponge on Tray* and *Can on Tray*-task, we select five sponge, five can, and seven tray meshes. For the *Paperroll upright*-task, we select five paperroll meshes. As for generating robot demonstrations (see Sec. III-B), during each evaluation episode, we randomly select meshes according to

²<https://motion-planning-lib.readthedocs.io/latest/>

| Method | Sponge on Tray | Coke on Tray | Paperroll upright | Mean SR (\uparrow) |
|------------------------|--------------------------------|--------------------------------|--------------------------------|--------------------------------|
| DITTO [12] | 6.3 \pm 2.1 | 26.0 \pm 3.6 | 0.3 \pm 0.5 | 10.9 \pm 2.1 |
| DITTO [12] w/ ZSP [21] | 4.3 \pm 1.2 | 19.7 \pm 3.8 | 0.7 \pm 0.6 | 8.2 \pm 1.8 |
| Real2Gen (ours) | 41.3\pm4.5 | 46.3\pm6.4 | 25.0\pm1.0 | 37.5\pm3.0 |

TABLE I: Performance comparison of Real2Gen against different baseline methods. We evaluate all methods on three different tasks and report per-task success rate as well as overall mean success rate [%] (\uparrow).

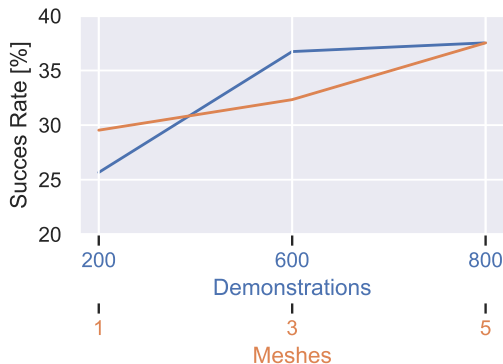


Fig. 3: Results of Ablation Study. We show the average success rate [%] (\uparrow) across all tasks. We either vary the number of demonstrations while using five meshes or we vary the number of meshes using 800 demonstrations.

the task and spawn them at random poses. To perform the evaluation, we select three random seeds. For each seed, we evaluate each method 100 times and record the successes of the roll-outs. We evaluate the successes as described in Sec. III-C. As done in [24], [51], we report the mean and standard deviation across all seeds.

Baseline Comparisons: We compare Real2Gen against two variants of DITTO [12]. First, the original variant using LoFTR [30] to match the live observation to the demonstration. This is an unfair comparison as DITTO [12] was designed to be faced with the same object instance as in the demonstration. Thus, we extend DITTO [12] by replacing the LoFTR [30] matching step with ZSP [21] which should enable DITTO [12] to work across instances from the same category.

The results of our experiment are reported in Tab. I. When comparing Real2Gen to DITTO [12] with and without ZSP [21], we see that the overall performance of Real2Gen is better. We assume the difference stems from DITTO [12] having only a single test-time image available, which increases difficulty for matching. To our surprise, the original DITTO [12] baseline outperforms the variant with ZSP [21]. Additionally, we observe that grasp and motion planning fail quite often, aligning with the reported results in DITTO [12], whereas our learned policy produces more robust results.

Ablation Studies: We additionally perform ablation studies of Real2Gen to study the effect of the number of generated meshes and demonstrations. Specifically, we change the number of demonstrations from 800 to 600 and 200, while still using five meshes, as well as changing the number of meshes to three and one with 800 demonstrations. We visualize the results in Fig. 3. As one would expect, the more meshes and demonstrations, the better. Nonetheless, the average performance increase diminishes rather soon when going from 600 to 800 (1.1%) demonstrations and three to

five meshes (5.2%), respectively.

B. Comparison of Mesh Generation

To highlight the effectiveness of our proposed automatic mesh generation, we compare the human effort needed to perform manual mesh retrieval from a large database, as performed, for instance, in Scaling Up And Distill Down [23]. To have representative results, we set the goal to generate 100 object meshes. For Real2Gen, this process is straightforward, and we run the asset generation a hundred times. For the object set comparison, we query Objaverse [25] to retrieve meshes with the tag of either sponge, coke/coca can, tray, or paperroll. If there are more than a hundred meshes available, we evaluate using the hundred most viewed meshes and a hundred random meshes. We additionally report the total meshes available. We then apply the same matching procedure as for the generated meshes. Lastly, we manually go through all of the generated and retrieved object set meshes and classify whether they are task relevant. For the retrieved meshes, this is particularly important as they are not anchored to the human demonstration, e.g., for sponge, sometimes the cartoon character *Sponge Bob* is returned. We report the results in Tab. II. Overall, compared to previous retrieval-based methods [23], our proposed way to generate meshes through a generative model results in almost 50% more available meshes on average, while also being able to generate an infinite amount.

C. VLMs as Size and Pose Estimators

We further evaluate the quality and robustness of using a VLM to estimate physical properties as proposed in Gen2Sim [38], as opposed to using our introduced foundational matching procedure (see Sec. III-B). We can not assume a known one-dimensional scale or canonical orientation for the generated meshes from Real2Gen nor for the Objaverse meshes. Thus, we evaluate inference performance for both physical properties in this experiment.

Data: We evaluate on a subset of 89 of all resulting matchable and semantically meaningful meshes from Sec. IV-B (essentially all last columns for each row in Tab. II)³. Both our proposed matching and the by Gen2Sim [38] inspired VLMs receive 10 rendered images as input, sampled from the full Fibonacci sphere as described in Sec. III-B.

Baselines: We compare our deterministic matching approach to a VLM baseline, as done in Gen2Sim [38]. For a fair comparison, we provide the VLM with the same inputs that the matching uses. For the rendered RGB images, this process is straightforward. Additionally, we provide the dimensions of the unscaled mesh when loaded from the file. To ground the prediction in the human demonstration, we compare providing the VLM with the full image I_0 or only with a segmented and cropped image of the object I_f .

³As for some sets the resulting amount of meshes is small, we test on all available meshes or on a maximum of ten meshes per set to prevent a strong bias in our results while still having sufficient and diverse meshes for statistic significance.

| Category | Source | Available Meshes | 100 Mesh Pre-Selection | Matching Successful [†] | Matching Successful and Task Relevant [†] |
|------------|---------------------|------------------|-------------------------------------|----------------------------------|--|
| Sponge | Point-E [19] (ours) | ∞ | Random | 99% | 61% |
| | Objaverse [25] | 351 | Most viewed Random | 58% 57% | 3% 7% |
| Tray | Point-E [19] (ours) | ∞ | Random | 35% | 18% |
| | Objaverse [25] | 278 | Most viewed Random | 21% 14% | 10% 5% |
| Coke Can | Point-E [19] (ours) | ∞ | Random | 67% | 39% |
| | Objaverse [25] | 41 | All | 82% | 73% |
| Paper Roll | Point-E [19] (ours) | ∞ | Random | 82% | 62% |
| | Objaverse [25] | 20 | All | 45% | 35% |
| All | Point-E [19] (ours) | ∞ | Random | 71% | 45% |
| | Objaverse [25] | 690 | Most viewed or All Random or All | 52% 50% | 30% 28% |

TABLE II: Comparison of Mesh Generation. We compare the number of resulting meshes when using our proposed generative way vs. using a large-scale object dataset. [†]We report the total percentage from the 100 selected meshes, or if fewer than 100 meshes are available from all available ones.

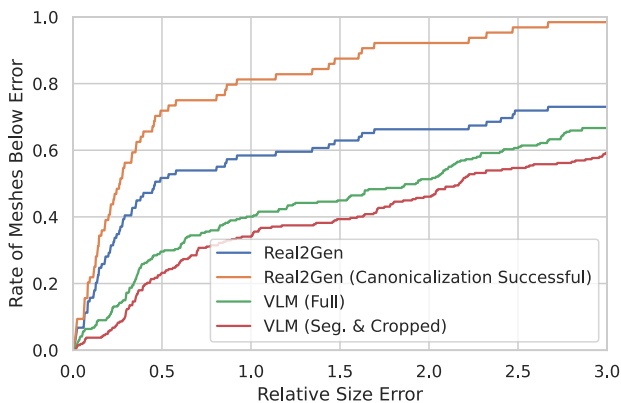


Fig. 4: Precision Curve for Scaling Factor. We plot the percentage of meshes below the relative size error ranging.

| Method | Reference Image | Canonicalization Successful (\uparrow) | Scaling mAP (\uparrow) |
|----------------------------------|-----------------|--|----------------------------|
| Real2Gen (ours) | Seg. & cropped | 71.9% | 0.586 |
| VLM (Gen2Sim [38]) | Full | 21.3% | 0.380 |
| | Seg. & cropped | 26.9% | 0.438 |
| Only Successful Canonicalization | | | |
| Real2Gen (ours) | Seg. & cropped | N/A | 0.810 |
| VLM (Gen2Sim [38]) | Full | N/A | 0.297 |
| | Seg. & cropped | N/A | 0.401 |

TABLE III: Comparison of Physical Property Inference. We compare our proposed matching against a VLM to infer the pose canonicalization success as well as calculate the mAP for relative errors between 0 and 3.

Lastly, we also provide both VLM variants with the point cloud dimensions of the segmented object in the reference depth image D_f . Given this information, the VLM has to infer a scaling factor for the mesh as well as determine the closest view of the rendered images to the reference image. As VLMs are non-deterministic, to further test the robustness, we query both variants three times.

Our matching process requires no modification for this experiment. If it fails, we will set the scaling result to identity.

Metrics: For size comparison, we measure real-world dimensions for each object category (see Fig. 6). Given these measurements, we optimize for the closest scaling factor δ of each mesh instance to the real-world reference size. The

optimized scaling factor δ for each mesh serves as our ground truth value, and we calculate the relative distance $(\hat{\delta} - \delta)/\delta$ to the inferred scaling factor $\hat{\delta}$. To compare the canonicalization of the object, we rely on proxies for both methods. For the matching approach, we directly use the information on whether the matching procedure succeeded. For the VLMs, we consider the canonicalization as successful if all three queries result in the same decision for the closest view.

Results: Overall, we observe that for dimensions estimation as well as object canonicalization, the matching procedure outperforms the VLM baseline. For the scaling factor, we plot the precision curve for relative size error thresholds ranging from 0 to 3 in Fig. 4. In Tab. III we compare the success of mesh canonicalization to the reference image as well as calculate the area under the scale factor precision curve.

D. Real-World Robotic Experiments


We also evaluate the effectiveness of Real2Gen on a real-world robot system. As done in PointFlowMatch [24], we apply additional augmentations through transforming the input point cloud and proprioceptive states as well as the robot actions using a sampled $SE(3)$ transform. In Sec. IV-C, we observed that the median error of the matching procedure scales objects 18% too large. We incorporate this finding through downscaling all meshes by 0.8 for training data generation. The remainder of the policy training for our real-world setup remains unchanged from our simulation experiments. For inference, we change the flow denoising steps to 30. As our experimental platform, we use a Franka Panda robot with 7 degrees of freedom, shown in Fig. 5a. We attach a close-view Intel Realsense D405⁴ to the wrist and use an Intel Realsense D435 as our external camera. We also align the simulation with our calibrated camera extrinsics.

We evaluate the *Sponge* on *Tray*-task on consecutive roll-outs for both 14 easy and 20 hard layouts. For the easier layouts, the average success rate achieved is 64% and 30% for the hard layouts. A successful execution is shown

⁴We apply a manual filtering of the depth image when either close ($< 5cm$) or far ($> 50cm$).



(a) Real-World Setup. As in simulation, we mount a wrist camera to the robot and place an external camera overlooking the scene.



(b) Successful Task Execution.

(c) Failed Task Executions.

Fig. 5: Real-World Robot Experiment. Failure cases include imperfect grasping and premature closing off the gripper.

in Fig. 5b. The most common failure case is missing the sponge by just a few centimeters because the real-world depth camera becomes inaccurate for close-up manipulation. We present failures in Fig. 5c. For the Coke on Tray and Paperroll upright-task we achieve an average success rate of 21% and 24% respectively across 29 roll-outs each. See the accompanying video for roll-outs of all tasks.

V. CONCLUSION

In this work, we investigated the use of 3D generative foundational models to transfer a single human demonstration to simulation. We showed that with Real2Gen, more robust policies are learned with an increase in success rate of 26.6% over baselines. We additionally analyzed the effort needed to generate task-relevant and usable simulation CAD models. We also showed the superior performance of our matching over VLM approaches and zero-shot transfer to a real-robot system. While we showed the applicability of Real2Gen to generate expert demonstration data, it can also be used without any modifications to serve as a reinforcement learning environment.

In the future, we see potential to extend Real2Gen to enable more tasks with constrained movements like interacting with articulated objects, e.g., sampled from an object prior like in CARTO [52]. Another promising extensions could combine Real2Gen with the capabilities of VLMs as described in Gen2Sim [38] to setup tasks different from the human demonstration but using the same generated object meshes.

APPENDIX

A. Task Overview

| | Sponge on Tray | | Coke on Tray | | Paperroll upright | |
|------------|----------------|-----|--------------|-----|-------------------|-------|
| | pos | rot | pos | rot | pos | rot |
| Generation | 15cm | No | 15cm | No | 5cm | 10deg |
| Evaluation | 15cm | No | 15cm | No | No | 10deg |

TABLE IV: Overview of the Used Success Criteria.

We evaluate Real2Gen on three different tasks, Sponge on Tray, Coke on Tray and Paperroll upright. Depending on the task we adapt the final success criterion in our demonstration generation and evaluation. Sponge on Tray and Coke on Tray both consist of grasping an object and placing it on a tray. For both tasks the poses of the objects are randomized and it is only important that the object is placed on the tray, thus we do not consider orientation. Paperroll upright differs as the task does not depend on a secondary object, but rather the paper roll has to be placed upright as the name suggests. Thus, during test time we only consider the orientation of the object but not the position. We given an overview of the criteria in Tab. IV.

B. Policy Hyper Parameters Setup

We adopt the training settings and hyperparameters from PointFlowMatch [24]. We randomly downsample the input point cloud to a size of 4096 points. The policy is trained using AdamW optimizer with a learning rate of $3e^{-5}$ and weight decay of $1e^{-6}$. We schedule the learning rate with cosine annealing and linear warmup of 5000 steps. The batch size is 128, and we apply EMA on the weights of the model. Throughout our experiments, we use 20 denoising steps for flow matching and predict actions with a horizon of $H = 32$.

C. Real Objects Overview

For evaluating the dimensions in Sec. IV-C and our real world robotic experiments in Sec. IV-D, we use the objects displayed in Fig. 6.




Fig. 6: Real-World Reference Objects for Size Estimation.

REFERENCES

- [1] C. Celemin, R. Pérez-Dattari, E. Chisari, G. Franzese, L. de Souza Rosa, R. Prakash, Z. Ajanović, M. Ferraz, A. Valada, J. Kober, et al., “Interactive imitation learning in robotics: A survey,” *Foundations and Trends® in Robotics*, vol. 10, no. 1-2, pp. 1–197, 2022.
- [2] S. Livanec, L. Londoño, M. Gorki, A. Röfer, A. Valada, and A. Kiesel, “Designing for difference: How human characteristics shape perceptions of collaborative robots,” *arXiv preprint arXiv:2507.16480*, 2025.
- [3] J. O. von Hartz, T. Welschhold, A. Valada, and J. Boedecker, “The art of imitation: Learning long-horizon manipulation tasks from few demonstrations,” *IEEE Robotics and Automation Letters*, 2024.
- [4] D. Honerkamp, H. Maheshka, J. O. von Hartz, T. Welschhold, and A. Valada, “Whole-body teleoperation for mobile manipulation at zero added cost,” *IEEE Robotics and Automation Letters*, 2025.
- [5] T. Welschhold, C. Dornhege, and W. Burgard, “Learning manipulation actions from human demonstrations,” in *Proc. IEEE Int. Conf. on Intel. Rob. and Syst.*, pp. 3772–3777, 2016.

- [6] L. Shao, T. Migimatsu, Q. Zhang, K. Yang, and J. Bohg, "Concept2robot: Learning manipulation concepts from instructions and human demonstrations," *Int. J. Rob. Res.*, vol. 40, no. 12-14, pp. 1419–1434, 2021.
- [7] D. Shan, J. Geng, M. Shu, and D. F. Fouhey, "Understanding human hands in contact at internet scale," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pp. 9869–9878, 2020.
- [8] G. Papagiannis, N. Di Palo, P. Vitiello, and E. Johns, "R+ x: Retrieval and execution from everyday human videos," *arXiv preprint arXiv:2407.12957*, 2024.
- [9] M. Lepert, J. Fang, and J. Bohg, "Phantom: Training robots without robots using only human videos," in *Proc. Conf. on Rob. Learn.*, 2025.
- [10] J. Ren, P. Sundaresan, D. Sadigh, S. Choudhury, and J. Bohg, "Motion tracks: A unified representation for human-robot transfer in few-shot imitation learning," in *Proc. IEEE Int. Conf. on Rob. and Auto.*, 2025.
- [11] H. Xiong, Q. Li, Y.-C. Chen, H. Bharadhwaj, S. Sinha, and A. Garg, "Learning by watching: Physical imitation of manipulation skills from human videos," in *Proc. IEEE Int. Conf. on Intel. Rob. and Syst.*, pp. 7827–7834, 2021.
- [12] N. Heppert, M. Argus, T. Welschehold, T. Brox, and A. Valada, "Ditto: Demonstration imitation by trajectory transformation," in *Proc. IEEE Int. Conf. on Intel. Rob. and Syst.*, pp. 7565–7572, 2024.
- [13] Y. Zhu, A. Lim, P. Stone, and Y. Zhu, "Vision-based manipulation from single human video with open-world object graphs," *arXiv preprint arXiv:2405.20321*, 2024.
- [14] J. Li, Y. Zhu, Y. Xie, Z. Jiang, M. Seo, G. Pavlakos, and Y. Zhu, "Okami: Teaching humanoid robots manipulation skills through single video imitation," in *Proc. Conf. on Rob. Learn.*, 2024.
- [15] X. Li, J. Li, Z. Zhang, R. Zhang, F. Jia, T. Wang, H. Fan, K.-K. Tseng, and R. Wang, "Robogsim: A real2sim2real robotic gaussian splatting simulator," *arXiv preprint arXiv:2411.11839*, 2024.
- [16] M. N. Qureshi, S. Garg, F. Yandun, D. Held, G. Kantor, and A. Silwal, "Splatsim: Zero-shot sim2real transfer of rgb manipulation policies using gaussian splatting," in *Proc. IEEE Int. Conf. on Rob. and Auto.*, 2025.
- [17] M. Torne, A. Simeonov, Z. Li, A. Chan, T. Chen, A. Gupta, and P. Agrawal, "Reconciling reality through simulation: A real-to-sim-to-real approach for robust manipulation," in *Proc. Rob.: Sci. and Syst.*, 2024.
- [18] J. Yu, L. Fu, H. Huang, K. El-Refai, R. A. Ambrus, R. Cheng, M. Z. Irshad, and K. Goldberg, "Real2render2real: Scaling robot data without dynamics simulation or robot hardware," in *Proc. Conf. on Rob. Learn.*, 2025.
- [19] A. Nichol, H. Jun, P. Dhariwal, P. Mishkin, and M. Chen, "Point-e: A system for generating 3d point clouds from complex prompts," *arXiv preprint arXiv:2212.08751*, 2022.
- [20] R. Liu, R. Wu, B. Van Hoorick, P. Tokmakov, S. Zakharov, and C. Vondrick, "Zero-1-to-3: Zero-shot one image to 3d object," in *Proc. Int. Conf. Comput. Vis.*, pp. 9298–9309, 2023.
- [21] W. Goodwin, S. Vaze, I. Havoutis, and I. Posner, "Zero-shot category-level object pose estimation," in *Proc. Springer Eur. Conf. Comput. Vis.*, pp. 516–532, 2022.
- [22] S. James, Z. Ma, D. Rovick Arrojo, and A. J. Davison, "Rlbench: The robot learning benchmark & learning environment," *IEEE Robotics and Automation Letters*, 2020.
- [23] H. Ha, P. Florence, and S. Song, "Scaling up and distilling down: Language-guided robot skill acquisition," in *Proc. Conf. on Rob. Learn.*, pp. 3766–3777, 2023.
- [24] E. Chisari, N. Heppert, M. Argus, T. Welschehold, T. Brox, and A. Valada, "Learning robotic manipulation policies from point clouds with conditional flow matching," *Proc. Conf. on Rob. Learn.*, 2024.
- [25] M. Deitke, D. Schwenk, J. Salvador, L. Weihs, O. Michel, E. VanderBilt, L. Schmidt, K. Ehsani, A. Kembhavi, and A. Farhadi, "Objaverse: A universe of annotated 3d objects," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pp. 13142–13153, 2023.
- [26] M. K. Srirama, S. Dasari, S. Bahl, and A. Gupta, "Hrp: Human affordances for robotic pre-training," in *Proc. Rob.: Sci. and Syst.*, 2024.
- [27] M. Xu, Z. Xu, Y. Xu, C. Chi, G. Wetzstein, M. Veloso, and S. Song, "Flow as the cross-domain manipulation interface," in *Proc. Conf. on Rob. Learn.*, 2024.
- [28] H. Bharadhwaj, R. Mottaghi, A. Gupta, and S. Tulsiani, "Track2act: Predicting point tracks from internet videos enables generalizable robot manipulation," in *Proc. Springer Eur. Conf. Comput. Vis.*, 2024.
- [29] K. Rana, J. Abou-Chakra, S. Garg, R. Lee, I. Reid, and N. Suenderhauf, "Affordance-centric policy learning: Sample efficient and generalisable robot policy learning using affordance-centric task frames," in *Proc. Conf. on Rob. Learn.*, 2024.
- [30] J. Sun, Z. Shen, Y. Wang, H. Bao, and X. Zhou, "LoFTR: Detector-free local feature matching with transformers," *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2021.
- [31] N. D. Palo and E. Johns, "Keypoint action tokens enable in-context imitation learning in robotics," in *Proc. Rob.: Sci. and Syst.*, 2024.
- [32] P. Ardon, E. Pairet, K. S. Lohan, S. Ramamoorthy, and R. P. A. Petrick, "Building affordance relations for robotic agents - a review," in *Proc. of the Int. Joint Conf. on Artificial Intelligence*, pp. 4302–4311, 8 2021.
- [33] Y. Kuang, J. Ye, H. Geng, J. Mao, C. Deng, L. Guibas, H. Wang, and Y. Wang, "RAM: Retrieval-based affordance transfer for generalizable zero-shot robotic manipulation," in *Proc. Conf. on Rob. Learn.*, 2024.
- [34] S. Bahl, R. Mendonca, L. Chen, U. Jain, and D. Pathak, "Affordances from human videos as a versatile representation for robotics," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pp. 13778–13790, 2023.
- [35] R. Mendonca, S. Bahl, and D. Pathak, "Structured world models from human videos," in *Proc. Rob.: Sci. and Syst.*, 2023.
- [36] L. Wang, Y. Ling, Z. Yuan, M. Shridhar, C. Bao, Y. Qin, B. Wang, H. Xu, and X. Wang, "Gensim: Generating robotic simulation tasks via large language models," in *Int. Conf. on Learn. Repr.*, 2024.
- [37] L. Downs, A. Francis, N. Koenig, B. Kinman, R. Hickman, K. Reymann, T. B. McHugh, and V. Vanhoucke, "Google scanned objects: A high-quality dataset of 3d scanned household items," in *Proc. IEEE Int. Conf. on Rob. and Auto.*, pp. 2553–2560, 2022.
- [38] P. Katara, Z. Xian, and K. Fragkiadaki, "Gen2sim: Scaling up robot learning in simulation with generative models," in *Proc. IEEE Int. Conf. on Rob. and Auto.*, pp. 6672–6679, 2024.
- [39] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, "High-resolution image synthesis with latent diffusion models," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2022.
- [40] A. Zook, F.-Y. Sun, J. Spjut, V. Blukis, S. Birchfield, and J. Tremblay, "Grs: Generating robotic simulation tasks from real-world images," *arXiv preprint arXiv:2410.15536*, 2024.
- [41] Y. Wang, Z. Xian, F. Chen, T.-H. Wang, Y. Wang, K. Fragkiadaki, Z. Erickson, D. Held, and C. Gan, "Robogen: Towards unleashing infinite data for automated robot learning via generative simulation," in *Int. Conf. on Mach. Learn.*, pp. 51936–51983, 2024.
- [42] J. Tang, "Stable-dreamfusion: Text-to-3d with stable-diffusion," 2022. <https://github.com/ashawkey/stable-dreamfusion>.
- [43] H. Jun and A. Nichol, "Shap-e: Generating conditional 3d implicit functions," *arXiv preprint arXiv:2305.02463*, 2023.
- [44] B. Wen, W. Yang, J. Kautz, and S. Birchfield, "Foundationpose: Unified 6d pose estimation and tracking of novel objects," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pp. 17868–17879, 2024.
- [45] S. Umeyama, "Least-squares estimation of transformation parameters between two point patterns," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 13, no. 04, pp. 376–380, 1991.
- [46] Á. González, "Measurement of areas on a sphere using fibonacci and latitude–longitude lattices," *Mathematical geosciences*, vol. 42, pp. 49–64, 2010.
- [47] L. Di Giammarino, B. Sun, G. Grisetti, M. Pollefeys, H. Blum, and D. Barath, "Learning where to look: Self-supervised viewpoint selection for active localization using geometrical information," in *Proc. Springer Eur. Conf. Comput. Vis.*, pp. 188–205, 2024.
- [48] F. Xiang, Y. Qin, K. Mo, Y. Xia, H. Zhu, F. Liu, M. Liu, H. Jiang, Y. Yuan, H. Wang, L. Yi, A. X. Chang, L. J. Guibas, and H. Su, "SAPIEN: A simulated part-based interactive environment," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2020.
- [49] H.-S. Fang, C. Wang, M. Gou, and C. Lu, "Graspnet-1billion: A large-scale benchmark for general object grasping," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pp. 11444–11453, 2020.
- [50] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, "Pointnet: Deep learning on point sets for 3d classification and segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pp. 652–660, 2017.
- [51] N. Funk, J. Uraim, J. Carvalho, V. Prasad, G. Chalvatzaki, and J. Peters, "Actionflow: Equivariant, accurate, and efficient policies with spatially symmetric flow matching," *arXiv preprint arXiv:2409.04576*, 2024.
- [52] N. Heppert, M. Z. Irshad, S. Zakharov, K. Liu, R. A. Ambrus, J. Bohg, A. Valada, and T. Kollar, "Carto: Category and joint agnostic reconstruction of articulated objects," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pp. 21201–21210, 2023.