

PROGRESSOR: A Perceptually Guided Reward Estimator with Self-Supervised Online Refinement

Tewodros W. Ayalew¹ Xiao Zhang^{*,1} Kevin Yuanbo Wu^{*,1} Tianchong Jiang²
Michael Maire¹ Matthew R. Walter²

¹University of Chicago ²Toyota Technological Institute at Chicago

<https://ripl.github.io/progressor>

Abstract

We present PROGRESSOR, a novel framework that learns a task-agnostic reward function from videos, enabling policy training through goal-conditioned reinforcement learning (RL) without manual supervision. Underlying this reward is an estimate of the distribution over task progress as a function of the current, initial, and goal observations that is learned in a self-supervised fashion. Crucially, PROGRESSOR refines rewards adversarially during online RL training by pushing back predictions for out-of-distribution observations, to mitigate distribution shift inherent in non-expert observations. Utilizing this progress prediction as a dense reward together with an adversarial push-back, we show that PROGRESSOR enables robots to learn complex behaviors without any external supervision. Pretrained on large-scale egocentric human video from EPIC-KITCHENS, PROGRESSOR requires no fine-tuning on in-domain task-specific data for generalization to real-robot offline RL under noisy demonstrations, outperforming contemporary methods that provide dense visual reward for robotic learning. Our findings highlight the potential of PROGRESSOR for scalable robotic applications where direct action labels and task-specific rewards are not readily available.

1. Introduction

Practical applications of reinforcement learning (RL) require that a domain expert design complex reward models that encourage desired behavior while simultaneously penalizing unwanted behavior [32, 35, 36]. However, manually constructing dense rewards that enable effective learning is difficult and can yield undesired behaviors [5, 15]. Sparse rewards require far less effort, however the lack of supervision typically results in a significant loss of sample efficiency. These issues are exacerbated for long-horizon

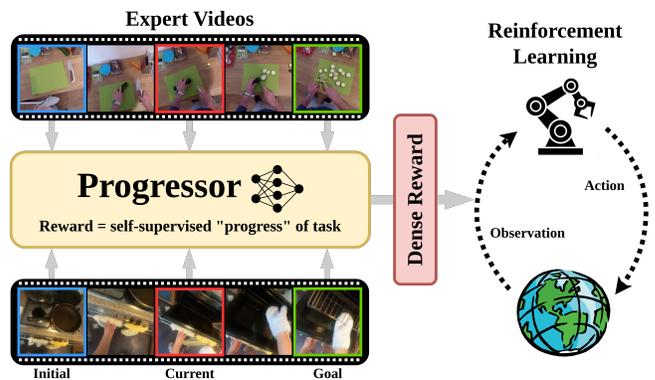


Figure 1. Trained in a self-supervised manner on expert videos, PROGRESSOR predicts an agent’s progress toward task completion, providing a reward signal for reinforcement learning. During online reinforcement learning, we employ an adversarial technique to refine this reward estimate, addressing the distribution shift between expert data and non-expert online rollouts.

tasks, where credit assignment is particularly difficult.

A promising alternative is to learn effective rewards from unlabeled (i.e., action-free) videos of task demonstrations [2, 8, 20, 24, 27, 40]. Large-scale video data is readily available on the internet, providing an effective alternative to costly, intensive task-specific data collection. An effective way to incorporate diverse video data into a *single* model without learning a new policy for each task is to condition the model on a goal image that visually specifies the desired environment changes upon task completion [9, 24].

To that end, we propose PROGRESSOR, a *self-supervised* framework that learns a *task-agnostic* reward function from videos for *goal-conditioned reinforcement learning* that only requires image observations, without the need for corresponding action labels. Following Yang et al. [40], we assume that expert demonstrations make monotonic progress towards a goal and employ a self-supervised temporal learning approach that utilizes a proxy objective aimed at esti-

mating task completion progress. PROGRESSOR predicts the distribution of progress from the current observation (i.e., the current image) relative to the initial and goal observations. This estimation provides dense supervision of an agent’s progression towards reaching the goal, and thus performing the task, guiding exploration in alignment with expert execution trajectories. However, during RL exploration, agents will often encounter states that were not visited as part of the expert trajectories. PROGRESSOR accounts for this distribution shift via an online adversarial refinement that pushes back pushing back predictions for out-of-distribution observations.

We evaluate the effectiveness of PROGRESSOR for robot manipulation tasks in both simulation and the real world. These results demonstrate the the benefits that PROGRESSOR’s self-supervised reward has on policy learning, particularly in settings where existing methods otherwise require intricate reward design. The key contributions of this paper are as follows:

- We present PROGRESSOR, a self-supervised reward model trained on unlabeled videos that guides a reinforcement learning agent by providing dense rewards based on predicted progress toward a goal.
- PROGRESSOR achieves state-of-the-art performance on six diverse tasks from the Meta-World benchmark [43] without the need for environment reward.
- Pretrained on large-scale egocentric human video from EPIC-KITCHENS [6], PROGRESSOR enables few-shot real-robot offline, *even when half of the demonstrations are unsuccessful*.

2. Related Work

Learning from expert videos has gained traction due to the potential to scale the acquisition of complex skills using the vast amount of video content readily available.

Inverse RL Inverse reinforcement learning (IRL) is a framework for learning reward functions from expert demonstrations as an alternative to specifying them manually. In contrast to traditional RL, where the reward function is predefined and the goal is to learn an optimal policy, IRL focuses on deducing the underlying reward structure that results in the expert’s behavior [1, 39, 46]. Recent IRL methods leverage adversarial techniques to learn reward functions [10, 11, 22], such as in Generative Adversarial Imitation Learning (GAIL) [19]. A key limitation of typical approaches to IRL is that they require action-labeled demonstrations, which are only available in small quantities. Another is the ambiguity that results from having multiple rewards that can explain the observed behavior [25], potentially leading to ill-shaped reward functions. In contrast, our notion of progress enables the learning of well-shaped dense rewards.

Imitation learning from video Several methods [4, 44]

learn from videos by inferring the inverse dynamics model, leveraging a small amount of annotated data. However, these methods depend on the presence and quality of an albeit limited action labels. Other methods have used existing hand pose estimation as intermediate annotation followed by standard imitation learning [3, 30]. However, these methods rely on the accuracy of the pose estimators and camera calibration.

Learning reward from video Recent methods learn visual representations useful for RL from video data using self-supervised temporal contrastive techniques [9, 24, 27, 34, 41]. Time-contrastive objectives treat observations close in time as positive pairs and those that are temporally distant as negative pairs. They then learn embeddings that encourage small distances between positive pairs and large distances between negative pairs. The learned embeddings can then be used to define reward functions, i.e., as the distance from the goal image in embedding space. Time-contrastive objectives can be sensitive to frame rate, and extracted rewards are assumed to be symmetric [24]. In comparison, our approach is agnostic to the sampling rate and can learn asymmetric dynamics depending on the specified initial and goal images.

Additionally, generative approaches have been used for reward learning. VIPER [8] estimates the log-likelihood of an observation as a reward function by training on an expert dataset using a VideoGPT-based autoregressive model. Similarly, Huang et al. [20] employ a diffusion-based approach to learn rewards by leveraging conditional entropy. However, these methods necessitate costly generative processes to estimate rewards.

Particularly relevant, Rank2Reward [40] demonstrates that learning the ranking of visual observations from a demonstration helps to infer well-shaped reward functions. By integrating ranking techniques with the classification of expert demonstratin data vs. non-expert data obtained through on-policy data collection, Rank2Reward guides robot learning from demonstration. A key limitation of the method is that it necessitates training distinct models for each task due to the task-dependent nature of ranking frames. In contrast, our approach employs a *single* reward model for all tasks. Additionally, we utilize an adversarial training approach that tackles the domain shift observed during online RL, rather than incorporating a separate classifier as used in Rank2Reward.

3. Preliminaries

We formulate our method in the context of a finite-horizon Markov decision process (MDP) \mathcal{M} defined by the tuple $(\mathcal{O}, \mathcal{A}, \mathcal{P}, \mathcal{R}, \gamma, \rho_0)$, which represents the observation space, action space, reward function, transition probability function, discount factor, initial state distribution, respectively. Typically, the agent starts from an initial state with

observation $\mathbf{o}_0 \sim \mathcal{O}$ and at each step performs an action $a \sim \mathcal{A}$. The fundamental objective in RL is to learn optimal policy π^* that maximizes the expected discounted reward, i.e., expressed as $\pi^* = \operatorname{argmax}_{\pi} \mathbb{E}_{\pi_{\theta}} [\sum_{t=1}^T (\gamma^{t-1} \mathcal{R}(\mathbf{o}_t))]$.

In this paper, we operate under the assumption that the environment is fully observable, wherein observations $\mathbf{o}_k \in \mathbb{R}^{H \times W \times 3}$ are high-dimensional image frames. Additionally, we assume access to a dataset \mathcal{D}_e that consists of expert demonstrations. These demonstrations are comprised of sets of trajectories $\{\tau_k\}_{k=1}^N$, each containing sequences of observations, $\tau_k = \{\mathbf{o}_0^{\tau_k}, \mathbf{o}_1^{\tau_k}, \dots, \mathbf{o}_n^{\tau_k}\}$.

However, we assume that neither the reward function \mathcal{R} nor the demonstrated actions $a \sim \mathcal{A}$ are known. This setup mirrors the real-world robot task-learning scenario in which we have a plethora of videos of human-performed tasks, without knowledge of the actions being performed or the rewards being received. The primary objective of this paper is to construct a model capable of predicting a well-structured dense reward from a dataset of expert demonstrations, without access to demonstrated actions or rewards. This model can subsequently serve as a reward estimating mechanism when training an agent using online RL in the absence of any environment rewards.

4. Method

We propose to learn a unified reward model via an encoder $E_{\theta}(\mathbf{o}_i, \mathbf{o}_j, \mathbf{o}_g)$ that estimates the relative progress of an observation \mathbf{o}_j with respect to an initial observation \mathbf{o}_i and a goal observation \mathbf{o}_g , all of which are purely pixel-based. Figure 2 illustrates PROGRESSOR’s framework for task progress estimation. Here, the objective is to learn the temporal ordering of observations that lead to the correct execution of a task. By learning this progress estimation, we can create a reward model that incentivizes progress toward the goal. Similar to Yang et al. [40], our method relies on the assumption that the values of states in optimal policies increase monotonically towards task completion.

4.1. Learning the Self-Supervised Reward Model

We measure an observation \mathbf{o}_j ’s progress towards a goal in terms of its relative position with respect to the initial observation \mathbf{o}_i and the goal observation \mathbf{o}_g . In order to learn this progress estimate, we first represent the relative position of an observation \mathbf{o}_j within an expert trajectory τ_k as

$$\delta(\mathbf{o}_i^{\tau_k}, \mathbf{o}_j^{\tau_k}, \mathbf{o}_g^{\tau_k}) = \frac{|j - i|}{|g - i|}. \quad (1)$$

Here, $\delta(\mathbf{o}_i^{\tau_k}, \mathbf{o}_j^{\tau_k}, \mathbf{o}_g^{\tau_k})$ calculates the ratio of the relative frame position differences between \mathbf{o}_j and \mathbf{o}_i and between \mathbf{o}_g and \mathbf{o}_i within a given expert trajectory τ_k . This progress label is always such that $\delta(\mathbf{o}_i^{\tau_k}, \mathbf{o}_j^{\tau_k}, \mathbf{o}_g^{\tau_k}) \in [0, 1]$.

Predicting progress based on real-world video is challenging due to factors like camera motion and repetitive

sub-trajectories. We address this by formulating the problem as one of estimating a Gaussian distribution over progress, allowing the model to estimate uncertainty as variance. With this approach, we approximate the ground-truth distribution for a triplet of frames $(\mathbf{o}_i^{\tau_k}, \mathbf{o}_j^{\tau_k}, \mathbf{o}_g^{\tau_k})$ within a video τ_k as a normal distribution with mean $\mu_{\tau_k} = \delta(\mathbf{o}_i^{\tau_k}, \mathbf{o}_j^{\tau_k}, \mathbf{o}_g^{\tau_k})$ and standard deviation $\sigma_{\tau_k} = \max(\frac{1}{g-i}, \epsilon)$:

$$p_{\text{target}}(\mathbf{o}_i^{\tau_k}, \mathbf{o}_j^{\tau_k}, \mathbf{o}_g^{\tau_k}) = \mathcal{N}(\mu_{\tau_k}, \sigma_{\tau_k}^2) \quad (2)$$

Our ϵ upper-bound of σ_{τ_k} downweights the penalty for triplets in which there are many frames between the initial and goal images, which we empirically find to improve robustness during training. Therefore, the predicted distribution of this form provides a distribution of the progress estimate. Sampling from this predicted distribution yields values that reflect the observed progress.

We optimize our reward model r_{θ} to predict the distribution of the progress on expert trajectory. We use a shared visual encoder to compute the per-frame representation, followed by several MLPs to produce the final estimation:

$$E_{\theta}(\mathbf{o}_i, \mathbf{o}_j, \mathbf{o}_g) = \mathcal{N}(\mu, \sigma^2) \quad (3)$$

Following Kingma and Welling [21], we estimate $\log \sigma^2$ and then compute σ as $\sigma = \exp(0.5 \log \sigma^2)$. We learn the parameters of the network by optimizing the Kullback-Leibler divergence loss between the ground-truth and the predicted distributions:

$$\mathcal{L}_{\text{expert}} = D_{\text{KL}}(p_{\text{target}}(\mathbf{o}_i^{\tau_k}, \mathbf{o}_j^{\tau_k}, \mathbf{o}_g^{\tau_k}) \| E_{\theta}(\mathbf{o}_i^{\tau_k}, \mathbf{o}_j^{\tau_k}, \mathbf{o}_g^{\tau_k})) \quad (4)$$

During the training of our reward model r_{θ} , we randomly select and rank three frames from a video sequence of expert trajectories as a training triplet. Additionally, since these triplets can be randomly drawn from any trajectory within the dataset of different robotics tasks, a single reward model suffices to handle a variety of tasks.

Using the Reward Model in Online RL: Due to the monotonic nature of task progress in our problem formulation, the estimates derived from the trained model can serve as dense rewards for training, replacing the unknown true reward. We create the reward model by defining a function derived from the model’s predicted outputs given a sample of frame triplet $(\mathbf{o}_i, \mathbf{o}_j, \mathbf{o}_g)$ of trajectory as:

$$r_{\theta}(\mathbf{o}_i, \mathbf{o}_j, \mathbf{o}_g) = \mu - \alpha \mathcal{H}(\mathcal{N}(\mu, \sigma^2)), \quad (5)$$

where r_{θ} denotes PROGRESSOR’s reward estimate and \mathcal{H} is the entropy. The first term μ in Eqn. 5 uses the mean prediction from the progress estimation model to determine the position of the current observation relative to the goal. The second term penalizes observations that yield high variance. Here, α is hyperparameter, which we set to $\alpha = 0.4$ in all our experiments. See S8 for an ablation of α .

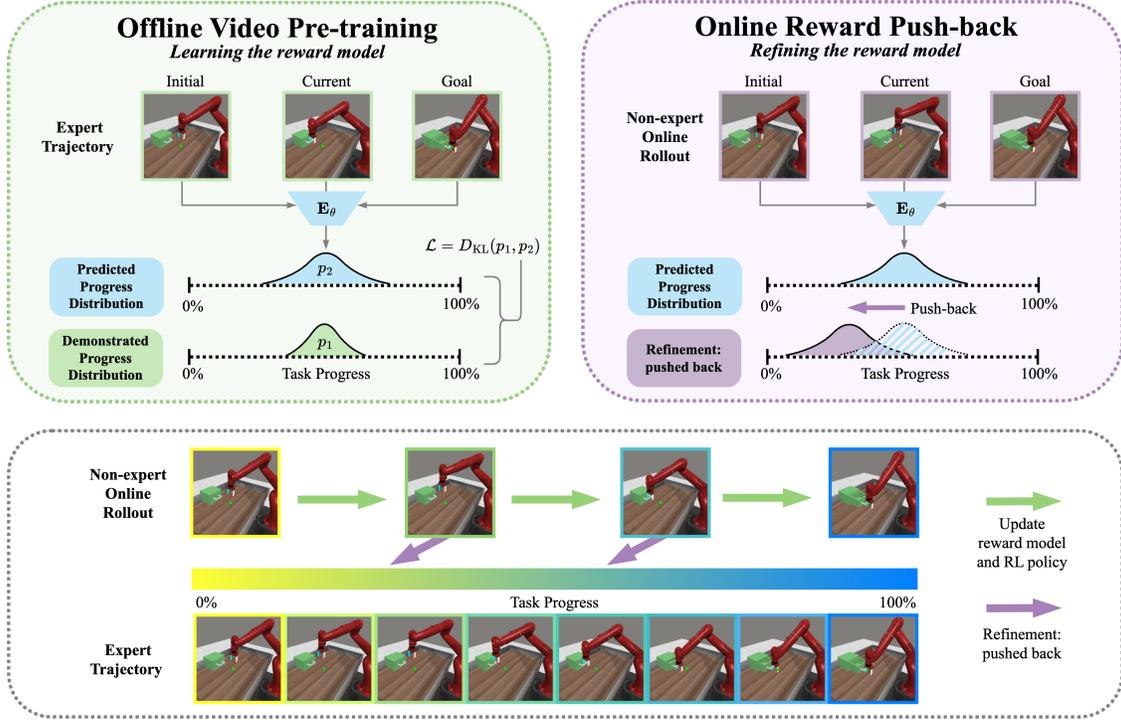


Figure 2. **Top Left:** Initial phase of reward model pretraining on expert data, where the model learns to predict the parameters of a Gaussian distribution centered on normalized progress, reflecting expected progress as demonstrated by experts. **Top Right:** In online reinforcement learning (RL) training, an adversarial online refinement (i.e., push-back) is applied to counteract non-expert predictions made by the reward model, effectively distinguishing expert from non-expert progress. **Bottom:** During online RL, the reward model is updated on expert and non-expert data.

4.2. Adversarial Online Refinement via Push-Back

Our reward model r_θ can be used to train a policy that maximizes the expected sum of discounted rewards as:

$$\pi^* = \arg \max_{\pi} \mathbb{E}_{\pi_\theta} \left[\sum_{t=1}^T \gamma^{t-1} r_\theta(\mathbf{o}_i^{\tau'}, \mathbf{o}_t^{\tau'}, \mathbf{o}_g^{\tau'}) \right] \quad (6)$$

The reward encourages the agent to take actions that align its observed trajectories τ' with those of the expert demonstrations τ , thereby promoting task completion. However, the all-but-random actions typical of learned policies early in training result in out-of-distribution observations relative to the expert trajectories. As a result, optimizing the policy using a reward r_θ that is only trained on an expert trajectory τ can result in unintended behaviors, since r_θ does not reflect the actual progress of τ' . To tackle this distribution shift, we implement an adversarial online refinement strategy [12, 13], which we refer to as “push-back”, that enables the reward model r_θ to differentiate between in- and out-of-distribution observations τ and τ' , providing better estimation. Specifically, for a frame triplet $(\mathbf{o}_i^{\tau'_k}, \mathbf{o}_j^{\tau'_k}, \mathbf{o}_g^{\tau'_k})$ sampled from τ'_k and the estimated progress $\mu_{\tau'_k}$ from E_θ (3), we update E_θ so that it learns to *push-back* the current

Algorithm 1 Training PROGRESSOR with Online RL

Require: Expert demonstration data $\mathcal{D}_e = \{\tau_k\}_{k=1}^N$

- 1: Initialize policy π , empty replay buffer \mathcal{D}_{RB}
 - 2: Initialize reward model r_θ .
 - 3: // Pretraining PROGRESSOR
 - 4: **for** step n in $\{1, \dots, N_{\text{pretrain}}\}$ **do**
 - 5: Sample triplets $(\mathbf{o}_i^{\tau_k}, \mathbf{o}_j^{\tau_k}, \mathbf{o}_g^{\tau_k})$ from $\tau_k \in \mathcal{D}_e$
 - 6: Learn r_θ with objective Eqn. 4
 - 7: **end for**
 - 8: // Policy optimization and reward model fine-tuning
 - 9: **for** step n in $\{1, \dots, N\}$ **do**
 - 10: With π , collect transitions $\{\tau_l\}_{l=1}^M$ and store in \mathcal{D}_{RB}
 - 11: **if** $n \% \text{reward_update_frequency} == 0$ **then**
 - 12: Update r_θ with Eqn. 4 using samples from \mathcal{D}_e
 - 13: Update r_θ with Eqn. 7 using samples from \mathcal{D}_{RB}
 - 14: **end if**
 - 15: Sample batch of transitions s_π from \mathcal{D}_{RB}
 - 16: Update π to maximize returns using Eqn. 5
 - 17: **end for**
-

estimation as $\beta \mu_{\tau'_k}$ with $\beta \in [0, 1]$ as the decay factor. In our experiments, we use $\beta = 0.9$. We formulate our online

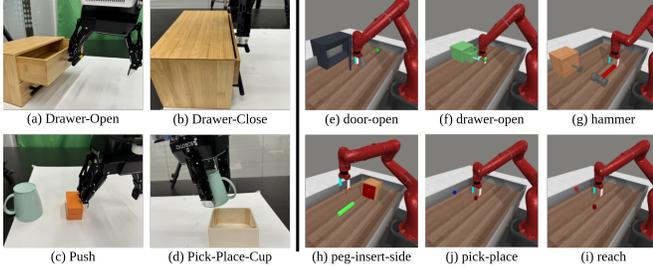


Figure 3. Visualization of the robotic tasks: (a-d) Real world environments with a UR5 arm. (e-j) Simulation environments for evaluation using the Meta-World [43] benchmark.

refinement objectives as:

$$p_{\text{push-back}} = \text{sg} \left(\mathcal{N} \left(\beta \mu_{\tau'_k}, \frac{1}{(g-i)^2} \right) \right)$$

$$\mathcal{L}_{\text{push-back}} = D_{\text{KL}} \left(p_{\text{push-back}} \parallel \mathbf{E}_{\theta} \left(\mathbf{o}_i^{\tau'_k}, \mathbf{o}_j^{\tau'_k}, \mathbf{o}_g^{\tau'_k} \right) \right), \quad (7)$$

where $\text{sg}(\cdot)$ denotes the stop gradient operator. With this approach, we differentiate between the two sources of trajectories by treating observations from the (inferior) online roll-out trajectories τ' as corresponding to lower progress compared to expert trajectories τ . In this case, optimizing for π^* would encourage the agent to consistently improve its actions until the policy trajectories are well aligned with the distribution of expert trajectories τ_k .

During online training, we fine-tune \mathbf{E}_{θ} (3) using hybrid objectives: we optimize $\mathcal{L}_{\text{push-back}}$ based on τ'_k sampled from the most recent buffers and at the same time, we optimize $\mathcal{L}_{\text{expert}}$ using samples from τ_k , preventing \mathbf{E}_{θ} from being biased by τ'_k .

5. Experimental Evaluation

We evaluate the effectiveness with which PROGRESSOR learns reward functions from visual demonstrations that enable robots to perform various manipulation tasks in simulation as well as the real world. The results demonstrate noticeable improvements over the previous state-of-the-art.

5.1. Simulated Experiments

In our simulated experiments, we used benchmark tasks from the Meta-World environment [43], selecting six table-top manipulation tasks (Figure 3 e-j): `door-open`, `drawer-open`, `hammer`, `peg-insert-side`, `pick-place`, and `reach`. For all tasks, we trained the model-free reinforcement learning algorithm DrQ-v2 [42] using the reward predicted by PROGRESSOR in place of the environment reward.

We compare our method to three relevant baselines: (1) **TCN** [33], a temporal contrastive method that optimizes embeddings for measuring task progress; (2) **GAIL** [18],

an adversarial approach that aligns the state-action distribution of a policy with that of an expert. GAIL’s discriminator is used as a reward function for our baseline following Yang et al. [40]; (3) **SOIL** [31] jointly trains an inverse dynamics model and policy by predicting actions from state-only demonstrations.; (4) **ROT** [16] leverages optimal transport to match the expert and agent’s trajectory to provide extra regularization for robotics learning; and (5) **Rank2Reward** [40], a state-of-the-art reward learning method that combines temporal ranking and classification of expert demonstrations to estimate rewards. Unlike our method, which uses a single model for all tasks, Rank2reward requires a separate model for each task and needs an additional classifier for online adaptation. For a fair comparison, we use their default hyperparameters and trained all methods under the same setup. All methods were trained for 1.5M steps, relying solely on the reward estimates provided by the model.

We evaluate the methods in terms of the episodic return, which reflects the cumulative reward received by the agent during each episode of the task. To ensure the robustness of our results, we average the episodic return over five different seeds for both the baseline methods and our approach, providing a comprehensive comparison of performance across all tested tasks. This setup allows us to analyze the impact of our proposed method in a controlled, simulated environment.

We present the results in Figure 4. Our method significantly outperforms the baseline methods across most tasks, with a notable advantage in efficiency. In the `drawer-open` and `hammer` tasks, PROGRESSOR requires only 10% of the total training budget to achieve higher rewards than the baselines. In the `door-open` and `peg-insert-side` tasks, where almost all other methods fail entirely, our approach demonstrates strong performance and a higher success rate.

Ablating Push-back We examine the impact of adversarial online refinement via push-back in our framework. Across various tasks (see Figure 4), we find that online push-back significantly enhances PROGRESSOR’s performance, particularly on challenging tasks such as `door-open` and `hammer`. These tasks demand more precise discrimination between expert behavior and online roll-outs to provide an informative reward signal. Notably, PROGRESSOR without push-back still outperforms every baseline on the `drawer-open` and `reach` tasks.

In the case of the `reach` task, we observe that PROGRESSOR with push-back shows a decrease in return as training progresses, a trend that is less pronounced in the version without push-back. We hypothesize that this occurs because the model initially achieves a high success rate—reaching a peak with only 10% of the total steps—leading the refinement process to overly penalize the agent, even

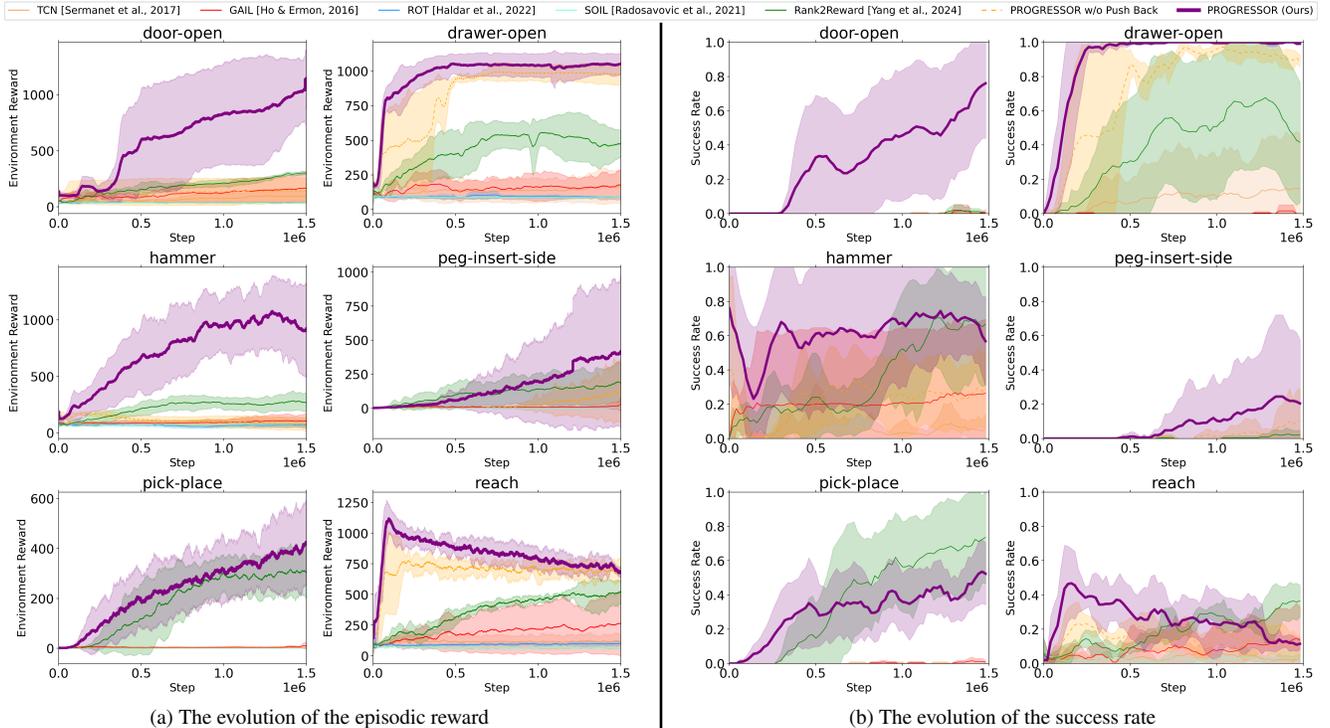


Figure 4. Visualization of policy learning in the Meta-World [43] simulation environment. We run PROGRESSOR and several baselines on six diverse tasks of various difficulties. We also run PROGRESSOR without online push-back as an ablation. We follow the standard evaluation protocol to report the environment reward during training (**left**) and the task success rate from 10 rollouts (**right**) averaged over five seeds. The line denotes the mean and the transparent area denotes standard deviation. PROGRESSOR demonstrates clear advantages in both metrics, especially at early stages of training.

when its behavior closely resembles the expert.

Online training with non-expert data is crucial for robust performance. TCN, which is only trained on expert demonstrations without online updates, fails at every task due to the domain shift between pre-training and online RL.

5.2. Real-World Robotic Experiments

In this section, we demonstrate how PROGRESSOR, pre-trained on the egocentric EPIC-KITCHENS dataset [6, 7], can efficiently learn robotic tasks from a limited number of demonstrations, even when some are unsuccessful. Our approach enhances sample efficiency and robustness to noisy data in offline RL, making it more effective than traditional behavior cloning (BC) methods.

5.2.1. Pretraining on Kitchen Dataset

Using ResNet34 [17] as a backbone, we first pretrain our encoder E_θ with Equation 4 taking P01-P07 sequences from the EPIC-KITCHENS dataset composed of approximately 1.29M frames. We randomly sample frame triplets (o_i, o_j, o_g) from the videos ensuring a maximal frame gap $\|i - g\| \leq 2000$. To improve the robustness of E_θ , we additionally train with distractor frames where we replace the current observation o_j with a frame randomly sampled from

Task	Behavior of Failed Demonstrations
Drawer-Open	The robot approaches the drawer but misses the handle by a few centimeters, failing to open it.
Drawer-Close	The robot misses the drawer by a few centimeters, failing to close it.
Push-Block	The robot moves the block halfway toward the target cup but stops before reaching it.
Pick-Place-Cup	The robot lowers to pick up the cup but misses it by a few centimeters and remains in that position.

Table 1. Failure modes in the collected suboptimal trajectories for the real-world robotic experiments. For each task we collected 20 failed demonstrations to go with the 20 expert demonstrations.

different video sequence $o_{j'}$ as negative examples. For negative triplets, we replace μ in Eqn. 3 with -1 as the label. We train our model with batch size 128 for 30000 iterations using the Adam optimizer with constant learning rate $2e-4$.

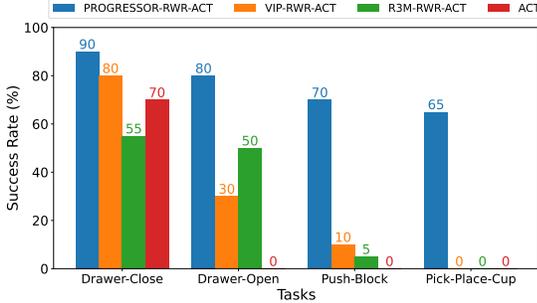


Figure 5. Success rates for four real-world tasks, where RWR-ACT is trained on a combination of correct and failed demonstrations using PROGRESSOR, R3M, and VIP as reward models.

5.2.2. Baselines

We compare PROGRESSOR with R3M [26] and VIP [23]. VIP and R3M are self-supervised visual representations shown to provide dense reward functions for robotic tasks. Both R3M and VIP are pretrained on Ego4D [14] with 4.3M frames from videos from 72000 clips. R3M is trained via time contrastive learning, where the distance between frames closer in time is smaller than for images farther in time. Additionally, they leverage L1 weight regularization and language embedding consistency losses. VIP uses a contrastive approach treating frames close in time as negative pairs and those further away as positive pairs towards learning visual goal-conditioned value functions.

5.2.3. Real-World Few-Shot Offline Reinforcement Learning with Noisy Demonstrations

Following the offline reinforcement learning experiments of Ma et al. [23], we leveraged the reward-weighted regression (RWR) method [28, 29]. Our aim in applying RWR is to show that a reward model trained on human videos can help robots learn to perform tasks even when the training data contains noisy, suboptimal trajectories. In scenarios with highly suboptimal trajectories, vanilla behavior cloning (BC) approaches often struggle to learn the correct behaviors. However, goal-conditioned reward weighting provides a signal that focuses learning from the correct subtrajectories, enabling the visual imitation learning model to effectively learn accurate action executions.

To this end, we start with a vanilla BC model and adapt the loss function to be weighted by a pretrained reward model. Specifically, we employed Action-Chunking Transformer (ACT) [45], a BC model designed for learning fine-grained manipulation tasks. We modified ACT’s reconstruction loss—originally defined as the mean absolute error between the predicted action and the ground-truth action—to incorporate RWR as follows:

$$\mathcal{L}_{\text{reconst}} = \exp(\omega \cdot \hat{r}) \cdot \|\pi_{\text{ACT}}(o_j) - a_t\|_1, \quad (8)$$

where \hat{r} is the reward prediction for the current observation o_j , ω is the temperature parameter, a_t is the ground-truth action, and $\pi_{\text{ACT}}(o_j)$ represents the action predicted by ACT by taking the current observation image (o_j) as an input. In this paper, we refer to ACT trained with the reconstruction loss replaced by Equation 8 as RWR-ACT, to distinguish it from the standard ACT.

We compare PROGRESSOR with R3M and VIP by freezing the pre-trained models and using them as reward prediction models to train RWR-ACT on downstream robotic learning tasks. The reward predictions from these models are used in place of \mathbf{u}_θ in Equation 8. In both VIP and R3M, the reward prediction is parameterized by the current observation, previous observation, and goal observation. Following Ma et al., the reward for a given transition tuple (o_t, o_{t+1}) is defined as $R(o_t, o_{t+1}; \phi, o_g) := S_\phi(o_{t+1}; o_g) - S_\phi(o_t; o_g)$. Where S_ϕ is the distance function in the ϕ -representation space extracted from VIP and R3M, defined as: $S_\phi(o_t; o_g) := -\|\phi(o_t) - \phi(o_g)\|_2$.

In contrast, our method parameterizes reward using the initial observation, current observation, and goal observation. Additionally, we include vanilla ACT as a baseline, which applies uniform weighting (i.e., $\omega = 0$) across all observations in the training trajectories.

We design four tabletop robotic manipulation tasks (see Figure 3 a-d): Drawer-Open, Drawer-Close, Pick-Place-Cup, and Push-Block. For each task, we collect 40 demonstrations, half of which are suboptimal and fail to complete the task. Including these failed demonstrations is crucial for evaluating whether the learned reward model can accurately signal progress toward a goal by assigning high reward to transitions that lead toward completion and low reward to those that do not. Table 1 summarizes the behaviors in failed demonstrations. Detailed task and data descriptions along with frame sequences from both successful and failed trajectories, are provided in S9. We categorize Drawer-Close as an easy task and the other three as hard tasks. This distinction in difficulty is based on the complexity of the tasks and the level of suboptimality in failed demonstrations.

We train all policies using the same hyperparameters employed by Zhao et al. [45] for training ACT in their real-world behavior cloning experiments (full details in S9.2). For all RWR experiments, we set $\omega = 0.1$. Each method is then evaluated over 20 rollouts, and the success rate is reported. The success criteria for each task are presented in S9.

The evaluation results presented in Figure 5 highlight the significant advantage of PROGRESSOR-RWR-ACT, which consistently outperforms all baseline methods across all tasks. This advantage is particularly pronounced in the more challenging tasks, such as Push-Block and Pick-Place-Cup, on which other methods struggle to

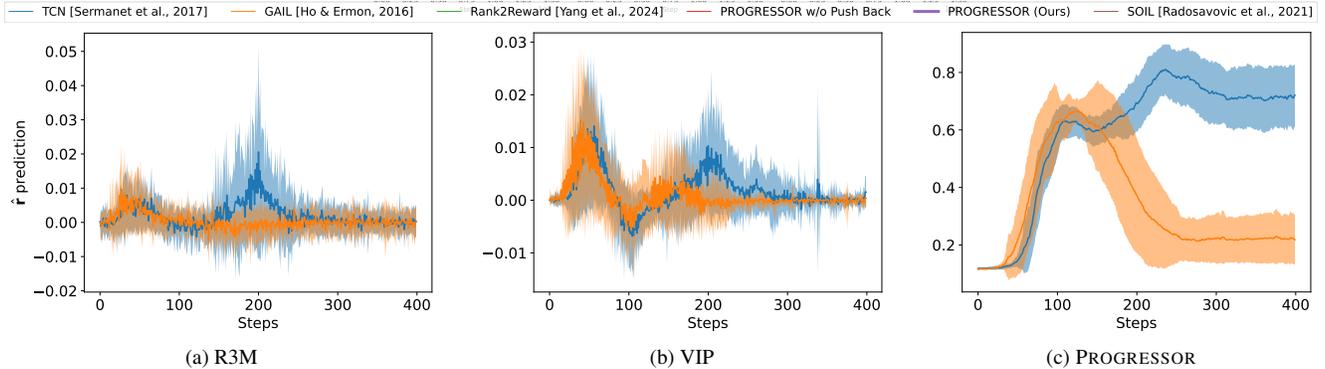


Figure 6. Mean reward predictions, \hat{r} , for **correct** and **incorrect** demonstrations in the *Drawer-Open* task from PROGRESSOR, R3M, and VIP. PROGRESSOR provides more distinct reward weighting between correct and failed trajectories compared to the baseline models.

achieve meaningful task performance. While R3M-RWR-ACT and VIP-RWR-ACT show moderate success in simpler tasks, such as *Drawer-Close*, their performance significantly drops in *Push-Block* and *Pick-Place-Cup* tasks. As can be seen in Figure 6, the gap in reward prediction between failed and correct trajectories for PROGRESSOR is notably larger compared to R3M and VIP, starting when the robot begins to deviate from the goal observation (around step 125). This supports our hypothesis that PROGRESSOR provides more distinct reward weighting between both correct and failed trajectories compared to the baseline models. Vanilla ACT performs poorly, especially in difficult tasks. Since ACT assigns equal weight to all sub-trajectories, it cannot prioritize more relevant transitions or filter out noise effectively, leading to poor task execution.

Figure 7 shows reward predictions for a sample frame sequence from EPIC-KITCHENS and a trajectory from our *Pick-Place-Cup* robotic task. The figure shows that, despite being pretrained solely on human videos, PROGRESSOR generates well-formed reward predictions for a robotic task in a zero-shot setting.

6. Conclusion and Limitations

In this work, we presented PROGRESSOR, a self-supervised framework that learns task-agnostic reward functions from video via progress estimation. By learning progress estimation between observations in expert trajectories, PROGRESSOR generates dense rewards that effectively guide policy learning. During online RL training, this progress estimation is further refined through an adversarial push-back strategy, helping the model handle non-expert observations and minimize distribution shift. Our method shows enhanced performance compared to previous state-of-the-art approaches across a range of simulated experiments. In real-robot experiments, we applied PROGRESSOR, pretrained on in-the-wild human videos, to learn policies with limited and noisy task demonstrations, outperforming other

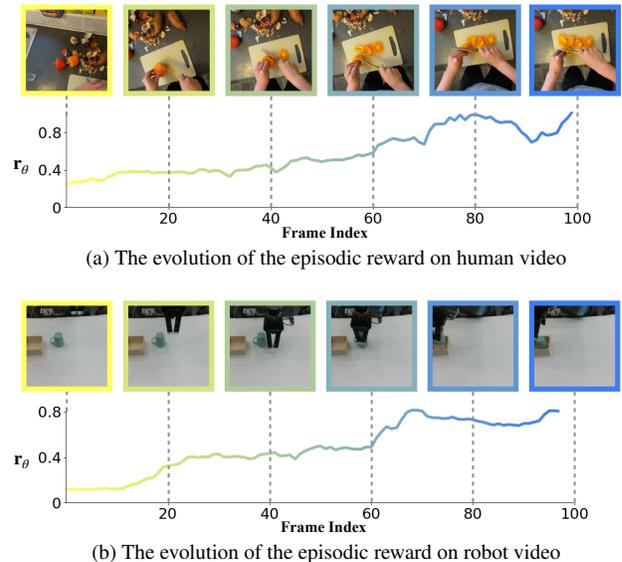


Figure 7. Visualization of the predicted rewards from Epic Kitchen pretrained PROGRESSOR, evaluated on a sample clip from EPIC-KITCHENS (**top**) and on sample frames from an unseen *Pick-Place-Cup* demonstration on our robot (**bottom**).

visual reward models.

Limitations and Future Work We acknowledge several limitations in our work. (1) Our method assumes a linear progression of tasks and models progress as a unimodal prediction, making it unsuitable for tasks with cyclic observations, such as those in the DeepMind Control Suite [37]. Future enhancements that model progress as a multimodal prediction could address this limitation more effectively. (2) While our online RL experiments demonstrate that our hyperparameters are robust across various tasks, incorporating a dynamic weighting factor β for refinement may further enhance performance.

References

- [1] Pieter Abbeel and Andrew Y. Ng. Apprenticeship learning via inverse reinforcement learning. In *Proceedings of the Twenty-First International Conference on Machine Learning*, page 1, New York, NY, USA, 2004. Association for Computing Machinery. 2
- [2] Minttu Alakuijala, Gabriel Dulac-Arnold, Julien Mairal, Jean Ponce, and Cordelia Schmid. Learning reward functions for robotic manipulation by observing humans, 2023. 1
- [3] Shikhar Bahl, Abhinav Gupta, and Deepak Pathak. Human-to-robot imitation in the wild. *arXiv preprint arXiv:2207.09450*, 2022. 2
- [4] Bowen Baker, Ilge Akkaya, Peter Zhokov, Joost Huizinga, Jie Tang, Adrien Ecoffet, Brandon Houghton, Raul Sampedro, and Jeff Clune. Video pretraining (VPT): Learning to act by watching unlabeled online videos. 2022. 2
- [5] Serena Booth, W Bradley Knox, Julie Shah, Scott Niekum, Peter Stone, and Alessandro Allievi. The perils of trial-and-error reward design: Misdesign through overfitting and invalid task specifications. In *AAAI*, 2023. 1
- [6] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Sanja Fidler, Antonino Furnari, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, and Michael Wray. Scaling egocentric vision: The EPIC-KITCHENS dataset. In *ECCV*, 2018. 2, 6, 1
- [7] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Antonino Furnari, Jian Ma, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, and Michael Wray. Rescaling egocentric vision: Collection, pipeline and challenges for EPIC-KITCHENS-100. *IJCV*, 130:33—55, 2022. 6
- [8] Alejandro Escontrela, Ademi Adeniji, Wilson Yan, Ajay Jain, Xue Bin Peng, Ken Goldberg, Youngwoon Lee, Danijar Hafner, and Pieter Abbeel. Video prediction models as rewards for reinforcement learning. 2024. 1, 2
- [9] Benjamin Eysenbach, Tianjun Zhang, Ruslan Salakhutdinov, and Sergey Levine. Contrastive learning as goal-conditioned reinforcement learning, 2023. 1, 2
- [10] Chelsea Finn, Sergey Levine, and Pieter Abbeel. Guided cost learning: Deep inverse optimal control via policy optimization. *arXiv preprint arXiv:1603.00448*, 2016. 2
- [11] Justin Fu, Katie Luo, and Sergey Levine. Learning robust rewards with adversarial inverse reinforcement learning. *arXiv preprint arXiv:1710.11248*, 2018. 2
- [12] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario March, and Victor Lempitsky. Domain-adversarial training of neural networks. 17(59):1–35, 2016. 4
- [13] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. 2014. 4
- [14] Kristen Grauman, Andrew Westbury, Eugene Byrne, Zachary Chavis, Antonino Furnari, Rohit Girdhar, Jackson Hamburger, Hao Jiang, Miao Liu, Xingyu Liu, et al. Ego4D: Around the world in 3,000 hours of egocentric video. In *CVPR*, pages 18995–19012, 2022. 7
- [15] Dylan Hadfield-Menell, Smitha Milli, Pieter Abbeel, Stuart J Russell, and Anca Dragan. Inverse reward design. 2017. 1
- [16] Siddhant Haldar, Vaibhav Mathur, Denis Yarats, and Lerrel Pinto. Watch and match: Supercharging imitation with regularized optimal transport. In *Conference on Robot Learning*, pages 32–43. PMLR, 2023. 5
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 6, 1
- [18] Jonathan Ho and Stefano Ermon. Generative adversarial imitation learning. 2016. 5
- [19] Jonathan Ho and Stefano Ermon. Generative adversarial imitation learning. *arXiv preprint arXiv:1606.03476*, 2016. 2
- [20] Tao Huang, Guangqi Jiang, Yanjie Ze, and Huazhe Xu. Diffusion reward: Learning rewards via conditional video diffusion. *arXiv preprint arXiv:2312.14134*, 2023. 1, 2
- [21] Diederik P Kingma and Max Welling. Auto-encoding variational Bayes. *arXiv preprint arXiv:1312.6114*, 2013. 3
- [22] Yunzhu Li, Jiaming Song, and Stefano Ermon. InfoGAIL: Interpretable imitation learning from visual demonstrations. *arXiv preprint arXiv:1703.08840*, 2017. 2
- [23] Yecheng Jason Ma, Shagun Sodhani, Dinesh Jayaraman, Osbert Bastani, Vikash Kumar, and Amy Zhang. VIP: Towards universal visual reward and representation via value-implicit pre-training. *arXiv preprint arXiv:2210.00030*, 2022. 7
- [24] Yecheng Jason Ma, Shagun Sodhani, Dinesh Jayaraman, Osbert Bastani, Vikash Kumar, and Amy Zhang. Vip: Towards universal visual reward and representation via value-implicit pre-training, 2023. 1, 2
- [25] Alberto Maria Metelli, Filippo Lazzati, and Marcello Restelli. Towards theoretical understanding of inverse reinforcement learning. *arXiv preprint arXiv:2304.12966*, 2023. 2
- [26] Suraj Nair, Aravind Rajeswaran, Vikash Kumar, Chelsea Finn, and Abhinav Gupta. R3M: A universal visual representation for robot manipulation. *arXiv preprint arXiv:2203.12601*, 2022. 7
- [27] Suraj Nair, Aravind Rajeswaran, Vikash Kumar, Chelsea Finn, and Abhinav Gupta. R3M: A universal visual representation for robot manipulation. *arXiv preprint arXiv:2203.12601*, 2022. 1, 2
- [28] Xue Bin Peng, Aviral Kumar, Grace Zhang, and Sergey Levine. Advantage-weighted regression: Simple and scalable off-policy reinforcement learning. *arXiv preprint arXiv:1910.00177*, 2019. 7
- [29] Jan Peters and Stefan Schaal. Reinforcement learning by reward-weighted regression for operational space control. pages 745–750, 2007. 7
- [30] Yuzhe Qin, Yueh-Hua Wu, Shaowei Liu, Hanwen Jiang, Ruihan Yang, Yang Fu, and Xiaolong Wang. DexMV: Imitation learning for dexterous manipulation from human videos. In *ECCV*, 2022. 2
- [31] Ilija Radosavovic, Xiaolong Wang, Lerrel Pinto, and Jitendra Malik. State-only imitation learning for dexterous manipulation. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 7865–7871. IEEE, 2021. 5

- [32] Stuart J Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall, Englewood Cliffs, NJ, USA, 1 edition, 1995. [1](#)
- [33] Pierre Sermanet, Corey Lynch, Yevgen Chebotar, Jasmine Hsu, Eric Jang, Stefan Schaal, and Sergey Levine. Time-contrastive networks: Self-supervised learning from video. pages 1134–1141, 2018. [5](#)
- [34] Pierre Sermanet, Corey Lynch, Yevgen Chebotar, Jasmine Hsu, Eric Jang, Stefan Schaal, and Sergey Levine. Time-contrastive networks: Self-supervised learning from video, 2018. [2](#)
- [35] Satinder Singh, Richard L. Lewis, and Andrew G. Barto. Where do rewards come from? In *Proceedings of the International Symposium on AI Inspired Biology*, pages 111–116, 2010. [1](#)
- [36] Richard S Sutton and Andrew G Barto. *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, MA, USA, 2 edition, 2018. [1](#)
- [37] Yuval Tassa, Yotam Doron, Alistair Muldal, Tom Erez, Yazhe Li, Diego de Las Casas, David Budden, Abbas Abdolmaleki, Josh Merel, Andrew Lefrancq, Timothy Lillicrap, and Martin Riedmiller. DeepMind control suite. *arXiv preprint arXiv:1801.00690*, 2018. [8](#)
- [38] Michael C Welle, Nils Ingelhart, Martina Lippi, Maciej K. Wozniak, Andrea Gasparri, and Danica Kragic. Quest2ROS: An app to facilitate teleoperating robots. In *Proceedings of the International Workshop on Virtual, Augmented, and Mixed-Reality for Human-Robot Interactions*, 2024. [1](#)
- [39] Markus Wulfmeier, Peter Ondruska, and Ingmar Posner. Maximum entropy deep inverse reinforcement learning, 2016. [2](#)
- [40] Daniel Yang, Davin Tjia, Jacob Berg, Dima Damen, Pulkit Agrawal, and Abhishek Gupta. Rank2Reward: Learning shaped reward functions from passive video. *arXiv preprint arXiv:2404.14735*, 2024. [1](#), [2](#), [3](#), [5](#)
- [41] Mengjiao Yang and Ofir Nachum. Representation matters: Offline pretraining for sequential decision making. *arXiv preprint arXiv:2102.05815*, 2021. [2](#)
- [42] Denis Yarats, Ilya Kostrikov, and Rob Fergus. Image augmentation is all you need: Regularizing deep reinforcement learning from pixels. In *International Conference on Learning Representations*, 2021. [5](#)
- [43] Tianhe Yu, Deirdre Quillen, Zhanpeng He, Ryan Julian, Karol Hausman, Chelsea Finn, and Sergey Levine. Meta-World: A benchmark and evaluation for multi-task and meta reinforcement learning. 2020. [2](#), [5](#), [6](#), [1](#)
- [44] Qihang Zhang, Zhenghao Peng, and Bolei Zhou. Learning to drive by watching YouTube videos: Action-conditioned contrastive policy pretraining. In *ECCV, 2022*. [2](#)
- [45] Tony Z Zhao, Vikash Kumar, Sergey Levine, and Chelsea Finn. Learning fine-grained bimanual manipulation with low-cost hardware. *arXiv preprint arXiv:2304.13705*, 2023. [7](#), [2](#)
- [46] Brian D. Ziebart, Andrew Maas, J. Andrew Bagnell, and Anind K. Dey. Maximum entropy inverse reinforcement learning. In *AAAI*, pages 1433–1438, 2008. [2](#)

PROGRESSOR: A Perceptually Guided Reward Estimator with Self-Supervised Online Refinement

Supplementary Material

7. PROGRESSOR Training Details

7.1. Architecture and Training

PROGRESSOR can, in principle, be trained with any visual encoding architecture, requiring only minor modifications to the final layer to predict Gaussian parameters. In our experiments, we utilize the standard ResNet34 model [17], replacing its final fully-connected layer with an MLP of size [512, 512, 128]. Given that our method processes triplets of inputs ($\mathbf{o}_i, \mathbf{o}_j, \mathbf{o}_g$), the resulting representation has a size of [128 × 3]. This representation is then fed into an MLP with layers of size [128 × 3, 2048, 256]. Finally, two prediction heads are derived from the 256-dimensional output, predicting μ and $\log \sigma^2$.

We pretrain PROGRESSOR for 30000 steps using the EPIC-KITCHENS dataset [6] for the real-world experiments, and for 10000 steps for experiments performed in simulation. In the pretraining steps for both our simulated and real-world experiments, we first sample a trajectory (i.e., a video clip) from the pretraining dataset. We then randomly select an initial frame \mathbf{o}_i as well as a goal frame \mathbf{o}_g from the selected trajectory such that $\|g-i\| \leq 2000$. Finally, we uniformly randomly select an intermediate frame \mathbf{o}_j , where $i < j < g$.

7.2. Hyperparameters

	Simulation	Real-World RWR
α	0.4	0.4
β	0.9	–

Table 2. Hyperparameters used by PROGRESSOR for experiments performed in simulation and the real world.

Table 2 lists the hyperparameters used for both the simulation online RL and real-world offline RL experiments. A consistent $\alpha = 0.4$ is used across all experiments. The push-back decay factor is employed exclusively in the simulated online RL experiments.

8. Simulation Experiment Details

In this section, we describe the tasks and the data generation process employed using the MetaWorld environment [43] for our simulation experiments.

8.1. Meta-World Tasks

We took six diverse tasks from the Meta-World environment [43], described in Table 3. In all tasks, the position of the target object, such as the drawer or hammer, is randomized between episodes.

Task	Task Description
door-open	Open a door with a revolving joint.
drawer-open	Open a drawer.
hammer	Hammer a screw on the wall.
peg-insert-side	Insert a peg sideways.
pick-place	Pick and place a puck to a goal.
reach	Reach a goal position.

Table 3. Meta-World [43] task descriptions.

8.2. Expert Data Generation

To collect expert trajectories for our simulated experiments, we execute Meta-World’s oracle policies. For each task, we generated 100 successful rollouts for training and 10 for testing. This dataset is subsequently used for pretraining PROGRESSOR in our simulated experiments, following the steps outlined in Section 7.1.

9. Real-World Robot Experiment Details

9.1. Robotic Experiment Setup

The real-robot experiments are performed using a Universal Robots UR5 robot arm equipped with a Robotiq 3-Finger Gripper (Figure 8). The setup includes two RealSense cameras: one mounted on the robot’s wrist that images the gripper, and the second on a fixed tripod facing the robot.

9.2. Robotic Demonstration Data Collection

We collect demonstrations by teleoperating the UR5 using a Meta Quest 3 controller [38]. We record each demonstration at 30 Hz for 400 steps. Figure 9 shows video frame sequences from correct and incorrect demonstrations for the four real-world tasks.

9.3. Task Descriptions

The tasks involve a variety of object manipulation challenges designed to test reward weighting in offline Reinforcement Learning. In the `Drawer-Close` task, the objective is to close a drawer starting from an open position. The `Drawer-Open` task requires the agent to pull a drawer open from a closed state. In the `Push-Block` task, the

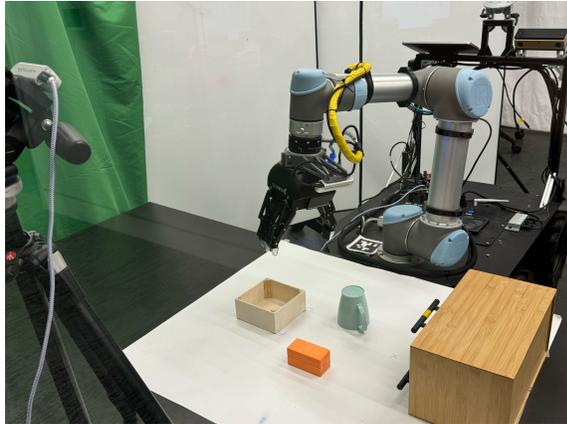


Figure 8. The real-world experiments were conducted using a Universal Robots UR5 robot arm equipped with a Robotiq 3-Finger Gripper. One RealSense camera is mounted to the end-effector and the other RealSense camera is fixed in the environment using a tripod (visible in the left of the image).

goal is to push a block toward a specified target, a cup, ensuring the block moves into proximity with the cup. Finally, `Pick-Place-Cup` involves picking up a cup and carefully placing it into a designated box. Figure 10 displays the goal frames representing the completion of each task.

9.4. Training and Evaluation Details

Task	Success Criterion
<code>Drawer-Open</code>	Drawer is open by more than 5 cm.
<code>Drawer-Close</code>	Drawer is within 1 cm of being fully closed.
<code>Push-Block</code>	Block is within 5 cm of the cup.
<code>Pick-Place-Cup</code>	Cup is placed inside the box.

Table 4. Success criterion for the real-robot experiments.

Our few-shot offline RL implementation builds upon the Action-Chunking Transformer (ACT) [45]. The inputs to the model include (i) two 640×480 RGB images from the RealSense cameras, and (ii) proprioceptive data consisting of the 6-DoF joint angles and the binary open or close state of the gripper. The action space consists of the 6-DoF translational and rotational velocity of the end-effector and a binary open or close command of the gripper.

In the reward-weighted regression (RWR) setup, rewards are computed by providing all reward models with the final frame of a correct demonstration from each task as the goal image. For all reward predictions, frames from the fixed RealSense camera were used. Figure 10 illustrates the goal images used for the four robotic tasks, which were consistently used across all reward predictions. The temperature scale in RWR was set to $\omega = 0.1$ for all tasks.

Hyperparameter	Value
prediction horizon	30
learning rate	10^{-5}
batch size	64
epochs	5000
ω	0.1

Table 5. Hyperparameters for reward-weighted ACT training.

Table 5 shows the hyperparameters used to train the {`PROGRESSOR`, `VIP`, `R3M`}-RWR-ACT models. For the vanilla ACT, we used the same hyperparameters as listed in Table 5, except for $\omega = 0$.

During inference, the model predicts a sequence of actions (a “chunk”) of length 30 and then executes each action before predicting the next chunk. We did not use a temporal ensemble as proposed by Zhao et al. [45], since we found that it causes the gripper to drift and negatively impacts performance.

For evaluation, we conduct 20 test rollouts for each task and report the success rate. The success criteria for each task are outlined in Table 4.

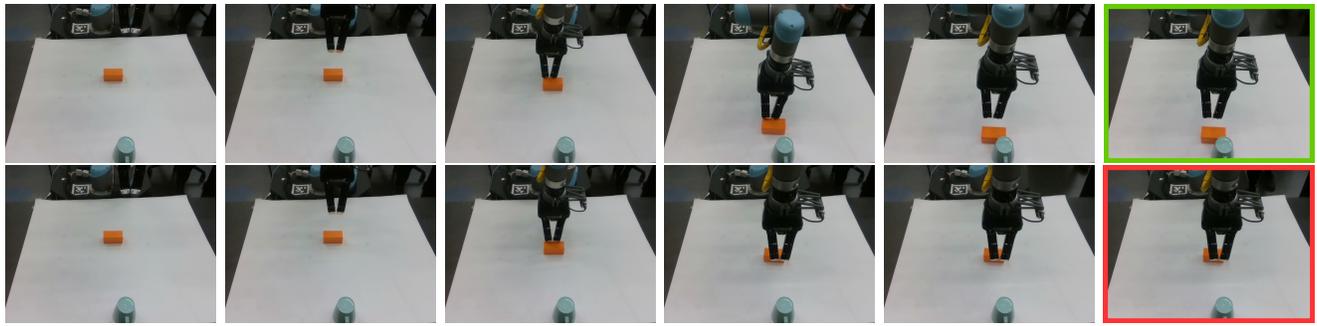
10. Ablation

In this section, we present an ablation study evaluating different values of the push-back decay factor (β) while training a DrQ-v2 agent on Meta-World’s `hammer` task, using a fixed seed of 121. The case of $\beta = 0$ (`PROGRESSOR` without Push-back) is discussed in the main paper. Figure 11 depicts the environment rewards accumulated during training. As shown in the figure, the agent achieves higher rewards with $\beta = 0.9$.

11. Qualitative Analysis

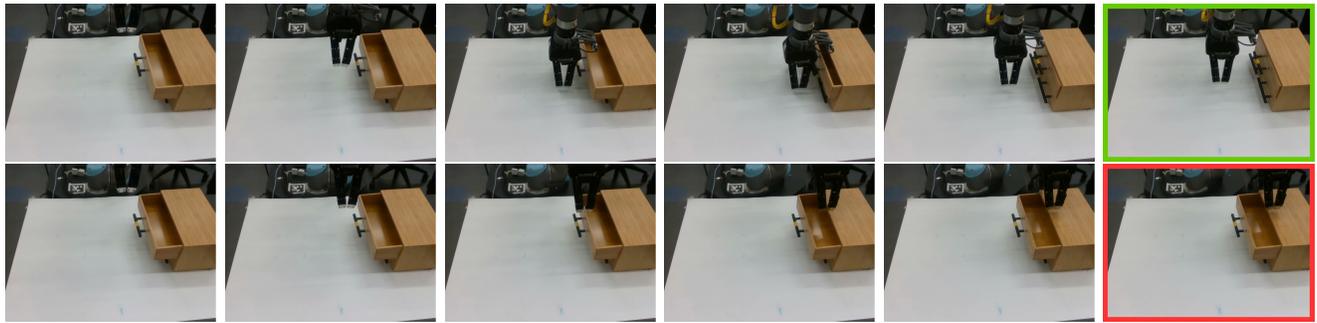
Figure 12 presents zero-shot reward predictions from `PROGRESSOR` pretrained on the EPIC-KITCHENS dataset. This figure serves as an extension to Figure 7 for completeness. It includes zero-shot reward predictions for sample correct trajectories from our collected real-robot demonstrations for the tasks `Drawer-Open`, `Drawer-Close`, and `Push-Block`.

For completeness, we have included plots similar to the mean reward prediction figure from the main paper (see Figure 7). The plots in Figure 13 provide a qualitative comparison of reward predictions from our robotic demonstration dataset for three additional real-world tasks: `Drawer-Close`, `Push-Block`, and `Pick-Place-Cup`. As illustrated in the figure, for all tasks, our reward model consistently predicts lower average rewards for the sub-trajectories in the incorrect demonstrations, where the failures occur.



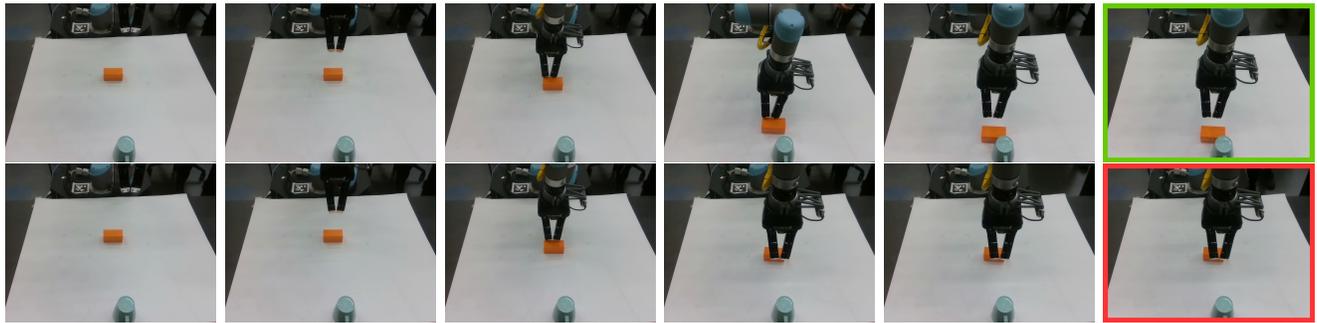
Time →

(a) Push-Block: (top) correct and (bottom) incorrect



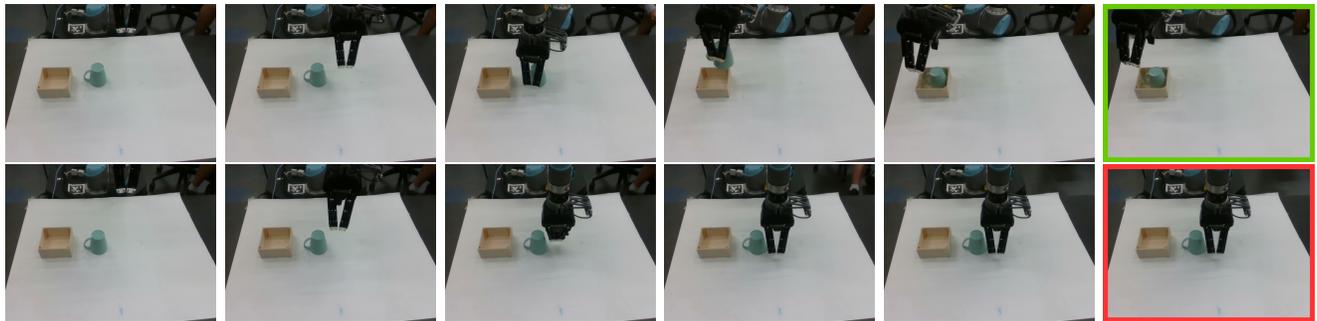
Time →

(b) Drawer-Close: (top) correct and (bottom) incorrect



Time →

(c) Push-Block: (top) correct and (bottom) incorrect



Time →

(d) Pick-Place-Cup: (top) correct and (bottom) incorrect

Figure 9. Correct and incorrect demonstrations (every 80th frame) for each real-robot task from a third-person camera view used in our experiments. To see how PROGRESSOR and the baselines differentiate between correct and incorrect trajectories, see Figure 6 (Drawer-Open) and Figure 13 (other three tasks).

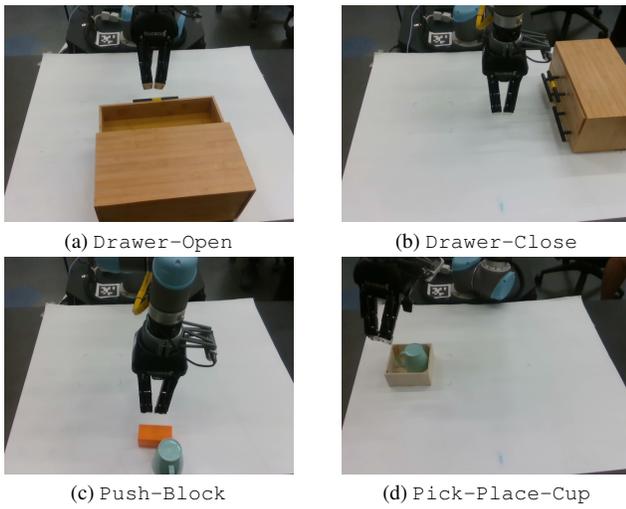


Figure 10. The goal images that we use for each task for reward weighting in the offline reinforcement learning experiments.

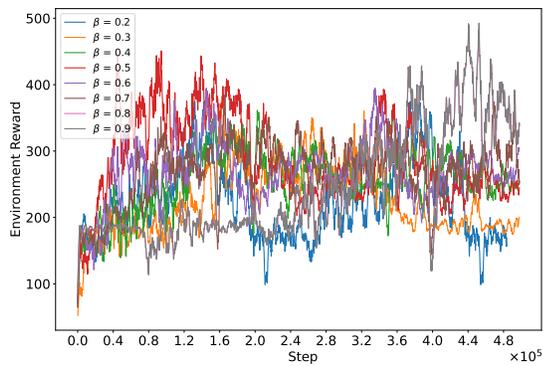
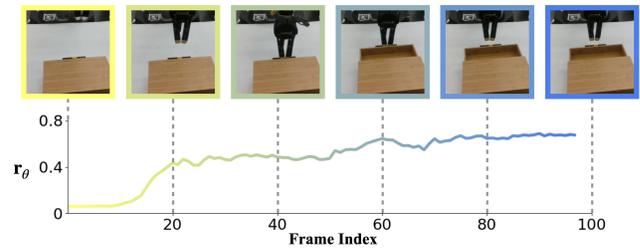
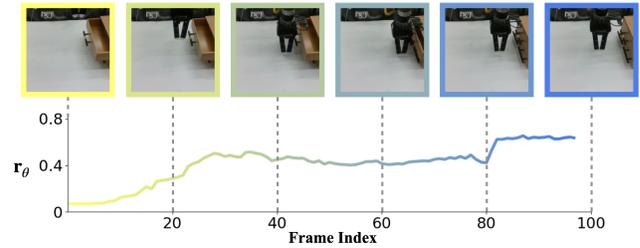


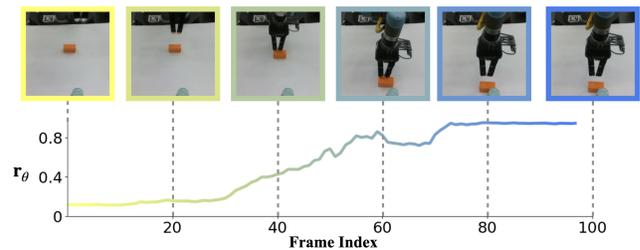
Figure 11. Visualization of policy learning performance in MetaWorld’s hammer environment, evaluating the effect of different values for the push-back decay factor β . The plot highlights the accumulated rewards over training, demonstrating how varying β values influences the agent’s ability to optimize performance.



(a) Reward plot for Drawer-Open task



(b) Reward plot for Drawer-Close task



(c) Reward plot for Push-Block task

Figure 12. Visualization of the predicted reward by PROGRESSOR, pretrained on EPIC-KITCHENS and evaluated zero-shot on correct robotic demonstrations for (a) Drawer-Open (b) Drawer-Close and (c) Push-Block tasks.

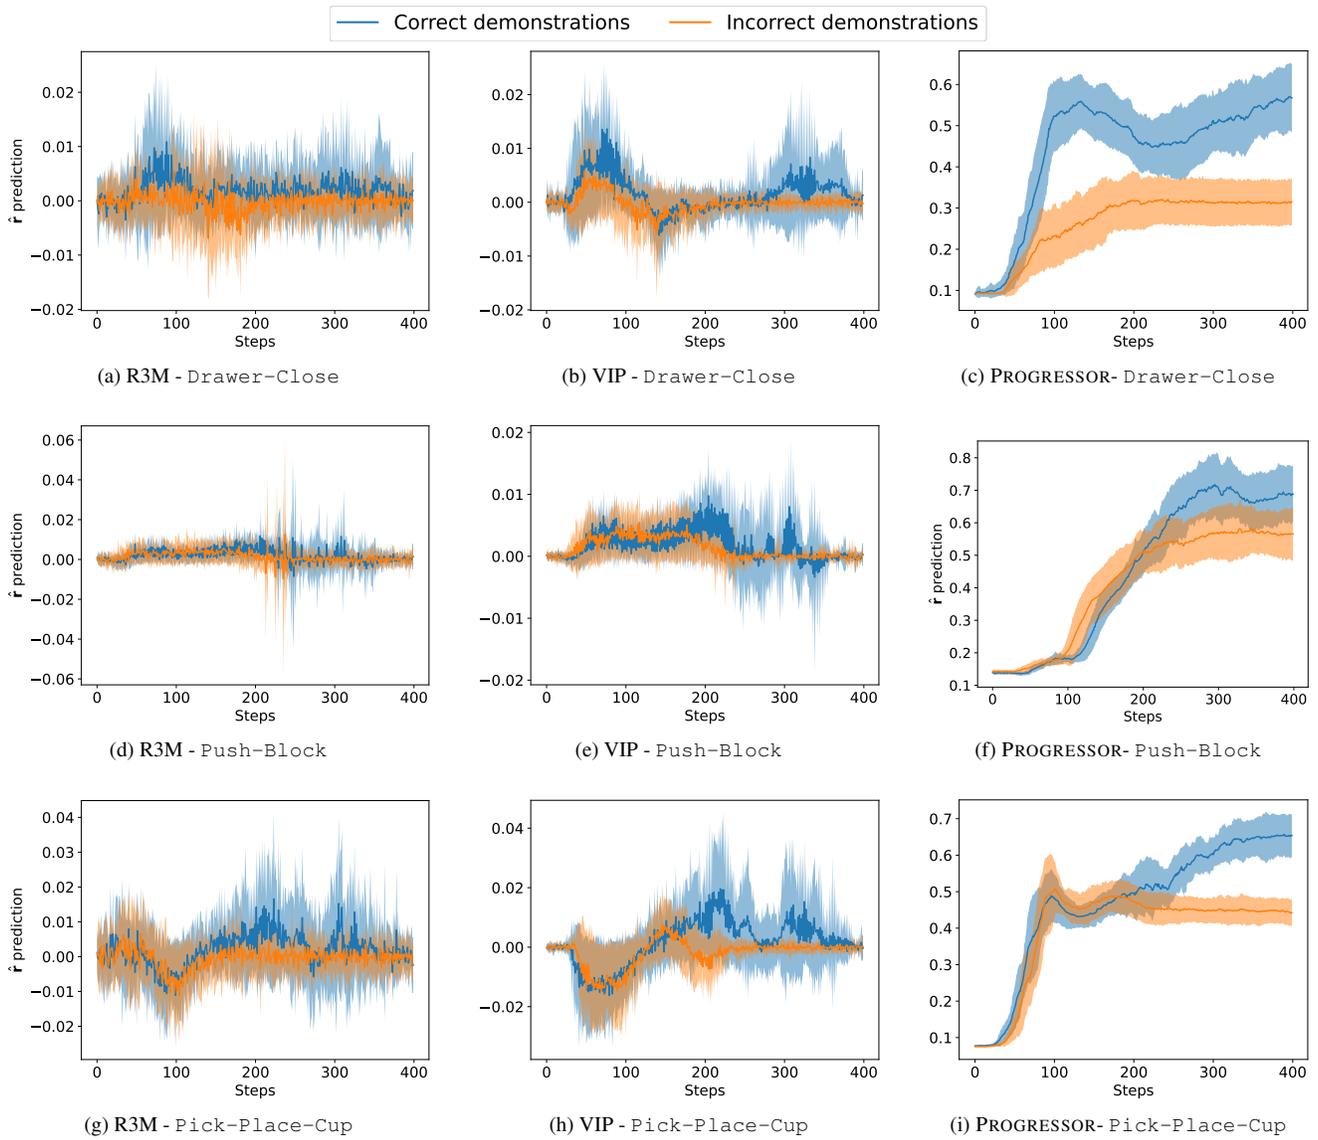


Figure 13. Mean reward predictions \hat{r} of (left column) R3M, (middle column) VIP, and (right column) PROGRESSOR for correct and incorrect demonstrations for the Drawer-Close, Push-Block, and Pick-Place-Cup tasks. PROGRESSOR provides reward predictions (weights) that better differentiate between correct and incorrect trajectories, consistently outperforming the baseline models.