

MoDem-V2: Visuo-Motor World Models for Real-World Robot Manipulation

Patrick Lancaster¹ Nicklas Hansen² Aravind Rajeswaran¹ Vikash Kumar¹

Abstract—Robotic systems that aspire to operate in *uninstrumented real-world environments* must perceive the world directly via onboard sensing. Vision-based learning systems aim to eliminate the need for environment instrumentation by building an implicit understanding of the world based on raw pixels, but navigating the contact-rich high-dimensional search space from solely sparse visual reward signals significantly exacerbates the challenge of exploration. The applicability of such systems is thus typically restricted to simulated or heavily engineered environments since agent exploration in the real-world without the guidance of explicit state estimation and dense rewards can lead to unsafe behavior and safety faults that are catastrophic. In this study, we isolate the root causes behind these limitations to develop a system, called MoDem-V2, capable of learning contact-rich manipulation directly in the uninstrumented real world. Building on the latest algorithmic advancements in model-based reinforcement learning (MBRL), demo-bootstrapping, and effective exploration, MoDem-V2 can acquire contact-rich dexterous manipulation skills directly in the real world. We identify key ingredients for leveraging demonstrations in model learning while respecting real-world safety considerations – exploration centering, agency handover, and actor-critic ensembles. We empirically demonstrate the contribution of these ingredients in four complex visuo-motor manipulation problems in both simulation and the real world. To the best of our knowledge, our work presents the first successful system for demonstration-augmented visual MBRL trained directly in the real world. Visit sites.google.com/view/modem-v2 for videos and more details.

I. INTRODUCTION

Robot agents learning manipulation skills directly from raw visual feedback avoid the need for explicit state estimation and extensive environment instrumentation for rewards, but face heightened exploration, and thereby safety, challenges in navigating the contact-rich high-dimensional search space purely based on sparse visual reward signals. These challenges are especially critical for agents operating in the real world where inefficiency can be expensive, and safety faults can be catastrophic. One approach to developing robot manipulation policies that avoid such safety restrictions is simulation to reality transfer [1], [2], [3], [4]. However, the creation and calibration of accurate physics simulations (from first principles) for contact-rich tasks is extremely challenging and time-consuming. In this work, we alternatively study the use of visual world model learning [5], [6], [7] for robot manipulation directly from real-world interaction.

¹Patrick Lancaster, Aravind Rajeswaran, and Vikash Kumar are with Meta AI {plancaster, aravraj}@meta.com, vikashplus@gmail.com

²Nicklas Hansen is with the University of California San Diego, nihansen@ucsd.edu




Fig. 1. We use MoDem-V2 to train the robot on four contact-rich manipulation tasks. These tasks cover a wide range of manipulation skills, namely non-prehensile pushing, object picking, and in-hand manipulation. In recognition of the difficulty of robust pose tracking and dense reward specification in the real world, the robot performs these tasks using only raw visual feedback, proprioceptive signals, and sparse rewards.

Model-Based Reinforcement Learning (MBRL) with visual world models involves the learning of dynamics models using real-world data, directly from visual observations. When applied to robot manipulation, visual MBRL can mitigate the need for detailed physics simulations from first principles, as well as the need for specialized sensor instrumentation and state estimation pipelines. However, visual MBRL for real-world robotics still has two major challenges: (a) sample inefficiency; and (b) sparse/weakly shaped rewards. While a number of recent algorithms such as RRL [8] and MoDem [9] circumvent these challenges by leveraging a small number of expert demonstrations to improve sample-efficiency, they rely on aggressive exploration to compensate for weak reward supervision that can result in unsafe behaviors, restricting their application to simulated or heavily engineered scenarios.

We indeed found MoDem to be infeasible for direct application in the real world due to excessively aggressive exploratory behavior. Even with significant engineering investments for safety, we found that the built-in low-level hardware controllers/drivers fault repetitively owing to ex-




Fig. 2. Agent performance on the inclined pushing task before failure due to safety violations. An asterisk indicates that agent training was terminated due to significant safety violations. *Left*: On a real-world robot, MoDem violates (robot manufacturer specified) torque limits at the onset of online interaction and is unable to learn, whereas MoDem-V2’s conservative exploration allows it to perfect the task. *Right*: Further evaluation in simulation reveals that simply penalizing the amount of torque exerted by the robot does not prevent termination due to significant safety violations. Other baseline agents are either terminated due to unsafe behavior or achieve significantly lower success than MoDem-V2.

cessive velocity, acceleration, and torque in the robot’s joints that exert dangerous amounts of force on the environment or the robot itself. While some level of safety can be imposed through hard-coded limits on velocity, acceleration, and torque (as done in all of our experiments), this is an insufficient solution for preventing unsafe behavior during online interaction (Figure 2 left). Tuning these task-specific limits for interaction in the real world is costly and faced with operational and safety challenges. Balancing the tightness of these limits presents a further challenge, as aggressive limits can prohibit the robot from exerting the minimum energy needed to solve the task and weak limits will fail to inhibit unsafe actions. Finally, these static limits fail to incorporate the need for dynamic risk sensitivities across relevant time scales (such as at the scale of a single episode, or even at the scale of training epochs). Furthermore, simply penalizing the amount of torque exerted by the agent is an ineffective, retrospective solution that does not prevent unsafe actions at the onset of exploration as shown in Figure 2. How can we get around these challenges?

Our key insight is that conservative exploration can respect the safety constraints of real-world environments while still allowing the agent to modulate its strategy according to task progress such that it is able to learn quickly and efficiently. We translate this insight into implementation in three steps. First, rather than sampling actions from the entire action space, **warm-starting** exploration with actions sampled from a policy learned via behavioral cloning (BC) [10] prevents the agent from straying far from the provided demonstrations at the onset of online learning. Second, as our world model gains better coverage through online exploration, **agency-**

transfer gradually shifts the agent from executing BC policy actions to short-horizon planning. Agency transfer provides a mechanism for increased exploration while stymying over-optimistic evaluation of regions of the observation-action space far from the agent’s previous experience. Third, we use **actor-critic ensembles** to estimate the epistemic uncertainty of the value estimations of these short-horizon trajectories, allowing the agent to avoid overly optimistic actions. We integrate these three enhancements into the recently proposed MoDem [9] algorithm to develop MoDem-V2. Despite the simplicity of the individual components, their combination, and the resulting effectiveness in transforming overly aggressive, fault-prone MoDem agents into MoDem-V2 agents that efficiently and safely learn manipulation behaviors in the real-world is quite unique. Our work is, to the best of our knowledge, the first successful demonstration of demonstration-augmented visual MBRL trained directly in the real world.

To evaluate the effectiveness of our approach, we study four robot manipulation tasks from visual feedback (Figure 1), and their simulated counterparts. Our main contributions are:

- We **identify unsafe exploration and over-optimism** as the key issues in leveraging visual MBRL algorithms for real-world applications.
- Based on this insight, we develop MoDem-V2 by integrating three ingredients into MoDem, namely **policy centering**, **agency transfer**, and **actor-critic ensembles**.
- We demonstrate that MoDem-v2’s conservative exploration significantly enhances its safety profile compared to other baselines, while still retaining the sample-efficient learning capability of MoDem.
- Finally, we demonstrate that MoDem-V2 can **quickly learn a variety of contact-rich manipulation skills**, such as pushing, picking, and in-hand manipulation directly in the real world.
- We contribute towards lowering the barrier of entry for RL in the real world by **open-sourcing our implementation of MoDem-V2** and discussing practical considerations for training on real hardware.

II. PRELIMINARIES

We begin by introducing notations and providing an overview of MBRL settings.

Notation: The general setting of an agent interacting with its environment can be formulated as a Markov Decision Process (MDP) described by the tuple $\mathcal{M} := (\mathcal{S}, \mathcal{A}, \mathcal{T}, R, \gamma)$. Here, \mathcal{S} denotes the state space, \mathcal{A} is the action space, the conditional probability distribution $\mathbf{s}_{t+1} \sim \mathcal{T}(\cdot | \mathbf{s}_t, \mathbf{a}_t)$ defines the dynamics of the MDP, and a scalar reward function is given by $r_t = R(\mathbf{s}_t, \mathbf{a}_t)$. Finally, $\gamma \in [0, 1)$ defines the discount factor for the MDP to trade-off future rewards to current ones. The goal for an agent is to learn a policy $\pi : \mathcal{S} \mapsto \mathcal{A}$ that can achieve high long term performance given by $\mathbb{E}_\pi [\sum_{t=0}^{\infty} \gamma^t r_t]$.

We specifically consider the problem of robotic manipulation from visual feedback on real hardware. We aim to

learn a control policy that controls a physical robot from RGB observations provided by cameras placed in the scene, and robot proprioception. We model this setting as a **high-dimensional** MDP with **sparse rewards**. This assumes that while the state space of the MDP (*e.g.* object poses) is not directly observable by the agent, a sufficient representation of the state can be well-approximated through the combination: $\mathbf{s} = (\mathbf{x}, \mathbf{q})$ where \mathbf{x} denotes stacked RGB observations from the robot’s camera(s) and \mathbf{q} denotes proprioception from the robot. Finally, we only assume access to a sparse task completion reward, which is much easier to obtain via visual inputs compared to a detailed well shaped reward function. The final goal is to learn a policy that achieves high performance, using minimal online interactions, while respecting the safety considerations of hardware.

MoDem - Model-Based Reinforcement Learning with Demonstrations: Our approach is based on *MoDem* [9] – a MBRL algorithm that combines (*i*) model predictive control (MPC) and the decoder-free world model of TD-MPC [7] with (*ii*) a small number of demonstrations to efficiently solve continuous control problems with limited online interaction. Concretely, MoDem learns the following five components:

$$\begin{array}{ll}
\text{State embedding} & \mathbf{z} = h_\theta(\mathbf{s}) \\
\text{Latent dynamics} & \mathbf{z}' = d_\theta(\mathbf{z}, \mathbf{a}) \\
\text{Reward predictor} & \hat{r} = R_\theta(\mathbf{z}, \mathbf{a}) \\
\text{Terminal value} & \hat{q} = Q_\theta(\mathbf{z}, \mathbf{a}) \\
\text{Policy guide} & \hat{\mathbf{a}} = \pi_\theta(\mathbf{z})
\end{array} \quad (1)$$

where $h_\theta, d_\theta, R_\theta$ and Q_θ are learned end-to-end using a combination of joint-embedding predictive learning [11], reward prediction, and Temporal Difference (TD) learning, and π_θ is a deterministic policy that learns to maximize Q_θ conditioned on a latent state \mathbf{z} . Throughout this work, we will refer to $(h_\theta, d_\theta, R_\theta, Q_\theta)$ as the *world model*, and π_θ as the *policy*. See [7], [9] for additional details on the world model. While *MoDem* has been shown to be effective in simulation, owing to aggressive exploration its applicability in domains where safety constraints can’t be overlooked is limited (Figure 2).

III. SAFETY

Agents that seek to learn via continuous operations must respect hardware and environmental safety constraints. Such constraints are diverse, obscure, and unobservable (intrinsic to low-level hardware details, or lack of appropriate sensing) and therefore cannot be directly accounted for by the agent during operations. Alternatives such as action penalization and user-defined safety are also insufficient (Figure 2) requiring extensive human intervention and monitoring for operations. In this work, we refer to these unobservable constraints as **safety violations**. For real-world operations, they are defined as hardware faults that require human intervention. In simulation, we define them as violations (unobserved by the policy) in either the robot’s torque limits (as defined by the robot manufacturer) or excessive contact force (100 N) applied by the robot’s end effector.

Algorithm 1 Planning procedure of MoDem-V2 (●Original MoDem ●MoDem-V2 modification)

Require: θ : learned network parameters
 μ, σ : initial parameters for \mathcal{N}
 N : num sample trajectories
 \mathbf{s}_0, h : current state, planning horizon
 τ : trajectory weighting temperature
 α : probability of using model rollouts
 w_1, w_2 : ensemble mixing weights

- 1: encode state $\mathbf{z}_0 \leftarrow h_\theta(\mathbf{s}_0)$ \triangleleft State embedding
- 2: **if** $\text{rand}() > \alpha$ **then**
- 3: $\Gamma := \{\mathbf{a}_0\}^N \sim \pi_\theta^{BC}$ \triangleleft Center actions around BC
- 4: $\phi_\Gamma = Q_\theta^{BC}(\mathbf{z}_0, \{\mathbf{a}_0\}^N)$ \triangleleft Critic evaluation
- 5: **else**
- 6: $\Gamma := \mathbf{a}_{t:t+h} \sim \mathcal{N}(\mu, \mathbf{I}\sigma^2)$ \triangleleft Prior sampling
- 7: $\Gamma := \mathbf{a}_{t:t+h} \sim \pi_\theta^{1:M}, d_\theta$ \triangleleft Policy ensemble sampling
- 8: **for all** N trajectories $\Gamma_i = (\mathbf{a}_t, \mathbf{a}_{t+1}, \dots, \mathbf{a}_{t+h})$ **do**
- 9: **for step** $t = 0..h - 1$ **do**
- 10: $\phi_\Gamma = \phi_\Gamma + \gamma^t R_\theta(\mathbf{z}_t, \mathbf{a}_t)$ \triangleleft Reward
- 11: $\mathbf{z}_{t+1} \leftarrow d_\theta(\mathbf{z}_t, \mathbf{a}_t)$ \triangleleft Latent transition
- 12: $\phi_\Gamma = \phi_\Gamma + \gamma^h Q_\theta(\mathbf{z}_H, \mathbf{a}_H)$ \triangleleft Terminal value
- 13: # Ensemble of terminal values
- 14: $\phi_\Gamma^{1:M} = \phi_\Gamma + \gamma^h Q_\theta^{1:M}(\mathbf{z}_H, \mathbf{a}_H)$
- 15: # Epistemic uncertainty estimation
- 16: $\phi_\Gamma = w_1 \text{mean}(\phi_\Gamma^{1:M}) + w_2 \text{std}(\phi_\Gamma^{1:M})$
- 17: $\Omega = e^{\tau(\phi_\Gamma)}, \mu = \frac{\sum_{i=1}^N \Omega_i \Gamma_i}{\sum_{i=1}^N \Omega_i}, \sigma = \sqrt{\frac{\sum_{i=1}^N \Omega_i (\Gamma_i - \mu)^2}{\sum_{i=1}^N \Omega_i}}$
- 18: **return** $\mathbf{a} \sim \mathcal{N}(\mu, \mathbf{I}\sigma^2)$

IV. METHOD

In this work, we aim to learn manipulation skills through environment interaction on *real* robots, from *visual* feedback, and with *minimal* human supervision and intervention. We first discuss two important metrics for any agent that aims to quickly and safely learn manipulation skills in the real world. We then address these limitations with three proposed enhancements to MoDem in order to develop MoDem-V2.

Strengths and Weaknesses of MoDem: MoDem accelerates learning through a three-stage framework in which h_θ, π_θ are first pretrained on a set of demonstrations using Behavior Cloning (BC), and the resulting policy $\pi_\theta \circ h_\theta(\cdot)$ is then used to *seed* the model, *i.e.* collect a small initial dataset for learning the model (with Gaussian noise injected into π_θ for exploration). After initializing the model on seeding data, the world model is iteratively used to collect new data via online interaction, and is optimized on all data: demonstrations, seeding data, and online interaction data. Using these insights, MoDem demonstrates accelerated model learning in a collection of simulated tasks. Yet, when exposed to the real world, MoDem faces a variety of challenges that were suppressed during its original simulation study.

We find that MoDem exerts excessive forces and torques far beyond those exemplified in the provided demonstrations (see Section VI). Even at the beginning of the online interaction phase at which the agent has only observed data close to the BC policy, MoDem relies on its world model and

value function to discriminate between high and low reward actions. Thus when MoDem samples actions across the entire action space, the world model and value function cannot (at least initially) provide good estimates. This can result in task failure due to poor action selection, or lead to the robot exerting unsafe forces and torques by choosing consecutive actions that are far apart in action space. Our primary contribution lies in eliminating these limitations while also improving the effectiveness of the original method.

A. MoDem-V2: Real-World Model-Based Reinforcement Learning with Demonstrations

MoDem generates agent actions by first mapping raw observations into a learned lower-dimensional space, and then performing efficient short horizon planning in that latent space with its learned dynamics model and value function. We propose the following three adaptations to MoDem’s planning procedure in order to improve its safety while maintaining its strengths in autonomy and data-efficiency (Algorithm 1).

Policy centered actions: Rather than sampling actions from across the entire action space, we propose to sample actions from our learned policy. This more conservative exploration strategy reduces the likelihood of world model and value function evaluation over unseen regions of the state-action space, enabling them to better discriminate the quality of the generated actions.

Agency transfer from BC actions to MPC: At the beginning of the online interaction phase, MoDem immediately begins using its learned world model and value function to do MPC. Yet both of these components have only seen limited data near the BC policy, so relying on them to choose actions for multiple consecutive timesteps at the beginning of interaction can quickly lead the agent into an unexplored region of the observation-action space from which it cannot recover. Our remedy is to *gradually* shift from executing actions sampled from the BC policy to actions computed by short horizon planning. We implement this in Algorithm 1 with a hyperparameter α that is initialized to 0 at the beginning of interaction and linearly increases to 1.0 over a fixed number of interaction steps.

Actor-Critic Ensembles for uncertainty aware planning: The use of actor-critic ensembles [12] improves the agent’s value estimations in two primary ways. First, note that each actor is trained by optimizing it to maximize its corresponding critic. While this provides a solution for efficiently finding the maximum value of Q over actions, it is subject to significant overestimation bias [13]. We mitigate this by only evaluating a critic with final trajectory actions produced by policies not directly optimized to maximize that particular critic. Actor-critic ensembles also improve value estimation by providing the agent with a pool of independently trained value functions, each of which computes its own value estimate. By estimating the epistemic uncertainty [14] of a trajectory, the agent can make uncertainty-aware decisions. We incorporate this into MoDem-V2 with weights $w_1 > 0$ and $w_2 < 0$ in Algorithm 1.

V. EXPERIMENTAL DESIGN

We design experiments to evaluate the design choices behind MoDem-V2 in enabling real-world contact-rich manipulation. Our investigation focuses on the following directions:

- How sample-efficient is MoDem-V2 relative to other methods?
- Is MoDem-V2 safer (*i.e.* fewer safety violations) than other methods including MoDem?
- Does MoDem-V2 actually enable physical robots to learn real-world manipulation tasks?

Our first step in answering these questions is to use simulation [15] to compare our method to both the original MoDem and other strong reinforcement learning baselines that also uses demonstrations to guide policy learning, Demonstration Augmented Policy Gradients (DAPG) [16] and the Framework for Efficient Robot Manipulation (FERM) [17]. We also provide further analysis of MoDem-V2 by ablating each of its three design decisions. Finally, we deploy MoDem-V2 onto a physical robot and evaluate its capability to learn four different manipulation tasks in the real world.

A. Hardware

We adopt a set of core hardware components that are common across all of our experimental tasks. Each task uses a Franka Panda arm. The pushing and picking tasks use a Robotiq two-fingered gripper, while the in-hand reorientation task uses a ten-degree of freedom D’Manus hand [18] from the *ROBEL* ecosystem [19]. For perception, three RealSense D435 cameras are mounted to the left, right and above the robot. Our hardware setup is depicted in Figure 3.



Fig. 3. A view of the in-hand reorientation task as an example of our hardware setup.

B. Task Suite

We evaluate MoDem-V2 on four manipulation tasks in both simulation and the real world. These tasks encompass a variety of manipulation skills, namely pushing, picking, and in-hand manipulation as shown in Figure 1. We briefly describe each task below.

Planar Pushing: This task requires the robot to push an oblong object towards a fixed goal position on a table top. This task is likely the easiest of all four tasks, and we view it as base case with which to compare the other tasks.

Inclined Pushing: This task requires the robot to push an object up an incline to reach a fixed goal position. During execution of the task, the robot must raise its gripper such that it can progress up the incline while also making sufficiently precise contact with the block to prevent it from slipping beneath or around the side of the gripper.

Bin Picking: To complete this task, the robot must grasp a juice container and then raise it out of the bin. This task requires accurate positioning of the gripper because the (mostly) non-deformable container has a primary width that is approximately 65% of the gripper’s maximum aperture.

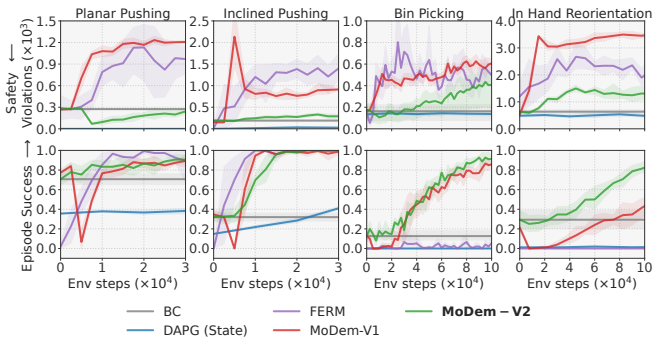


Fig. 4. The number of safety violations as defined in Section III (top row) and success rate (bottom row) for each of the four manipulation tasks in simulation. Lower is better for safety violations while higher is better for episode success. While both MoDem-V2 and MoDem achieve similar or better sample-efficiency than all of the baselines, MoDem-V2 exhibits significantly safer learning as evidenced by the drastically lower amount of safety violations.

This task also requires the robot to disambiguate spatially similar states; *e.g.* if the gripper is above the bin, the robot must understand whether or not the object is in its grasp so that it can decide to go down towards the bin to pick up the object or stay where it is in order to receive reward.

In-Hand Reorientation: This task requires the robot to grasp a water bottle laying on its side and then in-hand manipulate it to an upright position. Using the multi-fingered D’Manus hand more than doubles the dimensionality of the action space relative to the previous tasks.

VI. EXPERIMENTS

In this section, we evaluate MoDem-V2 against strong baselines in simulation and measure its performance in real world environments. Our simulation experiments measure both sample-efficiency and safety (defined in Section III), two important aspects for any learning method that is used in the real world. We then use MoDem-V2 to train a robot to perform all four manipulation tasks in the real world. All experiments train an initial policy with only ten demos, and each evaluation is aggregated over 30 trials.

A. Simulated Comparison to Baselines

With respect to sample-efficiency, MoDem-V2 and MoDem both significantly outperform DAPG (State), as shown in Figure 4 (bottom). This is despite DAPG having access to privileged state information (such as the object’s pose) and dense rewards (such as rewards based on the distance between the gripper and the object)! Although FERM achieves similar performance to both versions of MoDem on the easier pushing tasks, it is unable to learn the bin picking and in-hand re-orientation tasks. While MoDem achieves similar efficiency to MoDem-V2 on the first three tasks, it suffers a significant performance drop in the early stages of training across all tasks. This coincides with the point at which

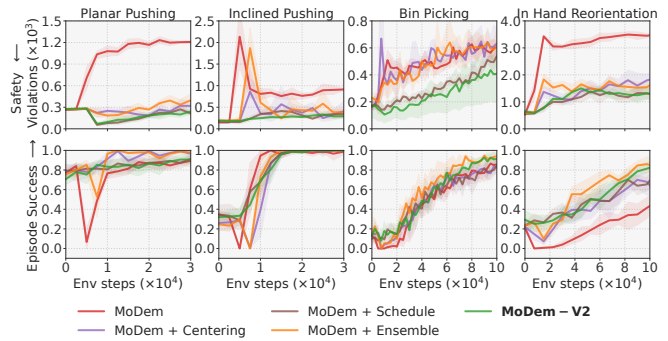


Fig. 5. Ablations of the three MoDem-V2 enhancements for all four tasks. Lower is better for safety violations (top row) while higher is better for episode success (bottom row). MoDem-V2 achieves both the higher sample-efficiency of *Ensemble* and the increased safety profile of *Schedule*.

MoDem begins to aggressively explore the action space, resulting in the robot applying excessive forces/torques beyond that observed in the provided demonstrations.

We measured the number of safety violations that occurred throughout training for MoDem and MoDem-V2, as shown in Figure 4 (top). While both methods initially commit few violations since their BC policies were trained from the same demonstrations, MoDem’s safety violations sharply increase as the interaction phase begins. Thanks to our design decisions, the number of violations exerted by MoDem-V2 is generally lower throughout online learning comparatively. Here MoDem-V2 demonstrates that it can achieve similar or better sample-efficiency than MoDem, while committing significantly fewer safety violations and thereby safer behavior.

B. Ablation of design choices

We perform ablations of MoDem-V2 by individually adding each improvement specified in Subsection IV-A to MoDem. As shown in Figure 5, we found that all of the ablations generally maintained or improved over the sample-efficiency of MoDem while significantly improving safety by committing fewer violations. The one exception to this is the bin picking task, for which the *Centering* and *Ensemble* ablations commit a greater number of safety violations whereas MoDem-v2 exhibits much superior performance. When comparing between ablations it is clear that each individual modification has both benefits and drawbacks. First, note that *Centering* is a necessary sub-component of *Schedule*. While *Schedule* is generally safer than *Ensemble*, *Ensemble* typically has better sample-efficiency. By combining these two ingredients, MoDem-V2 is able to achieve the improved sample-efficiency of *Ensemble* while maintaining the improved safety profile of *Schedule*.

C. Real World Results

As suggested by our simulation experiments, the original MoDem is unsafe for learning real-world manipulation tasks. When we did attempt to run MoDem on our real robot, we found that its aggressive exploration frequently violated

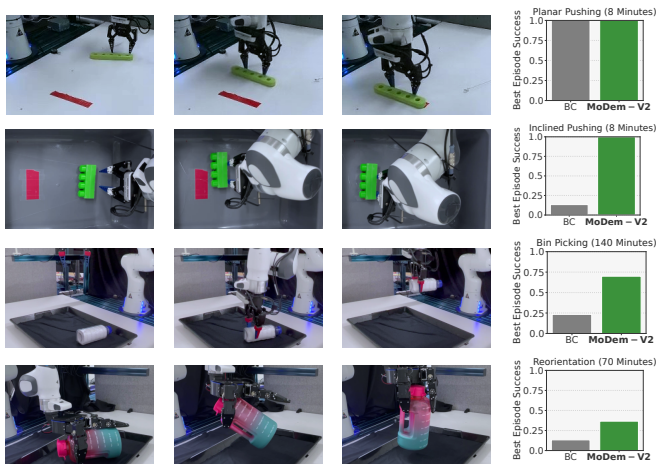


Fig. 6. Left: Example rollouts from our MoDem-V2 agents on real-world manipulation tasks. Right: The success rate of our best performing policy and its initial BC policy. MoDem-V2 is able to outperform its BC policy in hours or less.

safety limits at the beginning of online interaction. For example, on the inclined pushing task, the robot triggered a safety fault within the first two exploration episodes because it exerted excessive force/torque on the incline. Due to the unsafe behavior MoDem induces and the significant human intervention that would be required, it is not feasible to evaluate MoDem in the real world.

In contrast, MoDem-V2 is capable of safely learning manipulation tasks with minimal human intervention. MoDem-V2 enabled the robot to significantly exceed the performance of its initial BC policy with about two hours worth of online training data or less. Figure 6 shows the success rate of the initial policy cloned from just ten demonstrations and the best MoDem-V2 agent performance achieved throughout online training, as well as example trajectories from the MoDem-V2 agents. Please see sites.google.com/view/modem-v2 for additional videos and results.

VII. RELATED WORK

Visual MBRL: Improving sample-efficiency of visual RL by learning a model of the environment has been explored extensively in literature [20], [5], [6], [21], [22], [23], [7]. Here we focus on MBRL algorithms that leverage planning. Prior work typically learns a latent dynamics model from online interaction, and uses a sampling-based planning technique for action selection with candidate action sequences evaluated by the learned model. For continuous control, planning can be formalized as Model Predictive Control (MPC) [20], [6], [7], [9], whereas Monte-Carlo Tree Search (MCTS) is used for discrete action spaces [22], [23]. Regardless, the majority of work on visual MBRL focus on sample-efficiency in simulated tasks, where practicality and safety are of limited concern. Our work extends the MBRL algorithm of [7], [9] which has already been shown to be very sample-efficient in simulation, and instead focus on the challenges that arise when training MBRL in the real world.

Safe Reinforcement Learning: The field of safe RL encompasses a wide range of approaches; see [24] for a comprehensive review. A common framework for safe RL is to represent the task as a constrained markov decision process [25], [26], [27], but partial observability presents a challenge to applying such approaches to real world environments. Other methods encode safety as robustness through either domain randomization [28], [29] or adversarial perturbation [30], [31], [32]. Our work aligns with those previous that use ensemble methods to estimate model uncertainty for guiding the agent towards safer exploration [33], [34], [35]. Similar to our work, Thananjeyan et. al [34] uses demonstrations to limit policy exploration to be near known safe trajectories. However their method requires the user to provide a function indicating whether a given robot state is safe or not, which can be difficult to specify for high-dimensional or partially observed state spaces. In this work, we focus on proposing solutions for safe exploration that can be deployed on to real robots and do not diminish the high sample efficiency of the original MoDem.

Real-World Robot Learning: Researchers have explored a wide variety of approaches for robot learning on real hardware, most of which fall into one of three categories: learning from human demonstrations [36], [37], [38], [39], learning from large uncurated datasets [40], [41], [20], learning from online interaction via RL [42], [43], or any combination of them [17], [44], [45], [46]. Our work is most similar to Zhan et al. [17] in problem setting (RL with demonstrations) and experimental setup (robotic manipulation tasks in the real world). However, we focus on the unique challenges and opportunities of MBRL for real-world robot learning. Wu et al. [43] study real world MBRL, but consider simpler tasks with limited variation and do not leverage demonstrations.

VIII. DISCUSSION

In this work, we tackled the challenge of learning manipulation skills in the real world from only proprioceptive and visual feedback with sparse rewards. We developed MoDem-V2, a real-world ready adaptation of MoDem, by proposing to initially center rollouts around the BC policy, gradually increase the proportion of actions chosen by the learned world model, and implement uncertainty aware planning with actor-critic ensembles. We evaluated the sample-efficiency and safety of MoDem-V2 against strong baselines in simulation and found that it maintained the high sample-efficiency of MoDem while exhibiting significantly safer behavior through lower contact force exertion. We found that MoDem-V2 enabled a real, physical robot to learn a variety of manipulation skills, such as pushing, picking, and in-hand manipulation, from an hour or less worth of interaction data.

Limitations: One limitation of our work is that it requires a small number of demonstrations, which may not always be available or easy to collect. Also, this work assumes that the environment can be reset to a narrow set of starting states, which may not always be the case in the real-world. In future work, we hope to explore the reuse of our learned world model across changes in manipulated object and task goal.

REFERENCES

- [1] A. A. Rusu, M. Vecerík, T. Rothörl, N. M. O. Heess, R. Pascanu, and R. Hadsell, "Sim-to-real robot learning from pixels with progressive nets," *ArXiv*, vol. abs/1610.04286, 2016.
- [2] J. Tobin, R. Fong, A. Ray, J. Schneider, W. Zaremba, and P. Abbeel, "Domain randomization for transferring deep neural networks from simulation to the real world," *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 23–30, 2017.
- [3] A. Kumar, Z. Fu, D. Pathak, and J. Malik, "Rma: Rapid motor adaptation for legged robots," *ArXiv*, vol. abs/2107.04034, 2021.
- [4] A. Handa, A. Allshire, V. Makoviychuk, A. Petrenko, R. Singh, J. Liu, D. Makoviichuk, K. V. Wyk, A. Zhurkevich, B. Sundaralingam, Y. S. Narang, J.-F. Lafleche, D. Fox, and G. State, "Dextreme: Transfer of agile in-hand manipulation from simulation to reality," *ArXiv*, vol. abs/2210.13702, 2022.
- [5] D. Ha and J. Schmidhuber, "Recurrent world models facilitate policy evolution," in *Advances in Neural Information Processing Systems 31*. Curran Associates, Inc., 2018, pp. 2451–2463.
- [6] D. Hafner, T. Lillicrap, F. Fischer, R. Villegas, D. Ha, H. Lee, and J. Davidson, "Learning latent dynamics for planning from pixels," in *International Conference on Machine Learning*, 2019, pp. 2555–2565.
- [7] N. Hansen, X. Wang, and H. Su, "Temporal difference learning for model predictive control," *arXiv preprint arXiv:2203.04955*, 2022.
- [8] R. M. Shah and V. Kumar, "Rrl: Resnet as representation for reinforcement learning," in *International Conference on Machine Learning*. PMLR, 2021, pp. 9465–9476.
- [9] N. Hansen, Y. Lin, H. Su, X. Wang, V. Kumar, and A. Rajeswaran, "Modem: Accelerating visual model-based reinforcement learning with demonstrations," *arXiv preprint arXiv:2212.05698*, 2022.
- [10] C. G. Atkeson and S. Schaal, "Robot learning from demonstration," in *ICML*, 1997.
- [11] J.-B. Grill, F. Strub, F. Altché, C. Tallec, P. H. Richemond, E. Buchatskaya, C. Doersch, B. Á. Pires, Z. D. Guo, M. G. Azar, B. Piot, K. Kavukcuoglu, R. Munos, and M. Valko, "Bootstrap your own latent: A new approach to self-supervised learning," *ArXiv*, vol. abs/2006.07733, 2020.
- [12] Z. Huang, S. Zhou, B. Zhuang, and X. Zhou, "Learning to run with actor-critic ensemble," *arXiv preprint arXiv:1712.08987*, 2017.
- [13] S. Fujimoto, H. Hoof, and D. Meger, "Addressing function approximation error in actor-critic methods," in *International conference on machine learning*. PMLR, 2018, pp. 1587–1596.
- [14] K. Chua, R. Calandra, R. McAllister, and S. Levine, "Deep reinforcement learning in a handful of trials using probabilistic dynamics models," *Advances in neural information processing systems*, vol. 31, 2018.
- [15] "Robohive – a unified framework for robot learning," <https://sites.google.com/view/robohive>, 2020. [Online]. Available: <https://sites.google.com/view/robohive>
- [16] A. Rajeswaran, V. Kumar, A. Gupta, G. Vezzani, J. Schulman, E. Todorov, and S. Levine, "Learning Complex Dexterous Manipulation with Deep Reinforcement Learning and Demonstrations," in *Proceedings of Robotics: Science and Systems (RSS)*, 2018.
- [17] A. Zhan, P. Zhao, L. Pinto, P. Abbeel, and M. Laskin, "A framework for efficient robotic manipulation," *ArXiv*, vol. abs/2012.07975, 2020.
- [18] R. Bhirangi, A. DeFranco, J. Adkins, C. Majidi, A. Gupta, T. Hellebrekers, and V. Kumar, "All the feels: A dexterous hand with large area sensing," *arXiv preprint arXiv:2210.15658*, 2022.
- [19] M. Ahn, H. Zhu, K. Hartikainen, H. Ponte, A. Gupta, S. Levine, and V. Kumar, "Robel: Robotics benchmarks for learning with low-cost robots," in *Conference on robot learning*. PMLR, 2020, pp. 1300–1313.
- [20] F. Ebert, C. Finn, S. Dasari, A. Xie, A. X. Lee, and S. Levine, "Visual foresight: Model-based deep reinforcement learning for vision-based robotic control," *ArXiv*, vol. abs/1812.00568, 2018.
- [21] L. Kaiser, M. Babaeizadeh, P. Milos, B. Osinski, R. H. Campbell, K. Czechowski, D. Erhan, C. Finn, P. Kozakowski, S. Levine, R. Sepassi, G. Tucker, and H. Michalewski, "Model-based reinforcement learning for atari," *ArXiv*, vol. abs/1903.00374, 2020.
- [22] J. Schrittwieser, I. Antonoglou, T. Hubert, K. Simonyan, L. Sifre, S. Schmitt, A. Guez, E. Lockhart, D. Hassabis, T. Graepel, T. P. Lillicrap, and D. Silver, "Mastering atari, go, chess and shogi by planning with a learned model," *Nature*, vol. 588 7839, pp. 604–609, 2020.
- [23] W. Ye, S. Liu, T. Kurutach, P. Abbeel, and Y. Gao, "Mastering atari games with limited data," in *NeurIPS*, 2021.
- [24] L. Brunke, M. Greeff, A. W. Hall, Z. Yuan, S. Zhou, J. Panerati, and A. P. Schoellig, "Safe learning in robotics: From learning-based control to safe reinforcement learning," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 5, pp. 411–444, 2022.
- [25] Y. Ge, F. Zhu, X. Ling, and Q. Liu, "Safe q-learning method based on constrained markov decision processes," *IEEE Access*, vol. 7, pp. 165 007–165 017, 2019.
- [26] Y. Chow, O. Nachum, E. Duenez-Guzman, and M. Ghavamzadeh, "A lyapunov-based approach to safe reinforcement learning," *Advances in neural information processing systems*, vol. 31, 2018.
- [27] A. Wachi, Y. Sui, Y. Yue, and M. Ono, "Safe exploration and optimization of constrained mdps using gaussian processes," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 32, no. 1, 2018.
- [28] J. Tobin, R. Fong, A. Ray, J. Schneider, W. Zaremba, and P. Abbeel, "Domain randomization for transferring deep neural networks from simulation to the real world," in *2017 IEEE/RSJ international conference on intelligent robots and systems (IROS)*. IEEE, 2017, pp. 23–30.
- [29] G. D. Kontes, D. D. Scherer, T. Nisslbeck, J. Fischer, and C. Mutschler, "High-speed collision avoidance using deep reinforcement learning and domain randomization for autonomous vehicles," in *2020 IEEE 23rd international conference on Intelligent Transportation Systems (ITSC)*. IEEE, 2020, pp. 1–8.
- [30] B. Mehta, M. Diaz, F. Golemo, C. J. Pal, and L. Paull, "Active domain randomization," in *Conference on Robot Learning*. PMLR, 2020, pp. 1162–1176.
- [31] B. Yang, G. Habibi, P. Lancaster, B. Boots, and J. Smith, "Motivating physical activity via competitive human-robot interaction," in *Conference on Robot Learning*. PMLR, 2022, pp. 839–849.
- [32] B. Yang, L. Zheng, L. J. Ratliff, B. Boots, and J. R. Smith, "Stackelberg games for learning emergent behaviors during competitive autocurricula," *arXiv preprint arXiv:2305.03735*, 2023.
- [33] J. Zhang, B. Cheung, C. Finn, S. Levine, and D. Jayaraman, "Cautious adaptation for reinforcement learning in safety-critical settings," in *International Conference on Machine Learning*. PMLR, 2020, pp. 11 055–11 065.
- [34] B. Thananjeyan, A. Balakrishna, U. Rosolia, F. Li, R. McAllister, J. E. Gonzalez, S. Levine, F. Borrelli, and K. Goldberg, "Safety augmented value estimation from demonstrations (saved): Safe deep model-based rl for sparse cost robotic tasks," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 3612–3619, 2020.
- [35] M. Henaff, A. Canziani, and Y. LeCun, "Model-predictive policy learning with uncertainty regularization for driving in dense traffic," *arXiv preprint arXiv:1901.02705*, 2019.
- [36] E. Jang, A. Irpan, M. Khansari, D. Kappler, F. Ebert, C. Lynch, S. Levine, and C. Finn, "Bc-z: Zero-shot task generalization with robotic imitation learning," in *Conference on Robot Learning*, 2022.
- [37] S. Nair, A. Rajeswaran, V. Kumar, C. Finn, and A. Gupta, "R3m: A universal visual representation for robot manipulation," in *Conference on Robot Learning*, 2022.
- [38] Y. Zhu, A. Joshi, P. Stone, and Y. Zhu, "Viola: Imitation learning for vision-based manipulation with object proposal priors," *ArXiv*, vol. abs/2210.11339, 2022.
- [39] A. Brohan, N. Brown, J. Carbajal, Y. Chebotar, J. Dabis, C. Finn, K. Gopalakrishnan, K. Hausman, A. Herzog, J. Hsu, J. Ibarz, B. Ichter, A. Irpan, T. Jackson, S. Jesmonth, N. J. Joshi, R. C. Julian, D. Kalashnikov, Y. Kuang, I. Leal, K.-H. Lee, S. Levine, Y. Lu, U. Malla, D. Manjunath, I. Mordatch, O. Nachum, C. Parada, J. Peralta, E. Perez, K. Pertsch, J. Quiambao, K. Rao, M. S. Ryoo, G. Salazar, P. R. Sanketi, K. Sayed, J. Singh, S. A. Sontakke, A. Stone, C. Tan, H. Tran, V. Vanhoucke, S. Vega, Q. H. Vuong, F. Xia, T. Xiao, P. Xu, S. Xu, T. Yu, and B. Zitkovich, "Rt-1: Robotics transformer for real-world control at scale," *ArXiv*, vol. abs/2212.06817, 2022.
- [40] L. Pinto and A. K. Gupta, "Supersizing self-supervision: Learning to grasp from 50k tries and 700 robot hours," *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 3406–3413, 2015.
- [41] S. Levine, P. Pastor, A. Krizhevsky, and D. Quillen, "Learning hand-eye coordination for robotic grasping with deep learning and large-scale data collection," *The International Journal of Robotics Research*, vol. 37, pp. 421 – 436, 2016.
- [42] H. Zhu, A. Gupta, A. Rajeswaran, S. Levine, and V. Kumar, "Dexterous manipulation with deep reinforcement learning: Efficient, general, and low-cost," *2019 International Conference on Robotics and Automation (ICRA)*, pp. 3651–3657, 2018.

- [43] 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, pp. 2226–2240.
- [44] R. C. Julian, B. Swanson, G. S. Sukhatme, S. Levine, C. Finn, and K. Hausman, “Efficient adaptation for end-to-end vision-based robotic manipulation,” *ArXiv*, vol. abs/2004.10190, 2020.
- [45] D. Kalashnikov, J. Varley, Y. Chebotar, B. Swanson, R. Jonschkowski, C. Finn, S. Levine, and K. Hausman, “Mt-opt: Continuous multi-task robotic reinforcement learning at scale,” *ArXiv*, vol. abs/2104.08212, 2021.
- [46] A. Kumar, A. Singh, F. Ebert, Y. Yang, C. Finn, and S. Levine, “Pre-training for robots: Offline rl enables learning new tasks from a handful of trials,” *ArXiv*, vol. abs/2210.05178, 2022.