

MELD: Meta-Reinforcement Learning from Images via Latent State Models

Tony Z. Zhao^{*†}, Anusha Nagabandi^{*†}, Kate Rakelly^{*†}, Chelsea Finn[‡], Sergey Levine[†]

Abstract: Meta-reinforcement learning algorithms can enable autonomous agents, such as robots, to quickly acquire new behaviors by leveraging prior experience in a set of related training tasks. However, the onerous data requirements of meta-training compounded with the challenge of learning from sensory inputs such as images have made meta-RL challenging to apply to real robotic systems. Latent state models, which learn compact state representations from a sequence of observations, can accelerate representation learning from visual inputs. In this paper, we leverage the perspective of meta-learning as task inference to show that latent state models can *also* perform meta-learning given an appropriately defined observation space. Building on this insight, we develop meta-RL with latent dynamics (MELD), an algorithm for meta-RL from images that performs inference in a latent state model to quickly acquire new skills given observations and rewards. MELD outperforms prior meta-RL methods on several simulated image-based robotic control problems, and enables a real WidowX robotic arm to insert an Ethernet cable into new locations given a sparse task completion signal after only 8 hours of real world meta-training¹. To our knowledge, MELD is the first meta-RL algorithm trained in a real-world robotic control setting from images.

Keywords: meta-learning, reinforcement learning, robotic manipulation

1 Introduction

General purpose autonomous robots must be able to perform a wide variety of tasks and quickly acquire new skills. For example, consider a robot tasked with assembling electronics in a data center. This robot must be able to insert cables of varying shapes, sizes, colors, and weights into the correct ports with the appropriate amounts of force. While standard RL algorithms require hundreds or thousands of trials to learn a policy for each new setting, meta-RL methods hold the promise of drastically reducing the number of trials required. Given a task distribution, such as the variety of ways to insert cables described above, meta-RL algorithms leverage a set of training tasks to meta-learn a mechanism that can quickly learn unseen tasks from the same distribution. Despite promising results in simulation demonstrating that agents can learn new tasks in a handful of trials [1, 2, 3, 4], these algorithms remain largely unproven on real-world robotic systems.

Applying these algorithms to real-world robotic systems requires handling the raw sensory observations collected by a robot’s on-board sensors. In principle, deep reinforcement learning (RL) algorithms can directly map sensory inputs to actions. However, this automation comes at a steep cost in sample efficiency since the agent must learn to interpret observations from reward supervision alone. Fortunately, unsupervised learning of general-purpose latent state (or dynamics) models can serve as an additional training signal to help solve the representation learning problem [5, 6, 7, 8]. In this work, we seek to leverage the benefits of latent state models for representation learning to design a meta-RL algorithm that can acquire new skills quickly in the real world.

While meta-learning algorithms are often viewed as algorithms that learn to learn [9, 10, 3], an alternative viewpoint frames meta-learning as task inference [11, 12, 4]. From this perspective, the task is a hidden variable that can be inferred from experience consisting of observations and rewards. Our key insight is that the same latent dynamics models that greatly improve efficiency in end-to-end single-task RL can *also*, with minimal modification, be used for meta-RL by treating the

¹Results and videos at <https://sites.google.com/view/meld-lsm/home>

^{*} Equal contribution, [†] University of California Berkeley, [‡] Stanford University




Figure 1: At test time our algorithm MELD enables a 5-DoF WidowX robot to insert the ethernet cable into a novel insertion location and orientation within two episodes of experience, operating from image observations and a sparse task completion signal when the cable is correctly inserted. MELD achieves this result by meta-training a latent dynamics model to capture task and state information, as well a policy that conditions on this information to explore and identify the correct insertion point.

unknown task information as a hidden variable to be estimated from experience. We formalize the connection between latent state inference and meta-RL, and leverage this insight in our proposed algorithm MELD, Meta-RL with Latent Dynamics. To derive MELD, we cast meta-RL and latent state inference into a single partially observed Markov decision process (POMDP) in which task and state variables are aspects of a more general per-time step hidden variable. Concretely, we represent the agent’s belief over the hidden variable as the variational posterior in a sequential VAE latent state model that takes observations and rewards as input, and we condition the agent’s policy on this belief. During meta-training, the latent state model and policy are trained across a fixed set of training tasks sampled from the task distribution. The trained system can then quickly learn a new task from the distribution by inferring the posterior belief over the hidden variable and executing the conditional meta-learned policy.

We find in simulation that MELD substantially outperforms prior work on several challenging locomotion and manipulation problems, such as running at varying velocities, inserting a peg into varying targets, and putting away mugs of unknown weights to varying locations on a shelf. We then analyze MELD’s capability to meta-learn temporally-extended exploration strategies when only a sparse task completion signal is available. Finally, using a real WidowX robotic arm, we find that after eight hours of meta-training MELD successfully performs Ethernet cable insertion into ports at novel locations and orientations (Figure 1). This real-world experiment with the WidowX is to our knowledge the first demonstration of a meta-RL algorithm trained from images on a real robotic platform. Our open-source implementation of MELD can be found at <https://github.com/tonyzhaozh/meld>.

2 Related Work

Our approach can be viewed methodologically as a bridge between meta-RL methods for fast skill acquisition and latent state models for state estimation. In this section, we discuss these areas of work, as well as RL on real-world robotic platforms.

RL for Robotics. While prior work has obtained good results with geometric and force control approaches for a wide range of manipulation tasks [13, 14, 15], including insertion tasks [16, 17] such as those in our evaluation, such approaches typically require considerable manual design effort for each task. RL algorithms offer an automated alternative that has been demonstrated on a variety of robotic tasks [18] including insertion [19, 20, 21, 22, 23]. Although these policies learn impressive skills, they typically do not transfer to other tasks and must be re-trained from scratch for each task.

Meta-RL. Meta-RL algorithms learn to acquire new skills quickly by using experience from a set of meta-training tasks. Current meta-RL methods differ in how this acquisition procedure is represented, ranging from directly representing the learned learning process with a recurrent [1, 2] or recursive [24] deep network, to learning initial parameters for gradient-based optimization [3, 25, 26, 27], and learning tasks via variational inference [4, 28, 29]. Some of these works have formalized meta-RL as a special kind of POMDP in which the hidden state is constant throughout a task [4, 28, 30, 29]. Taking a broader view, we show that meta-RL can be tackled with a general POMDP algorithm that estimates a time-varying hidden state, rendering the same algorithm applicable to problems with both stationary and non-stationary sources of uncertainty.

Within this area, meta-learning approaches that enable few-shot adaptation have been studied with real systems for imitation learning [31, 32, 33, 34] and goal inference [35], but direct meta-RL in the real world has received comparatively little attention. Adapting to different environment parameters

has been explored in the sim2real setting for table-top hockey [36] and legged locomotion [37], and in the model-based RL setting for millirobot locomotion [38]. In Section 5.3, we demonstrate that our algorithm MELD can perform meta-RL trained from images in the real world.

Latent State Inference in RL. A significant challenge in real-world robotic learning is contending with the complex, high-dimensional, and noisy observations from the robot’s sensors. To handle general partial observability, recurrent policies can persist information over longer time horizons [39, 40, 41], while explicit state estimation approaches maintain a probabilistic belief over the current state of the agent and update it given experience [42, 43, 44, 45, 46, 47, 48]. In our experiments we focus on learning from image observations, which presents a state representation learning challenge that has been studied in detail. End-to-end deep RL algorithms can learn state representations implicitly, but currently suffer from poor sample efficiency due to the added burden of representation learning [49, 20, 50]. Pre-trained state estimation systems can predict potentially useful features such as object locations and pose [51, 52]; however, these approaches require ground truth supervision. On the other hand, unsupervised learning techniques can improve sample efficiency without access to additional supervision [53, 5, 54, 6, 55, 56, 57]. Latent dynamics models capture the time-dependence of observations and provide a learned latent space in which RL can be tractably performed [58, 59, 8, 60, 61, 7]. In our work, we generalize the learned latent variable to encode not only the state but also the task at hand, enabling efficient meta-RL from images.

3 Preliminaries: Meta-RL and Latent State Models

In this work, we leverage tools from latent state modeling to design an efficient meta-RL method that can operate in the real world from image observations. In this section, we review latent state models and meta-RL, developing a formalism that will allow us to derive our algorithm in Section 4.

3.1 POMDPs and Latent State Models

We first define the RL problem. A Markov decision process (MDP) consists of a set of states \mathcal{S} , a set of actions \mathcal{A} , an initial state distribution $p(\mathbf{s}_1)$, a state transition distribution $p(\mathbf{s}_{t+1}|\mathbf{s}_t, \mathbf{a}_t)$, a discount factor γ , and a reward function $r(\mathbf{s}_t, \mathbf{a}_t)$. We assume the transition and reward functions are unknown, but can be sampled by taking actions in the environment. The goal of RL is to learn a policy $\pi(\mathbf{a}_t|\mathbf{s}_t)$ that selects actions that maximize the sum of discounted rewards. However, robots operating in the real world do not have access to the underlying state \mathbf{s}_t , and must instead select actions using high-dimensional and often incomplete observations from cameras and other sensors. Such a system can be described as a partially observed Markov decision process (POMDP), where observations \mathbf{x}_t are a noisy or incomplete function of the unknown underlying state \mathbf{s}_t , and the policy is conditioned on a history of observations as $\pi(\mathbf{a}_t|\mathbf{x}_{1:t})$.

While end-to-end RL methods can acquire representations of observations from reward supervision alone [20, 21], the added burden of end-to-end representation learning limits sample efficiency and can make optimization more difficult. Methods that explicitly address this representation learning problem are more efficient and scale more efficiently to harder domains [47, 61, 7]. These approaches train latent state models to learn meaningful representations of the incoming observations by explicitly representing the unknown Markovian state as a hidden variable \mathbf{z}_t in a graphical model, as shown in Figure 2 (a). The parameters of these graphical models can be trained by approximately maximizing the log-likelihood of the observations: $\log p(\mathbf{x}_{1:T}|\mathbf{a}_{1:T-1}) = \log \int p(\mathbf{x}_T|\mathbf{z}_T)p(\mathbf{z}_T|\mathbf{z}_{T-1}, \mathbf{a}_{T-1})\dots p(\mathbf{z}_1)dz$. Given a history of observations and actions seen so far, the posterior distribution over the hidden variable captures the agent’s belief over the current underlying state, and can be written as $\mathbf{b}_t = p(\mathbf{z}_t|\mathbf{x}_{1:t}, \mathbf{a}_{1:t-1})$. Then, rather than conditioning the policy on raw observations, these methods learn a policy $\pi(\mathbf{a}_t|\mathbf{b}_t)$ as a function of this belief state.

3.2 Meta-Reinforcement Learning

In this work, we would like a robot to learn to acquire new skills quickly. We formalize this problem statement as meta-RL, where each task \mathcal{T} from a distribution of tasks $p(\mathcal{T})$ is a POMDP as described above, with initial state distribution $p_{\mathcal{T}}(\mathbf{s}_1)$, dynamics function $p_{\mathcal{T}}(\mathbf{s}_{t+1}|\mathbf{s}_t, \mathbf{a}_t)$, observation function $p_{\mathcal{T}}(\mathbf{x}_t|\mathbf{s}_t)$, and reward function $r_{\mathcal{T}}(r_t|\mathbf{s}_t, \mathbf{a}_t)$, as shown in Figure 2 (b). For example, a task distribution that varies both dynamics and rewards across tasks may consist of placing mugs of varying weights (dynamics) in different locations (rewards) on a kitchen shelf. The meta-training process leverages a set of training tasks sampled from $p(\mathcal{T})$ to learn an adaptation procedure that

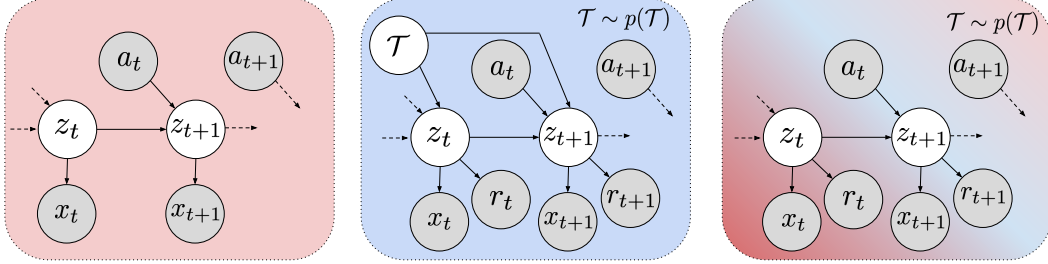


Figure 2: (a): When only partial observations of the underlying state are available, latent dynamics models can glean state information \mathbf{z}_t from a history of observations. (b): Meta-RL considers a task distribution where the current task \mathcal{T} is an unobserved variable that controls dynamics and rewards. (c): We interpret \mathcal{T} as part of \mathbf{z}_t , allowing us to leverage latent dynamics models for efficient image-based meta-RL.

can adapt to a new task from $p(\mathcal{T})$ using a small amount of experience. Similar to the framework in probabilistic inference-based and recurrence-based meta-RL approaches [28, 4, 2, 1], we formalize the adaptation procedure f_ϕ as a function of experience $(\mathbf{x}_{1:t}, r_{1:t}, \mathbf{a}_{1:t-1})$ that summarizes task-relevant information into the variable \mathbf{c}_t . The policy is conditioned on this updated variable as $\pi_\theta(\mathbf{a}_t | \mathbf{x}_t, \mathbf{c}_t)$ to adapt to the task. By training the adaptation mechanism f_ϕ and the policy π_θ end-to-end to maximize returns of the adapted policy, meta-RL algorithms can learn policies that effectively modulate and adapt their behavior with small amounts of experience in new tasks. We formalize this meta-RL objective as:

$$\max_{\theta, \phi} \mathbb{E}_{\mathcal{T} \sim p(\mathcal{T})} \mathbb{E}_{\substack{\mathbf{x}_t \sim p_{\mathcal{T}}(\cdot | \mathbf{s}_t) \\ \mathbf{a}_t \sim \pi_\theta(\cdot | \mathbf{x}_t, \mathbf{c}_t) \\ \mathbf{s}_{t+1} \sim p_{\mathcal{T}}(\cdot | \mathbf{s}_t, \mathbf{a}_t) \\ r_t \sim r_{\mathcal{T}}(\cdot | \mathbf{s}_t, \mathbf{a}_t)}} \left[\sum_{t=1}^T \gamma^t r_t \right] \quad \text{where } \mathbf{c}_t = f_\phi(\mathbf{x}_{1:t}, r_{1:t}, \mathbf{a}_{1:t-1}). \quad (1)$$

Meta-RL methods may differ in how the adaptation procedure f_ϕ is represented (e.g., as probabilistic inference [4, 28], as a recurrent update [2, 1], as a gradient step [3]), how often the adaptation procedure occurs (e.g., at every timestep [2, 28] or once per episode [4, 30]), and also in how the optimization is performed (e.g., on-policy [2], off-policy [4]). Differences aside, these methods all typically optimize this objective end-to-end, creating a representation learning bottleneck when learning from image inputs that are ubiquitous in real-world robotics. In the following section, we show how the latent state models discussed in Section 3.1 can be re-purposed for joint representation and task learning, and how this insight leads to a practical algorithm for image-based meta-RL.

4 MELD

In this section, we present MELD: an efficient algorithm for meta-RL from images. We first develop the algorithm in Section 4.1 and then describe its implementation in Section 4.2.

4.1 Meta-RL with Latent Dynamics Models

To see how task inference in meta-RL can be cast as latent state inference, consider the graphical models depicted in Figure 2. Panel (a) illustrates a standard POMDP with underlying latent state \mathbf{z}_t and observations \mathbf{x}_t , and panel (b) depicts standard meta-RL, where the hidden task variable \mathcal{T} is assumed constant throughout the episode. In the meta-RL setting, the policy must then be conditioned on both the observation and the task variable in order to adapt to a new task (see Equation 1). Casting this task variable as part of the latent state, panel (c) illustrates our graphical model, where the states \mathbf{z}_t now contain both state and task information. In effect, we cast the task distribution over POMDPs as a POMDP itself, where the state variables now additionally capture task information. This meld allows us to draw on the rich literature of latent state models discussed in Section 3.1, and use them here to tackle the problem of meta-RL from sensory observations. Note that the task is not explicitly handled since it is simply another hidden state variable, providing a seamless integration of meta-RL with learning from sensory observations.

Concretely, we learn a latent state model over hidden variables by optimizing the log-likelihood of the evidence (observations and rewards) in the graphical model in Figure 2c:

$$\max_{\phi} \mathbb{E}_{\mathcal{T} \sim p(\mathcal{T})} \mathbb{E}_{\substack{\mathbf{x}_t \sim p_{\mathcal{T}}(\cdot | \mathbf{s}_t) \\ \mathbf{a}_t \sim \pi_{\theta}(\cdot | \mathbf{b}_t) \\ \mathbf{s}_{t+1} \sim p_{\mathcal{T}}(\cdot | \mathbf{s}_t, \mathbf{a}_t) \\ r_t \sim r_{\mathcal{T}}(\cdot | \mathbf{s}_t, \mathbf{a}_t)}} [\log p_{\phi}(\mathbf{x}_{1:T}, r_{1:T} | \mathbf{a}_{1:T-1})]. \quad (2)$$

Note that the only change from the latent state model from Section 3.1 is the inclusion of rewards as part of the observed evidence. While this change appears simple, it enables meta-learning by allowing the hidden state to capture task information. Posterior inference in this model then gives the agent’s belief $\mathbf{b}_t = p(\mathbf{z}_t | \mathbf{x}_{1:t}, r_{1:t}, \mathbf{a}_{1:t-1})$ over latent state and task variables \mathbf{z}_t . Conditioned on this belief, the policy $\pi_{\theta}(\mathbf{a}_t | \mathbf{b}_t)$ can learn to adapt its behavior to the task. Prescribing the adaptation procedure f_{ϕ} from Equation 1 to be posterior inference in our latent state model, the meta-training objective in MELD is:

$$\max_{\theta} \mathbb{E}_{\mathcal{T} \sim p(\mathcal{T})} \mathbb{E}_{\substack{\mathbf{x}_t \sim p_{\mathcal{T}}(\cdot | \mathbf{s}_t) \\ \mathbf{a}_t \sim \pi_{\theta}(\cdot | \mathbf{b}_t) \\ \mathbf{s}_{t+1} \sim p_{\mathcal{T}}(\cdot | \mathbf{s}_t, \mathbf{a}_t) \\ r_t \sim r_{\mathcal{T}}(\cdot | \mathbf{s}_t, \mathbf{a}_t)}} \left[\sum_{t=1}^T \gamma^t r_t \right] \quad (3)$$

$$\text{where } \mathbf{b}_t = p(\mathbf{z}_t | \mathbf{x}_{1:t}, r_{1:t}, \mathbf{a}_{1:t-1}). \quad (4)$$

By melding state and task inference, MELD inherits the same representation learning mechanism as latent state models discussed in Section 3.1 to enable efficient meta-RL with images.

4.2 Implementing MELD

Exactly computing the posterior distribution over the latent state variable is intractable, so we take a variational inference approach to maximize a lower bound on the log-likelihood objective [62] in Equation 2. We factorize the variational posterior as $q(\mathbf{z}_{1:T} | \mathbf{x}_{1:T}, r_{1:T}, \mathbf{a}_{1:T-1}) = q(\mathbf{z}_T | \mathbf{x}_T, r_T, \mathbf{z}_{T-1}, \mathbf{a}_{T-1}) \dots q(\mathbf{z}_2 | \mathbf{x}_2, r_2, \mathbf{z}_1, \mathbf{a}_1) q(\mathbf{z}_1 | \mathbf{x}_1, r_1)$. With this factorization, we implement each component as a deep neural network and optimize the evidence lower bound of the joint objective, where $\mathbb{E}_{\mathbf{z}_{1:t} \sim q_{\phi}} [\log p(\mathbf{x}_{1:T}, r_{1:T} | \mathbf{a}_{1:T-1})] \geq \mathcal{L}_{model}$, with \mathcal{L}_{model} defined as:

$$\begin{aligned} \mathcal{L}_{model}(\mathbf{x}_{1:T}, r_{1:T}, \mathbf{a}_{1:T-1}) = & \mathbb{E}_{\mathbf{z}_{1:T} \sim q_{\phi}} \sum_{t=1}^T \log p_{\phi}(\mathbf{x}_t | \mathbf{z}_t) + \log p_{\phi}(r_t | \mathbf{z}_t) \\ & - D_{KL}(q_{\phi}(\mathbf{z}_1 | \mathbf{x}_1, r_1) || p(\mathbf{z}_1)) - \sum_{t=2}^T D_{KL}(q_{\phi}(\mathbf{z}_t | \mathbf{x}_t, r_t, \mathbf{z}_{t-1}, \mathbf{a}_{t-1}) || p_{\phi}(\mathbf{z}_t | \mathbf{z}_{t-1}, \mathbf{a}_{t-1})). \quad (5) \end{aligned}$$

The first two terms encourage a rich latent representation \mathbf{z}_t by requiring reconstruction, while the last term keeps the inference network consistent with latent dynamics. The first timestep posterior $q_{\phi}(\mathbf{z}_1 | \mathbf{x}_1, r_1)$ is modeled separately from the remaining steps, and $p(\mathbf{z}_1)$ is chosen to be a fixed unit Gaussian $\mathcal{N}(0, I)$. The learned inference networks $q_{\phi}(\mathbf{z}_1 | \mathbf{x}_1, r_1)$ and $q_{\phi}(\mathbf{z}_t | \mathbf{x}_t, r_t, \mathbf{z}_{t-1}, \mathbf{a}_{t-1})$, decoder networks $p_{\phi}(\mathbf{x}_t | \mathbf{z}_t)$ and $p_{\phi}(r_t | \mathbf{z}_t)$, and dynamics $p_{\phi}(\mathbf{z}_t | \mathbf{z}_{t-1}, \mathbf{a}_{t-1})$ are all fully connected networks that output parameters of Gaussian distributions. We follow the architecture of the latent variable model from SLAC [7] and provide the remaining implementation details in Appendix A.

Algorithm 1 MELD Meta-training

Require: Training tasks $\{\mathcal{T}_i\}_{i=1 \dots J}$ from $p(\mathcal{T})$, learning rates η_1, η_2, η_3

- 1: Init. model p_{ϕ}, q_{ϕ} , actor π_{θ} , critic Q_{ζ}
- 2: Init. replay buffers \mathcal{B}_i for each train task
- 3: **while** not done **do**
- 4: **for** each task \mathcal{T}_i **do** ▷ collect data
- 5: Infer $\mathbf{b}_1 = q_{\phi}(\mathbf{z}_1 | \mathbf{x}_1, r_1)$
- 6: Step with $\mathbf{a}_1 \sim \pi_{\theta}(\mathbf{a} | \mathbf{b}_1)$, get \mathbf{x}_2, r_2
- 7: **for** $t = 2, \dots, T-1$ **do**
- 8: Infer $\mathbf{b}_t = q_{\phi}(\mathbf{z}_t | \mathbf{x}_t, r_t, \mathbf{z}_{t-1}, \mathbf{a}_{t-1})$
- 9: Step $\mathbf{a}_t \sim \pi_{\theta}(\mathbf{a} | \mathbf{b}_t)$, get $\mathbf{x}_{t+1}, r_{t+1}$
- 10: **end for**
- 11: Add data $\{\mathbf{x}_{1:T}, r_{1:T}, \mathbf{a}_{1:T-1}\}$ to \mathcal{B}^i
- 12: **end for**
- 13: **for** step in train steps **do**
- 14: **for** each task \mathcal{T}_i **do**
- 15: Sample $\{\mathbf{x}_{1:T}, r_{1:T}, \mathbf{a}_{1:T-1}\} \sim \mathcal{B}_i$
- 16: Infer beliefs $\mathbf{b}_{1:T} = \{q_{\phi}(\mathbf{z}_t | \dots)\}_{1:T}$
- 17: Predict reconstructions $\{\hat{\mathbf{x}}_t, \hat{r}_t\}_{1:T}$
- 18: $\mathcal{L}_m^i = \mathcal{L}_{model}(\{\mathbf{x}_t, \hat{\mathbf{x}}_t, r_t, \hat{r}_t, \mathbf{a}_t\}_{1:T})$
- 19: **end for**
- 20: $\phi \leftarrow \phi - \eta_1 \nabla_{\phi} \sum_i \mathcal{L}_m^i$ ▷ train model
- 21: Update θ, ζ with SAC(η_2, η_3) ▷ train AC
- 22: **end for**
- 23: **end while**

Algorithm 2 MELD Meta-testing

Require: Test task $\mathcal{T} \sim p(\mathcal{T})$

- 1: Infer $\mathbf{b}_1 = q_{\phi}(\mathbf{z}_1 | \mathbf{x}_1, r_1)$
- 2: Step with $\mathbf{a}_1 \sim \pi_{\theta}(\mathbf{a} | \mathbf{b}_1)$, get \mathbf{x}_2, r_2
- 3: **for** $t = 2, \dots, T-1$ **do**
- 4: Infer $\mathbf{b}_t = q_{\phi}(\mathbf{z}_t | \mathbf{x}_t, r_t, \mathbf{z}_{t-1}, \mathbf{a}_{t-1})$
- 5: Step with $\mathbf{a}_t \sim \pi_{\theta}(\mathbf{a} | \mathbf{b}_t)$, get $\mathbf{x}_{t+1}, r_{t+1}$
- 6: **end for**

We use the soft actor-critic (SAC) [63] RL algorithm in this work, due to its high sample efficiency and performance. The actor $\pi_\theta(\mathbf{a}_t|\mathbf{b}_t)$ and the critic $Q_\psi(\mathbf{b}_t, \mathbf{a}_t)$ are conditioned on the posterior belief \mathbf{b}_t , modeled as fully connected neural networks, and trained as prescribed by the SAC algorithm. During meta-training, MELD alternates between collecting data with the current policy, training the model by optimizing \mathcal{L}_{model} , and training the policy with the current model. Meta-training and meta-testing are described in Algorithm 1 and 2 respectively.




Figure 3: MELD meta-training alternates between collecting data with π_θ and training the latent state model p_ϕ , inference networks q_ϕ , actor π_θ , and critic Q_ψ .

5 Experiments

In our experiments, we aim to answer the following: **(1)** How does MELD compare to prior meta-RL methods in enabling fast acquisition of new skills at test time in challenging simulated control problems? **(2)** Can MELD meta-learn effective exploration when only sparse task completion rewards are available at meta-test time? **(3)** Can MELD enable real robots to quickly acquire skills via meta-RL from images?

5.1 MELD in Simulated Environments

In this section, we evaluate MELD on the four simulated image-based continuous control problems in Figure 4. In (a) Cheetah-vel, each task is a different target running velocity for the 6-DoF legged robot. Reward is the difference in robot velocity from the target. The remainder of the problems use a 7-DOF Sawyer robotic arm. In (b) Reacher, each task is a different goal position for the end-effector. In (c) Peg-insertion, the robot must insert the peg into the correct box, where each task varies the goal box as well as locations of all four boxes. In (d) Shelf-placing, each task varies the weight (dynamics change) and target location (reward change) of a mug that the robot must move to the shelf. For the Sawyer environments, the reward function is the negative distance between the robot end-effector and the desired end location. For all environments, we train with 30 meta-training tasks and evaluate on 10 meta-test tasks from the same distribution that are not seen during training.




Figure 4: Simulated locomotion and manipulation meta-RL environments in the MuJoCo simulator [64]; goals illustrated for visualization purposes.

We compare MELD to two representative state-of-the-art meta-RL algorithms, PEARL [4] and RL² [2]. PEARL models a belief over a probabilistic latent task variable as a function of un-ordered batches of transitions, and conditions the policy on both the current observation and this inferred task belief. Unlike MELD, this algorithm assumes an exploration phase of several trajectories in the new task to gather information before adapting, so to get its best performance, we evaluate *only after* this exploration phase. RL² models the policy as a recurrent network that directly maps observations, actions, and rewards to actions. To apply PEARL and RL² with image observations, we augment them with the same convolutional encoder architecture used by MELD. Finally, to verify the need for task inference to solve new tasks, we compare to SLAC [7], which infers state information from a sequence of observations but does *not* perform meta-learning.

In Figure 5 we plot average performance on meta-test tasks over the course of meta-training, across 3 random seeds. See Appendix B for the the definitions of metrics used for each task. MELD achieves the highest performance in each environment and is the only method to fully solve Cheetah-vel, Peg-insertion, and Shelf-placing. The SLAC baseline fails in this meta-RL setting, as expected,




Figure 5: Rewards on test tasks versus meta-training environment steps, comparing MELD to prior methods. Black dashed line indicates consistent task success (see Appendix B for definition).

with qualitative behavior of always executing a single “average” motion, such as reaching toward a mean goal location and running at a medium speed. PEARL aggregates task information over time in its latent task variable, but relies on the current observation alone for state information. Its poor performance on Cheetah, Reacher, and Shelf reflect the need for state estimation from a sequence of observations to perform control in these environments. While RL² is capable of propagating both state and task information over time, we observe that it overfits heavily to training tasks and struggles on evaluation tasks.

5.2 Temporally-Extended Exploration

The previous section assumed a shaped reward function that is the negative distance between the current robot position and the desired one at every timestep. In the real world, this type of reward function is typically not available to the agent, since it requires information that may be difficult or impossible to obtain. For example, in the Ethernet cable insertion problem, the location of the insertion point is unknown, but the agent might receive a sparse task completion reward upon making the correct electrical connection. To quickly succeed at a new task given only this sparse signal at meta-test time, it is critical for MELD to reason over multiple episodes and acquire temporally-extended strategies during meta-training. Because RL with sparse rewards is very inefficient, we follow prior work [65, 4] and assume access to a shaped reward function during meta-training to help learn these strategies. We detail our particular approach to making use of shaped rewards during meta-training in Appendix C, and at meta-test time assume access to only the sparse reward signal.

To evaluate MELD in this setting, we design a simulated button-pressing environment with the Sawyer robot, where the button to push and the location of the panel changes with each task. Sparse reward is given only when the correct button is pushed, while shaped reward (used only in meta-training) is the negative distance from the robot’s end-effector to the insertion point. In Figure 6, we analyze the qualitative behavior of MELD when learning a new task at test time. Though the shaped reward is not used at test time, we plot the shaped reward reconstruction mean and variance to gain an insight into the contents of the learned latent state. In the first episode, predicted reward error and variance are high until the robot presses the correct button. MELD’s latent state model persists the task information to the second episode (predicted reward error and variance are very low) and the robot navigates immediately to the correct button. In Figure 7, we compare MELD to the same baselines introduced in the previous section and find that it is the only method able to press the correct button within two trajectories of experience. In the next section, we test MELD’s capability to perform such exploration and exploitation in the real world.

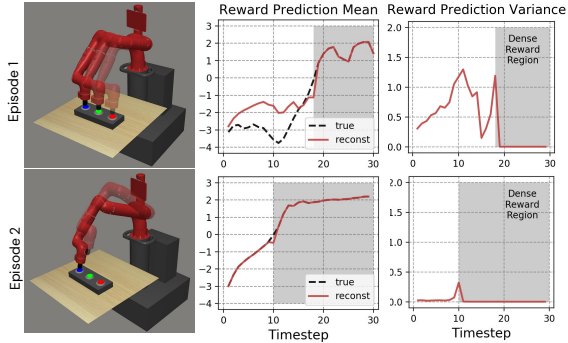


Figure 6: Button press with reward given only upon pressing correct button. The robot explores each button until it finds the correct one (top) and returns to that button immediately in the next episode (bottom). See text for discussion.




Figure 7: Comparison to approaches from Figure 5 on the button-pressing environment, showing test task performance versus meta-training environment steps.

5.3 MELD in the Real World

We now evaluate MELD on a real-world 5-DoF WidowX arm performing Ethernet cable insertion. The task distribution consists of different ports in a router that also varies in location and orientation (see Figure 8). To instrument these tasks in the real world, we build an automatic reset mechanism that moves and rotates the router, as detailed in Appendix E. At meta-test time the reward is a sparse signal given when the robot inserts the cable in the correct port. As in the previous section, during meta-training we make use of a shaped reward function that is the sum of the L2-norms of translational and rotational distances between the pose of the object in the end-effector and a goal pose. The agent’s observations are concatenated images from two webcams (Figure 9): one fixed view and one first-person view from a wrist-mounted camera. The policy sends joint velocity controls over a ROS interface to a low-level PID controller to move the joints of each robot.

We compare MELD to SLAC as described in Section 5.1 as well as a random policy, and plot the results in Figure 10. After training across 20 meta-training tasks using a total of 8 hours (80,000 samples at 3.3Hz) worth of data, MELD achieves a success rate of 90% over three rounds of evaluation in each of the 10 randomly sampled evaluation tasks that were not seen during training. To our knowledge, this experiment is the first demonstration of meta-RL trained entirely in the real world from image observations. We also conducted experiments with the Sawyer robot, finding that MELD enables the Sawyer to insert a peg into the correct hole given a per-timestep reward of distance to the hole (see Appendix F). Due to lab access restrictions as a result of COVID-19, we could not evaluate adaptation to new tasks on this platform. Videos of all experiments can be found on our project website.²




Figure 8: Ethernet cable insertion problem. The automatic reset mechanism changes the position and orientation of the router across tasks.




Figure 9: 64x128 image observations seen by WidowX.

6 Discussion

In this paper, we leverage the insight that meta-RL can be cast into the framework of latent state inference. This allows us to combine the fast skill-acquisition capabilities of meta-learning with the efficiency of unsupervised latent state models when learning from raw image observations. Based on this principle, we design MELD, a practical algorithm for meta-RL with image observations. MELD outperforms prior methods on simulated locomotion and manipulation tasks, and is efficient enough to perform meta-RL directly from images in the real world. However, neither our approach nor prior meta-learning works have shown convincing generalization to wider task distributions of qualitatively distinct manipulation tasks. Additionally, due to the difficulty of defining and instrumenting shaped reward functions for such tasks, it is important that meta-RL algorithms be able to learn from less structured signals and other forms of supervision, such as demonstrations and natural language. We view these directions as exciting avenues of future work that will broaden the applicability of robotic learning to less structured environments.




Figure 10: **(left)** Training task rewards versus meta-training environment steps. **(right)** Task success (defined as full insertion into correct hole) on held-out tasks after meta-training.

²<https://sites.google.com/view/meld-lsm/home>

References

- [1] J. X. Wang, Z. Kurth-Nelson, D. Tirumala, H. Soyer, J. Z. Leibo, R. Munos, C. Blundell, D. Kumaran, and M. Botvinick. Learning to reinforcement learn. *arXiv preprint arXiv:1611.05763*, 2016.
- [2] Y. Duan, J. Schulman, X. Chen, P. L. Bartlett, I. Sutskever, and P. Abbeel. RL2: Fast reinforcement learning via slow reinforcement learning. *arXiv preprint arXiv:1611.02779*, 2016.
- [3] C. Finn, P. Abbeel, and S. Levine. Model-agnostic meta-learning for fast adaptation of deep networks. *arXiv preprint arXiv:1703.03400*, 2017.
- [4] K. Rakelly, A. Zhou, D. Quillen, C. Finn, and S. Levine. Efficient off-policy meta-reinforcement learning via probabilistic context variables. In *ICML*, 2019.
- [5] C. Finn, X. Y. Tan, Y. Duan, T. Darrell, S. Levine, and P. Abbeel. Deep spatial autoencoders for visuomotor learning. In *ICRA*, pages 512–519. IEEE, 2016.
- [6] A. Ghadirzadeh, A. Maki, D. Kragic, and M. Björkman. Deep predictive policy training using reinforcement learning. In *IROS*, pages 2351–2358. IEEE, 2017.
- [7] A. X. Lee, A. Nagabandi, P. Abbeel, and S. Levine. Stochastic latent actor-critic: Deep reinforcement learning with a latent variable model. *arXiv preprint arXiv:1907.00953*, 2019.
- [8] M. Zhang, S. Vikram, L. Smith, P. Abbeel, M. Johnson, and S. Levine. Solar: Deep structured representations for model-based reinforcement learning. In *ICML*, pages 7444–7453, 2019.
- [9] J. Schmidhuber. *Evolutionary principles in self-referential learning, or on learning how to learn: the meta-meta-... hook*. PhD thesis, Technische Universität München, 1987.
- [10] S. Thrun, W. Burgard, D. Fox, et al. Probabilistic robotics, vol. 1, 2005.
- [11] J. B. Tenenbaum. *A Bayesian framework for concept learning*. PhD thesis, Massachusetts Institute of Technology, 1999.
- [12] L. Fei-Fei et al. A bayesian approach to unsupervised one-shot learning of object categories. In *Proceedings Ninth IEEE International Conference on Computer Vision*, pages 1134–1141. IEEE, 2003.
- [13] A. Bicchi and V. Kumar. Robotic grasping and contact: A review. In *ICRA*, volume 1, pages 348–353. IEEE, 2000.
- [14] G. A. Pereira, M. F. Campos, and V. Kumar. Decentralized algorithms for multi-robot manipulation via caging. *IJRR*, 23(7-8):783–795, 2004.
- [15] D. Henrich and H. Wörn. *Robot manipulation of deformable objects*. Springer Science & Business Media, 2012.
- [16] K. Kronander, E. Burdet, and A. Billard. Task transfer via collaborative manipulation for insertion assembly. In *WHRI*. Citeseer, 2014.
- [17] W. S. Newman, Y. Zhao, and Y.-H. Pao. Interpretation of force and moment signals for compliant peg-in-hole assembly. In *ICRA*, volume 1, pages 571–576. IEEE, 2001.
- [18] J. Kober, J. A. Bagnell, and J. Peters. Reinforcement learning in robotics: A survey. *IJRR*, 32(11):1238–1274, 2013.
- [19] V. Gullapalli, J. A. Franklin, and H. Benbrahim. Acquiring robot skills via reinforcement learning. *IEEE CSM*, 14(1):13–24, 1994.
- [20] S. Levine, C. Finn, T. Darrell, and P. Abbeel. End-to-end training of deep visuomotor policies. *JMLR*, 17(1):1334–1373, 2016.
- [21] A. Zeng, S. Song, S. Welker, J. Lee, A. Rodriguez, and T. Funkhouser. Learning synergies between pushing and grasping with self-supervised deep rl. *arXiv preprint arXiv:1803.09956*, 2018.
- [22] M. A. Lee, Y. Zhu, K. Srinivasan, P. Shah, S. Savarese, L. Fei-Fei, A. Garg, and J. Bohg. Making sense of vision and touch: Self-supervised learning of multimodal representations for contact-rich tasks. *arXiv preprint arXiv:1810.10191*, 2018.
- [23] G. Schoettler, A. Nair, J. Luo, S. Bahl, J. A. Ojea, E. Solowjow, and S. Levine. Deep reinforcement learning for industrial insertion tasks with visual inputs and natural reward signals. In *ICLR*, 2019.
- [24] N. Mishra, M. Rohaninejad, X. Chen, and P. Abbeel. Meta-learning with temporal convolutions. *arXiv:1707.03141*, 2017.
- [25] J. Rothfuss, D. Lee, I. Clavera, T. Asfour, and P. Abbeel. Promp: Proximal meta-policy search. *arXiv preprint arXiv:1810.06784*, 2018.
- [26] R. Houthoofd, R. Y. Chen, P. Isola, B. C. Stadie, F. Wolski, J. Ho, and P. Abbeel. Evolved policy gradients. In *NeurIPS*, 2018.
- [27] R. Mendonca, A. Gupta, R. Kralew, P. Abbeel, S. Levine, and C. Finn. Guided meta-policy search. *arXiv preprint arXiv:1904.00956*, 2019.
- [28] L. Zintgraf, K. Shiarlis, M. Igl, S. Schulze, Y. Gal, K. Hofmann, and S. Whiteson. Varibad: A very good method for bayes-adaptive deep rl via meta-learning. *arXiv preprint arXiv:1910.08348*, 2019.
- [29] C. F. Perez, F. P. Such, and T. Karaletsos. Generalized hidden parameter mdps transferable model-based rl in a handful of trials. *arXiv preprint arXiv:2002.03072*, 2020.
- [30] J. Humplik, A. Galashov, L. Hasenclever, P. A. Ortega, Y. W. Teh, and N. Heess. Meta reinforcement learning as task inference. *arXiv preprint arXiv:1905.06424*, 2019.
- [31] C. Finn, T. Yu, T. Zhang, P. Abbeel, and S. Levine. One-shot visual imitation learning via meta-learning. In *CoRL*, pages 357–368, 2017.

- [32] T. Yu, C. Finn, A. Xie, S. Dasari, T. Zhang, P. Abbeel, and S. Levine. One-shot imitation from observing humans via domain-adaptive meta-learning. In *ICLR*, 2018.
- [33] S. James, M. Bloesch, and A. J. Davison. Task-embedded control networks for few-shot imitation learning. In *CoRL*, pages 783–795, 2018.
- [34] A. Bonardi, S. James, and A. J. Davison. Learning one-shot imitation from humans without humans. *arXiv preprint arXiv:1911.01103*, 2019.
- [35] A. Xie, A. Singh, S. Levine, and C. Finn. Few-shot goal inference for visuomotor learning and planning. In *CoRL*, pages 40–52, 2018.
- [36] K. Arndt, M. Hazara, A. Ghadirzadeh, and V. Kyrki. Meta reinforcement learning for sim-to-real domain adaptation. *arXiv preprint arXiv:1909.12906*, 2019.
- [37] X. Song, Y. Yang, K. Choromanski, K. Caluwaerts, W. Gao, C. Finn, and J. Tan. Rapidly adaptable legged robots via evolutionary meta-learning. *arXiv preprint arXiv:2003.01239*, 2020.
- [38] A. Nagabandi, I. Clavera, S. Liu, R. S. Fearing, P. Abbeel, S. Levine, and C. Finn. Learning to adapt in dynamic, real-world environments through meta-rl. *arXiv preprint arXiv:1803.11347*, 2018.
- [39] N. Yadaiah and G. Sowmya. Neural network based state estimation of dynamical systems. In *IJCNN*, pages 1042–1049. IEEE, 2006.
- [40] N. Heess, J. J. Hunt, T. P. Lillicrap, and D. Silver. Memory-based control with recurrent neural networks. *arXiv preprint arXiv:1512.04455*, 2015.
- [41] M. Hausknecht and P. Stone. Deep recurrent q-learning for partially observable mdps. In *2015 AAAI Fall Symposium Series*, 2015.
- [42] L. P. Kaelbling, M. L. Littman, and A. R. Cassandra. Planning and acting in partially observable stochastic domains. *AI*, 101(1-2):99–134, 1998.
- [43] J. Pineau, G. Gordon, S. Thrun, et al. Point-based value iteration: An anytime algorithm for pomdps. In *IJCAI*, volume 3, pages 1025–1032, 2003.
- [44] S. Ross, J. Pineau, B. Chaib-draa, and P. Kreitmann. A bayesian approach for learning and planning in partially observable markov decision processes. *JMLR*, 12(May):1729–1770, 2011.
- [45] M. P. Deisenroth and J. Peters. Solving nonlinear continuous state-action-observation pomdps for mechanical systems with gaussian noise. 2012.
- [46] P. Karkus, D. Hsu, and W. S. Lee. Qmdp-net: Deep learning for planning under partial observability. In *NeurIPS*, pages 4694–4704, 2017.
- [47] M. Igl, L. Zintgraf, T. A. Le, F. Wood, and S. Whiteson. Deep variational reinforcement learning for pomdps. In *ICML*, pages 2122–2131, 2018.
- [48] K. Gregor and F. Besse. Temporal difference variational auto-encoder. *arXiv preprint arXiv:1806.03107*, 2018.
- [49] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.
- [50] A. Singh, L. Yang, K. Hartikainen, C. Finn, and S. Levine. End-to-end robotic reinforcement learning without reward engineering. *environment (eg, by placing additional sensors)*, 34:44, 2020.
- [51] J. Tremblay, T. To, B. Sundaralingam, Y. Xiang, D. Fox, and S. Birchfield. Deep object pose estimation for semantic robotic grasping of household objects. In *CoRL*, pages 306–316, 2018.
- [52] V. Kumar, T. Herman, D. Fox, S. Birchfield, and J. Tremblay. Contextual reinforcement learning of visuo-tactile multi-fingered grasping policies. *arXiv preprint arXiv:1911.09233*, 2019.
- [53] S. Lange, M. Riedmiller, and A. Voigtländer. Autonomous reinforcement learning on raw visual input data in a real world application. In *IJCNN*, pages 1–8. IEEE, 2012.
- [54] T. Schmidt, R. Newcombe, and D. Fox. Self-supervised visual descriptor learning for dense correspondence. *IEEE Robotics and Automation Letters*, 2(2):420–427, 2016.
- [55] P. Florence, L. Manuelli, and R. Tedrake. Self-supervised correspondence in visuomotor policy learning. *IEEE Robotics and Automation Letters*, 5(2):492–499, 2019.
- [56] D. Yarats, A. Zhang, I. Kostrikov, B. Amos, J. Pineau, and R. Fergus. Improving sample efficiency in model-free reinforcement learning from images. *arXiv preprint arXiv:1910.01741*, 2019.
- [57] A. Sax, B. Emi, A. R. Zamir, L. J. Guibas, S. Savarese, and J. Malik. Mid-level visual representations improve generalization and sample efficiency for learning visuomotor policies. In *CoRL*, 2019.
- [58] M. Watter, J. Springenberg, J. Boedecker, and M. Riedmiller. Embed to control: A locally linear latent dynamics model for control from raw images. In *NeurIPS*, pages 2746–2754, 2015.
- [59] M. Karl, M. Soelch, J. Bayer, and P. van der Smagt. Deep variational bayes filters: Unsupervised learning of state space models from raw data. In *ICLR*, 2016.
- [60] D. Hafner, T. Lillicrap, I. Fischer, R. Villegas, D. Ha, H. Lee, and J. Davidson. Learning latent dynamics for planning from pixels. In *ICML*, pages 2555–2565, 2019.
- [61] C. Gelada, S. Kumar, J. Buckman, O. Nachum, and M. G. Bellemare. Deepmdp: Learning continuous latent space models for representation learning. In *ICML*, pages 2170–2179, 2019.
- [62] M. J. Wainwright and M. I. Jordan. *Graphical models, exponential families, and variational inference*. Now Publishers Inc, 2008.
- [63] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. *arXiv preprint arXiv:1801.01290*, 2018.

- [64] E. Todorov, T. Erez, and Y. Tassa. Mujoco: A physics engine for model-based control. In *IRROS*, pages 5026–5033. IEEE, 2012.
- [65] A. Gupta, R. Mendonca, Y. Liu, P. Abbeel, and S. Levine. Meta-reinforcement learning of structured exploration strategies. In *NeurIPS*, pages 5302–5311, 2018.

A MELD Implementation Details

Here, we expand on Section 4.2 to provide further implementation details of our algorithm MELD. Please also refer to our open-source code: <https://github.com/tonyzhaozh/meld>.

As discussed in Section 4.2, the latent state model is comprised of latent dynamics distributions, posterior inference distributions, and generative distributions of observations and rewards. While most timesteps are processed by the same time-invariant dynamics model $p(\mathbf{z}_t|\mathbf{z}_{t-1}, \mathbf{a}_{t-1})$ and posterior inference network $q_\phi(\mathbf{z}_t|\mathbf{z}_{t-1}, \mathbf{x}_t, r_t, \mathbf{a}_{t-1})$, we use separate distributions to model the first time step of a trial: $p(\mathbf{z}_1|\mathbf{z}_0)$ and $q_\phi(\mathbf{z}_1|\mathbf{z}_0, \mathbf{x}_1, r_1)$. Note that in Sections 5.2 and 5.3, we assume trials contain 2 episodes, while in Section 5.1 trials contain 1 episode. Following SLAC, we implement two layers of latent variables (please refer to the SLAC paper [7], for more details). The network $q_\phi(\mathbf{z}_t|\mathbf{z}_{t-1}, \mathbf{x}_t, r_t, \mathbf{a}_{t-1})$ encodes image observations and the network $p(\mathbf{x}_t|\mathbf{z}_t)$ decodes them for the reconstruction loss. Both of these networks include the same convolutional architecture (the decoder simply the transpose of the encoder) that consists of five convolutional layers. The layers have 32, 64, 128, 256, and 256 filters and the corresponding filter sizes are 5, 3, 3, 3, 4. For environments in which the robot observes two images (such as scene and wrist camera), we concatenate the images and apply rectangular filters. All other model networks are fully connected and consist of 2 hidden layers of 32 units each. We use ReLU activations after each layer.

As discussed, we train the actor and critic using the SAC RL algorithm [63] with the belief state as input. The SAC algorithm maximizes discounted returns as well as policy entropy via policy iteration. The critic (Q-function) is trained to minimize the soft Bellman error, which takes the entropy of the policy into account in the backup. We instantiate the actor and critic as fully connected networks with 2 hidden layers of 256 units each. We follow the implementation of SAC, including the use of 2 Q-networks and the tanh actor output activation (please see the SAC paper [63] for more details).

Meta-training alternates between collecting data with the current model and policy, and training the model and actor-critic. Gradients from the actor-critic optimization do not flow into the latent state model. Before beginning this alternating scheme, we first train the latent state model on 60 trajectories of data collected with a random policy. Relevant hyper-parameters for meta-training can be found in Table 1.

Table 1: Meta-training hyper-parameters

Parameter	Value
num. training tasks	30
num. eval tasks	10
actor, critic learning rates	3e-4
model learning rate	1e-4
size of per-task replay buffers \mathcal{B}_i	1e5
num. tasks collect data	20
num. rollouts per task	1
num. train steps per epoch	640
num. tasks sample for update	20
model batch size	512
actor, critic batch size	512

B Simulation Experiment Details and Ablation Study

In this section we provide further details on the simulated experiments in Section 5.1. We also perform an ablation study of several design decisions in the MELD algorithm.

B.1 Success Metrics

We define a success metric for each environment that correlates with qualitatively solving the task: Cheetah-vel: within .2m/s of target velocity, Reacher: within 10cm of goal, Peg: complete insertion with 5cm variation possible inside the site, Shelf: mug within 5cm of goal. We use these success metrics because task reward is often misleading when averaged across a distribution of tasks; in peg-insertion, for example, the numerical difference between always inserting the peg correctly versus never inserting it can be as low as 0.1, since the distance between the center of the goal distribution and each goal is quite small and accuracy is required. In Figure 5, we plot this success threshold as a dashed black line.

B.2 Environment Details

In the Cheetah-vel environment, we control the robot by commanding the torques on the robot’s 6 joints. The reward function consists of the difference between the target velocity v_{target} and the current velocity v_x of the cheetah’s center of mass, as well a small control cost on the torques sent to the joints:

$$r_{\text{cheetal-vel}} = -|v_x - v_{\text{target}}| + 0.01\|a_t\|_2. \tag{6}$$

The episode length is 50 time steps, and the observation consists of a single 64x64 pixel image from a tracking camera (as shown in Fig. 11a), which sees a view of the full cheetah.

In all three Sawyer environments, we control the robot by commanding joint delta-positions for all 7 joints. The reward function indicates the difference between the current end-effector pose x_{ee} and a goal pose x_{goal} , as follows:

$$r_{\text{sawyer-envs}} = -(d^2 + \log(d + 1e-5)), \quad \text{where } d = \|x_{\text{ee}} - x_{\text{goal}}\|_2. \tag{7}$$

This reward function encourages precision near the goal, which is particularly important for the peg insertion task. We impose a maximum episode length of 40 time steps for these environments. The observations for all three of these environments consist of two images concatenated to form a 64x128 image: one from a fixed scene camera, and one from a wrist-mounted first-person view camera. These image observations for each environment are shown in Fig. 11b-d. The simulation time step and control frequency for each of these simulated environments is listed in Table 2.

Table 2: Simulation Environments

Environment	Sim. time step	Control freq.
Cheetah-vel	0.01	10Hz
Reacher	0.0025	4Hz
Peg-insert	0.0025	4Hz
Shelf-placing	0.0025	4Hz

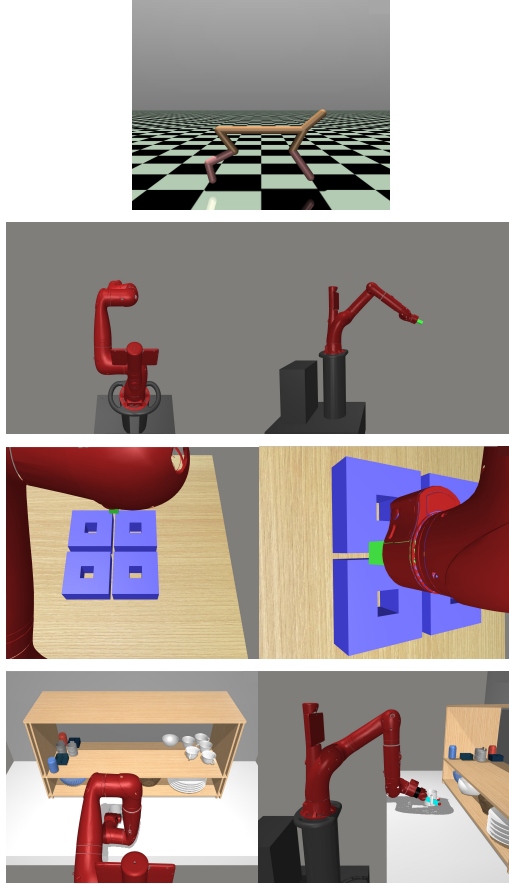


Figure 11: 64x64 and 64x128 image observations, seen as input by MELD for (a) Cheetah-vel, (b) Reacher, (c) Peg-insertion, and (d) Shelf-placing.

B.3 Ablation Study

In this section we ablate several hyper-parameters of the algorithm to understand which components affect its performance. First, we examine the effect of the number of meta-training tasks on MELD’s performance on test tasks. In Figure 12 we plot the performance on test tasks for different numbers of train tasks on the Reacher problem. We find that while generalization fails completely with only a single meta-training task, there are diminishing returns for increasing the number of training tasks beyond 20. This result demonstrates promising generalization, since a relatively low number of training tasks is actually needed in order to be able to solve the held-out test tasks. This result also has a practical benefit in that it precludes the need for instrumenting a large task distribution for meta-training in the real world, which can be cumbersome.

Although the reference implementation of SAC calls for each data collection step to be interleaved with a training step, this process of stopping after each environment step in order to perform training is not possible in the real world. Instead, in MELD, we collect batches of data that consist of full episodes, and we interleave these collection phases with training phases that consist of many gradient steps. Here, we examine the affect of this choice on the performance of the algorithm in simulation. From the results shown in Figure 13, we see that MELD is not too sensitive to scaling between this type of batched off-policy training and the original SAC-style on-policy training, which is essential for training in the real world.

In Figure 14, we examine the effect of the dimension of the latent variable z_t . These experiments show that a larger latent dimension performs better; we use dimension 256 in all our experiments.




Figure 12: **Number of meta-training tasks** Generalization to evaluation tasks requires training across set of tasks, but increasing beyond 20 tasks does not improve.




Figure 13: **Data / gradients ratio** MELD is not very sensitive to scaling the batch size of data collected, along with the amount of actor-critic training, at each iteration.




Figure 14: **Latent variable dimension** Latent state dimension affects task performance, larger is better.

C Sparse Reward Method and Experiment Details

As described in Section 5.2, when reward is given only upon completion of the task, efficient exploration is required to identify a new task within a few trials. The agent can acquire these exploration strategies during meta-training by learning strategies tailored to the task distribution. For example, in the button-pressing problem presented in Section 5.2, a learned exploration strategy might try pushing each button in succession, but would not try to e.g., pick up the control panel.

While in principle these behaviors can be acquired by MELD as described in Section 4.2, in practice performing RL with sparse rewards at meta-training time presents a significant exploration challenge. In effect, to learn useful exploration strategies for meta-test time, the agent must first explore effectively during meta-training. Because RL with sparse rewards is very inefficient, we follow prior work [65, 4] and assume access to a shaped reward function during meta-training to help learn these strategies. This setup corresponds to a setting where meta-training is performed in a laboratory with access to instrumentation to calculate the shaped reward, while meta-testing occurs outside the lab where only sparse rewards are available.

To make use of the shaped reward during meta-training time, we follow a two-stage procedure. First, we perform meta-training using the shaped reward as prescribed in Section 4.2. We then add data collected by this agent to the replay buffer of a second agent, which is trained with a small modification made to the model training loss from Equation 5. Here, the latent state model takes the **sparse** reward signal as input (to match our desired meta-testing setup), but we still use the shaped reward for the reconstruction target. We denote the shaped reward as \tilde{r} and the sparse reward as r , and highlight the difference from Equation 5 in blue.

$$\begin{aligned} \mathcal{L}_{model}(\mathbf{x}_{1:T}, r_{1:T}, \tilde{r}_{1:T}, \mathbf{a}_{1:T-1}) &= \mathbb{E}_{\mathbf{z}_{1:T} \sim q_\phi} \sum_{t=1}^T \log p_\phi(\mathbf{x}_t | \mathbf{z}_t) + \log p_\phi(\tilde{r}_t | \mathbf{z}_t) \\ &\quad - D_{\text{KL}}(q_\phi(\mathbf{z}_1 | \mathbf{x}_1, r_1) \| p(\mathbf{z}_1)) - \sum_{t=2}^T D_{\text{KL}}(q_\phi(\mathbf{z}_t | \mathbf{x}_t, r_t, \mathbf{z}_{t-1}, \mathbf{a}_{t-1}) \| p_\phi(\mathbf{z}_t | \mathbf{z}_{t-1}, \mathbf{a}_{t-1})). \end{aligned} \quad (8)$$

Additionally, we use the shaped reward to train the critic, since this is also not required at meta-test time.

Note that MELD does *not* simply learn to copy the trajectories from the shaped-reward training that were used to warmstart the sparse-reward training. The former (“expert”) trajectories move from the starting position directly to the correct button, since the shaped reward contains information to almost immediately identify the correct task. The MELD trajectories that result from receiving only sparse rewards as input, however, demonstrate systematic exploration of visiting different buttons in order to determine the correct one (see Figure 6). Finally, note that we use this same approach for the real-world WidowX experiments in Section 5.3.

D Importance of Performing Inference at Every Timestep

In this section, we discuss the benefit of the time-varying nature of MELD’s latent belief. We emphasize that this design decision is useful in the standard meta-RL setting, as well as in other realistic settings of the underlying task itself changing within an episode.

D.1 Fast Adaptation

We first present a didactic image-based 2D navigation problem to illustrate how MELD can learn extended exploration strategies to adapt to a new task, similar to the results shown in Figure 6. Here, the task distribution consists of goals located along a semi-circle around the start state. The agent receives inputs in the form of 64x64 image observations and rewards that are non-zero only upon reaching the correct goal. We use the same approach to using shaped reward for meta-training as described in Appendix C.

As shown in Figure 15a, the agent learns an efficient exploration strategy of traversing the semi-circle goal region until the goal is found. By updating the posterior belief at each step, MELD is able to find the goal within 10 – 20 steps, instead of multiple episodes as required by methods that explore via posterior sampling [4, 65] that hold the task variable constant across each episode. Furthermore, note that once the goal is found, MELD can navigate directly to it (Figure 15b) in the next episode.

This experiment demonstrates that even in standard meta-RL setting where the underlying task remains constant throughout an episode, updating task information at each timestep can enable faster adaptation. We argue that this behavior has safety benefits in the real world, since the agent need not complete full episodes of potentially hazardous exploration before incorporating task information.




Figure 15: Image-based 2D navigation with reward given after finding the right goal: (Left) trajectory traces of meta-learned exploration finding goal in Episode-1 (red) and going directly there in Episode-2 (blue). (Right) higher rewards in episode-2 show preservation of information across episodes.

D.2 Adapting to Task Changes within Episodes

The experiments in the main paper as well as the section above consider the standard meta-RL paradigm, where the agent adapts to one test task at a time. However, many realistic scenarios consist of a sequence of tasks. For example, consider a robot moving a mug filled with liquid; if some of the liquid spills, the robot must adapt to the new dynamics of the lighter mug to finish the job. Because MELD updates the belief over the hidden variables at each time step, it can be directly applied to this setting without modification. We evaluate MELD in the Cheetah-vel environment on a sequence of 3 different target velocities within a single episode and observe in Figure 16 that MELD adapts to track each velocity within a few time steps.




Figure 16: MELD tracking changing velocity targets for Cheetah-vel.

E Task Reset Mechanism for WidowX Experiments

Since meta-training requires training across a distribution of tasks, we build an automatic task reset mechanism for the real-world experiments with the WidowX robot performing ethernet cable insertion. This mechanism controls the translational and rotational displacement of the network switch. The network switch(A) is mounted to a 3D printed housing(B) with gear attached. We control the rotation of the housing through motor 1. This setup is then mounted on top of a linear rail(C) and motor 2 controls its translational displacement through a timing pulley. In our experiments, the

training task distribution consisted of 20 different tasks, where each task was randomly assigned from a rotational range of 16 degrees and a translational range of 2cm.




Figure 17: **Automatic task reset mechanism** The network switch is rotated and translated by a series of motors in order to generate different tasks for meta-learning. This allows our meta-learning process to be entirely automated, without needing human intervention to reset either the robot or the task at the beginning of each rollout.

F Sawyer Multi-task Peg Insertion

In this experiment, we demonstrate that MELD can reason jointly about state and task information to perform real-world peg insertion with a 7-DoF Sawyer robot (Figure 18, left). On the Sawyer robot, we learn precise peg insertion where the task distribution consists of three tasks, each corresponding to a different target box. Note that the goal is not provided to the agent, but must be inferred from its history of observations and rewards. The reward function is the sum of the L2-norms of translational and rotational distances between the pose of the object in the end-effector and a goal pose. The agent’s observations are concatenated images from two webcams (Figure 18, right): one fixed view and one first-person view from a wrist-mounted camera. The robot succeeds on all three tasks after training on 4 hours of data (60,000 samples at 4Hz), as shown in Figure 19. Videos of the experiment may be found on the project website.




Figure 18: **(left)** Performing precise peg insertion in the real world with a 7-DoF Sawyer robot. **(right)** 64x128 image observations seen as input by the Sawyer.




Figure 19: Rewards on train tasks during meta-training for Sawyer peg insertion.