

# LUMOS: Language-Conditioned Imitation Learning with World Models

Iman Nematollahi<sup>1</sup>, Branton DeMoss<sup>2</sup>, Akshay L Chandra<sup>1</sup>, Nick Hawes<sup>2</sup>, Wolfram Burgard<sup>3</sup>, Ingmar Posner<sup>2</sup>

**Abstract**—We introduce LUMOS, a language-conditioned multi-task imitation learning framework for robotics. LUMOS learns skills by practicing them over many long-horizon rollouts in the latent space of a learned world model and transfers these skills zero-shot to a real robot. By learning on-policy in the latent space of the learned world model, our algorithm mitigates policy-induced distribution shift which most offline imitation learning methods suffer from. LUMOS learns from unstructured play data with fewer than 1% hindsight language annotations but is steerable with language commands at test time. We achieve this coherent long-horizon performance by combining latent planning with both image- and language-based hindsight goal relabeling during training, and by optimizing an intrinsic reward defined in the latent space of the world model over multiple time steps, effectively reducing covariate shift. In experiments on the difficult long-horizon CALVIN benchmark, LUMOS outperforms prior learning-based methods with comparable approaches on chained multi-task evaluations. To the best of our knowledge, we are the first to learn a language-conditioned continuous visuomotor control for a real-world robot within an offline world model. Videos, dataset and code are available at <http://lumos.cs.uni-freiburg.de>.

## I. INTRODUCTION

Building systems that can perform long-horizon tasks specified by natural language is a long-standing goal in robotics [1]. When complete task specifications are available, reinforcement learning (RL) methods can be used to produce policies by trial and error. However, typical RL algorithms can require millions of episodes to learn good policies, which is infeasible for many robotics tasks. Furthermore, many tasks are difficult to specify as reward functions, requiring substantial engineering effort to elicit reasonable learned skills from RL alone.

Imitation learning offers an alternative by learning directly from task demonstrations. Naïve approaches like behavior cloning [2] predict actions from states based on demonstration data and apply them during deployment. However, these methods overlook the sequential nature of decision-making. Since each observation depends on the previous action, the independence assumption in standard statistical machine learning is broken. Ross and Bagnell [3] showed that a behavior-cloned policy with error  $\epsilon$  has expected regret of  $\mathcal{O}(T^2\epsilon)$ , growing quadratically with decision horizon  $T$ . Intuitively, because the behavior-cloned policy is not trained under its own distribution, small prediction errors compound, pulling the policy out of distribution to unseen states, which is especially problematic for long-horizon tasks due to the quadratic regret growth.

<sup>1</sup>University of Freiburg. <sup>2</sup>University of Oxford. <sup>3</sup>University of Technology Nuremberg.



Fig. 1: LUMOS learns a general-purpose language-conditioned visuomotor policy within the latent space of a learned world model. By optimizing an intrinsic reward to match expert performance, it recovers from its own mistakes across multiple time steps and reduces covariate shift. As a result, LUMOS can handle complex, long-horizon robot tasks from abstract language instructions in real-world scenarios, without online learning or fine-tuning.

An approach to alleviating distribution shift in off-policy learning is to train on-policy in a simulation of the environment. However, simulators may be unavailable or suffer from the sim2real problem, where policies that perform well in simulation fail to generalize to the real world due to the “reality gap”—a mismatch between real and simulated dynamics [4].

World models offer a promising alternative to hand-crafted simulators [5], approximating dynamics by predicting states or observations conditioned on actions from real trajectories over multiple steps. By compressing observations into latent representations that capture temporal context, contemporary world models [6], [7] achieve high-fidelity predictions over thousands of time steps. Learning within the world model’s latent space enables both on-policy and long-horizon learning, mitigating distribution shift and avoiding the reality gap from mis-specified simulators.

Taking steps towards achieving language-conditioned, long-horizon, multi-task policies, we introduce LUMOS, a language-conditioned imitation learning framework. LUMOS employs world models to learn a multi-task policy that can be guided by free-form natural language instructions entirely offline and transferred to the real world in a zero-shot fashion. It leverages the latent-matching intrinsic reward proposed by DITTO [8] within the world model’s latent space, enabling its actor-critic agent to recover expert performance without any online interaction. To further enhance

the long-horizon performance of the actor-critic agent, we build upon recent advancements in imitation learning [9], [10], and augment LUMOS with a latent planning network. This network uses the benefit of hindsight during training to distill predicted information about the future trajectory into a “latent plan”. We demonstrate that LUMOS can learn imitation policies solely within the world model’s latent space, achieving zero-shot transfer to the real world. Extensive tests on the CALVIN benchmark and in real-world scenarios show that our algorithm excels in handling multi-objective, long-horizon tasks specified in natural language, outperforming previous learning-based methods.

Our contributions are as follows: 1) **Improved Long-Horizon Imitation Learning:** LUMOS reduces covariate shift and enhances long-horizon performance by combining world-model-based exploration with latent planning, enabling agents to practice skills over simulated long horizons using full demonstration trajectories. Ablation studies show that multi-step practice in latent space significantly outperforms single-step behavioral cloning. 2) **Language Goal-Conditioned Learning:** Our language-conditioned imitation learning framework outperforms prior methods with comparable approaches on the challenging CALVIN benchmark, testing long-horizon manipulation tasks guided by language instructions. 3) **Zero-Shot Real-World Transfer:** We achieve successful zero-shot transfer of policies learned entirely offline in the world model with language conditioning to the real environment.

## II. RELATED WORK

**Language-Conditioned Imitation Learning for Robotics:** Lynch *et al.* [11] introduce MCIL, which trains a behavior cloning agent using unstructured robot play data. They annotate 1% of trajectories with human language, making the policy steerable at test-time by natural language commands. Mees *et al.* [1] introduced the open-source CALVIN environment, a benchmark for evaluating language-conditioned, long-horizon multi-task robotic manipulation. It includes a simulator with  $\sim 24$  hours of teleoperated play data and 20K language task annotations. HULC [10] introduces a framework for hierarchical language-conditioned imitation learning using contrastive learning to ground language instructions to robotic actions, as well as other improvements like using a multimodal transformer encoder (discussed below). HULC outperforms MCIL on the CALVIN benchmark [12]. HULC++ [13] further enhances long-horizon performance by integrating affordance prediction and motion planning into the control loop, only deferring to the learned HULC controller when the end-effector is near the manipulation target. SPIL [12] improves upon the state-of-the-art by labeling all actions in the dataset and probabilistically assigning them to a base skill (translation, rotation, grasping) based on their magnitudes. In this work, we do not incorporate any prior knowledge about the semantic relationship between action dimensions and particular skills, as we aim to maintain a 7-DoF continuous action space. Diffusion Generative Models have gained traction for policy

representation in robotics, diffusing actions from Gaussian noise [14]–[16]. They can acquire diverse behaviors conditioned on language-goals [17]–[19]. The latest, Multimodal Diffusion Transformer (MDT) [20], uses Masked Generative Foresight and Contrastive Latent Alignment to boost performance in long-horizon tasks. Although this approach is orthogonal to our work, we acknowledge the potential of diffusion models in both world modeling and behavior learning. In this paper, however, we aim to investigate how to train policies in the latent space of a learned world model.

**Improving Behavior Cloning with Latent Planning:** Lynch *et al.* [9] introduce GCBC, which augments standard behavior cloning with image-based goal conditioning. GCBC further introduces a planning module which trains a plan-predictor conditioned only on the current and goal states to predict a latent variable produced by a planning module that has privileged access to the entire expert trajectory. At test time, latent variables sampled from the plan predictor are used to condition the policy to achieve improved long-horizon performance. One limitation of this work is the reliance on goal images for conditioning the plan predictor. Zhang *et al.* [21] improve upon the latent planner in GCBC by replacing the discrete latent plan embedding model with a diffusion model. They learn to iteratively denoise the latent plan via a language-conditioned U-Net, achieving superior long-horizon performance over GCBC.

**Learning with World Models:** World models have recently emerged as a promising approach to data-driven simulation. The seminal work of Ha and Schmidhuber [5] demonstrated the effectiveness of learning world models directly from pixel observations to predict future observations, and then learning a policy inside the latent space of the learned model. Building on this, the Dreamer family of works [22], [23] introduced a world-model-based RL algorithm that achieved SOTA performance on the Atari benchmark, and recently became the first algorithm to produce diamonds in Minecraft, which requires acting coherently across thousands of time-steps [7]. DayDreamer [24] applied the Dreamer algorithm to learning quadrupedal locomotion online on a real robot, demonstrating the robustness of policies trained in the world model to transfer to the real environment.

Our work builds primarily on DITTO [8], a world-model-based imitation learning algorithm that penalizes an on-policy divergence from expert trajectories in the world model latent space. Specifically, it defines an *intrinsic reward* in the latent space that measures the divergence between agent rollouts and expert demonstrations and then optimizes this intrinsic reward using actor-critic RL. Optimizing this intrinsic reward induces imitation learning that is robust to errors over long-horizon rollouts, effectively mitigating covariate shifts.

## III. PROBLEM FORMULATION

We investigate goal-conditioned imitation learning in a partially observable Markov decision process (POMDP) with continuous actions and high-dimensional observations from an unknown environment. We model the interaction between the environment and the goal-conditioned policy




Fig. 2: **LUMOS** learns a language-guided general-purpose policy within the latent space of a world model. (1) The world model, comprising an image encoder, a Recurrent State-Space Model (RSSM) for dynamics, and an image decoder, transforms play dataset experience into a predictive model that enables behavior learning in the latent state space. (2) The goal-conditioned policy learning and uses *either* a language annotation *or* the final latent state as the goal, with plan recognition and proposal networks being trained to identify and organize behaviors in a latent plan space. The action decoder is intrinsically rewarded by matching the expert’s latent trajectory. (3) During inference, the policy acts based on the latent state inferred by the world model from the current observation and is guided by a user’s language command.

using a goal-augmented POMDP  $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathcal{R}, \mathcal{T}, \mathcal{G}, \gamma)$ , where  $\mathcal{S}$  is the state-observation space,  $\mathcal{A}$  is the action space,  $\mathcal{R}(s, a)$  is the reward function,  $\mathcal{T}(s'|s, a)$  defines state-transition dynamics,  $\mathcal{G}$  is the goal space, and  $\gamma \in (0, 1)$  is the discount factor. The agent receives visual observations instead of direct state access and aims to maximize the expected discounted sum of extrinsic rewards  $\mathbb{E}[\sum_t \gamma^t r_t]$ , which it cannot directly observe. We conduct offline learning using a large, unlabeled, and undirected, fixed play dataset  $\mathcal{D} = \{(s_1, a_1), \dots, (s_T, a_T)\}$ . This dataset is relabeled into long temporal state-action streams [25]. Each visited state is treated as a “reached goal state”, resulting in a trajectory dataset  $\mathcal{D} = \{\tau_i = (s_t, a_t)_{t=0}^k\}_{i=1}^N$ . This transformation creates  $\mathcal{D}_{\text{play}} = \{(\tau, s_g)_i\}_{i=0}^{|\mathcal{D}_{\text{play}}|}$ , pairing each goal state  $s_g$  with the corresponding demonstration trajectory  $\tau$ . We adopt the multi-context imitation learning approach from play data by Lynch *et al.* [11] to utilize language annotations. This method shows that combining a few random windows with language-based retrospective instructions helps learn a unified, language-conditioned visuomotor policy for various robotic tasks. Free-form language instructions  $l \in \mathcal{L}$  guide the policy, providing flexible and natural task descriptions.

#### IV. LUMOS

In this section, we introduce LUMOS. Training consists of two phases: first, a world model is learned from the unlabeled play dataset  $\mathcal{D}$ . Next, an actor-critic agent is trained within this learned world model to acquire a goal-conditioned policy by guiding imagined sequences of latent model states to match the latent trajectory of the expert demonstrations. During inference, the language-conditioned policy  $\pi_\theta(a_t | s_t, l)$ , trained entirely in the latent space of the world model, successfully transfers to the real environment.

Fig. 2 shows an overview of the approach.

##### A. World Model Learning

World models capture an agent’s experience into a predictive model, enabling behavior learning without direct interaction with the environment. By converting high-dimensional images into compact state representations, they allow efficient, parallel long-term predictions in latent space. While our method is agnostic regarding the choice of the world-model architecture, we follow DITTO by adopting the DreamerV2 architecture [23] as the backbone for LUMOS, which includes an image encoder, a recurrent state-space model (RSSM) for learning transition dynamics, and a decoder for reconstructing observations from latent states. To learn robotic manipulation skills, we use observations from static and gripper-mounted cameras, with separate CNN encoders and transposed CNN decoders for each. These encodings are concatenated before being passed to the RSSM, which generates a sequence of deterministic recurrent states  $(h_t)_{t=0}^T$ . Each state defines two distributions over stochastic hidden states: the stochastic posterior state  $z_t$ , determined by the current observation  $x_t$  and recurrent state  $h_t$ , and the stochastic prior state  $\hat{z}_t$ , trained to approximate the posterior without the current observation. By predicting  $\hat{z}_t$ , the model learns the environment dynamics. The combined model state, comprising deterministic and stochastic components  $\hat{s}_t = (h_t, z_t)$ , reconstructs the observation. The RSSM includes the following modules:

$$\text{RSSM} \begin{cases} \text{Recurrent state:} & h_t = f_\phi(\hat{s}_{t-1}, a_{t-1}) \\ \text{Representation model:} & z_t \sim q_\phi(z_t | h_t, x_t) \\ \text{Dynamics predictor:} & \hat{z}_t \sim p_\phi(\hat{z}_t | h_t) \\ \text{Image decoder:} & \hat{x}_t \sim p_\phi(\hat{x}_t | \hat{s}_t) \end{cases} \quad (1)$$

The prior and posterior models predict categorical distributions, optimized using straight-through gradient estimation [26]. All modules are implemented as neural networks parameterized by  $\phi$ , and are optimized jointly by minimizing the negative variational lower bound [27]:

$$\min_{\phi} \mathbb{E}_{q_{\phi}(z_{1:T}|a_{1:T},x_{1:T})} \left[ \sum_{t=1}^T -\log p_{\phi}(x_t | \hat{s}_t) + \beta D_{\text{KL}}(q_{\phi}(z_t | \hat{s}_t) \| p_{\phi}(\hat{z}_t | h_t)) \right]. \quad (2)$$

After training, the world model can operate without input observations by using the prior  $\hat{z}$  instead of the posterior  $z$ . This allows the model to generate unlimited imagined trajectories of the form  $\{(h_t, \hat{z}_t, a_t)_{t=0}^H\}$ , where  $H$  is the time horizon for imagination.

### B. Behavior Learning

LUMOS employs an actor-critic agent to learn long-horizon, language-conditioned behaviors within the latent space of its world model. It learns a goal-conditioned policy  $\pi_{\theta}(a_t | s_t, g)$  and a value function  $v_{\psi}(s_t, g)$  from the dataset  $\mathcal{D}_{\text{play}}$ , where  $s_t$  denotes the current latent state and  $g \in \mathcal{G}$  represents a free-form language instruction  $l$  or a latent goal state  $s_g$ . In this section, we outline the design choices of our model, explain how to intrinsically reward the agent for matching expert latent state-action pairs over a trajectory, and detail the optimization process for the actor and critic.

**Observation and Action Spaces** Our goal-reaching policy builds on the HULC architecture by Mees *et al.* [10], enhancing language-conditioned imitation learning from unstructured data. Unlike HULC, which uses high-dimensional image observations, our method utilizes the latent state space of the world model for its compact representations. In addition to encoding RGB images from static and gripper cameras, LUMOS encodes compact states of latent trajectories from the current to the goal state. The action space is a 7-DoF continuous space, including relative XYZ positions, Euler angles, and a gripper action.

**Latent Plan Encoding** Learning control with free-form imitation data faces the challenge of multiple valid trajectories connecting the same  $(s_t, s_g)$  pairs. To address this, we use a sequence-to-sequence conditional variational auto-encoder (seq2seq CVAE) [9] to encode contextual latent trajectories into a latent “plan” space. This space is determined by two stochastic encoders: one for recognizing plans during training and the other for proposing plans during inference. The plan recognition encoder identifies the performed behavior from the entire sequence, while the plan proposal encoder generates possible behaviors from the initial and final states (see the second column of Fig. 2). Minimizing the KL divergence between these encoders, referred to as  $\mathcal{L}_{\text{KL}}$ , ensures the plan proposal encoder accurately reflects observed behaviors. Similar to HULC, we use a multimodal transformer encoder to develop a contextualized representation of latent trajectories, mapping them into a vector of multiple latent categorical variables, optimized with straight-through gradients.

**Semantic Alignment of Latent Trajectories and Language** Addressing the symbol grounding problem [28] requires relating language instructions to the robot’s perception and actions. Aligning instructions with visual observations is crucial for distinguishing similar objects like colored blocks. Unlike HULC, which aligns visual features with language features, we align the world model’s latent features with language features. This involves maximizing cosine similarity between latent and language features while minimizing similarity with unrelated instructions, using a contrastive loss similar to CLIP [29] and incorporating in-batch negatives for efficient training. We denote this objective as  $\mathcal{L}_{\text{contrast}}$ .

**Reward** Unlike behavior cloning approaches, which rely exclusively on supervised action labels for training, LUMOS leverages a simple intrinsic reward defined within the latent space of the world model to address the problem of covariate shift. To align the agent’s behavior with the expert, we employ a reward function proposed by DeMoss *et al.* [8] that encourages the agent to match the expert’s state-action pairs in the latent space. This intrinsic reward is defined as a modified inner product on the latent state representations, encouraging similarity without requiring exact matches. This lets the agent explore different trajectories in the world-model latent space, while aligning its behavior to the demonstrator over the entire training horizon. The reward at each step  $t$  is given by:

$$r_t^{\text{int}}(s_t^E, s_t^{\pi}) = \frac{s_t^E \cdot s_t^{\pi}}{\max(\|s_t^E\|, \|s_t^{\pi}\|)^2} \quad (3)$$

**Action and value models** Our stochastic actor, parameterized by  $\theta$ , samples actions based on the current latent state  $s_t$  and a latent goal  $g$ , aiming to maximize expected rewards (Eqn. 3). The deterministic critic, parameterized by  $\psi$ , predicts the expected discounted future rewards, providing value estimates for the same latent state and goal. The models are defined as:

$$\begin{aligned} \text{Actor: } & a_t \sim \pi_{\theta}(a_t | s_t, g), \\ \text{Critic: } & v_{\psi}(s_t, g) \approx \mathbb{E}_{\pi_{\theta}, p_{\phi}} \left[ \sum_{t=0}^H \gamma^t r_t^{\text{int}} \right]. \end{aligned} \quad (4)$$

Fully-connected neural networks are used for the actor and critic. The action model outputs a tanh-transformed Gaussian [30], facilitating reparameterized sampling for backpropagation. The world model is fixed during behavior learning, ensuring that the agent gradients do not alter its representations.

**Learning Objective** The objective in training the critic is to regress the computed value estimate to the  $\lambda$ -target [31], denoted  $V^{\lambda}$  (refer to [23] for the definition):

$$\mathcal{L}(\psi) \doteq \mathbb{E}_{\pi_{\theta}, p_{\phi}} \left[ \sum_{t=1}^{H-1} \frac{1}{2} (v_{\psi}(\hat{s}_t) - sg(V_t^{\lambda}))^2 \right] \quad (5)$$

where  $sg(\cdot)$  is the stop-gradient operator. Moreover, a target network, which updates its weights slowly, is used to provide stable value bootstrap targets [32]. The actor’s objective function uses backpropagation through the learned dynamics

to maximize value estimates while minimizing the KL divergence between plan encoders and the contrastive loss for aligning latent trajectories with language. For sampled windows without language annotations, the contrastive loss is omitted:

$$\mathcal{L}(\theta) \doteq \mathbb{E}_{\pi_{\theta}, p_{\phi}} \left[ \sum_{\tau=t}^{t+H} (-V_{\lambda}(s_{\tau}) + \alpha_1 \mathcal{L}_{KL} + \alpha_2 \mathcal{L}_{\text{contrast}}) \right] \quad (6)$$

## V. EXPERIMENTAL RESULTS

We evaluate LUMOS in learning language-conditioned imitation learning in both simulated and real-world settings. The objectives of our experiments are threefold: (i) to determine the model’s performance in executing long-horizon language-specified tasks, (ii) to ablate components of the objective function and analyze their impact, and (iii) to assess its scalability to real-world robotics tasks.

### A. Experiments in Simulation

We conduct our simulation experiments in environment  $D$  of the CALVIN Challenge [1], which offers six hours of unstructured tele-operated play data to train a 7-DoF Franka Emika Panda robot arm. CALVIN features 34 distinct sub-tasks and evaluates 1000 unique instruction chain sequences. The agent’s objective is to sequentially solve up to five language instructions using only onboard sensors, requiring it to transition between subgoals based solely on language cues. Notably, only 1% of this dataset is annotated with language instructions through crowd-sourcing.

**Evaluation Protocol** We compare our model with the following language-conditioned robotic imitation learning approaches over unstructured data: **GCBC** [9]: Goal-conditioned behavior cloning does not model latent plans but instead maximizes the likelihood of each action as reconstruction loss at each time step. **MCIL** [11]: Multi-context imitation learning employs a sequential CVAE to predict the next actions from image or language-based goals by modeling reusable latent sequences of states as latent plans. **HULC** [10]: Hierarchical universal language conditioned policy provides a strong baseline by building upon MCIL, introducing several improvements, such as adding a self-supervised contrastive loss for alignment of video and language representation. Our approach differs from the above models by not regressing action labels. Instead, we learn a world model from the data, before training an agent to match the expert’s latent trajectory in the world model. To reduce the computational burden of LUMOS’ world model, we standardize the input resolution of all images to  $64 \times 64$ . For fair

comparison, all baselines are trained at this resolution, which sometimes lowers the original implementations’ resolutions (e.g.,  $200 \times 200$  for static and  $84 \times 84$  for the gripper camera).

Table I contains quantitative results of the average success rate of each model over 1,000 instruction chains across three seeded runs. LUMOS successfully performs multi-stage, long-horizon, language-conditioned robotic manipulation tasks, and outperforms all the baselines. This experiment shows that our agent can optimize its intrinsic reward, which measures its deviation from expert latent trajectories. This capability allows the agent to recover from its mistakes over multiple time steps and successfully chain more tasks together in sequence.

In our first ablation, we examine the impact of our latent matching reward function. Specifically, we supervise our agent using behavior cloning with mean squared error instead of optimizing for intrinsic rewards. This ablation demonstrates a significant drop in our model’s performance under these conditions and indicates that our reward formulation effectively reduces covariate shift. For the second ablation study, we explore the effect of learning a latent plan space on long-horizon behavior. For this test, we removed the plan proposal and recognition components from the actor’s architecture. The results show that, although the agent starts the trials relatively well, it struggles to complete longer-horizon tasks. Lastly, in our final ablation, we omit the alignment objective between language and latent state features. While this variant maintains some long-horizon performance, it performs slightly worse at each step compared to our full model. This is evident in its occasional manipulation of blocks of the wrong color due to a suboptimal association between instructions and the scene.

### B. Experiments in Real-World

To assess the performance of LUMOS at handling various real-world tasks, we conduct experiments using a Franka Emika Panda robot arm within a 3D tabletop environment. This environment includes a table with a stove, a bowl, a cabinet, a carrot, and an eggplant. We collected a play dataset by recording three hours of teleoperated interactions with the robot arm, using a VR controller to guide its movements. During recording, we captured RGB observations from both static and gripper cameras. We used language annotation for less than 1% of the data ( $\sim 2800$  random windows). Annotators described behaviors in randomly sampled portions of the interaction data. This dataset encompasses over 22 unique tasks for manipulating the available objects.

Method	No. Instructions in a Row (1000 chains)					
	1	2	3	4	5	Avg. Len.
GCBC [9]	56.7% (3.4)	20.9% (1.9)	7.3% (2.9)	1.6% (0.5)	0.04% (0.0)	0.63 (0.4)
MCIL [11]	70.3% (3.1)	40.2% (2.9)	21.6% (4.1)	11.1% (1.9)	5.4% (1.0)	1.48 (0.3)
HULC [10]	77.6% (1.1)	58.05% (0.3)	41.6% (1.3)	29.8% (0.9)	20.0% (1.5)	2.27 (0.05)
LUMOS (ours)	<b>80.7%</b> (1.0)	<b>59.3%</b> (1.4)	<b>42.6%</b> (1.8)	<b>30.7%</b> (1.2)	<b>21.1%</b> (0.8)	<b>2.34</b> (0.05)
No DITTO	71.8% (1.3)	45.6% (1.5)	27.8% (1.9)	14.2% (1.8)	8.7% (0.7)	1.68 (0.02)
No latent plan	76.5% (2.5)	50.5% (0.9)	28.7% (1.3)	15.9% (0.8)	11.0% (0.7)	1.81 (0.01)
No alignment	73.3% (1.5)	52.6% (0.9)	39.8% (2.0)	27.3% (1.3)	17.2% (0.9)	2.05 (0.06)

TABLE I: Performance of all models on environment  $D$  of the CALVIN Challenge with an ablation study of key components evaluated over three seeds. All models receive  $64 \times 64$  resolution RGB images from both a static and gripper camera. Standard deviation in parentheses.




Fig. 3: **World Model Long-Horizon Predictions.** Using the first five images of a hold-out trajectory as context, the world model predicts the next 195 steps using its latent dynamics, given only the actions. Trained on sequences of horizon 50, our model makes precise long-term predictions, aiding efficient behavior learning in a compact latent space. Only gripper camera reconstructions are visualized.

First, we used the offline play dataset to learn a generative world model. As shown in Fig. 3 our model demonstrates competence in long-horizon predictions. To assess its long-term prediction abilities, we applied the representation model to the first five images of hold-out trajectories to establish context. Given an action trajectory, we then tasked the model with predicting forward for 195 steps using its latent dynamics, despite being trained only with a horizon length of 50. The accuracy of the reconstructed images decoded from the latent trajectory of our recurrent state-space model indicates that our world model is well-suited for learning behaviors within its compact latent space.

We trained our language-conditioned policy within the latent space of the world model and subsequently assessed each task through 20 rollouts, comparing its performance to HULC [10] under the same conditions. We utilized various neutral starting positions to avoid biasing the policy toward any specific task. This neutral setup eliminates the correlation between the robot’s initial state and the queried task and ensures that the agent depends entirely on language cues to comprehend and execute the tasks. The success rates are reported in Table II. For brevity, we combined the success rates for tasks involving either the carrot or eggplant under the category “vegetable”. However, each task was evaluated individually (i.e., if the robot picks up the carrot when the instruction was to pick up the eggplant, it does not count as a success). Moreover, we assessed our model and HULC on 20 unique chains of 5 instructions in a row to evaluate their multi-stage long-horizon capabilities. We reported the average number of successfully completed sequential tasks, as shown in Table II. Our findings indicate that LUMOS outperforms HULC in both individual and multi-stage long-horizon tasks, demonstrating superior performance in recov-




Fig. 4: **Real-world Tasks.** Examples from left to right are: placing the carrot onto the pan, lifting the carrot from the table, dropping the carrot into the bowl, storing the carrot in the cabinet, placing the pan into the cabinet, setting the pan onto the stove, adding the eggplant to the pan, picking up the pan from the table, taking the eggplant from the cabinet, and putting the eggplant into the bowl.

Task / Method	LUMOS (ours)	HULC [10]	World Model
Lift the vegetable from the table	67.5%	62.5%	72.5%
Lift the vegetable from the pan	75%	72.5%	77.5%
Lift the vegetable from the bowl	72.5%	67.5%	75%
Lift the vegetable from the cabinet	55%	50%	70%
Lift the pan from the table	65%	75%	70%
Lift the pan from the stove	80%	70%	75%
Lift the pan from the cabinet	35%	35%	60%
Place the vegetable onto the table	90%	85%	92.5%
Place the vegetable onto the pan	62.5%	65%	77.5%
Place the vegetable into the bowl	85%	75%	80%
Place the vegetable into the cabinet	70%	60%	77.5%
Place the pan onto the table	70%	65%	80%
Place the pan onto the stove	65%	55%	85%
Place the pan into the cabinet	55%	50%	70%
Average over tasks	<b>67.68%</b>	63.39%	75.89%
Average no. of sequential tasks	<b>2.05</b>	1.90	-

TABLE II: The average success rate of language-conditioned general-purpose policies in the real world setup.

ering from its own mistakes across multiple time steps and reducing covariate shift.

Finally, we analyzed LUMOS’s predicted action trajectories by simulating their outcomes in the learned world model. The performance gap between real and simulated environments stems from inaccuracies in the world model, particularly for tasks with lower real-world success rates. We believe that training the world model with more comprehensive data could improve its robustness and reduce this gap.

## VI. CONCLUSION AND LIMITATIONS

In this paper, we propose *Language-Conditioned Imitation Learning via World Models* (LUMOS), which leverages a learned world model’s latent space for fully offline, multi-step on-policy learning. Practicing skills in long-horizon model-based rollouts allows our model to recover expert performance without any online interaction. Our approach demonstrates superior performance on the CALVIN benchmark, achieving robust long-horizon manipulation guided solely by language instructions. LUMOS serves as a proof of concept, demonstrating that dynamics and policies learned offline within a world model can successfully transfer zero-shot to real environments.

Although LUMOS performs well, it has several aspects that warrant future research. One difficulty with world-model-based approaches is that it is hard to know a priori whether the world model’s representations and dynamics are of sufficient quality to produce strong policies. Additionally, it has been observed that increasing image input resolution substantially improves the final performance. However, training world models at higher resolution significantly increases the compute resources required to run experiments.

## REFERENCES

- [1] O. Mees, L. Hermann, E. Rosete-Beas, and W. Burgard, "CALVIN: A Benchmark for Language-Conditioned Policy Learning for Long-Horizon Robot Manipulation Tasks," *IEEE Robotics and Automation Letters (RA-L)*, 2022.
- [2] D. A. Pomerleau, "ALVINN: An Autonomous Land Vehicle in a Neural Network," in *NeurIPS*, 1989.
- [3] S. Ross and J. A. Bagnell, "Efficient Reductions for Imitation Learning," in *AISTATS*, 2010.
- [4] J. Hwangbo, J. Lee, A. Dosovitskiy, D. Bellicoso, V. Tsounis, V. Koltun, and M. Hutter, "Learning Agile and Dynamic Motor Skills for Legged Robots," *Science Robotics*, 2019.
- [5] D. Ha and J. Schmidhuber, "World Models," *arXiv preprint arXiv:1803.10122*, 2018.
- [6] J. Schrittwieser, I. Antonoglou, T. Hubert, K. Simonyan, L. Sifre, S. Schmitt, A. Guez, E. Lockhart, D. Hassabis, T. Graepel, et al., "Mastering atari, go, chess and shogi by planning with a learned model," *Nature*, 2020.
- [7] D. Hafner, J. Pasukonis, J. Ba, and T. Lillicrap, "Mastering diverse domains through world models," *arXiv preprint arXiv:2301.04104*, 2023.
- [8] B. DeMoss, P. Duckworth, N. Hawes, and I. Posner, "Ditto: Offline imitation learning with world models," *arXiv preprint arXiv:2302.03086*, 2023.
- [9] C. Lynch, M. Khansari, T. Xiao, V. Kumar, J. Tompson, S. Levine, and P. Sermanet, "Learning latent plans from play," in *CoRL*, 2019.
- [10] O. Mees, L. Hermann, and W. Burgard, "What matters in language conditioned robotic imitation learning over unstructured data," *IEEE Robotics and Automation Letters*, 2022.
- [11] C. Lynch and P. Sermanet, "Language Conditioned Imitation Learning Over Unstructured Data," in *RSS*, 2021.
- [12] H. Zhou, Z. Bing, X. Yao, X. Su, C. Yang, K. Huang, and A. Knoll, "Language-conditioned imitation learning with base skill priors under unstructured data," *arXiv preprint arXiv:2305.19075*, 2023.
- [13] O. Mees, J. Borja-Diaz, and W. Burgard, "Grounding Language with Visual Affordances over Unstructured Data," in *ICRA*, London, UK, 2023.
- [14] C. Chi, S. Feng, Y. Du, Z. Xu, E. Cousineau, B. Burchfiel, and S. Song, "Diffusion Policy: Visuomotor Policy Learning via Action Diffusion," in *Proceedings of Robotics: Science and Systems (RSS)*, 2023.
- [15] M. Reuss, M. Li, X. Jia, and R. Lioutikov, "Goal-conditioned imitation learning using score-based diffusion policies," *arXiv preprint arXiv:2304.02532*, 2023.
- [16] Z. Xian, N. Gkanatsios, T. Gervet, T.-W. Ke, and K. Fragkiadaki, "Chaineddiffuser: Unifying trajectory diffusion and keypose prediction for robotic manipulation," in *Conference on Robot Learning*, 2023.
- [17] T.-W. Ke, N. Gkanatsios, and K. Fragkiadaki, "3d diffuser actor: Policy diffusion with 3d scene representations," *arXiv preprint arXiv:2402.10885*, 2024.
- [18] K. Black, M. Nakamoto, P. Atreya, H. Walke, C. Finn, A. Kumar, and S. Levine, "Zero-shot robotic manipulation with pretrained image-editing diffusion models," *arXiv preprint arXiv:2310.10639*, 2023.
- [19] H. Ha, P. Florence, and S. Song, "Scaling up and distilling down: Language-guided robot skill acquisition," in *Conference on Robot Learning*, 2023.
- [20] M. Reuss, Ö. E. Yağmurlu, F. Wenzel, and R. Lioutikov, "Multimodal Diffusion Transformer: Learning Versatile Behavior from Multimodal Goals," in *First Workshop on Vision-Language Models for Navigation and Manipulation at ICRA 2024*, 2024.
- [21] E. Zhang, Y. Lu, S. Huang, W. Y. Wang, and A. Zhang, "Language Control Diffusion: Efficiently Scaling through Space, Time, and Tasks," in *The Twelfth International Conference on Learning Representations*, 2024.
- [22] D. Hafner, T. Lillicrap, J. Ba, and M. Norouzi, "Dream to Control: Learning Behaviors by Latent Imagination," in *International Conference on Learning Representations*, 2019.
- [23] D. Hafner, T. Lillicrap, M. Norouzi, and J. Ba, "Mastering Atari with Discrete World Models," *arXiv preprint arXiv:2010.02193*, 2020.
- [24] P. Wu, A. Escontrela, D. Hafner, P. Abbeel, and K. Goldberg, "Daydreamer: World models for physical robot learning," in *Conference on Robot Learning*. PMLR, 2023.
- [25] M. Andrychowicz, F. Wolski, A. Ray, J. Schneider, R. Fong, P. Welinder, B. McGrew, J. Tobin, P. Abbeel, and W. Zaremba, "Hindsight experience replay," in *NeurIPS*, 2017.
- [26] Y. Bengio, N. Léonard, and A. Courville, "Estimating or Propagating Gradients Through Stochastic Neurons for Conditional Computation," *arXiv preprint arXiv:1308.3432*, 2013.
- [27] D. P. Kingma and M. Welling, "Auto-Encoding Variational Bayes," *arXiv preprint arXiv:1312.6114*, 2013.
- [28] S. Harnad, "The Symbol Grounding Problem," *Physica D: Nonlinear Phenomena*, 1990.
- [29] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, et al., "Learning transferable visual models from natural language supervision," *arXiv preprint arXiv:2103.00020*, 2021.
- [30] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, "Soft Actor-Critic: Off-Policy Maximum Entropy Deep Reinforcement Learning with a Stochastic Actor," *arXiv preprint arXiv:1801.01290*, 2018.
- [31] R. S. Sutton and A. G. Barto, "Reinforcement Learning: An Introduction," *IEEE Transactions on Neural Networks*, 2005.
- [32] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, et al., "Human-level control through deep reinforcement learning," *nature*, 2015.
- [33] N. Reimers and I. Gurevych, "Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks," in *EMNLP*, 2019.
- [34] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *NeurIPS*, 2017.
- [35] T. Salimans, A. Karpathy, X. Chen, and D. P. Kingma, "Pixelcnn++: Improving the pixelcnn with discretized logistic mixture likelihood and other modifications," *arXiv preprint arXiv:1701.05517*, 2017.
- [36] R. S. Sutton, "Learning to predict by the methods of temporal differences," *Machine learning*, 1988.
- [37] D. Yarats, R. Fergus, A. Lazaric, and L. Pinto, "Mastering Visual Continuous Control: Improved Data-Augmented Reinforcement Learning," *arXiv preprint arXiv:2107.09645*, 2021.

## VII. APPENDIX

### A. Hyperparameters and Training Details of the World Model

To develop the world model of LUMOS, we used DreamerV2 [23] as our foundation. While DreamerV2 models Atari game environments, our world model is specifically tailored for robotic manipulation. As a result, LUMOS incorporates distinct encoders and decoders for the static and gripper cameras. We concatenate their encodings and employ a fully-connected layer to fuse them before passing the features to the RSSM. This results in our model comprising approximately 40 million parameters. As described in (Eqn. 2), all components of the world model are trained jointly using a modified ELBO objective. Following [23], we incorporate KL balancing in this training process, which is used to control the regularization of the prior and posterior relative to each other using a parameter  $\delta$ :

$$D_{\text{KL}}(q \parallel p) = \underbrace{\delta D_{\text{KL}}(q \parallel sg(p))}_{\text{posterior regularizer}} + (1 - \delta) \underbrace{D_{\text{KL}}(sg(q) \parallel p)}_{\text{prior regularizer}} \quad (7)$$

where  $sg(\cdot)$  is the stop-gradient operator. KL balancing is necessary because the prior and posterior should not be regularized at the same rate; the prior must update more quickly towards the posterior, which encodes more information. This approach ensures the prior better approximates the aggregate posterior, enhancing the model’s performance.

To achieve uniform coverage of various segments of play data while modeling the world, we randomly sample the start index of each training sequence within the episode and clip it to prevent exceeding the episode length minus the training sequence length.

Since our approach focuses on learning from offline, unstructured tele-operated play data, it includes very long-horizon episodes and consequently very few reset states. To address this, we randomly reset the hidden state of the RSSM during training with a probability  $\zeta$ . This ensures that the world model learns to attend to both the current observation and the trajectory’s history while learning from long-horizon streams of data.

In all simulated and real-world experiments, we use similar hyperparameters to learn the world model, as outlined in Table III.

The latent state consists of deterministic and stochastic components,  $\hat{s}_t = (h_t, z_t)$ . Flattening the stochastic sample from 32 categorical distributions, each with 32 classes, results in a sparse binary vector of length 1024. Consequently, the latent size for each time step becomes  $k = 2048$ .

### B. Hyperparameters and Training Details of the Agent

Since our goal-reaching policy is inspired by the HULC architecture [10], we maintain similar design choices wherever possible for a fair comparison. Any deviations from the original design are specified in this section. Notably, HULC operates on raw observations, whereas our model first infers the latent state trajectories by the learned world model and then takes latent trajectories as input.

**Perceptual Encoder** While HULC encodes the sequence of visual observations from both static and gripper cameras  $X_{\{static, gripper\}} \in \mathbb{R}^{T \times H \times W \times 3}$  using separate convolutional encoders, LUMOS first leverages its learned world model to infer the latent state trajectories from the camera observations and then encodes the sequence of latent states  $S \in \mathbb{R}^{T \times k}$  using a fully connected encoder with two layers, featuring a hidden size of 1024 and an output size of 128 units with ReLU activations. Both methods map to the same output space, forming a perceptual representation

Name	Symbol	Value
Batch size	$B$	50
Sequence length	$L$	50
Deterministic latent state dimensions	—	1024
Discrete latent state dimensions	—	32
Discrete latent state classes	—	32
Latent dimensions	$k$	2048
KL loss scale	$\beta$	0.3
KL balancing	$\delta$	0.8
RSSM reset probability	$\zeta$	0.01
World model learning rate	—	$3 \cdot 10^{-4}$
Gradient clipping	—	100
Adam epsilon	$\epsilon$	$10^{-5}$
Weight decay (decoupled)	—	$5 \cdot 10^{-2}$

TABLE III: hyperparameters for learning the world model.

$V \in \mathbb{R}^{T \times d}$ , where  $T$  represents the sequence length and  $d$  the feature dimension.

**Goal Encoder** The goal-conditioned policy samples latent trajectories and uses *either* a language annotation *or* the final latent state as the goal. For each goal type, we employ a goal encoder similar to HULC. Notably, to transform raw text into a semantically meaningful vector space, we use the paraphrase-MiniLM-L3-v2 model [33]. This model is trained on paraphrase datasets primarily sourced from Wikipedia. It supports a vocabulary of 30,522 words and converts any length of sentence into a 384-dimensional vector. Both goal encoders map their respective inputs to a latent goal representation in  $\mathbb{R}^g$ , where  $g$  is the latent goal dimension.

**Plan Recognition and Proposal** While MCIL [11] proposed using bidirectional recurrent neural networks (RNNs) as the plan recognition network to encode randomly sampled play sequences into a latent Gaussian distribution, HULC demonstrated improved performance with a multimodal transformer encoder [34], incorporating positional embeddings to capture temporal information, and representing latent plans as vectors of multiple categorical variables [23]. We adopt HULC’s method by passing the encoded perceptual representation  $V$  to the transformer, leveraging its scaled dot-product attention mechanism, which allows elements within a sequence to attend to each other. This transformer encoder consists of two blocks, eight self-attention heads, and a hidden size of 2048 units. The plan proposal network is implemented using four fully connected layers, each with 2048 units and ReLU activations. To minimize the KL divergence between these encoders, referred to as  $\mathcal{L}_{\text{KL}}$ , we employ the KL balancing technique (Eqn. 7) again, ensuring that the KL loss decreases faster with respect to the prior than the posterior.

**Alignment of Latent Trajectories and Language** LUMOS maximizes the cosine similarity between the latent trajectory features of sequence  $i$  and its corresponding language features while minimizing the similarity with unrelated language instructions within the same batch. This is achieved using a contrastive loss similar to CLIP [29]. First, the latent trajectory features ( $x_i$ ) and language features ( $y_j$ ) are normalized, and the cosine similarity between each pair is computed. The logits for each batch form an  $M \times M$  matrix, where each entry is defined as:

$$\text{logit}(x_i, y_j) = \cos.\text{sim}(x_i, y_j) \cdot \exp(\tau), \quad \forall (i, j) \in \{1, 2, \dots, M\},$$

with  $\tau$  being a trainable temperature parameter. Only the diagonal entries of this matrix, which represent true pairs, are considered positive examples. The contrastive loss is computed as follows:

$$\mathcal{L}_{\text{contrast}} = \frac{1}{2} \left( \sum_{i=1}^M \text{CrossEntropy}(\text{logits\_per\_latent}(i), i) + \sum_{i=1}^M \text{CrossEntropy}(\text{logits\_per\_text}(i), i) \right) \quad (8)$$

where  $\text{logits\_per\_latent}$  represents the matrix of logits with rows corresponding to latent features and columns to language features, and  $\text{logits\_per\_text}$  is the transpose of this matrix.

**Action Decoder** The actor decodes the actions to be executed at each time step, conditioned on the current and goal latent states as well as

Name	Symbol	Value
Batch size	$B$	512
Sequence length	$T$	32
Minimum window length	—	20
Maximum window length	—	32
Feature dimensions	$d$	128
Latent goal dimensions	$g$	32
Discrete latent plan dimensions	—	32
Discrete latent plan classes	—	32
KL loss scale	$\alpha_1$	0.1
Contrastive loss scale	$\alpha_2$	3.0
KL balancing	$\delta$	0.8
Discount factor	$\gamma$	0.995
$TD(\lambda)$ parameter	$\lambda$	0.95
Optimizer	—	Adam
Actor learning rate	—	$2 \cdot 10^{-4}$
Critic learning rate	—	$3 \cdot 10^{-4}$
Slow critic update interval	—	100

TABLE IV: hyperparameters for training the actor and critic of LUMOS.




Fig. 5: Visualization of the complete real-world robot setup. This figure highlights the static and gripper cameras, as well as the VR controller and tracking system.

the sampled latent plan. These three inputs are concatenated and passed as input to the action decoder. While HULC parameterizes the action decoder as a recurrent neural network (RNN) that outputs parameters of a discretized logistic mixture distribution [35], our action decoder is significantly different: it is not recurrent and it outputs parameters of a tanh-transformed Gaussian distribution [30]. HULC’s action decoder consists of 2 RNN layers, each with a hidden dimension of 2048, utilizing 10 distributions to predict 10 classes per dimension. This design results in their action decoder having approximately 15 million parameters. In contrast, our action decoder is composed of eight fully connected layers, each with a hidden size of 256 units. This architecture results in a more compact model with approximately 1 million parameters. Unlike HULC, which needs a recurrent network to capture temporal patterns due to its reliance on action labels at each time step, our action decoder leverages the learned world model that encapsulates the temporal structure of latent state-action trajectories. This allows LUMOS to maximize intrinsic rewards and recover from mistakes across multiple time steps by matching the expert latent trajectory, thereby eliminating the need for a recurrent architecture.

**Critic** The critic of LUMOS consists of eight fully connected layers, each with a hidden size of 1024 units. It predicts the discounted sum of future rewards (state value) achieved by the actor, given the current and goal latent states. As detailed in Section 4, we use temporal-difference (TD) learning [36], training the critic towards a value target constructed from intermediate rewards and future state predictions. While 1-step targets are common, our model leverages the ability to generate multi-step on-policy trajectories, allowing us to use  $n$ -step targets for faster reward incorporation. This is achieved using the  $\lambda$ -target [31]:

$$V_t^\lambda = r_t + \gamma \left( (1 - \lambda)v_\psi(\hat{s}_{t+1}) + \lambda V_{t+1}^\lambda \right), \quad V_{t+H}^\lambda = v_\psi(\hat{s}_{t+H}).$$

The hyperparameter  $\lambda$  determines the horizon for TD learning. When  $\lambda = 0$ , it results in 1-step TD learning. A value of  $\lambda = 1$  corresponds to unbiased Monte Carlo returns. Intermediate values of  $\lambda$  provide an exponentially weighted sum of  $n$ -step returns. We set  $\lambda = 0.95$  to prioritize long-horizon targets over short-horizon targets.

As reported in Table IV, during training, we randomly sample windows with lengths between 20 and 32 and pad them to a maximum length of 32 to set the training horizon for the agent.

### C. Modifications to Baselines for Fair Comparisons

To efficiently handle the computational demands associated with increasing the resolution of LUMOS’ world model, we standardize the input resolution of all images to  $64 \times 64$ . Consequently, all baselines are trained at this same resolution for a fair comparison. This reduces the resolution from the original implementations (e.g.,  $200 \times 200$  for the static camera and  $84 \times 84$  for the gripper camera). To adapt the vision networks of our baselines to  $64 \times 64$  images, we made the following changes to their convolutional layers:

- The first convolutional layer was changed from a  $8 \times 8$  kernel with stride 4 to a  $4 \times 4$  kernel with stride 2.
- The second convolutional layer was adjusted from a  $4 \times 4$  kernel with stride 2 to a  $2 \times 2$  kernel with stride 2.
- The third convolutional layer was modified from a  $3 \times 3$  kernel with stride 1 to a  $2 \times 2$  kernel with stride 2.



Fig. 6: Visualization of the real-world environment observations across two time steps. At each time step, the RGB image from the static camera is presented on the left, and the RGB image from the gripper camera is shown on the right.

While our model does not use any image augmentation techniques, our baselines require data augmentation of image observations to facilitate learning. Specifically, during training, we apply stochastic shifts of 0-3 pixels to both the gripper and static camera images, modifying the original shifts of 0-4 pixels and 0-10 pixels, respectively, as in Yarats *et al.* [37]. Additionally, we apply bilinear interpolation to the shifted images, replacing each pixel with the average of the nearest pixels. This augmentation technique is used consistently across all baselines.

### D. 7-DoF Action Framework

In all our experiments, including simulated and real-world scenarios, we employ a 7-dimensional action framework:

$$[\delta x, \delta y, \delta z, \delta \phi, \delta \theta, \delta \psi, gripperAction]$$

The first three dimensions,  $\delta x$ ,  $\delta y$ , and  $\delta z$ , represent translations along the  $x$ ,  $y$ , and  $z$  axes, respectively, to control the end-effector’s position in 3D space. The following three dimensions,  $\delta \phi$ ,  $\delta \theta$ , and  $\delta \psi$ , are Euler angles which specify changes in the end-effector’s orientation relative to the robot’s base frame. All six of these parameters accept continuous values within the range  $[-1, 1]$ . The final dimension, *gripperAction*, controls the gripper state. Unlike the others, it accepts discrete values:  $-1.0$  for opening and  $1.0$  for closing. LUMOS treats *gripperAction* as a continuous variable between  $-1.0$  and  $1.0$ . To ensure compatibility with the environment, we apply thresholding before sending this value to the robot. Values greater than 0 trigger an open action, while values less than or equal to 0 cause the gripper to close.

### E. Data Collection Details

For our real-world experiment, we collected three hours of play data by teleoperating a Franka Emika Panda robot with a HTC VIVE Pro headset in a 3D tabletop environment (Figure 5). The environment contained various everyday objects for interaction, including a stove, a bowl, a cabinet, and toy vegetables like a carrot and an eggplant. During data collection, we recorded a comprehensive set of sensor measurements from the robot. These measurements included proprioceptive data, which provide information about the robot’s joint positions and the pose of its end-effector. Additionally, we captured static RGB images ( $200 \times 200$  resolution) using an Azure Kinect camera to capture the surrounding environment. We also acquired high-resolution RGB images ( $200 \times 200$  resolution) from a FRAMOS Industrial Depth D435e mounted on the robot’s gripper, providing detailed information about the objects being manipulated (Figure 6). Lastly, we logged the commanded absolute actions sent to the robot for control purposes.

For training, we derived relative actions by computing the differences in absolute actions between successive time steps. Additionally, we resized the observations to a resolution of  $64 \times 64$ . Given the minimal movement between frames, the original recording frequency of 30 Hz was downsampled to 15 Hz.