

# SIMPACT: Simulation-Enabled Action Planning using Vision-Language Models

Haowen Liu<sup>\*,1</sup> Shaoxiong Yao<sup>\*,2</sup> Haonan Chen<sup>3</sup> Jiawei Gao<sup>3</sup>  
 Jiayuan Mao<sup>4,5</sup> Jia-Bin Huang<sup>1</sup> Yilun Du<sup>3</sup>  
<sup>1</sup>UMD, <sup>2</sup>UIUC, <sup>3</sup>Harvard, <sup>4</sup>Amazon FAR, <sup>5</sup>UPenn



Figure 1. **Simulation-Enabled VLM Action Planning.** Given a single RGB-D image and a language task description (*left*), our method efficiently constructs a physics simulator that enables test-time VLM reasoning with physical grounding. This physically grounded reasoning allows the robot to succeed in fine-grained manipulation tasks (*bottom*), outperforming a vanilla VLM planner (*top*) that lacks awareness of physical dynamics.

## Abstract

Vision-Language Models (VLMs) exhibit remarkable common-sense and semantic reasoning capabilities. However, they lack a grounded understanding of physical dynamics. This limitation arises from training VLMs on static internet-scale visual-language data that contain no causal interactions or action-conditioned changes. Consequently, it remains challenging to leverage VLMs for fine-grained robotic manipulation tasks that require physical understanding, reasoning, and corresponding action planning. To overcome this, we present **SIMPACT**, a test-time, **SIM**ulation-enabled **ACT**ion **P**lanning framework that equips VLMs with physical reasoning through simulation-in-the-loop world modeling, without requiring any additional training. From a single RGB-D observation, SIMPACT efficiently constructs physics simulations, enabling the VLM to propose informed actions, observe simulated rollouts, and iteratively refine its reasoning. By integrating language reasoning with physics prediction, our simulation-enabled VLM can understand contact dynamics and action outcomes in a physically grounded way. Our method demonstrates state-of-the-art performance on seven challenging, real-world rigid-body and deformable

manipulation tasks that require fine-grained physical reasoning, outperforming existing general-purpose robotic manipulation models. Our results demonstrate that embedding physics understanding via efficient simulation into VLM reasoning at test time offers a promising path towards generalizable embodied intelligence. Project webpage can be found at <https://simpact-bot.github.io>.

## 1. Introduction

General-purpose robots hold significant promise for handling complex, labor-intensive tasks in unstructured environments, but realizing this potential requires advanced scene perception and robust action planning. Vision-Language Models (VLMs), trained on static internet-scale visual and language data, offer a promising solution by equipping robots to understand scenes and respond to diverse queries. These models can understand object semantics, infer task goals, and generate action descriptions aligned with human intent [10, 12, 46]. However, despite their remarkable commonsense and semantic reasoning capabilities, VLMs lack a **grounded understanding of physical dynamics**. They can describe what to do, but often fall short in predicting how actions will unfold when executed in the physical world.

\*Equal contribution

As such, VLMs have shown limited capabilities in robotic manipulation, particularly for tasks involving rich physical interactions, such as turning an object in place or carefully stacking objects. These tasks require reasoning about how objects behave under forces and constraints, where small variations in contact or timing can lead to drastically different outcomes. Lacking physical understanding, VLMs often propose plans that appear reasonable in language but fail during execution.

To address this limitation, we propose a framework that augments VLMs with physical simulation rollouts as contextual feedback, enabling test-time physical reasoning for action planning. Our approach begins with a novel simulation generation pipeline that leverages pretrained visual foundation models—including segmentation, 3D generation, and pose estimation models—to efficiently build a physical simulator directly from a single-view RGB-D image. In addition, we use VLMs to automate the setup of a multi-physics simulator, enabling it to model the behavior of both rigid and deformable objects across diverse material properties. The resulting physical simulation characterizes intricate contact dynamics that are difficult to infer from static images and language alone, providing VLMs with physical insights for manipulation planning.

Powered by the generated simulation, we introduce a planning framework driven by VLMs’ reasoning capabilities. Our key idea is to leverage the rich prior knowledge of VLMs to generate action sequence proposals, and to use simulated rollouts as context for the VLM to iteratively refine these proposals. This test-time reasoning paradigm, inspired by model-based control frameworks [53, 64], enables VLMs to reason not only about the world through language but also about its dynamics through simulated interaction. By augmenting VLMs with physical simulation, our framework enables them to anticipate action consequences, evaluate predicted outcomes, and iteratively adjust their decisions at test time, without any task-specific training. This process unlocks significantly stronger physical reasoning, enabling more reliable and robust real-world performance than state-of-the-art general-purpose manipulation models. In summary, this paper makes the following contributions:

- We introduce a test-time, zero-shot framework enabling VLMs to plan physics-aware embodied actions;
- We present a pipeline for automatically generating multi-physics simulations from a single RGB-D observation using visual foundation models and VLM;
- We propose a novel in-context learning approach for robot action generation, where physics simulation serves as context, enabling a new form of test-time reasoning in robotics.

## 2. Related Works

**Vision-Language Models for Robotics.** VLMs excel at scene understanding and language interpretation [33, 34, 46, 51, 52], making them promising for natural language-based robot control in open-world environments [5, 14, 15, 23, 61]. Many existing works directly fine-tune an action prediction head, i.e., vision-language-action (VLA) models, but these methods require large amounts of action-labeled training data, and their generalization ability remains limited [2, 3, 43, 70, 73]. Other works adopt carefully designed 3D geometric representations to enable VLMs to reason about actions. Examples include volumetric value maps [25], motion arrows [44], keypoint affordances [16, 69], and keypoint constraints [24, 26]. While these spatial representations advance VLM reasoning capacity through 3D grounding, they critically lack temporal dynamics, which are essential for tasks involving physical interaction and sequential manipulation. Early works have explored using physics simulation to augment reasoning in VLMs [37], and physical grounding for VLMs has also been investigated [18]. However, these efforts focus on question answering rather than the more challenging task of manipulation planning, which requires a robot to generate and refine continuous action sequences.

**Model-based Planning in Robotics.** Model-based planning has long been studied in robotics as a generalizable way to automatically synthesize long-horizon, complex action sequences [19, 22, 30–32, 42]. However, existing planning frameworks face limitations in open-world settings, where we need to build dynamic models from real-world perception and plan long-horizon actions from language instructions. With advances in deep learning, neural dynamics models have been developed to capture physical dynamics through image-space prediction [14, 17, 68], latent-space dynamics modeling [1, 20], and structured world representations [7, 8, 71, 72]. To improve planning efficiency, methods have been proposed to learn sampling distributions [49, 50] or to increase optimization efficiency using energy-based approaches [13, 27]. Nevertheless, these extensions still require training within a specific problem domain.

We argue that existing works do not fully address the open-world manipulation challenge. Task-specific models struggle with the diversity of real-world scenes. In contrast, pre-trained VLMs offer general scene understanding and reasoning capabilities, so we leverage them to support each component of our framework.

Our work also advances the construction of simulations from real-world observations. Compared to real-to-sim-to-real approaches such as digital twins [21, 29, 48, 59, 60, 65] and cousin creation [11], our approach constructs simulations more efficiently from a single-view RGB-D image. The recent method Prompting-with-the-Future [45] uses

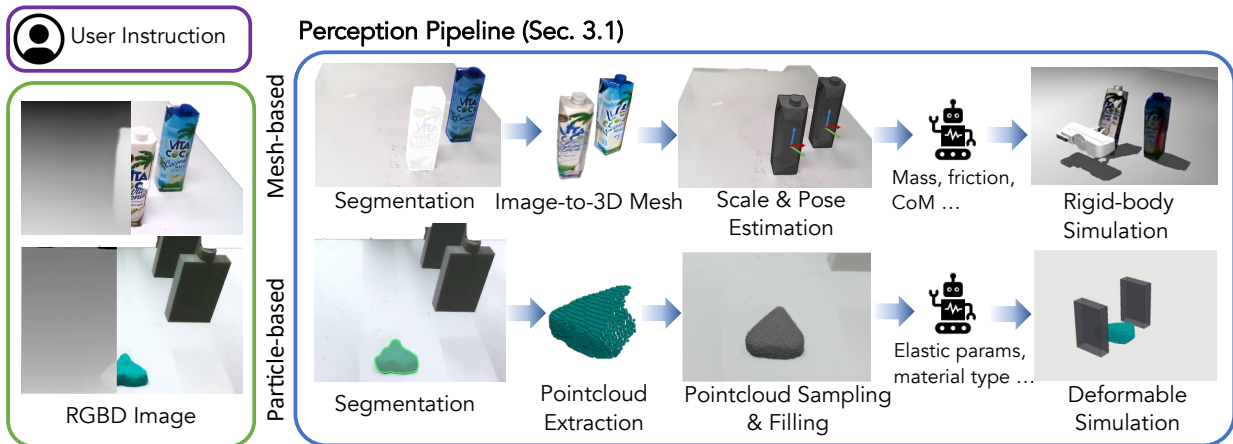


Figure 2. **Simulation construction from a single RGBD image.** Given an RGB-D image and a language task description, our pipeline automatically generates either a mesh-based simulation (*top*) for rigid objects or a particle-based simulation (*bottom*) for deformables. After segmenting objects-of-interest via GroundedSAM2 [55], we reconstruct either the 3D shape, scale, and pose of the object for rigid-body simulation, or perform dense sampling of particles within the volumes between the object surface and the table for the particle-based simulation pipeline. In both cases, we prompt the VLM to infer the relevant physical parameters required for simulation.

rigid-body simulation and a VLM solely as a reward signal in a model-predictive-control setup. In contrast, our method integrates multi-physics simulation and exploits VLMs for both informed action sampling and in-context learning optimization, resulting in substantial performance gains as demonstrated by our experiments.

### 3. Method

Our framework enables zero-shot robotic manipulation action generation from a single RGB-D image input  $I_0$  and natural language instruction  $\ell_{\text{task}}$  and outputs robot action sequence  $\mathbf{a} = \{a_t\}_{1 \leq t \leq T}$ , where  $a_t \in \text{SE}(3) \times \mathbb{R}$ , defining end-effector pose and gripper open width. For each task, the natural-language specification  $\ell_{\text{task}}$  defines the task requirements, along with potential success and failure conditions, to guide the VLM in proposing plausible actions.

Our simulation-enabled VLM planning framework operates as illustrated in Fig. 3. First, we construct a physical simulator SIM using an automated perception pipeline that reconstructs complete 3D geometries and configures appropriate simulation parameters as shown in Fig. 2. Next, we instantiate a manipulation planner that integrates the simulator with a VLM as its core reasoning module. The planner begins by generating a scene context from an initial visual observation, which is augmented with robot proprioceptive data and object states. Based on this context and prior knowledge, the VLM proposes action sequences, which are evaluated through simulation rollouts. The resulting visual observations and object states from each rollout are then fed back to the VLM as additional context for iterative refinement. This process continues until a rollout is validated as

successful. Finally, the optimized action sequence is executed as end-effector commands on the real robot system.

#### 3.1. Simulation Construction

Our approach employs a physics-based simulator to predict the consequences of actions for manipulation planning. The simulation follows the discrete-time state transition:

$$s_t = \text{SIM}(s_{t-1}, a_t; \theta) \quad (1)$$

where  $s_t$  denotes the state at time step  $t$ ,  $a_t$  represents the applied action, and  $\theta$  comprises time-invariant simulation parameters. The state space captures all task-relevant information: rigid objects are represented by a 6DoF pose in  $\text{SE}(3)$ , while deformable objects are described by  $N$  particle positions in  $\mathbb{R}^{3 \times N}$ . We initialize the state as  $s_0$ , assuming objects remain static prior to interaction, and construct parameters via  $\theta = \text{CreateSim}(I_0)$  from the initial RGBD image  $I_0$ . Here, the simulation parameters are defined as  $\theta = (\theta_{\text{geom}}, \theta_{\text{phys}})$ , where  $\theta_{\text{geom}}$  specifies the object shape and pose, and  $\theta_{\text{phys}}$  characterizes its mechanical properties.

Our geometry pipeline begins by prompting a VLM to generate object labels based on the user’s instructions, as shown in Fig. 2. We first apply a pretrained segmentation model, GroundedSAM2 [54, 55], to segment each identified object in  $I_0$ . We prompt the VLM to automatically select different physics engines based on object characteristics: MuJoCo [58] for rigid bodies, a variant of the projective dynamics [4] solver for stiff deformable objects that ensures numerical stability, and the Material Point Method [28] solver for soft objects to handle potential topological changes. We automate physical parameters  $\theta_{\text{phys}}$  inference by prompting the VLM to leverage its common-

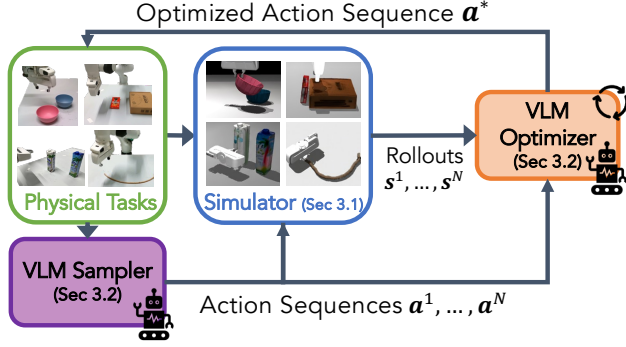


Figure 3. **Method overview.** Our method first instantiates a physics simulator given the real-world scene. Next, a VLM-based action sampler and optimizer iteratively refine the action sequence towards task success using simulated rollouts as context. The final optimized actions are then executed in the real world.

sense reasoning for plausible predictions, following prior works [6, 66, 67], with more details in Suppl. Sec. A.1.

**Mesh-based Rigid Body Simulation.** For rigid bodies, we define the geometry parameters as  $\theta_{\text{geom}} = \{(\mathcal{M}_i, X_i)\}_{i=1}^{N_{\text{obj}}}$ , where  $\mathcal{M}_i$  denotes the triangle mesh and  $X_i$  represents the initial 6DoF pose of object  $i$ . Using the segmented RGB image, we reconstruct complete triangle meshes for each object using a pretrained image-to-3D model [57], denoted as the unscaled mesh  $\hat{\mathcal{M}}_i$ . Each reconstructed mesh is then centered and scaled according to the size of its corresponding real-world bounding box obtained from point cloud segmentation, yielding  $\mathcal{M}_i = \alpha_i(\hat{\mathcal{M}}_i - \beta_i)$ , where  $\alpha_i$  denotes the ratio between the diagonal length of the real-world bounding box and that of the unscaled mesh, and  $\beta_i$  represents the 3D centroid of the unscaled mesh. Finally, we estimate the 6DoF pose  $X_i$  for each object using its triangle mesh  $\mathcal{M}_i$  in model-based mode and the RGB-D observation  $I_0$ , employing FoundationPose [63]. The physical parameters  $\theta_{\text{phys}}$  include mass, friction, and the center of mass for each rigid body.

**Particle-based Deformable Object Simulation.** For deformable objects, we define  $\theta_{\text{geom}} = \{P_i\}_{i=1}^{N_{\text{obj}}}$ , where each  $P_i \subset \mathbb{R}^3$  denotes the point set representing object  $i$ . We first back-project the segmented object mask from the depth image to obtain 3D surface points. To construct the full particle representation, we uniformly sample points within the volume bounded by the object surface and the supporting table surface, as illustrated in the bottom row of Fig. 2. Deformable bodies have  $\theta_{\text{phys}}$  defined by elasticity and plasticity parameters; see Suppl. Sec. A.1 for details.

### 3.2. Action Planning via Simulation-enabled VLM

Given the constructed simulator SIM, our action planning framework follows an iterative refinement process, as outlined in Fig. 3. As shown in Alg. 1, our planner takes as input the initial RGB-D observation  $I_0$ , the initial simulator state  $s_0$ , task description  $\ell_{\text{task}}$ , VLM, and SIM. The plan-

#### Algorithm 1: Action Planning Algorithm

```

1 Input: VLM, SIM,  $I_0, \ell_{\text{task}}, s_0$ ;
2  $\mathcal{A} = \emptyset, \mathcal{S} = \emptyset$ ; // Action sequences and rollouts
  // Initial action sampling and simulation
3 for  $k = 1..K$  do
4    $\mathcal{A} \leftarrow \mathcal{A} \cup \{\mathbf{a}^i \leftarrow \text{SAMPLE}(I_0, \ell_{\text{task}}, s_0; \text{VLM})\}$ ;
5    $\mathcal{S} \leftarrow \mathcal{S} \cup \{\mathbf{s}^i \leftarrow \text{SIMROLLOUT}(s_0, \mathbf{a}^i; \text{SIM})\}$ ;
  // Iterative action optimization
6 for  $k = K+1$  to  $K_{\text{max}}$  do
7    $\mathbf{a}^k \leftarrow \text{OPTIMIZE}(\mathcal{A}, \mathcal{S}, \ell_{\text{task}}; \text{VLM})$ ;
8    $\mathbf{s}^k \leftarrow \text{SIMROLLOUT}(s_0, \mathbf{a}^k; \text{SIM})$ ;
9   if  $\text{TASKSUCCESS}(\mathbf{s}^k; \text{VLM})$  then
10    break; // Stop when successful
11  else
12     $\mathcal{A} \leftarrow \mathcal{A} \cup \{\mathbf{a}^k\}, \mathcal{S} \leftarrow \mathcal{S} \cup \{\mathbf{s}^k\}$ ;
13 return  $\mathbf{a}^k$ 

```

ner begins by sampling an initial set of action sequences  $\mathcal{A}$  from the VLM prior. For each action sequence  $\mathbf{a}^i \in \mathcal{A}$ , the SIMROLLOUT procedure iteratively applies each action  $\mathbf{a}_t^i$  and uses the SIM function to obtain the next state  $s_{t+1}^i$ , adding simulation rollouts  $\mathbf{s}^i \in \mathcal{S}$ .

After initialization, each iteration proceeds as follows. Using both  $\mathcal{A}$  and  $\mathcal{S}$ , a VLM-based optimizer refines the proposed action sequences and produces a new action sequence  $\mathbf{a}^k$ . Based on the simulated rollout  $\mathbf{s}^k$ , the VLM model then evaluates whether  $\mathbf{a}^k$  achieves the task goal. If successful, the corresponding action sequence  $\mathbf{a}^k$  is executed on the real robot, and the process terminates. Otherwise, the planner proceeds to another round of action optimization by adding a newly generated  $\mathbf{a}^k$  to  $\mathcal{A}$  and  $\mathbf{s}^k$  to  $\mathcal{S}$ , until either a successful plan is found or the maximum iteration limit  $K_{\text{max}}$  is reached.

At the heart of our planning framework is the VLM, which uses its pretrained knowledge to instantiate the SAMPLE, OPTIMIZE, and TASKSUCCESS modules. For each role, we define a corresponding system prompt  $\ell_*$ , where  $*$  denotes sample, opt, or eval, specifying the function that the VLM performs.

**VLMs for Action Proposal Generation.** To instantiate SAMPLE using a VLM, we build upon two key ideas: (1) constructing an informed contextual description of the environment, and (2) leveraging hierarchical action generation.

(1) *Contextual representation.* We begin by constructing a comprehensive context that includes the initial visual observation  $I_0$  and the robot’s proprioceptive state. For the manipulated objects, we further incorporate their 6-DoF poses along with key geometric attributes, such as bounding box dimensions. Full details are shown in Suppl. Sec. A.2.

(2) *Hierarchical action generation.* Directly prompting a VLM to generate continuous 6-DoF end-effector poses for long-horizon tasks is challenging, as such representations lack clear semantic meaning and are difficult for VLMs to reason about. In contrast, we find that VLMs are highly effective at producing high-level symbolic action sequences




Figure 4. **Action optimization process.** We show a representative example from the *non-toppling push* task. The left three images show simulation rollouts from initial VLM-sampled action sequence proposals, all of which fail due to insufficient/overshooting push, or because the bottle topples. From these proposals, the VLM optimizer reasons a non-trivial action update that pushes the bottle for the correct distance without toppling in both simulation and real-world execution.

Table 1. **Definition of tasks.** For each manipulation task, we list the corresponding instruction and success criteria.

Tasks	Instruction	Success Condition
Non-toppling push	Push the white carton forward to align horizontally with the others.	The bottle does not topple and aligns within $\pm 2$ cm.
Bowl stacking	Grasp the pink bowl at its edge and stack it with the blue bowl.	The pink bowl stably lies inside the blue bowl.
Pivoting	Make the red pocky box lean vertically against the brown box.	The red pocky box reaches vertical pose.
Shape rope	Grab the free end of the rope and arrange the rope to a U shape.	The deformed rope has an opening ratio in $[0.5, 2.0]$ .
Shape playdoh	Squeeze the Play-Doh to a square shape with equal sides.	Two sides of the Play-Doh have a ratio within 1.5.
Avoid obstacle	Push the orange bottle around the box to the image’s bottom-right.	The orange bottle reaches the other side without colliding.
Sweeping	Sweep the coffee beans into the purple box.	All beans are pushed into the box.

that align with patterns seen in their pretraining data. Accordingly, we define a compact set of symbolic actions e.g., MOVE, GRASP, and RELEASE, to better exploit the semantic reasoning capabilities of VLMs. Each symbolic action is further parameterized by continuous control variables, enabling fine-grained and precise motion execution.

Formally, we represent a high-level action at time  $t$  as  $A_t = (\tau_t, u_t)$  where  $\tau_t$  denotes the high-level action type and  $u_t$  represents the continuous control parameters. A deterministic mapping, ACTION2POSE, translates a sequence of high-level actions into continuous 6-DoF control trajectories. Within the SAMPLE function, let  $A^i$  denote the  $i$ -th high-level action sequence. The VLM then samples an action sequence as

$$\mathbf{a}^i = \text{ACTION2POSE}(A^i = \text{VLM}(I_0, \ell_{\text{task}}, s_0; \ell_{\text{sample}})).$$

where  $\ell_{\text{sample}}$  denotes the system prompt that specifies the sampling behavior of the VLM (refer to Suppl. Sec. A.2 for details).

**VLMs for Action Optimization.** Given sampled action sequences  $\mathcal{A} = \{\mathbf{a}^i\}_{i=1}^K$ , we first perform simulation rollouts to obtain their corresponding state trajectories  $\mathcal{S} = \{\mathbf{s}^i\}_{i=1}^K$ . Next, we instantiate the OPTIMIZE function using the VLM via in-context learning. For each action sequence, we construct an optimization context  $c^i$  by subsampling time steps and gathering intermediate information. In particular, at each selected time step  $t$ , we render a simulator observation image  $I_t^i$  and include the numerical action  $a_t^i$  and state  $s_t^i$  in the context. This provides the VLM with both visual and state-based evidence to guide optimization.

$$\mathbf{a}^k = \text{VLM}(c^1, \dots, c^K; \ell_{\text{opt}}) \quad (2)$$

This optimization process is not restricted to local updates as in numerical optimizers; instead, it can perform reasoning and even learn from all failure examples. Fig. 4 illustrates a case where the VLM-based optimizer learns from failed attempts, produces a successful action sequence through its internal reasoning.

**VLMs for Success Evaluation.** Given the simulation rollout  $s^k$ , we render the final simulation state and extract both an observation image  $I_T^k$  and the simulator state  $s_T^k$ . These are used as contextual inputs for the VLM to assess whether the task is successfully completed. The evaluation is formulated as

$$\text{TASKSUCCESS}(\mathbf{s}^k) = \text{VLM}(I_T^k, s_T^k, \ell_{\text{task}}; \ell_{\text{eval}}).$$

If the VLM determines that the proposed action sequence achieves the task objective, the sequence is executed in the real environment. Otherwise, the system continues to optimize actions until the iteration limit is reached.

## 4. Experiments

To evaluate the effectiveness of our framework, we design seven challenging, real-world, physics-aware, fine-grained manipulation tasks. We assess whether our method enables zero-shot planning on these tasks, comparing it against other state-of-the-art zero-shot methods. We validate our design choices through systematic ablation studies.

### 4.1. Experimental Setup

We evaluate our system using a Franka Research 3 robot arm with a parallel-jaw gripper. For the Play-Doh manipulation task, we use a custom 3D-printed end effector to

Table 2. **Success rates of our method and baselines.** For each task, we run 10 trials per method. Our approach consistently achieves a substantially higher success rate than baselines, highlighting the effectiveness of simulation-enabled VLMs for action planning.

Method	Non-toppling push	Bowl stacking	Pivoting	Shape rope	Shape dough	Avoid obstacle	Sweeping
$\pi_{0.5}$ [2]	0%	0%	0%	0%	0%	0%	0%
VoxPoser[25]	0%	20%	0%	0%	0%	0%	20%
MOKA[16]	0%	10%	0%	20%	0%	0%	0%
Ours	<b>80%</b>	<b>60%</b>	<b>40%</b>	<b>90%</b>	<b>80%</b>	<b>80%</b>	<b>70%</b>



Figure 5. **Qualitative results.** The figure shows the initial state, execution progress, and final state for six tasks in both the real world (top) and the simulation (bottom). By leveraging VLM’s powerful generalization, rendered simulation images can guide VLM’s test-time reasoning for action planning despite the visual sim2real gap. Fig. 10 shows the results of the remaining task.

achieve a sufficiently large contact area. A single calibrated Intel RealSense D435i RGBD camera is used.

**Tasks and Metrics.** We design diverse tasks requiring fine-grained, physics-aware manipulation planning. The objects span rigid bodies (cartons, bowls, boxes) to deformable materials (rope, Play-Doh), enabling evaluation across different physical properties and manipulation strategies, including pushing, grasping, pivoting, squeezing, and sweeping. Success rate is our primary evaluation metric. Task instructions and success criteria are detailed in Table 1.

**Baselines.** We compare our approach against the following baselines: (1) *VLA models* that are trained on large-scale robot action datasets to directly predict joint velocities from visual observations and language instructions. We use  $\pi_{0.5}$  [2], a recent open-source VLA model pretrained on a large robot manipulation dataset, as a representative base-

line. (2) *VLM-based methods* that leverage geometric representations to augment VLM for manipulation planning. We compare against VoxPoser [25], which uses volumetric value maps to represent spatial affordances in 3D, and MOKA [16], which predicts keypoints and affordance regions to generate manipulation actions. For our pushing and squeezing scenarios, we extend MOKA, which originally supports only grasping, to represent contact location with a target contact point and infer contact direction from pre- and post-contact positions.

**Implementation details.** For simulation, we implement the projective dynamics variant solver using PyTorch [47] and the MPM simulator using Warp [39]. We use Google Gemini 2.5 Pro as the default VLM [10] and generate  $K = 10$  initial action proposals, setting  $K_{\max} = 13$ , corresponding to a maximum of 3 action-optimization iterations.

Table 3. **Ablation.** Success rates (%) over 10 trials for each task after removing each component of our method. Results demonstrate the importance of VLM-conditioned sampling and the VLM’s simulation-enabled test-time reasoning capabilities.

Method	Non-toppling push	Bowl stacking	Pivoting	Shape rope	Shape dough	Avoid obstacle	Sweeping
w/o VLM sampler	0%	10%	0%	0%	0%	0%	20%
w/o simulation rollout	20%	0%	0%	30%	30%	0%	20%
w/o VLM optimizer	30%	50%	30%	40%	80%	20%	70%
Ours	<b>80%</b>	<b>60%</b>	<b>40%</b>	<b>90%</b>	<b>80%</b>	<b>80%</b>	<b>70%</b>

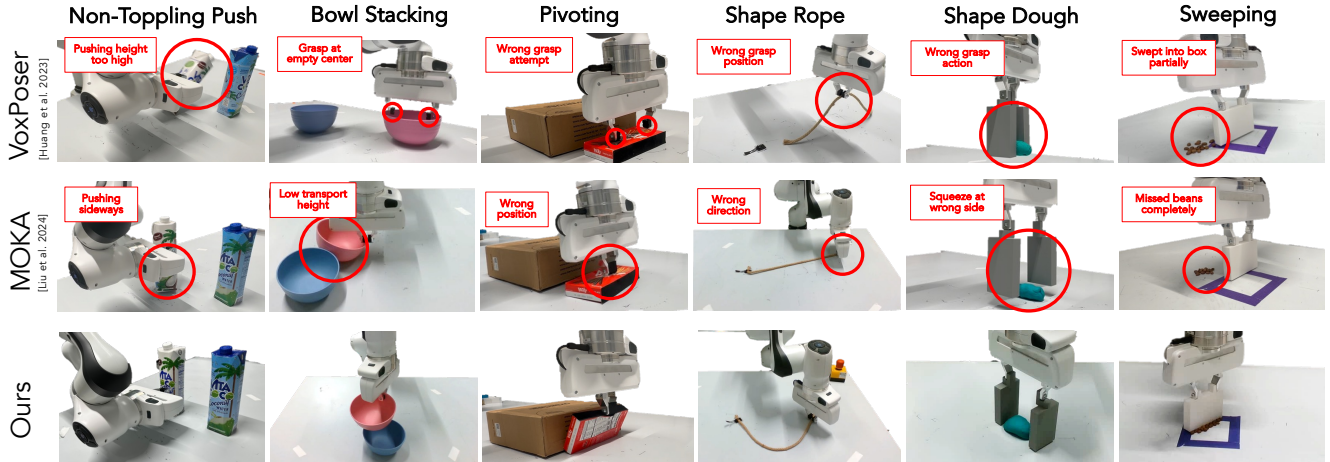


Figure 6. **Qualitative comparison with baseline methods.** We show representative failures from baseline methods that lack simulation-enabled reasoning. These methods often choose incorrect action parameters, using improper pushing heights that cause toppling, or attempting to grasp the bowl at its center. Pivoting tasks fail because the baselines do not maintain contact with the box’s side face. For the rope task, baselines place the rope in the wrong direction due to missing deformation reasoning; for the dough-shaping task, baselines fail to plan the perpendicular squeezes needed to form a square; for the sweeping task, baselines fail to push all beans into the box.

## 4.2. Results

Table 2 shows the success rates of our method compared to baseline approaches. Overall, our method consistently outperforms baseline methods across all evaluated tasks, highlighting its strong performance on challenging, physics-aware, fine-grained manipulation tasks. Fig. 5 shows simulation and real-world rollouts of six of our seven tasks.

From the table, the VLA model  $\pi_{0.5}$  consistently fails on all tasks. While we observe that  $\pi_{0.5}$  can sometimes generate actions that approach the target object, it fails to complete the manipulation. This is because while VLA models can perform zero-shot on tasks similar to those seen during training, they generalize poorly to out-of-domain, challenging tasks used in our experiment. VLM-based methods, VoxPoser and MOKA, leveraging VLM’s strong scene-understanding and reasoning capabilities, achieve non-zero success rates on tasks such as *bowl stacking*, *shape rope* and *sweeping*. However, they struggle with tasks that require precise action planning, where small errors, such as pushing the wrong part of an object (in *non-toppling push*) or squeezing an incorrect region of deformable materials (in *shape dough*) lead to failures, as shown in Fig. 6. In contrast, our method integrates simulation-enabled reasoning with VLM, enabling the robot to iteratively refine its

action plan using simulation rollouts as context. This enables the system to identify and avoid physically unstable or ineffective strategies. For example, in *non-toppling push*, the simulation shows that pushing near the top of the carton would cause toppling, so the system adapts by pushing from a more stable point, as shown in Fig. 4.

## 4.3. Ablation study

We consider three ablated variants of our method. (1) *Removing the VLM sampler*: To assess the importance of VLM-guided action sampling, we replace them with uninformed sampling from a Gaussian distribution over gripper pose deltas. To ensure fairness, we increase the sample size by  $5\times$ . This variant resembles the Prompting-with-the-Future approach [45], but uses a VLM-based optimizer rather than the cross-entropy method (CEM). The VLM optimizer is more effective due to its reasoning capability, as shown in Fig. 4, rather than being limited to the local action updates assumed by CEM. (2) *Removing simulation rollout context*: We evaluate whether current VLMs can reason effectively without simulation rollouts. Following a proposer-verifier structure, the VLM generates and evaluates multiple action proposals using only its internal reasoning. (3) *Removing the VLM optimizer*: We disable iterative refinement and let the VLM select the best action from the initial

Table 4. **Sampling length ablation.** Success rates (%) over 10 trials varying numbers of in-context examples for tasks *non-toppling push*, *bowl stacking*, *shape rope*.

#Samples	Non-toppling push	Bowl stacking	Shape rope
3 samples	50%	50%	40%
10 samples	80%	60%	90%
20 samples	90%	20%	80%

proposals based on simulation outcomes, testing whether a naive optimization process is sufficient.

From Table 3, removing the VLM sampling module causes a significant performance drop. For fine-grained manipulation tasks, purely random sampling often yields actions far from feasible solutions, providing no useful guidance for subsequent VLM reasoning. This underscores the importance of VLM-conditioned action sampling in generating reasonable action proposals. Removing simulation-rollout validation also substantially degrades performance, particularly in tasks such as *bowl stacking*, *pivoting*, and *avoiding obstacles*. This indicates that language-based reasoning without physical grounding cannot reliably infer successful action. However, the variant still outperforms baseline methods, largely due to the hierarchical action sampling strategy introduced in Sec. 3. Finally, disabling the VLM optimizer also results in a performance decrease. This decline is especially pronounced in tasks such as *non-toppling pushing*, *shape rope*, and *avoid obstacle*, where the initial VLM-generated samples are often inadequate for completing the tasks and require iterative refinement.

In addition, we investigate how the number of initial VLM-sampled actions ( $K$  in Alg. 1) affects performance, reporting success rates over 10 trials in Table 4: using only 3 samples degrades performance, as limited rollouts fail to provide sufficient task information, leading to poor optimization; Increasing samples also does not always help, as performance drops for *bowl stacking* and *shape rope* when increasing samples from 10 to 20, particularly for task *bowl stacking*. This likely occurs because longer contexts reduce VLM reasoning efficacy [62], which could be mitigated by pre-selecting informative samples. In contrast, *non-toppling push* slightly improves, as its shorter action horizon limits context growth.

#### 4.4. Failure Case Analysis

Fig. 7 shows the failure distribution across tasks, categorized into three failure types: perception, planning, and execution failures. *Perception failures* mainly stem from errors in single-view 3D reconstruction, which could be reduced by using better image-to-3D models or changing the observation view. *Planning failures* occur when the robot fails to generate a feasible action sequence even after multiple rounds of action optimization. These are the most frequent failure cases, especially in the pivoting task, where finding a




Figure 7. **Failure case decomposition graph.** Failures are categorized as perception, planning, or execution.

successful action sequence is particularly challenging. *Execution failures* arise when kinematic or dynamic discrepancies between simulation and reality cause actions that succeed in simulation to fail in real-world execution.

#### 4.5. Limitations

There are several limitations to our method. First, simulation quality depends heavily on the underlying image-to-3D reconstruction in the rigid-body case, where single-view methods struggle with occlusions. Incorporating inpainting or generative 3D completion models [38] may help alleviate this issue. Recent works on articulated object reconstruction [9, 36] can also be integrated into our framework. Next, since we estimate physical parameters via VLM prompting, inaccuracies can lead to sim-to-real discrepancies and affect downstream planning. Integrating system identification using real-world interaction data could help refine these estimates. Finally, our current system performs only open-loop execution, making it vulnerable to compounding errors. We include an optional replanning mechanism for recovery that updates the simulator using real-world feedback (see Suppl. Sec. B.9), and future work will explore tighter integration of real-time feedback for improved robustness.

### 5. Conclusion

We introduce **SIMPACT**, a novel action-planning framework that leverages simulation-enabled VLM to enable zero-shot robotic manipulation without any task-specific training. Our approach is made possible by a foundation-model-enabled simulation construction pipeline and a test-time VLM reasoning framework that together unlock the rich commonsense knowledge and reasoning capabilities of VLMs for physics-aware, fine-grained robotic manipulation. Real-world experiments demonstrate that SIMPACT provides substantial improvements over state-of-the-art general-purpose manipulation models. Additional ablation studies further highlight the importance of both simulation construction and test-time reasoning in achieving generalizability and high performance.

## Acknowledgements

We would like to thank Sicong Pan and Hanxiao Jiang for insightful discussions on experimental design, Xiaoshen Han for assistance with hardware setup preparation, Wei-Cheng Huang and Kaifeng Zhang for valuable suggestions on real2sim implementations, and Fangchen Liu for MOKA implementation. This work was partly supported by the Kempner Institute, Amazon, and Pickle Robot.

## References

- [1] Pulkit Agrawal, Ashvin V Nair, Pieter Abbeel, Jitendra Malik, and Sergey Levine. Learning to poke by poking: Experiential learning of intuitive physics. *Advances in neural information processing systems*, 29, 2016. 2
- [2] Kevin Black, Noah Brown, James Darpinian, Karan Dhabalia, Danny Driess, Adnan Esmail, Michael Robert Equi, Chelsea Finn, Niccolo Fusai, Manuel Y. Galliker, Dibya Ghosh, Lachy Groom, Karol Hausman, brian ichter, Szymon Jakubczak, Tim Jones, Liyiming Ke, Devin LeBlanc, Sergey Levine, Adrian Li-Bell, Mohith Mothukuri, Suraj Nair, Karl Pertsch, Allen Z. Ren, Lucy Xiaoyang Shi, Laura Smith, Jost Tobias Springenberg, Kyle Stachowicz, James Tanner, Quan Vuong, Homer Walke, Anna Walling, Haohuan Wang, Lili Yu, and Ury Zhilinsky.  $\pi_{0.5}$ : a vision-language-action model with open-world generalization. In *Proceedings of The 9th Conference on Robot Learning*, pages 17–40. PMLR, 2025. 2, 6
- [3] Kevin Black, Noah Brown, Danny Driess, Adnan Esmail, Michael Robert Equi, Chelsea Finn, Niccolo Fusai, Lachy Groom, Karol Hausman, Brian Ichter, Szymon Jakubczak, Tim Jones, Liyiming Ke, Sergey Levine, Adrian Li-Bell, Mohith Mothukuri, Suraj Nair, Karl Pertsch, Lucy Xiaoyang Shi, Laura Smith, James Tanner, Quan Vuong, Anna Walling, Haohuan Wang, and Ury Zhilinsky.  $\pi_0$ : A Vision-Language-Action Flow Model for General Robot Control. In *Proceedings of Robotics: Science and Systems*, Los Angeles, CA, USA, 2025. 2
- [4] Sofien Bouaziz, Sebastian Martin, Tiantian Liu, Ladislav Kavan, and Mark Pauly. Projective dynamics: fusing constraint projections for fast simulation. *ACM Trans. Graph.*, 33(4), 2014. 3
- [5] Boyuan Chen, Zhuo Xu, Sean Kirmani, Brain Ichter, Dorsa Sadigh, Leonidas Guibas, and Fei Xia. SpatialVlm: Endowing vision-language models with spatial reasoning capabilities. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14455–14465, 2024. 2
- [6] Boyuan Chen, Hanxiao Jiang, Shaowei Liu, Saurabh Gupta, Yunzhu Li, Hao Zhao, and Shenlong Wang. Physgen3d: Crafting a miniature interactive world from a single image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6178–6189, 2025. 4
- [7] Haonan Chen, Yilong Niu, Kaiwen Hong, Shuijing Liu, Yixuan Wang, Yunzhu Li, and Katherine Rose Driggs-Campbell. Predicting object interactions with behavior primitives: An application in stowing tasks. In *7th Annual Conference on Robot Learning*, 2023. 2
- [8] Haonan Chen, Jiaming Xu, Lily Sheng, Tianchen Ji, Shuijing Liu, Yunzhu Li, and Katherine Driggs-Campbell. Learning coordinated bimanual manipulation policies using state diffusion and inverse dynamics models. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5644–5651, 2025. 2
- [9] Zoey Chen, Aaron Walsman, Marius Memmel, Kaichun Mo, Alex Fang, Karthikeya Vemuri, Alan Wu, Dieter Fox, and Abhishek Gupta. Urdformer: A pipeline for constructing articulated simulation environments from real-world images. In *Robotics: Science and Systems*, 2024. 8
- [10] Gheorghe Comanici, Eric Bieber, Mike Schaekermann, Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blstein, Ori Ram, Dan Zhang, Evan Rosen, et al. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*, 2025. 1, 6
- [11] Tianyuan Dai, Josiah Wong, Yunfan Jiang, Chen Wang, Cem Gokmen, Ruohan Zhang, Jiajun Wu, and Li Fei-Fei. Automated creation of digital cousins for robust policy learning. In *8th Annual Conference on Robot Learning*, 2024. 2
- [12] Danny Driess, Fei Xia, Mehdi S. M. Sajjadi, Corey Lynch, Aakanksha Chowdhery, Brian Ichter, Ayzaan Wahid, Jonathan Tompson, Quan Vuong, Tianhe Yu, Wenlong Huang, Yevgen Chebotar, Pierre Sermanet, Daniel Duckworth, Sergey Levine, Vincent Vanhoucke, Karol Hausman, Marc Toussaint, Klaus Greff, Andy Zeng, Igor Mordatch, and Pete Florence. Palm-e: An embodied multimodal language model. In *arXiv preprint arXiv:2303.03378*, 2023. 1
- [13] Yilun Du, Toru Lin, and Igor Mordatch. Model based planning with energy based models. In *Conference on Robot Learning*, 2019. 2
- [14] Yilun Du, Sherry Yang, Pete Florence, Fei Xia, Ayzaan Wahid, brian ichter, Pierre Sermanet, Tianhe Yu, Pieter Abbeel, Joshua B. Tenenbaum, Leslie Pack Kaelbling, Andy Zeng, and Jonathan Tompson. Video language planning. In *The Twelfth International Conference on Learning Representations*, 2024. 2
- [15] Jiafei Duan, Wentao Yuan, Wilbert Pumacay, Yi Ru Wang, Kiana Ehsani, Dieter Fox, and Ranjay Krishna. Manipulate-anything: Automating real-world robots using vision-language models. In *8th Annual Conference on Robot Learning*, 2024. 2
- [16] Kuan Fang, Fangchen Liu, Pieter Abbeel, and Sergey Levine. Moka: Open-world robotic manipulation through mark-based visual prompting. *Robotics: Science and Systems (RSS)*, 2024. 2, 6
- [17] Chelsea Finn and Sergey Levine. Deep visual foresight for planning robot motion. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2786–2793, 2017. 2
- [18] Jiahui Gao, Bikram Sarkar, Fei Xia, Tony Xiao, Jiajun Wu, Brian Ichter, Anirudha Majumdar, and Dorsa Sadigh. Physically grounded vision-language models for robotic manipulation. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 12462–12469. IEEE, 2024. 2

- [19] Caelan Reed Garrett, Rohan Chitnis, Rachel Holladay, Beomjoon Kim, Tom Silver, Leslie Pack Kaelbling, and Tomás Lozano-Pérez. Integrated task and motion planning. *Annual Review of Control, Robotics, and Autonomous Systems*, 4(Volume 4, 2021):265–293, 2021. 2
- [20] David Ha and Jürgen Schmidhuber. Recurrent world models facilitate policy evolution. In *Advances in Neural Information Processing Systems*. Curran Associates, Inc., 2018. 2
- [21] Xiaoshen Han, Junqiu Yu, Minghuan Liu, Yilun Chen, Xiaoyang Lyu, Yang Tian, Bolun Wang, Weinan Zhang, Weinan Zhang, and Jiangmiao Pang. Re<sup>3</sup>sim: Generating high-fidelity simulation data via 3d-photorealistic real-to-sim for robotic manipulation. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2026. 2
- [22] Taylor Howell, Nimrod Gileadi, Saran Tunyasuvunakool, Kevin Zakka, Tom Erez, and Yuval Tassa. Predictive Sampling: Real-time Behaviour Synthesis with MuJoCo. 2022. 2
- [23] Yingdong Hu, Fanqi Lin, Tong Zhang, Li Yi, and Yang Gao. Look before you leap: Unveiling the power of GPT-4v in robotic vision-language planning. In *First Workshop on Vision-Language Models for Navigation and Manipulation at ICRA 2024*, 2024. 2
- [24] Haoxu Huang, Fanqi Lin, Yingdong Hu, Shengjie Wang, and Yang Gao. Copa: General robotic manipulation through spatial constraints of parts with foundation models. In *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 9488–9495. IEEE, 2024. 2
- [25] 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. In *7th Annual Conference on Robot Learning*, 2023. 2, 6
- [26] Wenlong Huang, Chen Wang, Yunzhu Li, Ruohan Zhang, and Li Fei-Fei. Rekep: Spatio-temporal reasoning of relational keypoint constraints for robotic manipulation. In *8th Annual Conference on Robot Learning*, 2024. 2
- [27] Michael Janner, Yilun Du, Joshua B. Tenenbaum, and Sergey Levine. Planning with diffusion for flexible behavior synthesis. In *International Conference on Machine Learning*, 2022. 2
- [28] Chenfanfu Jiang, Craig Schroeder, Joseph Teran, Alexey Stomakhin, and Andrew Selle. The material point method for simulating continuum materials. In *ACM SIGGRAPH 2016 Courses*, New York, NY, USA, 2016. Association for Computing Machinery. 3
- [29] Hanxiao Jiang, Hao-Yu Hsu, Kaifeng Zhang, Hsin-Ni Yu, Shenlong Wang, and Yunzhu Li. Phystwin: Physics-informed reconstruction and simulation of deformable objects from videos. *ICCV*, 2025. 2
- [30] Shucheng Kang, Xiaoyang Xu, Jay Sarva, Ling Liang, and Heng Yang. Fast and certifiable trajectory optimization. In *International Workshop on the Algorithmic Foundations of Robotics*, 2024. 2
- [31] Shucheng Kang, Guorui Liu, and Heng Yang. Global contact-rich planning with sparsity-rich semidefinite relaxations. In *Robotics: Science and Systems*, 2025.
- [32] Steven M. LaValle. *Planning Algorithms*. Cambridge University Press, USA, 2006. 2
- [33] Junnan Li, Dongxu Li, Caiming Xiong, and Steven CH Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International Conference on Machine Learning*, pages 12888–12900. PMLR, 2022. 2
- [34] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: bootstrapping language-image pre-training with frozen image encoders and large language models. In *Proceedings of the 40th International Conference on Machine Learning*. JMLR.org, 2023. 2
- [35] Xinghang Li, Minghuan Liu, Hanbo Zhang, Cunjun Yu, Jie Xu, Hongtao Wu, Chilam Cheang, Ya Jing, Weinan Zhang, Huaping Liu, Hang Li, and Tao Kong. Vision-language foundation models as effective robot imitators. In *The Twelfth International Conference on Learning Representations*, 2024. 5
- [36] Jiayi Liu, Denys Iliash, Angel X Chang, Manolis Savva, and Ali Mahdavi Amiri. SINGAPO: Single image controlled generation of articulated parts in objects. In *The Thirteenth International Conference on Learning Representations*, 2025. 8
- [37] Ruibo Liu, Jason Wei, Shixiang Shane Gu, Te-Yen Wu, Soroush Vosoughi, Claire Cui, Denny Zhou, and Andrew M. Dai. Mind’s eye: Grounded language model reasoning through simulation. In *The Eleventh International Conference on Learning Representations*, 2023. 2
- [38] Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van Gool. Repaint: Inpainting using denoising diffusion probabilistic models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11461–11471, 2022. 8
- [39] Miles Macklin. Warp: A high-performance python framework for gpu simulation and graphics, 2022. NVIDIA GPU Technology Conference (GTC). 6
- [40] Oier Mees, Lukas Hermann, and Wolfram Burgard. What matters in language conditioned robotic imitation learning over unstructured data. *IEEE Robotics and Automation Letters (RA-L)*, 7(4):11205–11212, 2022. 5
- [41] Oier Mees, Lukas Hermann, Erick Rosete-Beas, and Wolfram Burgard. Calvin: A benchmark for language-conditioned policy learning for long-horizon robot manipulation tasks. *IEEE Robotics and Automation Letters (RA-L)*, 7(3):7327–7334, 2022. 5
- [42] Igor Mordatch, Emanuel Todorov, and Zoran Popović. Discovery of complex behaviors through contact-invariant optimization. *ACM Trans. Graph.*, 31(4), 2012. 2
- [43] Yao Mu, Qinglong Zhang, Mengkang Hu, Wenhai Wang, Mingyu Ding, Jun Jin, Bin Wang, Jifeng Dai, Yu Qiao, and Ping Luo. Embodiedgpt: Vision-language pre-training via embodied chain of thought. *Advances in Neural Information Processing Systems*, 36:25081–25094, 2023. 2
- [44] Soroush Nasiriany, Fei Xia, Wenhao Yu, Ted Xiao, Jacky Liang, Ishita Dasgupta, Annie Xie, Danny Driess, Ayzaan Wahid, Zhuo Xu, Quan Vuong, Tingnan Zhang, Tsang-Wei Edward Lee, Kuang-Huei Lee, Peng Xu, Sean Kirmani, Yuke Zhu, Andy Zeng, Karol Hausman, Nicolas Heess, Chelsea Finn, Sergey Levine, and Brian Ichter. Pivot: iterative visual prompting elicits actionable knowledge for vlms.

- In *Proceedings of the 41st International Conference on Machine Learning*. JMLR.org, 2024. 2
- [45] Chuanruo Ning, Kuan Fang, and Wei-Chiu Ma. Prompting with the future: Open-world model predictive control with interactive digital twins. In *RSS*, 2025. 2, 7, 1
- [46] OpenAI. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023. 1, 2
- [47] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019. 6
- [48] Shivansh Patel, Xinchen Yin, Wenlong Huang, Shubham Garg, Hooshang Nayyeri, Li Fei-Fei, Svetlana Lazebnik, and Yunzhu Li. A real-to-sim-to-real approach to robotic manipulation with vlm-generated iterative keypoint rewards. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, 2025. 2
- [49] Thomas Power and Dmitry Berenson. Learning a generalizable trajectory sampling distribution for model predictive control. *IEEE Transactions on Robotics*, 40:2111–2127, 2024. 2
- [50] Ahmed H Qureshi, Anthony Simeonov, Mayur J Bency, and Michael C Yip. Motion planning networks. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 2118–2124. IEEE, 2019. 2
- [51] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021. 2
- [52] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *International Conference on Machine Learning*, pages 8821–8831. PMLR, 2021. 2
- [53] James Blake Rawlings, David Q Mayne, Moritz Diehl, et al. *Model predictive control: theory, computation, and design*. Nob Hill Publishing Madison, WI, 2020. 2
- [54] Tianhe Ren, Shilong Liu, Ailing Zeng, Jing Lin, Kunchang Li, He Cao, Jiayu Chen, Xinyu Huang, Yukang Chen, Feng Yan, Zhaoyang Zeng, Hao Zhang, Feng Li, Jie Yang, Hongyang Li, Qing Jiang, and Lei Zhang. Grounded sam: Assembling open-world models for diverse visual tasks, 2024. 3
- [55] Tianhe Ren, Shuo Shen, and the IDEA-Research team. Grounded-sam-2: Ground and track anything in videos, 2025. Accessed: 2025-11-12. 3
- [56] Moritz Reuss, Hongyi Zhou, Marcel R’uhle, ”Omer Erdinç Yağmurlu, Fabian Otto, and Rudolf Lioutikov. FLOWER: Democratizing generalist robot policies with efficient vision-language-flow models. In *9th Annual Conference on Robot Learning*, 2025. 5
- [57] Tencent Hunyuan3D Team. Hunyuan3d 2.1: From images to high-fidelity 3d assets with production-ready pbr material, 2025. 4
- [58] Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5026–5033, 2012. 3
- [59] Marcel Torne, Anthony Simeonov, Zechu Li, April Chan, Tao Chen, Abhishek Gupta, and Pulkit Agrawal. Reconciling reality through simulation: A real-to-sim-to-real approach for robust manipulation. In *Robotics: Science and Systems*, 2024. 2
- [60] Maggie Wang, Stephen Tian, Aiden Swann, Ola Shorinwa, Jiajun Wu, and Mac Schwager. Phys2real: Fusing vlm priors with interactive online adaptation for uncertainty-aware sim-to-real manipulation. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2026. Accepted to ICRA 2026. 2
- [61] Yanwei Wang, Tsun-Hsuan Wang, Jiayuan Mao, Michael Hagenow, and Julie Shah. Grounding language plans in demonstrations through counterfactual perturbations. In *The Twelfth International Conference on Learning Representations*, 2024. 2
- [62] Zhaowei Wang, Wenhao Yu, Xiyu Ren, Jipeng Zhang, Yu Zhao, Rohit Saxena, Liang Cheng, Ginny Wong, Simon See, Pasquale Minervini, Yangqiu Song, and Mark Steedman. Mmlongbench: Benchmarking long-context vision-language models effectively and thoroughly. In *The 39th (2025) Annual Conference on Neural Information Processing Systems*, 2025. 8
- [63] Bowen Wen, Wei Yang, Jan Kautz, and Stan Birchfield. Foundationpose: Unified 6d pose estimation and tracking of novel objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17868–17879, 2024. 4
- [64] Grady Williams, Nolan Wagener, Brian Goldfain, Paul Drews, James M. Rehg, Byron Boots, and Evangelos A. Theodorou. Information theoretic mpc for model-based reinforcement learning. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1714–1721, 2017. 2
- [65] Fei Xia, Amir R Zamir, Zhiyang He, Alexander Sax, Jitendra Malik, and Silvio Savarese. Gibson env: Real-world perception for embodied agents. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 9068–9079, 2018. 2
- [66] Hongchi Xia, Zhi-Hao Lin, Wei-Chiu Ma, and Shenlong Wang. Video2game: Real-time interactive realistic and browser-compatible environment from a single video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4578–4588, 2024. 4
- [67] William Xie, Maria Valentini, Jensen Lavering, and Nikolaus Correll. Deligrasp: Inferring object properties with LLMs for adaptive grasp policies. In *8th Annual Conference on Robot Learning*, 2024. 4
- [68] Sherry Yang, Yilun Du, Seyed Kamyar Seyed Ghasemipour, Jonathan Tompson, Leslie Pack Kaelbling, Dale Schuurmans, and Pieter Abbeel. Learning interactive real-world simulators. In *The Twelfth International Conference on Learning Representations*, 2024. 2

- [69] Wentao Yuan, Jiafei Duan, Valts Blukis, Wilbert Pumacay, Ranjay Krishna, Adithyavairavan Murali, Arsalan Mousavian, and Dieter Fox. Robopoint: A vision-language model for spatial affordance prediction in robotics. In *8th Annual Conference on Robot Learning*, 2024. [2](#)
- [70] Michał Zawalski, William Chen, Karl Pertsch, Oier Mees, Chelsea Finn, and Sergey Levine. Robotic control via embodied chain-of-thought reasoning. In *8th Annual Conference on Robot Learning*, 2024. [2](#)
- [71] Kaifeng Zhang, Baoyu Li, Kris Hauser, and Yunzhu Li. Adaptigraph: Material-adaptive graph-based neural dynamics for robotic manipulation. In *Proceedings of Robotics: Science and Systems (RSS)*, 2024. [2](#)
- [72] Kaifeng Zhang, Baoyu Li, Kris Hauser, and Yunzhu Li. Particle-grid neural dynamics for learning deformable object models from rgb-d videos. In *Proceedings of Robotics: Science and Systems (RSS)*, 2025. [2](#)
- [73] Brianna Zitkovich, Tianhe Yu, Sichun Xu, Peng Xu, Ted Xiao, Fei Xia, Jialin Wu, Paul Wohlhart, Stefan Welker, Ayzaan Wahid, et al. Rt-2: Vision-language-action models transfer web knowledge to robotic control. In *Conference on Robot Learning*, pages 2165–2183. PMLR, 2023. [2](#)