

Self-Supervised Multisensory Pretraining for Contact-Rich Robot Reinforcement Learning

Rickmer Krohn^{1,2,3}, Vignesh Prasad^{1,2,3}, Gabriele Tiboni¹ and Georgia Chalvatzaki^{1,2,3}

Abstract—Effective contact-rich manipulation requires robots to synergistically leverage vision, force, and proprioception. However, Reinforcement Learning agents struggle to learn in such multisensory settings, especially amidst sensory noise and dynamic changes. We propose MultiSensory Dynamic Pretraining (MSDP), a novel framework for learning expressive multisensory representations tailored for task-oriented policy learning. MSDP is based on masked autoencoding and trains a transformer-based encoder by reconstructing multisensory observations from only a subset of sensor embeddings, leading to cross-modal prediction and sensor fusion. For downstream policy learning, we introduce a novel asymmetric architecture, where a cross-attention mechanism allows the critic to extract dynamic, task-specific features from the frozen embeddings, while the actor receives a stable pooled representation to guide its actions. Our method demonstrates accelerated learning and robust performance under diverse perturbations, including sensor noise, and changes in object dynamics. Evaluations in multiple challenging, contact-rich robot manipulation tasks in simulation and the real world showcase the effectiveness of MSDP. Our approach exhibits strong robustness to perturbations and achieves high success rates on the real robot with as few as 6,000 online interactions, offering a simple yet powerful solution for complex multisensory robotic control. Website: <https://msdp-pearl.github.io/>

Index Terms—Reinforcement Learning; Representation Learning; Sensorimotor Learning

I. INTRODUCTION

REINFORCEMENT Learning (RL) has shown impressive successes in learning complex tasks ranging from Atari [1], locomotion [2], vision-based manipulation [3] to multisensory peg insertion [4], [5]. However, incorporating multiple sensor modalities, especially in complex contact-rich robotic manipulation tasks, remains a challenge for RL, due to the heterogeneous dynamics of different sensor modalities. Additionally, the importance of each input modality changes during the execution of a manipulation task, e.g., coarse scene understanding from visual input to fine-grained force feedback when in contact. Thus, robotic agents need to learn how to dynamically focus on the most relevant sensory information while adapting to perturbations and dynamic changes in the




Fig. 1: Multisensory Dynamic Pretraining fuses multiple sensors, like human senses, to solve complex contact-rich manipulation tasks.

environment. This field of *sensor fusion* is a long-studied problem in the field of robotics for various tasks ranging from control [6], [7], manipulation [8], [9], localization, and navigation [10], but remains underexplored in RL settings.

To this end, Imitation Learning approaches [11], [12], [13] have shown promise in utilizing multisensory data for learning skills, but require experts collecting informative data, thus limiting their usability, robustness and generalization. The challenge of data collection becomes particularly pronounced in tasks involving objects with varying or uncertain properties, such as mass or friction. In such cases, substantial and often expensive data acquisition efforts are required to develop a robust policy capable of generalizing effectively across diverse contextual variations [14]. Self-directed exploration in RL, on the other hand, facilitates the development of adaptable strategies that can be generalized across diverse object properties and contexts. These advantages position RL as a compelling approach for learning multisensory contact-rich manipulation [4], [15], [16]. Recent advances in multimodal learning [17], [18] have shown the advantages of reconstructing masked or noised inputs to learn expressive cross-modal representations for downstream tasks. Such masking-based self-supervision approaches, whether at the sensory or embedding level, have also been shown to improve the network’s robustness [19], [20], [21].

In this paper, we propose *MultiSensory Dynamic Pretraining* (MSDP), a novel RL framework to learn expressive multisensory representations for contact-rich manipulation tasks. MSDP first learns to extract and fuse sensor features via an offline, pretraining phase based on masked autoencoding and cross-sensor prediction; then, an online RL agent leverages the features from the pretrained multisensory encoder through a new combination of cross-attention and pooling mechanisms applied to the critic and actor, respectively. In turn, this design enables RL policies to synergistically fuse multisensory inputs and inherently cope with sensor noise or missing modalities.

Manuscript received: October 24, 2025; Revised: February 11, 2026; Accepted: March 20, 2026. This paper was recommended for publication by Editor Jens Kober upon evaluation of the Associate Editor and Reviewers comments. Corresponding author: rickmer.krohn@tu-darmstadt.de.

¹ Interactive Robot Perception & Learning (PEARL) Lab, TU Darmstadt, Germany, ² Hessian.AI; ³ Robotics Institute Germany (RIG); This research is funded by the German Research Foundation (DFG) Emmy Noether Programme (CH 2676/1-1), the EU’s Horizon Europe project “ARISE” (Grant no.: 101135959), the German Federal Ministry of Education and Research (BMBF) project “RiG” (Grant no.: 16ME1001) and the European Research Council (ERC) project “SIREN” (Grant No.: 101163933).

Digital Object Identifier (DOI): see top of this page.

Our experimental results demonstrate that MSDP yields effective representations that accelerate RL on a variety of contact-rich manipulation tasks while being robust to sensor noise and changing object dynamics. The FT-sensor boosts MSDP’s performance in two challenging real robot tasks by 14 % leading to near-optimal performance. Notably, the policy is trained directly in the real world on MSDP’s multisensory latent representation, without any sim-to-real transfer. We achieve near-optimal performance in only 6,000 online interactions, which takes less than 55 minutes including data collection and pretraining. Beyond these results, MSDP scales naturally to an increasing number of diverse input modalities and only introduces a few learnable parameters in downstream training of RL agents.

To summarize, the key contributions of our work are twofold: (i) we develop an effective pretraining strategy based on masked autoencoding to form a rich multisensory representation, and (ii) we introduce a novel architecture that allows task-specific feature extraction to achieve efficient robot Reinforcement Learning. Our findings pave the way for adaptive and resilient RL agents to handle multiple input modalities to master complex contact-rich manipulation tasks.

II. RELATED WORK

A. Reinforcement Learning for Contact-Rich Manipulation

Tasks where a robotic manipulator has to interact with its environment, either directly or indirectly via a tool, require a good understanding of the interaction forces that shape the task at hand. Solving contact-rich tasks often relies on an accurate estimate of force and dynamics. When such estimates are available, classical control approaches can be adapted to solve the task [22], [23]. To better handle unseen or difficult contact dynamics in unstructured environments, RL presents itself as an ideal candidate that can learn directly via interactions [24], without accurate state estimation [25]. Various works have employed RL to learn contact-rich manipulation policies on a variety of tasks [26], [27], [28].

B. Multimodal Self-Supervision for Robotics

Lee et al. [4] learns a multimodal representation from vision, force torque, and proprioception via MLP-fusion and multiple self-supervised objectives. The frozen representation leads to a robust representation for RL to solve multiple peg insertion tasks. [15] extended the Vision Transformer [29] with a force torque sensor to solve a variety of contact-rich tasks. They additionally incorporate the self-supervised objectives from [4] to shape a representation using SLAC [30]. [31] pretrains an audio-encoder [32] to combine vision and audio via a transformer decoder for Imitation Learning. The audio signal from the contact microphone provides rich feedback for various manipulation tasks. To also account for different sensor frequencies [13] developed a multi-resolution policy based on pretrained Vision Language Models to improve inference time using proprioception and force.

C. Deep Sensor Fusion

A robust representation is essential to integrate and process various sensory inputs to ensure stable and efficient learning in multisensory RL. Previous architectures often focused on straightforward latent fusion approaches by concatenating the various representations for downstream tasks [33], [34], [35], [36], [37]. Feng et al. [38] take this a step further by adding a subgoal-aware weighting for learning the stage-wise importance of difference sensors. Alternatively, contrastive learning approaches [39], [40] emphasize feature alignment, rather than fusion, to learn a shared latent representation between multiple modalities. However, distilling task-relevant features while maintaining sensor-specific information can be challenging. In contrast, inspired by the success of masked token prediction [17], recent works have also explored how masked multisensory pretraining can enhance representation learning for contact-rich manipulation [5], [16].

III. PRELIMINARIES

We define a multisensorial POMDP as a tuple $\mathcal{M} = \langle \mathcal{S}, \mathcal{O}^{MS}, \mathcal{A}, \mathcal{R}, \mathcal{P}, \gamma \rangle$ [41], where \mathcal{S} is the true state of the environment. The agent does not have access to \mathcal{S} and observes it through multiple sensor observations, which we specify as Proprioception $\mathcal{O}^P \in \mathbb{R}^{14}$ (joint position and velocities), Force Torque (FT) $\mathcal{O}^{FT} \in \mathbb{R}^{4 \times 6}$ (FT-sensor at robot wrist) and Vision $\mathcal{O}^V \in \mathbb{R}^{64 \times 64 \times 3}$ (external RGB camera). The sensors build up the multisensory Observation-space $\mathcal{O}^{MS} = [\mathcal{O}^P, \mathcal{O}^{FT}, \mathcal{O}^V]$. \mathcal{A} is the action space (cartesian control), $\mathcal{R} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}$ is the reward function, $\mathcal{P} : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$ is the transition kernel, and $\gamma \in [0, 1)$ is the discount factor. Our goal is to learn a policy $\pi : \mathcal{O}^{MS} \rightarrow \mathcal{A}$ that maximizes the discounted sum of rewards $\mathbb{E}_\pi [\sum_t^\infty \gamma^t r(s_t, a_t, s_{t+1})]$.

We aim to extract an expressive, fused representation from all sensors in order to solve contact-rich manipulation tasks with RL. We only provide the current multisensory observation and no history. We use the off-policy algorithm SAC [42] in simulation and RLPD [43] for efficiency in the real world, however we highlight that our approach is compatible with any actor-critic algorithm.

IV. MULTISENSORY DYNAMIC PRETRAINING

We present **MultiSensory Dynamic Pretraining (MSDP)**, a novel framework for representation learning that builds upon Masked Autoencoders [44], tailored to enhance contact-rich robot reinforcement learning tasks that require perception through multiple sensor modalities. MSDP introduces a modular architecture to seamlessly handle diverse input sensors, and a multisensory masking scheme to promote rich cross-modal representations, i.e., to retain knowledge about the task and the environment dynamics even in the absence or perturbation of one or more modalities.

Furthermore, we shed light on the different ways to map pretrained transformer embeddings to input states for downstream RL tasks, referred to as *latent bridging*, a problem often overlooked in practice by the community. In this context, we propose employing a simple cross-attention layer to obtain expressive task-specific features from frozen multisensory

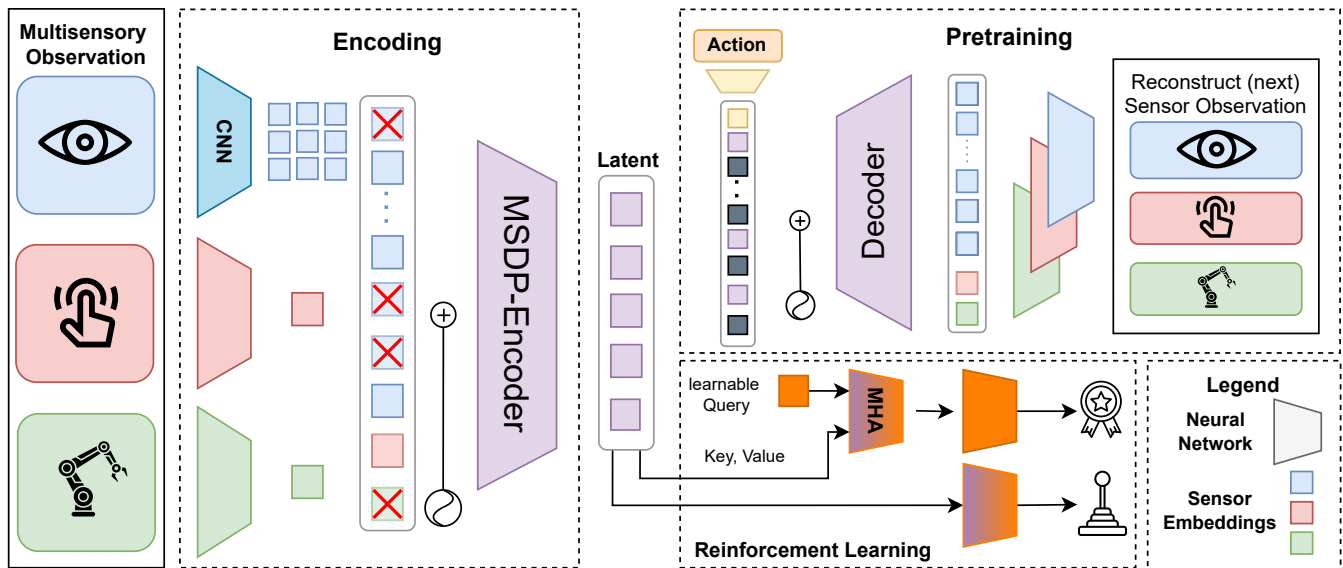


Fig. 2: The MSDP framework with MSDP-Encoder (left), Pretraining (top right) and downstream RL (bottom right): The current multisensory observation gets projected with a CNN-stem and linear layers to the embedding space. The MSDP-encoder fuses all sensor embeddings to form our expressive multisensory latent representation. The encoder is trained via the decoder and (next) sensor observation reconstruction from a subset of sensor embeddings. This pretraining results in dynamic cross-sensor prediction, shaping and fusing sensor representations. For downstream RL we extract multisensory task-specific features via a single cross-attention layer for the critic and via pooling for the actor. Sensor embeddings are only masked during pretraining. Our Framework offers an expressive and robust multisensory representation for complex contact-rich manipulation tasks in simulation and the real world.

embeddings for the critic representation, which is also referred to as attention pooling [45]. The actor on the other hand, receives a special pooling of the sensor embeddings to ensure a stable representation. Both mechanisms offer a fixed low-dimensional observation for the downstream RL agent.

Overall, our framework decouples representation learning and downstream RL to ensure rich feature extraction and fusion, while offering stable, robust and compact downstream representations. Notably, our architecture is scalable w.r.t. the number and types of input modalities and also supports pretraining using additional sensors, which may not be available during policy training. As a result, our pretraining phase on limited offline data leads to rich representations that may be directly used to solve complex contact-rich RL tasks from high-dimensional input (including images) in less than 500k environment steps, whereas in the real world only 6k online samples are needed.

A. MSDP Architecture

Given the inputs from multiple sensor modalities, such as vision, force torque readings and proprioception, we encode each modality using a separate network referred to as *sensor encoders*. Particularly for the vision input, we use a 4 layer CNN-stem similar to [46], which allows us to mask at embedding- rather than pixel-level, as masked object positions cannot be retrieved in pixel-space. The CNN-stem introduces redundancy as vision embeddings have overlapping receptive fields and, most importantly, stabilizes training [47]. Furthermore, compared to a patchify-stem, it reduces the burden of the multisensory encoder to extract vision features, while fusing with other modalities. We use a linear projection to encode force torque readings and proprioception to the 128-dim embedding space given their low dimensionalities.

Subsequently, we incorporate position encoding for the embedded sensor features to contextualize the position and modality of each embedding. Rather than using fixed positional/modality encoding [48] we use learnable embeddings to encode both together. Once we obtain embedding representations from all sensors, we randomly mask out a subset (70 %) of sensor tokens and feed the remaining to our 2 layer transformer encoder with 4 attention heads, pre-normalization and a mlp ratio of 2. The attention mechanism of the encoder leads to dynamic multisensory fusion promoted by the representation objective described in Section IV-B. A learnable mask embedding is added to the randomly masked out embedding positions, next to the multisensory embeddings from the encoder, to generate the original number of embeddings. We reapply our positional/modality encoding before feeding all embeddings to the decoder. During decoding, embeddings exchange information, especially the mask embedding, before being fed to separate decoder heads to fulfill the representation learning objective. We use a shared linear projection for all vision embeddings to reconstruct their corresponding patch.

B. Representation Learning

Our representation learning objective is based on multimodal masked autoencoding [44], [17], [18]. The objective is to either reconstruct the current \mathbf{O}_t^{MS} or next observation \mathbf{O}_{t+1}^{MS} from a random subset of sensor embeddings. Vision has a high number of embeddings and needs to extract information about other, potentially masked, sensors e.g. identifying contact to estimate force. Non-vision sensors on the other hand, are beneficial to reconstruct the vision observation as e.g. proprioception defines the robot position. Vision as a global sensing modality never gets fully masked out, whereas low-dimensional sensors

are not available when masked. This representation objective results in cross-sensor prediction thus leading to fusion of all modalities. Furthermore, the decoder is conditioned on action in the prediction objective and needs to extract dynamic action-related features important for downstream RL. We denote our method **MSDP-P** (Prediction) wherein we reconstruct the next observation $\|\mathbf{O}_{t+1}^{MS} - \Phi(\mathbf{O}_t^{MS}, \mathbf{A}_t)\|^2$ and **MSDP-R** (Reconstruction) where we reconstruct the current observation $\|\mathbf{O}_t^{MS} - \Phi(\mathbf{O}_t^{MS})\|^2$. $\Phi(\cdot)$ denotes the network prediction conditioned on the current observation \mathbf{O}_t^{MS} and action \mathbf{A}_t . $\|\cdot\|^2$ denotes the Mean Squared Error. Sensor embeddings are only masked during pretraining.

C. Policy Learning

From the pretrained encoder, we receive expressive multi-sensory embeddings given the available sensors. Extracting a compact representation from those embeddings for downstream task-solving is a crucial and often under-discussed design choice. We name the mapping between embeddings and a compact representation *latent bridging*, which has a considerable impact on performance. To address this, we depart from works that naively extract the "CLS"-token from the high-dimensional embeddings, analogous to the Vision Transformer [21], [29], [49], [50] and, instead, propose an asymmetric *latent bridging* strategy between actor and critic to obtain a compact representation. The critic uses a single cross-attention layer with a learnable query and the multisensory embeddings from the MSDP encoder as keys and values (see Figure 2). It offers dynamic task-specific feature extraction (e.g. object positions, robot state, contact) over the task-solving process. Fine-grained understanding of the environment leads to faster convergence compared to a global representation. The policy, on the other hand, does not profit from a cross-attention layer as it may destabilize training. Instead, we mean pool all embeddings originating from the vision sensor before pooling all embeddings, to account for the uneven number of sensor embeddings. Pooling sensor tokens results in a stable and parameter-free latent bridging similar to [5], [46]. This asymmetric representation between actor and critic, follows [51], where the actor benefits from a stable representation over task stages and the critic from a dynamic-specific representation. The cross-attention layer is trained by the critic, resulting in task-specific feature extraction. We train with the off-policy algorithm SAC in simulation and RLPD in the real world, where we incorporate offline data for pretraining in our replay buffer. For pretraining and policy training we use a batch size of 64 and a learning rate of 3×10^{-4} . Our approach is working with any actor-critic RL algorithm.

V. EXPERIMENTS AND RESULTS

In this section, we present our experimental setting, three competitive baselines, training details and multisensory environments. Furthermore, we ablate the impact of sensor settings and latent bridging mechanisms. Finally, we investigate how pretraining with multiple sensors can enrich the vision representation for downstream RL and showcase our task-specific feature extraction.




Fig. 3: Multisensory contact-rich robot environments

A. Baselines and Training Details

We compare our methods against one transformer and two non transformer-based baselines. The latter models extract sensor-specific features with a CNN for vision and an MLP for proprioception and force torque readings. The **Concat** model fuses the concatenated features with a 2-layer MLP to form the multisensory latent representation [4]. The **PoE** model generates separate means and variances for each sensor and fuses them with a product of experts approach. **VTT** [15] fuses all sensors via a transformer-encoder and compresses the features via multiple linear layers to a compact latent. As the baseline models introduce a bottleneck, we neglect masking and pretrain the encoders with observation reconstruction to form the latent. All models use a similar decoder architecture for the representation objectives, where proprioception and force torque signals are reconstructed using separate two-layer MLPs and vision via a deconvolutional CNN.

We collect 30,000 random samples from the environment and train each model with its respective representation learning objective (see Section IV-B) for 30,000 update steps. After pretraining, we freeze the encoder and train the downstream task for 100 epochs. Each epoch consists of 5,000 interaction steps, totaling 500,000 RL update steps. We report the mean and 95 % confidence interval over 12 runs.

B. Multisensory Environments

Our multisensory environments are based on panda-gym [52] and Pybullet [53]. The tasks Peg Insertion, Push Cube, Close Drawer Gently and Dual Arm Peg Insertion are challenging contact-rich manipulation tasks with sensor noise, varying dynamics and number of sensors. Each task has a dense reward function and comes with multisensory observations consisting of proprioception, vision and force torque readings. The observation of the Dual Arm environment has an additional FT-sensor and proprioception of the second robot arm. The action space consists of the 3-dimensional endeffector displacement. Task details are provided below and an overview of the environments can be found in Figure 3.




Fig. 4: Performance comparison between MSDP-P and MSDP-R to the baselines in Peg Insertion, Push Cube, Close Drawer Gently and Dual Arm Peg Insertion. Our method significantly accelerates RL training and achieves the highest final success rate across all tasks.

Peg Insertion: Insert the triangular peg into the hole. Next to the position, the z-orientation of the peg needs to match the hole, leading to a 4-dimensional action space and complex insertion dynamics. We randomize the robot initial position, the hole position and orientation. Furthermore, we add gaussian noise to the vision observation to mimic sensor noise and encourage the usage of the FT-sensor.

Push Cube: Push a cube to a target location, with a round peg held by the endeffector. Robot and cube initial position are randomized. Additionally, we vary the cube’s mass and center of mass, leading to changing object dynamics.

Close Drawer Gently: Instead of fully closing a drawer, the task is successful when the drawer is nearly closed, while having a minimal velocity. It requires fine-grained control to be solved, compared to standard close drawer formulations, where undesirable, forceful closing behavior may occur. We randomize the position and orientation of the drawer and the friction of its prismatic joint.

Dual Arm Peg Insertion: extends the Peg Insertion task to a bi-manual setting. We randomize both base positions of the robot and use a similar reward function. Both robot arms need to align the peg and hole positions to solve the task.

C. Simulation Experiments

Figure 4 shows the performance of our MSDP-P and MSDP-R models compared to the baselines. We obtain superior performance in Peg Insertion and Push Cube, while also outperforming in Close Drawer Gently and Dual Arm Peg Insertion, indicating an expressive representation for contact-rich tasks. Especially in Peg Insertion, our sensor fusion provides fine-grained features to solve the task around 80 % of the time after only 200,000 environment steps. The baselines struggle to combine sensor knowledge to consistently insert the peg. Baseline models need an extended amount of training in Push Cube to be able to account for varying object dynamics, like mass and center of mass of the cube, while our representations allow quick adaptation to those changes, leading to a steep learning curve and optimal performance. The Close Drawer Gently task offers no clear results, mainly due to simulation bottlenecks. We observed sudden changes in the drawer during contact that limit the agent’s ability to close

it smoothly. Still, our method can learn the task faster with better final performance. For the Dual Arm Peg Insertion task additional sensor encoders are used to map the proprioceptive state of the second robot and its force torque reading to the respective latent. Next to PoE we achieve optimal performance over all runs. Our embedding-based latent would allow for more advanced mapping to actor/critic, which we omit for a fair comparison.




Fig. 5: Peg Insertion Sensor Ablation: Proprioception is crucial to identify the precise peg pose under vision noise, while the force torque sensor allows for precise exploration around the hole, resulting in consistent insertions. Vision only is not able to achieve a high success rate.

Sensor Ablation: To evaluate the role of individual sensors, we conduct sensor ablation studies on Peg Insertion and Push Cube using the MSDP-P model. Figure 5 compares the success rate in various sensor settings, highlighting the performance gains compared to a vision only approach. Proprioception is essential to identify the peg pose under vision noise, while the force torque sensor guides exploration and allows for consistent insertions, indicated by the lower confidence interval. Figure 6 shows the episode lengths in the Push Cube task as all sensor combinations achieve near-optimal performance. Vision is sufficient to locate the cube and solve the task. The usage of all sensor allows the agent to make and keep contact with the cube to solve the task in an efficient manner, indicated by the fast completion time (30 % faster compared to vision only). In addition to the performance gains, we observe smoother object interactions with access to the force torque modality. Most prominent in the Peg Insertion task, where the agent without




Fig. 6: Push Cube Sensor Ablation: Best performance is achieved using all sensors as the agent is able to maintain contact with the cube in order to push it fast to the goal.

having access to the FT-sensor asserts high forces on the hole, indicating unwanted exploration actions which may damage the robot or environment in the real world.

Latent Bridging: The *latent bridging* mechanism to obtain a compact representation from multisensory embeddings, described in Section IV-C, can have a substantial impact on performance. We compare our cross-attention extraction with common approaches. *CLS* is using the CLS embedding of the encoder, commonly used in the Vision Transformer [29] or Imitation Learning [13]. In *Pooling* we mean-pool all sensor embeddings similar to [5], [46]. To account for the different number of embeddings between sensors, we first take the mean of all vision embeddings. To avoid possible dilution *Cat* provides the Concatenation of sensor embeddings, where we again pool the vision embeddings. Common latent bridging mechanisms do not introduce new learnable parameters, while our cross-attention only adds minimal parameters to extract task-specific features.




Fig. 7: Push Cube latent bridging ablation: Our mechanism extracts fine-grained multisensory features resulting in efficient policy training.

Figure 7 shows different latent bridging mechanisms on our MSDP-P model in the Push Cube task. Our approach successfully extracts task-specific features over an episode in an efficient manner. Our mechanism can extract more fine-grained vision features and combine them with other sensor information, resulting in superior performance. The *CLS* embedding doesn't contain rich features for downstream task learning, while *Pooling* takes longer to solve the task consistently, as the agent needs to extract environment details from a potentially diluted representation.




Fig. 8: Pretraining the representation for Peg Insertion with multiple sensors improves performance, when only vision is available during policy training.

Enriched Representation for vision-based RL: Our MSDP transformer encoder allows for varying input length, offering us the option to pretrain the representation on more sensors and only use a subset, for downstream task learning. We pretrain our MSDP-P model with different sensor configurations and only provide the vision modality during task solving. Figure 8 shows how different pretraining settings can enhance vision-based RL in the Peg Insertion task. The encoder stores information about proprioception and force torque readings promoted by the masked autoencoding pretraining objective, thus enriching the vision representation, leading to increased performance.

Task-specific features extraction: The cross-attention map in figure 9 highlights how the critic focuses on task-relevant features e.g. cube and target position. We want to stress that all embeddings from the pretrained encoder are multisensorial, even if we can assign their origin from a specific sensor. As shown in the previous paragraph, additional sensor features are encoded in the vision modality.




Fig. 9: Critic's cross-attention maps in the Push Cube task. The critic is focusing on the multisensory embeddings originated at the cube position, indicating task-specific feature extraction.

D. Real World Experiments

We conduct the Peg Insertion and Push Cube tasks in the real world using a Franka robot arm with a wrist-mounted FT-sensor and a custom endeffector. The observation space consists of endeffector position and velocity as proprioception, four force torque readings and an downsampled 64 by 64 RGB image from the 3rd person RealSense camera for both tasks. Peg and hole are both 3d printed. The robot needs to align the orientation of the triangular peg and place it precisely to fully insert the peg and complete the task. In Push Cube the agent needs to push the 7.5 cm block to the 15 cm away goal location. Successful episodes are detected with the endeffector position in Peg Insertion and an aruco marker and a second camera in the Push Cube task. The starting position of the robot is randomized and we use a sparse reward of +1 upon success in both tasks. An overview of the setup and results can be found in Figure 10a and 10b.

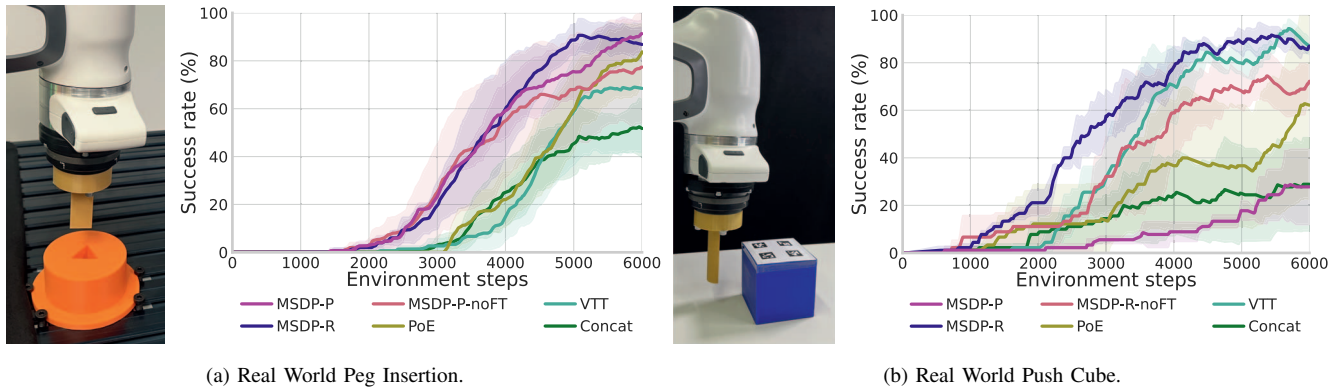


Fig. 10: Real world setup and experimental results. Our MSDP framework enables training RL policies directly in the real world, with first successful episodes after only 2,000 / 1,000 online interactions, outperforming various baselines. Task success is detected via endeffector position or the aruco marker on the cube. Force torque readings are essential to consistently push the cube to the goal or to insert the peg under occlusion improving task success by 14 % (cf. MSDP-noFT). Policies are learned directly on the pretrained multisensory representation, without any sim-to-real transfer, with only 6,000 online interactions.

We build upon the SERL-package [54] and use a cartesian impedance controller for safe interaction. The multisensory encoder is pretrained with data from 20 demonstrations and $\sim 2,000$ samples of play data, resulting in a total of $\sim 3,000$ samples. Downstream tasks are trained directly on the real system for only 6,000 environment interactions with the proposed RLPD [43] algorithm from SERL. We pretrain for 6,000 steps in Peg Insertion, updating critic and actor twice per online sample. For Push Cube, due to higher vision variance, we use 10,000 update steps and apply four critic and one actor update step per online interaction to obtain meaningful representations and efficient training. The complete training pipeline, including data collection, pretraining, and online RL, learns the final policy in less than 55 minutes for both tasks. Compared to prior real-world multisensory RL methods [4], [15], our approach requires only a fraction of the data (3 % pretraining data in Peg Insertion) and training time (< 20 %) to learn the final policy.

Figures 10a and 10b show results for Peg Insertion (5 runs) and Push Cube (3 runs), respectively. Both our pretraining objectives achieve superior performance in Peg Insertion, whereas baselines suffer from vision noise and occlusion. Particularly, we notice that our synergistic use of force torque readings enables consistent insertion behavior, while reducing the asserted force between the peg and the hole during exploration. On the other hand, the Push Cube task introduces several challenges as the agent needs to make and maintain contact with the cube without pushing it out of the workspace. Here, MSDP-R and VTT are able to obtain a high final success rate (see Figure 10b), while MSDP-P is not able to extract a suitable representation. We attribute this shortcoming to the complexity of learning forward dynamics in vision-demanding tasks on limited real world data and lack of observation histories. Overall, we highlight that our expressive multisensory representation, in combination with task-specific feature extraction via latent bridging, is able to obtain first success with less than 2,000 online interactions. FT-sensor usage particularly improves the performance of MSDP by 14 % in both real-world tasks, underscoring the significance of



Fig. 11: Visual disturbances of the robustness evaluation. From left to right: back light, front light, occlusion, and disco lights. MSDP achieves a high success rate across all disturbances, which have not been seen during training.

multi-sensor integration in contact-rich manipulation.

Robustness Evaluation: We evaluate the final policy of MSDP-P in the Peg Insertion task under various disturbances to showcase the robustness of our pretrained multisensory encoder and policy. We evaluate each change that hasn't been observed during training for 20 trials. Trained on $K_c = 2000$ cartesian-stiffness the policy achieves 90 % success rate with decreased ($K_c = 1500$) and 100 % with increased cartesian-stiffness ($K_c = 2500$). MSDP shows remarkable robustness against changed light settings, e.g., back light (100 %), front light (100 %), disco lights (100 %) and visual occlusion (partly blocked camera view, 95 %) and external forces (80 %). All visual disturbances are shown in Figure 11.

This work assumes the presence of all sensor modalities at the start of training to obtain an expressive representation from limited multisensory data. The usage of MSDP to enhance existing representations seems promising to leverage common available datasets for generalization.

VI. CONCLUSION AND FUTURE WORK

In this work, we propose MultiSensory Dynamic Pretraining (MSDP), a novel pretraining framework to learn multisensory representations for contact-rich manipulation tasks. Specifically, MSDP learns to reconstruct varying sensory information from a subset of input sensor embeddings, leading to cross-sensor predictions and sensor fusion. MSDP captures the interplay between different sensor modalities to learn a rich multisensory representation, which in combination with task-specific feature extraction through cross-attention, leads to a superior performance in challenging contact-rich manipulation

tasks in simulation and the real world. Future work may investigate MSDP's capabilities for sensor modalities beyond vision and force-torque, such as tactile or audio, as well as its use for fine-tuning existing representations with multisensory data. Finally, we reckon that extending MSDP's pretraining stage to incorporate play data from environments at large scale may yield representations that generalize over new tasks.

REFERENCES

- [1] V. Mnih *et al.*, "Playing Atari with Deep Reinforcement Learning," 2013, arXiv:1312.5602 [cs].
- [2] Z. Zhuang *et al.*, "Robot Parkour Learning," 2023, arXiv:2309.05665 [cs].
- [3] F. Zhang *et al.*, "Towards Vision-Based Deep Reinforcement Learning for Robotic Motion Control," 2015, arXiv:1511.03791 [cs].
- [4] M. A. Lee *et al.*, "Making Sense of Vision and Touch: Self-Supervised Learning of Multimodal Representations for Contact-Rich Tasks," *CoRR*, vol. abs/1810.10191, 2018, eprint: 1810.10191.
- [5] C. Sferrazza *et al.*, "The Power of the Senses: Generalizable Manipulation from Vision and Touch through Masked Multimodal Learning," 2023, arXiv:2311.00924 [cs].
- [6] J. G. García *et al.*, "Sensor fusion for compliant robot motion control," *IEEE Transactions on Robotics*, vol. 24, no. 2, pp. 430–441, 2008.
- [7] F. F. Khalil and P. Payeur, "Dexterous robotic manipulation of deformable objects with multi-sensory feedback—a review," *Robot Manipulators Trends and Development*, no. March 2010, 2010.
- [8] Y. Hu *et al.*, "Development of sensory-motor fusion-based manipulation and grasping control for a robotic hand-eye system," *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, vol. 47, no. 7, pp. 1169–1180, 2016.
- [9] Z. Xia *et al.*, "A review on sensory perception for dexterous robotic manipulation," *International Journal of Advanced Robotic Systems*, 2022.
- [10] M. B. Alatise and G. P. Hancke, "A review on challenges of autonomous mobile robot and sensor fusion methods," *IEEE Access*, vol. 8, pp. 39 830–39 846, 2020.
- [11] H. Li *et al.*, "See, Hear, and Feel: Smart Sensory Fusion for Robotic Manipulation," 2022, arXiv:2212.03858 [cs].
- [12] Y. Hao *et al.*, "Masked Imitation Learning: Discovering Environment-Invariant Modalities in Multimodal Demonstrations," in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2023, pp. 1–7, iISSN: 2153-0866.
- [13] S. Saxena, M. Sharma, and O. Kroemer, "MResT: Multi-Resolution Sensing for Real-Time Control with Vision-Language Models," 2024, arXiv:2401.14502 [cs].
- [14] S. Belkhale, Y. Cui, and D. Sadigh, "Data quality in imitation learning," *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [15] Y. Chen *et al.*, "Visuo-Tactile Transformers for Manipulation," 2022, arXiv:2210.00121 [cs].
- [16] Q. Liu *et al.*, *Masked Visual-Tactile Pre-training for Robot Manipulation*, 2024.
- [17] R. Bachmann *et al.*, "MultiMAE: Multi-modal Multi-task Masked Autoencoders," 2022.
- [18] X. Geng *et al.*, "Multimodal Masked Autoencoders Learn Transferable Representations," 2022.
- [19] S. Skand *et al.*, "Simple Masked Training Strategies Yield Control Policies That Are Robust to Sensor Failure," 2024.
- [20] G.-H. Liu *et al.*, "Learning End-to-end Multimodal Sensor Policies for Autonomous Navigation," 2017.
- [21] I. Radosavovic *et al.*, "Real-World Robot Learning with Masked Visual Pre-training," 2022, arXiv:2210.03109 [cs].
- [22] D. E. Whitney, "Historical perspective and state of the art in robot force control," *The International Journal of Robotics Research*, vol. 6, no. 1, pp. 3–14, 1987.
- [23] D. E. Whitney *et al.*, "Quasi-static assembly of compliantly supported rigid parts," *Journal of Dynamic Systems, Measurement, and Control*, vol. 104, no. 1, pp. 65–77, 1982.
- [24] M. Nuttin and H. Van Brussel, "Learning the peg-into-hole assembly operation with a connectionist reinforcement technique," *Computers in Industry*, vol. 33, no. 1, pp. 101–109, 1997.
- [25] S. Levine *et al.*, "End-to-end training of deep visuomotor policies," *Journal of Machine Learning Research*, vol. 17, no. 39, pp. 1–40, 2016.
- [26] Í. Elguea-Aguinaco *et al.*, "A review on reinforcement learning for contact-rich robotic manipulation tasks," *Robotics and Computer-Integrated Manufacturing*, vol. 81, p. 102517, 2023.
- [27] M. Suomalainen, Y. Karayiannidis, and V. Kyrki, "A survey of robot manipulation in contact," *Robotics and Autonomous Systems*, vol. 156, p. 104224, 2022.
- [28] R. Liu *et al.*, "Deep reinforcement learning for the control of robotic manipulation: a focussed mini-review," *Robotics*, vol. 10, no. 1, p. 22, 2021.
- [29] A. Dosovitskiy *et al.*, "An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale," *CoRR*, vol. abs/2010.11929, 2020, eprint: 2010.11929.
- [30] A. X. Lee *et al.*, "Stochastic Latent Actor-Critic: Deep Reinforcement Learning with a Latent Variable Model," 2020, arXiv:1907.00953 [cs].
- [31] J. Mejia *et al.*, "Hearing Touch: Audio-Visual Pretraining for Contact-Rich Manipulation," 2024, arXiv:2405.08576 [cs].
- [32] P. Morgado, N. Vasconcelos, and I. Misra, "Audio-Visual Instance Discrimination with Cross-Modal Agreement," 2021, arXiv:2004.12943 [cs].
- [33] M. A. Lee *et al.*, "Making sense of vision and touch: Learning multimodal representations for contact-rich tasks," *IEEE Transactions on Robotics*, vol. 36, no. 3, pp. 582–596, 2020.
- [34] H. Li *et al.*, "See, hear, and feel: Smart sensory fusion for robotic manipulation," in *Conference on Robot Learning*. PMLR, 2023, pp. 1368–1378.
- [35] A. Li *et al.*, "Reinforcement learning strategy based on multimodal representations for high-precision assembly tasks," in *Intelligent Robotics and Applications: 14th International Conference, ICIRA 2021, Yantai, China, October 22–25, 2021, Proceedings, Part I 14*. Springer, 2021, pp. 56–66.
- [36] I. Guzey *et al.*, "Dexterity from touch: Self-supervised pre-training of tactile representations with robotic play," *arXiv preprint arXiv:2303.12076*, 2023.
- [37] Y. Han *et al.*, "Learning generalizable vision-tactile robotic grasping strategy for deformable objects via transformer," *IEEE/ASME Transactions on Mechatronics*, 2024.
- [38] R. Feng *et al.*, "Play to the score: Stage-guided dynamic multi-sensory fusion for robotic manipulation," in *8th Annual Conference on Robot Learning*, 2024.
- [39] V. Dave, F. Lygerakis, and E. Rueckert, "Multimodal Visual-Tactile Representation Learning through Self-Supervised Contrastive Pre-Training," 2024, arXiv:2401.12024 [cs].
- [40] F. Lygerakis, V. Dave, and E. Rueckert, "M2CURL: Sample-Efficient Multimodal Reinforcement Learning via Self-Supervised Representation Learning for Robotic Manipulation," 2024, arXiv:2401.17032 [cs].
- [41] H. Kurniawati, "Partially Observable Markov Decision Processes (POMDPs) and Robotics," 2021, arXiv:2107.07599 [cs].
- [42] T. Haarnoja *et al.*, "Soft Actor-Critic Algorithms and Applications," 2019, eprint: 1812.05905.
- [43] P. J. Ball *et al.*, "Efficient Online Reinforcement Learning with Offline Data," 2023, arXiv:2302.02948 [cs].
- [44] K. He *et al.*, "Masked Autoencoders Are Scalable Vision Learners," 2021, arXiv:2111.06377 [cs].
- [45] X. Chen *et al.*, "Context Autoencoder for Self-Supervised Representation Learning," 2023, arXiv:2202.03026 [cs].
- [46] Y. Seo *et al.*, "Masked World Models for Visual Control," 2023, arXiv:2206.14244 [cs].
- [47] T. Xiao *et al.*, "Early Convolutions Help Transformers See Better," 2021, arXiv:2106.14881 [cs].
- [48] A. Vaswani *et al.*, "Attention is All you Need," in *Advances in Neural Information Processing Systems*, I. Guyon *et al.*, Eds., vol. 30. Curran Associates, Inc., 2017.
- [49] T. Xiao *et al.*, "Masked Visual Pre-training for Motor Control," 2022, arXiv:2203.06173.
- [50] A. Majumdar *et al.*, "Where are we in the search for an Artificial Visual Cortex for Embodied Intelligence?" 2024, arXiv:2303.18240 [cs].
- [51] S. Garcin *et al.*, "Studying the Interplay Between the Actor and Critic Representations in Reinforcement Learning," 2025, arXiv:2503.06343 [cs].
- [52] Q. Gallouédec *et al.*, "panda-gym: Open-Source Goal-Conditioned Environments for Robotic Learning," *4th Robot Learning Workshop: Self-Supervised and Lifelong Learning at NeurIPS*, 2021.
- [53] E. Coumans and Y. Bai, "PyBullet, a Python module for physics simulation for games, robotics and machine learning," 2016.
- [54] J. Luo *et al.*, "Serl: A software suite for sample-efficient robotic reinforcement learning," 2024.