

---

# Decoupling Exploration and Exploitation for Meta-Reinforcement Learning without Sacrifices

---

Evan Zheran Liu<sup>1</sup> Aditi Raghunathan<sup>1</sup> Percy Liang<sup>1</sup> Chelsea Finn<sup>1</sup>

## Abstract

The goal of meta-reinforcement learning (meta-RL) is to build agents that can quickly learn new tasks by leveraging prior experience on related tasks. Learning a new task often requires both exploring to gather task-relevant information and exploiting this information to solve the task. In principle, optimal exploration and exploitation can be learned end-to-end by simply maximizing task performance. However, such meta-RL approaches struggle with local optima due to a chicken-and-egg problem: learning to explore requires good exploitation to gauge the exploration’s utility, but learning to exploit requires information gathered via exploration. Optimizing separate objectives for exploration and exploitation can avoid this problem, but prior meta-RL exploration objectives yield suboptimal policies that gather information irrelevant to the task. We alleviate both concerns by constructing an exploitation objective that automatically identifies task-relevant information and an exploration objective to recover only this information. This avoids local optima in end-to-end training, without sacrificing optimal exploration. Empirically, DREAM substantially outperforms existing approaches on complex meta-RL problems, such as sparse-reward 3D visual navigation. Videos of DREAM: <https://ezliu.github.io/dream/>

## 1. Introduction

A general-purpose agent should be able to perform multiple related tasks across multiple related environments. Our goal is to develop agents that can perform a variety of tasks in novel environments, based on previous experience and only a small amount of experience in the new environment. For example, we may want a robot to cook a meal (a new task) in

---

<sup>1</sup>Department of Computer Science, Stanford University. Correspondence to: Evan Zheran Liu <evanliu@cs.stanford.edu>.

a new kitchen (the environment) after it has learned to cook other meals in other kitchens. To adapt to a new kitchen, the robot must both explore to find the ingredients, and use this information to cook. Existing meta-reinforcement learning (meta-RL) methods can adapt to new tasks and environments, but, as we identify in this work, struggle when adaptation requires complex exploration strategies.

In the meta-RL setting, the agent is presented with a set of meta-training problems, each in an environment (e.g., a kitchen) with some task (e.g., make pizza); at meta-test time, the agent is given a new, but related environment and task. It is allowed to gather information in a few initial exploration episodes, and its goal is to then maximize returns on all subsequent exploitation episodes, using this information. A common meta-RL approach is to learn to explore and exploit *end-to-end* by training a policy and updating exploration behavior based on how well the policy later exploits using the information discovered from exploration (Duan et al., 2016; Wang et al., 2016a; Stadie et al., 2018; Zintgraf et al., 2019; Humplik et al., 2019). With enough model capacity, such approaches can express optimal exploration and exploitation, but they create a chicken-and-egg problem that leads to bad local optima and poor sample efficiency: Learning to explore requires good exploitation to gauge the exploration’s utility, but learning to exploit requires information gathered via exploration. Therefore, with only final performance as signal, one cannot be learned without already having learned the other. For example, a robot chef is only incentivized to explore and find the ingredients if it already knows how to cook with those ingredients, but the robot can only learn to cook if it can already find the ingredients by exploration.

To avoid this chicken-and-egg problem, we propose to optimize separate objectives for exploration and exploitation by leveraging the *problem ID*—an easy-to-provide unique one-hot for each training meta-training task and environment. Such a problem ID can be realistically available in real-world meta-RL tasks: e.g., in a robot chef factory, each training kitchen (problem) can be easily assigned a unique ID, and in a personalized recommendation system, each user (problem) is typically identified by a unique username. Some prior works (Humplik et al., 2019; Kamienny et al., 2020) also use these problem IDs, but not in a way that avoids the chicken-and-egg problem. Others (Rakelly et al.,

2019; Zhou et al., 2019b; Gupta et al., 2018; Gurumurthy et al., 2019; Zhang et al., 2020) also optimize separate objectives, but their exploration objectives learn suboptimal policies that needlessly gather task-irrelevant information.

Instead, we propose an exploitation objective that automatically identifies task-relevant information, and an exploration objective to recover only this information. We learn an exploitation policy without the need for exploration, by conditioning on a learned representation of the problem ID, which provides task-relevant information (e.g., by memorizing the locations of the ingredients for each ID / kitchen). We apply an information bottleneck to this representation to encourage discarding of any information not required by the exploitation policy (i.e., task-irrelevant information). Then, we learn an exploration policy to only discover task-relevant information by training it to produce trajectories containing the same information as the learned ID representation (Section 4). Crucially, unlike prior work, we prove that our separate objectives are *consistent*: optimizing them yields optimal exploration and exploitation, assuming expressive-enough policy classes and enough meta-training data (Section 5.1).

Overall, this work’s main contribution is a consistent decoupled meta-RL algorithm, called **DREAM: Decoupling exploration and Exploitation in Meta-RL**, which overcomes the chicken-and-egg problem (Section 4.2). Theoretically, in a simple tabular example, we show that addressing the coupling problem with DREAM provably improves sample complexity over existing end-to-end approaches by a factor exponential in the horizon (Section 5). Empirically, we stress test DREAM’s ability to learn sophisticated exploration strategies on 3 challenging, didactic benchmarks and a sparse-reward 3D visual navigation benchmark. On these, DREAM learns to optimally explore and exploit, achieving 90% higher returns than existing state-of-the-art approaches (PEARL, E-RL<sup>2</sup>, IMPORT, VARIBAD), which struggle to learn an effective exploration strategy (Section 6).

## 2. Related Work

We draw on a rich literature on learning to adapt to related tasks (Schmidhuber, 1987; Thrun & Pratt, 2012; Naik & Mammone, 1992; Bengio et al., 1991; 1992; Hochreiter et al., 2001; Andrychowicz et al., 2016; Santoro et al., 2016). Many meta-RL works focus on adapting efficiently to a new task from few samples without optimizing the sample collection process, via updating the policy parameters (Finn et al., 2017; Agarwal et al., 2019; Yang et al., 2019; Houthoof et al., 2018; Mendonca et al., 2019), learning a model (Nagabandi et al., 2018; Sæmundsson et al., 2018; Hiraoka et al., 2020), multi-task learning (Fakoor et al., 2019), or leveraging demonstrations (Zhou et al., 2019a). In contrast, we focus on problems where targeted exploration is critical for few-shot adaptation.

Approaches that specifically explore to obtain the most informative samples fall into two main categories: *end-to-end* and *decoupled* approaches. End-to-end approaches optimize exploration and exploitation end-to-end by updating exploration behavior from returns achieved by exploitation (Duan et al., 2016; Wang et al., 2016a; Mishra et al., 2017; Rothfuss et al., 2018; Stadie et al., 2018; Zintgraf et al., 2019; Humplik et al., 2019; Kamienny et al., 2020; Dorfman & Tamar, 2020). These approaches can represent the optimal policy (Kaelbling et al., 1998), but they struggle to escape local optima due to a chicken-and-egg problem between learning to explore and learning to exploit (Section 4.1). Several of these approaches (Humplik et al., 2019; Kamienny et al., 2020) also leverage the problem ID during meta-training, but they still learn end-to-end, so the chicken-and-egg problem remains.

Decoupled approaches instead optimize separate exploration and exploitation objectives, via, e.g., Thompson-sampling (TS) (Thompson, 1933; Rakelly et al., 2019), obtaining exploration trajectories predictive of dynamics or rewards (Zhou et al., 2019b; Gurumurthy et al., 2019; Zhang et al., 2020), or exploration noise (Gupta et al., 2018). While these works do not identify the chicken-and-egg problem, decoupled approaches coincidentally avoid it. However, existing decoupled approaches, including those (Rakelly et al., 2019; Zhang et al., 2020) that leverage the problem ID, do not learn optimal exploration: TS (Rakelly et al., 2019) explores by guessing the task and executing a policy for that task, and hence cannot represent exploration behaviors that are different from exploitation (Russo et al., 2017). Predicting the dynamics (Zhou et al., 2019b; Gurumurthy et al., 2019; Zhang et al., 2020) is inefficient when only a small subset of the dynamics are relevant to solving the task. In contrast, we propose a separate mutual information objective for exploration, which both avoids the chicken-and-egg problem and yields optimal exploration when optimized (Section 5). Past work (Gregor et al., 2016; Houthoof et al., 2016; Eysenbach et al., 2018; Warde-Farley et al., 2018) also optimize mutual information objectives, but not for meta-RL.

Beyond meta-RL, learning a policy in the general RL setting (i.e., learning from scratch) also requires targeted exploration to gather informative samples. In contrast to exploration algorithms for general RL (Bellemare et al., 2016; Pathak et al., 2017; Burda et al., 2018; Leibfried et al., 2019), which must visit many novel states to find regions with high reward, exploration in meta-RL can be substantially more targeted by leveraging prior experience from related problems during meta-training. As a result, DREAM can learn new tasks in just *two* episodes (Section 6), while learning from scratch can require millions of episodes to learn a new task.




Figure 1. Meta-RL setting: Given a new environment and task, the agent is allowed to first explore and gather information (exploration episode), and then must use this information to solve the task (in subsequent exploitation episodes).

### 3. Preliminaries

**Meta-reinforcement learning.** The meta-RL setting considers a family of Markov decision processes (MDPs)  $\langle \mathcal{S}, \mathcal{A}, \mathcal{R}_\mu, \mathcal{T}_\mu \rangle$  with states  $\mathcal{S}$ , actions  $\mathcal{A}$ , rewards  $\mathcal{R}_\mu$ , and dynamics  $\mathcal{T}_\mu$ , indexed by a one-hot *problem ID*  $\mu \in \mathcal{M}$ , drawn from a distribution  $p(\mu)$ . Colloquially, we refer to the dynamics as the *environment*, the rewards as the *task*, and the entire MDP as the *problem*. Borrowing terminology from Duan et al. (2016), meta-training and meta-testing both consist of repeatedly running *trials*. Each trial consists of sampling a problem ID  $\mu \sim p(\mu)$  and running  $N+1$  episodes on the corresponding problem. Following prior evaluation settings (Finn et al., 2017; Rakelly et al., 2019; Rothfuss et al., 2018; Fakoor et al., 2019), we designate the first episode in a trial as an *exploration* episode consisting of  $T$  steps for gathering information, and define the goal as maximizing the returns in the subsequent  $N$  *exploitation* episodes (Figure 1). Following Rakelly et al. (2019); Humplik et al. (2019); Kamienny et al. (2020), the easy-to-provide problem ID is available for meta-training, but not meta-testing trials.

We formally express the goal in terms of an exploration policy  $\pi^{\text{exp}}$  used in the exploration episode and an exploitation policy  $\pi^{\text{task}}$  used in exploitation episodes, but these policies may be the same or share parameters. Rolling out  $\pi^{\text{exp}}$  in the exploration episode produces an exploration trajectory  $\tau^{\text{exp}} = (s_0, a_0, r_0, \dots, s_T)$ , which contains information discovered via exploration. The exploitation policy  $\pi^{\text{task}}$  may then condition on  $\tau^{\text{exp}}$  and optionally, its history across all exploitation episodes in a trial, to maximize exploitation episode returns. The goal is therefore to maximize:

$$\mathcal{J}(\pi^{\text{exp}}, \pi^{\text{task}}) = \mathbb{E}_{\mu \sim p(\mu), \tau^{\text{exp}} \sim \pi^{\text{exp}}} [V^{\text{task}}(\tau^{\text{exp}}; \mu)], \quad (1)$$

where  $V^{\text{task}}(\tau^{\text{exp}}; \mu)$  is the expected returns of  $\pi^{\text{task}}$  conditioned on  $\tau^{\text{exp}}$ , summed over the  $N$  exploitation episodes in a trial with problem ID  $\mu$ .

**End-to-end meta-RL.** A common meta-RL approach (Wang et al., 2016a; Duan et al., 2016; Rothfuss et al., 2018; Zintgraf et al., 2019; Kamienny et al., 2020; Humplik et al., 2019) is to learn to explore and exploit *end-to-end* by directly optimizing  $\mathcal{J}$  in (1), updating both from rewards achieved during exploitation. These approaches typically learn a single recurrent policy  $\pi(a_t | s_t, \tau_t)$  for both exploration and exploitation (i.e.,  $\pi^{\text{task}} =$

$\pi^{\text{exp}} = \pi$ ), which takes action  $a_t$  given state  $s_t$  and history of experiences spanning all episodes in a trial  $\tau_t = (s_0, a_0, r_0, \dots, s_{t-1}, a_{t-1}, r_{t-1})$ . Intuitively, this policy is learned by rolling out a trial, producing an exploration trajectory  $\tau^{\text{exp}}$  and, conditioned on  $\tau^{\text{exp}}$  and the exploitation experiences so far, yielding some exploitation episode returns. Then, credit is assigned to both exploration (producing  $\tau^{\text{exp}}$ ) and exploitation by backpropagating the exploitation returns through the recurrent policy. Directly optimizing the objective  $\mathcal{J}$  this way can learn optimal exploration and exploitation strategies, but optimization is challenging, which we detail in the next section.

### 4. Decoupling Exploration and Exploitation

In this section, we first illustrate how end-to-end optimization approaches face a chicken-and-egg problem between learning exploration and exploitation, leading to local optima and poor sample complexity (Section 4.1). Next, in Section 4.2, we propose DREAM to sidestep this chicken-and-egg problem by optimizing separate objectives for exploration and exploitation. Finally, we describe a practical implementation of DREAM in Section 4.3. Prior decoupled approaches also optimize separate exploration and exploitation objectives. However, crucially, as we show in the next section, the optimum of DREAM’s objectives maximizes returns, while the optimum of prior objectives does not.

#### 4.1. The Problem with Coupling Exploration and Exploitation

We begin by showing that end-to-end optimization struggles with local optima due to a chicken-and-egg problem, illustrated in Figure 2. Learning  $\pi^{\text{exp}}$  relies on gradients passed through  $\pi^{\text{task}}$ . If  $\pi^{\text{task}}$  cannot effectively solve the task, then these gradients will be uninformative. However, to learn to efficiently solve the task,  $\pi^{\text{task}}$  needs good exploration data (trajectories  $\tau^{\text{exp}}$ ) from a good exploration policy  $\pi^{\text{exp}}$ . This results in bad local optima as follows: if our current (sub-optimal)  $\pi^{\text{task}}$  obtains low rewards with a good informative trajectory  $\tau_{\text{good}}^{\text{exp}}$ , the low reward would cause  $\pi^{\text{exp}}$  to learn to *not* generate  $\tau_{\text{good}}^{\text{exp}}$ . This causes  $\pi^{\text{exp}}$  to instead generate trajectories  $\tau_{\text{bad}}^{\text{exp}}$  that lack information required to obtain high reward, further preventing the exploitation policy  $\pi^{\text{task}}$  from learning. Typically, early in training, both  $\pi^{\text{exp}}$  and  $\pi^{\text{task}}$  are




Figure 2. Coupling between the exploration policy  $\pi^{\text{exp}}$  and exploitation policy  $\pi^{\text{task}}$ . These policies are illustrated separately for clarity, but may be a single policy. Since the two policies depend on each other (for gradient signal and the  $\tau^{\text{exp}}$  distribution), it is challenging to learn one when the other policy has not learned.

suboptimal and hence will likely reach this local optimum.

More succinctly, estimates of the expected exploitation returns  $V^{\text{task}}(\tau^{\text{exp}}; \mu)$  in (1) (e.g., from value-function approximation) form the learning signal for exploration. Escaping the local optima requires accurately estimating  $V^{\text{task}}$ , which requires many episodes, leading to sample inefficiency. In Section 5.2, we illustrate this in a simple example.

## 4.2. DREAM: Decoupling Exploration and Exploitation in Meta-Reinforcement Learning

While we can sidestep the local optima of end-to-end training by optimizing separate objectives for exploration and exploitation, the challenge is to construct objectives that yield the same optimal solution as the end-to-end approach. We now discuss how we can use the problem IDs during meta-training to do so. Intuitively, a good exploration objective should encourage discovering task-relevant distinguishing attributes of the problem (e.g., ingredient locations), and ignoring task-irrelevant attributes (e.g., wall color). To create this objective, the key idea behind DREAM is to *learn* to extract only the task-relevant information from the problem ID, which encodes all information about the problem. Then, DREAM’s exploration objective seeks to recover only this task-relevant information.

Concretely, DREAM extracts only the task-relevant information from the problem ID  $\mu$  via a stochastic encoder  $F_\psi(z | \mu)$ . To learn this encoder, we train an exploitation policy  $\pi^{\text{task}}(a | s, z)$  to maximize rewards, conditioned on samples  $z \sim F_\psi(z | \mu)$ , while simultaneously applying an information bottleneck to  $z$  to discard information not needed by  $\pi^{\text{task}}$  (i.e., task-irrelevant information). Then, DREAM learns an exploration policy  $\pi^{\text{exp}}$  to produce trajectories with high mutual information with  $z$ . In this approach, the exploitation policy  $\pi^{\text{task}}$  no longer relies on effective exploration from  $\pi^{\text{exp}}$  to learn, and once  $F_\psi(z | \mu)$  is learned, the exploration policy also learns independently from  $\pi^{\text{task}}$ , decoupling the two optimization processes. During meta-testing,  $\mu$  is either unavailable or uninformative because it is simply a novel one-hot ID. However, the two policies can be easily combined, since the trajectories generated by  $\pi^{\text{exp}}$  are optimized to contain the same information as the encodings  $z \sim F_\psi(z | \mu)$  that the exploitation policy  $\pi^{\text{task}}$  trained on. Next we describe each of these components in detail.

**Learning the problem ID encodings and exploitation policy.** We first learn a stochastic encoder  $F_\psi(z | \mu)$  parametrized by  $\psi$  and exploitation policy  $\pi_\theta^{\text{task}}(a | s, z)$  parametrized by  $\theta$ , which conditions on  $z$ , by solving the following constrained optimization problem:

$$\begin{aligned} &\text{minimize} && I(z; \mu) \\ &\text{subject to} && \mathbb{E}_{z \sim F_\psi(z | \mu)} [V^{\pi_\theta^{\text{task}}}(z; \mu)] = V^*(\mu) \text{ for all } \mu, \end{aligned} \quad (2)$$

where  $V^{\pi_\theta^{\text{task}}}(z; \mu)$  is the expected returns of  $\pi_\theta^{\text{task}}$  on problem  $\mu$ , given encoding  $z$ , and  $V^*(\mu)$  is the maximum expected returns achievable by any policy on problem  $\mu$ . Intuitively, optimizing this problem discards any (task-irrelevant) information from  $z$  (the objective) that does not help maximize returns (the constraint), and importantly, is independent of exploration.

In practice, we solve this problem (without knowing  $V^*(\mu)$ ), by maximizing the Lagrangian, with dual variable  $\lambda^{-1}$ :

$$\begin{aligned} &\text{maximize}_{\psi, \theta} && \underbrace{\mathbb{E}_{\mu \sim p(\mu), z \sim F_\psi(z | \mu)} [V^{\pi_\theta^{\text{task}}}(z; \mu)]}_{\text{Returns}} - \lambda \underbrace{I(z; \mu)}_{\text{Information bottleneck}}. \end{aligned} \quad (3)$$

We maximize the returns via standard RL and minimize the mutual information  $I(z; \mu)$  by minimizing a variational upper bound on it (Alemi et al., 2016),  $\mathbb{E}_\mu [\text{D}_{\text{KL}}(F_\psi(z | \mu) || j(z))]$ , where  $j$  is any prior and  $z$  is distributed as  $p_\psi(z) = \int_\mu F_\psi(z | \mu) p(\mu) d\mu$ . Note that the returns are optimized with respect to both the exploitation policy  $\pi_\theta^{\text{task}}$  and the encoder  $F_\psi$ , while the information bottleneck only depends on and is only optimized with respect to  $F_\psi$ .

**Learning an exploration policy given problem ID encodings.** Once we’ve obtained an encoder  $F_\psi(z | \mu)$  to extract only the necessary task-relevant information required to optimally solve each task, we can optimize the exploration policy  $\pi^{\text{exp}}$  to produce trajectories that contain this same information by maximizing their mutual information  $I(\tau^{\text{exp}}; z)$ . We slightly abuse notation to use  $\pi^{\text{exp}}$  to denote the probability distribution over the trajectories  $\tau^{\text{exp}}$ . Then, the mutual information  $I(\tau^{\text{exp}}; z)$  can be efficiently maximized by maximizing a variational lower bound (Barber & Agakov, 2003) as follows:

$$\begin{aligned} I(\tau^{\text{exp}}; z) &= H(z) - H(z | \tau^{\text{exp}}) \\ &\geq H(z) + \mathbb{E}_{\mu, z \sim F_\psi, \tau^{\text{exp}} \sim \pi^{\text{exp}}} [\log q_\omega(z | \tau^{\text{exp}})] \\ &= H(z) + \mathbb{E}_{\mu, z \sim F_\psi} [\log q_\omega(z)] + \\ &\quad \mathbb{E}_{\mu, z \sim F_\psi, \tau^{\text{exp}} \sim \pi^{\text{exp}}} \left[ \sum_{t=1}^T \log \frac{q_\omega(z | \tau_{:t}^{\text{exp}})}{q_\omega(z | \tau_{:t-1}^{\text{exp}})} \right], \end{aligned} \quad (4)$$

where  $q_\omega$  is any distribution parametrized by  $\omega$ , and the last line comes from expanding a telescoping series. We

**Algorithm 1** DREAM meta-training trial

- 
- 1: Sample a problem  $\mu \sim p(\mu)$
  - 2: Compute problem ID encoding  $z \sim F_\psi(z | \mu)$
  - 3: // Exploration episode
  - 4: Roll out exploration policy  $\tau^{\text{exp}} \sim \pi_\phi^{\text{exp}}(a_t | s_t, \tau^{\text{exp}})$
  - 5: Update  $\pi_\phi^{\text{exp}}$  and  $q_\omega$  to maximize  $I(\tau^{\text{exp}}; z)$  via rewards in (5)
  - 6: // Exploitation episode
  - 7: Every other episode, choose  $z \sim q_\omega(z | \tau^{\text{exp}})$
  - 8: Roll out exploitation policy  $\pi_\theta^{\text{task}}(a | s, z)$
  - 9: Update  $\pi_\theta^{\text{task}}$  and  $F_\psi$  to maximize (3)
- 

maximize the above expression over  $\tau^{\text{exp}}$  and over  $\omega$  to learn  $q_\omega$  that approximates the true conditional distribution  $p(z | \tau^{\text{exp}})$ , which makes this bound tight. In addition, we do not have access to the problem  $\mu$  at test time and hence cannot sample from  $F_\psi(z | \mu)$ . Therefore,  $q_\omega$  serves as a *decoder* that generates the encoding  $z$  from the exploration trajectory  $\tau^{\text{exp}}$ .

Recall, our goal is to maximize (4) w.r.t., trajectories  $\tau^{\text{exp}}$  from the exploration policy  $\pi_\phi^{\text{exp}}$ . Only the third term depends on  $\tau^{\text{exp}}$ , so we train  $\pi_\phi^{\text{exp}}$  on rewards set to be this third term, the information gain:

$$r_t^{\text{exp}}(a_t, r_t, s_{t+1}, \tau_{t-1}^{\text{exp}}; \mu) = \mathbb{E}_{z \sim F_\psi(z | \mu)} \left[ \log \frac{q_\omega(z | \tau_{t-1}^{\text{exp}} = [s_{t+1}; a_t; r_t; \tau_{t-1}^{\text{exp}}])}{q_\omega(z | \tau_{t-1}^{\text{exp}})} \right] - c. \quad (5)$$

Intuitively, the exploration reward for taking action  $a_t$  and transitioning to state  $s_{t+1}$  is high if this transition encodes more information about the problem (and hence the encoding  $z \sim F_\psi(z | \mu)$ ) than was already present in the trajectory  $\tau_{t-1}^{\text{exp}} = (s_0, a_0, r_0, \dots, s_{t-2}, a_{t-2}, r_{t-2})$ . We also include a small penalty  $c$  to encourage exploring efficiently in as few timesteps as possible. This reward is attractive because (i) it is independent from the exploitation policy and hence avoids the local optima described in Section 4.1, and (ii) it is dense, so it helps with credit assignment. It is also non-Markov, since it depends on  $\tau^{\text{exp}}$ , so we maximize it with a recurrent  $\pi_\phi^{\text{exp}}(a_t | s_t, \tau_{t-1}^{\text{exp}})$ , parametrized by  $\phi$ .

### 4.3. A Practical Implementation of DREAM

Altogether, DREAM learns four components. We summarize each component and detail practical choices for parametrizing them as neural networks below.

1. Encoder  $F_\psi(z | \mu)$ : The encoder learns to extract only task-relevant information from the problem ID  $\mu$  via Equation 3. Then, DREAM learns to efficiently explore by recovering the extracted information. For simplicity, we parametrize the stochastic encoder by learning a deterministic encoding  $f_\psi(\mu)$  and apply Gaussian noise, i.e.,  $F_\psi(z | \mu) = \mathcal{N}(f_\psi(\mu), \rho^2 I)$ . We choose a conve-

nient prior  $j(z)$  to be a unit Gaussian with same variance  $\rho^2 I$ , which makes the information bottleneck take the form of simple  $\ell_2$ -regularization  $I(z; \mu) = \|f_\psi(\mu)\|_2^2$ .

2. Decoder  $q_\omega(z | \tau^{\text{exp}})$ : The decoder learns to map exploration trajectories  $\tau^{\text{exp}}$  to encodings  $z$ , used by the exploitation policy during meta-test time, via maximizing Equation 4. Similar to the encoder, we parametrize the decoder  $q_\omega(z | \tau^{\text{exp}})$  as a Gaussian centered around a deterministic encoding  $g_\omega(\tau^{\text{exp}})$  with variance  $\rho^2 I$ . Then,  $q_\omega$  maximizes  $\mathbb{E}_{\mu, z \sim F_\psi(z | \mu)} \left[ \|z - g_\omega(\tau^{\text{exp}})\|_2^2 \right]$  w.r.t.,  $\omega$  (Equation 4), and the exploration rewards take the form:  $r^{\text{exp}}(a, r, s', \tau^{\text{exp}}; \mu) = \|f_\psi(\mu) - g_\omega([\tau^{\text{exp}}; a; r; s'])\|_2^2 - \|f_\psi(\mu) - g_\omega(\tau^{\text{exp}})\|_2^2 - c$  (Equation 5).
3. Exploitation policy  $\pi_\theta^{\text{task}}$  and 4. Exploration policy  $\pi_\phi^{\text{exp}}$ : We learn both policies with double deep Q-learning (van Hasselt et al., 2016), treating  $(s, z)$  as the state for  $\pi_\theta^{\text{task}}$ .

In practice, we jointly learn all components by following Algorithm 1 each meta-training trial. Overall, this avoids the chicken-and-egg problem in Section 4.1 by learning exploitation and the encoder (lines 6–9) independently from exploration. This enables the encoder to learn quickly, and once it is learned, it forms a learning signal for exploration separate from the expected exploitation returns (lines 3–5), which improves sample efficiency (Section 5.2).

During meta-testing,  $\mu$  is unavailable, but since  $\pi_\phi^{\text{exp}}$  learns to produce exploration trajectories  $\tau^{\text{exp}}$  containing the same information as  $z \sim F_\psi(z | \mu)$ , we can generate  $z$  from  $q_\omega(z | \tau^{\text{exp}})$  instead of from  $F_\psi(z | \mu)$  for the exploitation policy  $\pi_\theta^{\text{task}}(a | s, z)$ . Since the exploitation policy conditions on  $z \sim q_\omega(z | \tau^{\text{exp}})$  from the decoder during meta-testing, we also train the exploitation policy conditioned on  $z \sim q_\omega(z | \tau^{\text{exp}})$  every other episode during meta-training (line 7), which improves stability. See Appendix A for detailed pseudocode and other training details.

## 5. Analysis of DREAM

### 5.1. Theoretical Consistency of the DREAM Objective

A key property of DREAM is that it is *consistent*: maximizing our decoupled objective also maximizes expected returns (Equation 1). This contrasts prior decoupled approaches (Zhou et al., 2019b; Rakelly et al., 2019; Gurusurthy et al., 2019; Zhang et al., 2020), which also decouple exploration from exploitation, but do not recover the optimal policy even with infinite meta-training trials. Formally,

**Proposition 1.** *Assume  $\langle \mathcal{S}, \mathcal{A}, \mathcal{R}_\mu, \mathcal{T}_\mu \rangle$  is ergodic for all problems  $\mu \in \mathcal{M}$ . Let  $V^*(\mu)$  be the maximum expected returns achievable by any exploitation policy with access to the problem ID  $\mu$ , i.e., with complete information. Let  $\pi_\star^{\text{task}}, \pi_\star^{\text{exp}}, F_\star$  and  $q_\star(z | \tau^{\text{exp}})$  be the optimizers of the*

## Decoupling Exploration and Exploitation for Meta-Reinforcement Learning

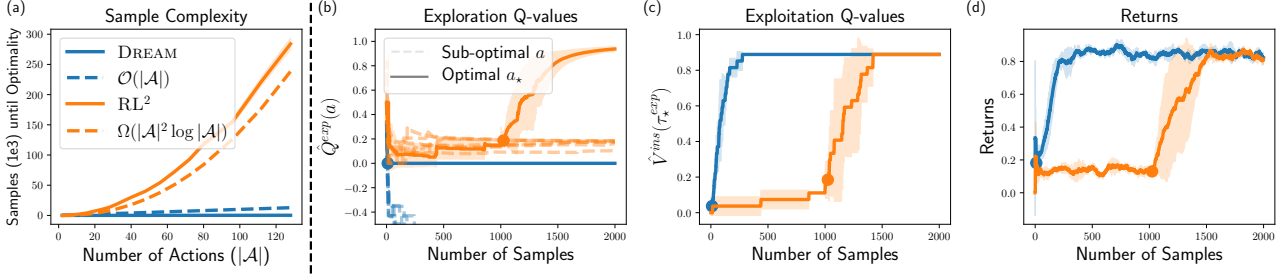


Figure 3. (a) Sample complexity of learning the optimal exploration policy as the action space  $|\mathcal{A}|$  grows (1000 seeds). (b) Exploration Q-values  $\hat{Q}^{\text{exp}}(a)$ . The policy  $\arg \max_a \hat{Q}^{\text{exp}}(a)$  is optimal after the dot. (c) Exploitation values given optimal trajectory  $\hat{V}^{\text{task}}(\tau_*^{\text{exp}})$ . (d) Returns achieved on a tabular MDP with  $|\mathcal{A}| = 8$  (3 seeds).




Figure 4. Simple bandit domain. In problem ID  $\mu$ , action  $a_\mu$  obtains reward 1; all other actions obtain no reward. In all problems, action  $a_* = A$  reveals  $\mu$  and hence, is optimal for exploration.

*DREAM objective.* Then, if the function classes DREAM optimizes over are well-specified, there exists a finite  $T$  such that if the length of the exploration episode is at least  $T$ ,

$$\mathbb{E}_{\mu \sim p(\mu), \tau^{\text{exp}} \sim \pi_*^{\text{exp}}, z \sim q_*(z | \tau^{\text{exp}})} [V^{\pi_*^{\text{task}}}(z; \mu)] = \mathbb{E}_{\mu \sim p(\mu)} [V^*(\mu)].$$

At the optimum of DREAM’s objective, the exploration and exploitation policies together achieve the maximal returns  $V^*(\mu)$  without access to  $\mu$  during meta-testing (proof in Appendix C.1). This result depends on an ergodicity assumption, which ensures that it is possible to recover all task-relevant information in a single exploration episode. However, this can be easily removed by increasing the number of exploration episodes. Furthermore, DREAM empirically achieves near-optimal returns even on non-ergodic MDPs in the experiments (Section 6).

### 5.2. Illustrating the Effect of Coupling on Sample Complexity

With enough capacity, end-to-end approaches can also learn the optimal policy, but can be highly sample inefficient due to the coupling problem in Section 4.1. We highlight this in a simple tabular example to remove the effects of function approximation, illustrated in Figure 4: Each episode is a one-step bandit problem with action space  $\mathcal{A} = \{1, \dots, A\}$ . Taking action  $a_* = A$  in the exploration episode leads to a trajectory  $\tau_*^{\text{exp}}$  that reveals the problem ID  $\mu$ ; all other actions  $a$  reveal no information and lead to  $\tau_a^{\text{exp}}$ . The ID  $\mu$  identifies a unique action  $a_\mu$  that receives reward 1 during exploitation; all other actions receive reward 0. Therefore, taking  $a_*$  during exploration is necessary and suffi-

cient to obtain optimal reward 1. We now study the number of samples required for RL<sup>2</sup> (the canonical end-to-end approach) and DREAM to learn the optimal exploration policy with  $\epsilon$ -greedy tabular Q-learning. We precisely describe a more general setup in Appendix C.2 and prove that DREAM learns the optimal exploration policy in  $\Omega(|\mathcal{A}|^H |\mathcal{M}|)$  times fewer samples than RL<sup>2</sup> in this simple setting with horizon  $H$ . Figure 3a empirically validates this result and we provide intuition below.

In the tabular analog of RL<sup>2</sup>, the exploitation Q-values form targets for the exploration Q-values:  $\hat{Q}^{\text{exp}}(a) \leftarrow \hat{V}^{\text{task}}(\tau_a^{\text{exp}}) := \max_{a'} \hat{Q}^{\text{task}}(\tau_a^{\text{exp}}, a')$ . We drop the fixed initial state from notation. This creates the local optimum in Section 4.1. Initially  $\hat{V}^{\text{task}}(\tau_*^{\text{exp}})$  is low, as the exploitation policy has not learned to achieve reward, even when given  $\tau_*^{\text{exp}}$ . This causes  $\hat{Q}^{\text{exp}}(a_*)$  to be small and therefore  $\arg \max_a \hat{Q}^{\text{exp}}(a) \neq a_*$  (Figure 3b), which then prevents  $\hat{V}^{\text{task}}(\tau_*^{\text{exp}})$  from learning (Figure 3c) as  $\tau_*^{\text{exp}}$  is roughly sampled only once per  $\frac{|\mathcal{A}|}{\epsilon}$  episodes. This effect is mitigated only when  $\hat{Q}^{\text{exp}}(a_*)$  becomes higher than  $\hat{Q}^{\text{exp}}(a)$  for the other uninformative  $a$ ’s (the dot in Figure 3b-d). Then, learning both the exploitation and exploration Q-values accelerates, but getting there takes many samples.

In DREAM, the exploration Q-values regress toward the decoder  $\hat{q}$ :  $\hat{Q}^{\text{exp}}(a) \leftarrow \log \hat{q}(\mu | \tau_a^{\text{exp}}(a))$ . This decoder learns much faster than  $\hat{Q}^{\text{task}}$ , since it does not depend on the exploitation actions. Consequently, DREAM’s exploration policy quickly becomes optimal (dot in Figure 3b), which enables quickly learning the exploitation Q-values and achieving high reward (Figures 3c and 3d).

In general, DREAM learns in far fewer samples than end-to-end approaches, since in end-to-end approaches like RL<sup>2</sup>, exploration is learned from a quantity requiring many samples to accurately estimate (i.e., the exploitation Q-values in this case). Initially, this quantity is estimated poorly, so the signal for exploration can erroneously "down weight" good exploration behavior, leading to the chicken-and-egg problem. In contrast, in DREAM, the exploration policy learns from the decoder, which requires far fewer samples to accu-




Figure 5. Didactic grid worlds to stress test exploration. (a) Navigation. (b) Cooking.

rately estimate, avoiding the chicken-and-egg problem.

## 6. Experiments

Many real-world problem distributions (e.g., cooking) require exploration (e.g., locating ingredients) that is distinct from exploitation (e.g., cooking these ingredients). Therefore, we desire benchmarks that require distinct exploration and exploitation to stress test aspects of exploration in meta-RL, such as if methods can: (i) efficiently explore, even in the presence of distractions; (ii) leverage informative objects (e.g., a map) to aid exploration; (iii) learn exploration and exploitation strategies that generalize to unseen problems; (iv) scale to challenging exploration problems with high-dimensional visual observations. Existing benchmarks (e.g., MetaWorld (Yu et al., 2019) or MuJoCo tasks like HalfCheetahVelocity (Finn et al., 2017; Rothfuss et al., 2018)) were not designed to test exploration and are unsuitable for answering these questions. These benchmarks mainly vary the rewards (e.g., the speed to run at) across problems, so naively exploring by exhaustively trying different exploitation behaviors (e.g., running at different speeds) is optimal. They further don’t include visual states, distractors, or informative objects, which test if exploration is efficient. We therefore design new benchmarks meeting the above criteria, testing (i-iii) with didactic benchmarks, and (iv) with a sparse-reward 3D visual navigation benchmark, based on Kamienny et al. (2020), that combines complex exploration with high-dimensional visual inputs. To further deepen the exploration challenge, we make our benchmarks goal-conditioned. This requires exploring to discover information relevant to *any* potential goal, rather than just a single task (e.g., locating all ingredients for *any* meal vs. just the ingredients for pasta).

**Comparisons.** We compare DREAM with state-of-the-art end-to-end (E-RL<sup>2</sup> (Stadie et al., 2018), VARIBAD (Zintgraf et al., 2019), and IMPORT (Kamienny et al., 2020)) and decoupled approaches (PEARL-UB, an upper bound on the final performance of PEARL (Rakelly et al., 2019)). For PEARL-UB, we analytically compute the expected rewards achieved by optimal Thompson sampling (TS) exploration,

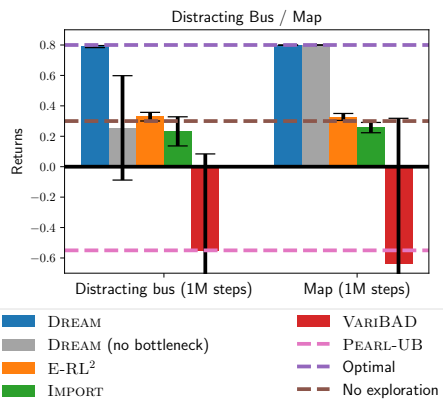


Figure 6. Navigation results. Only DREAM optimally explores all buses and the map.

assuming access to the optimal problem-specific policy and true posterior problem distribution. Like DREAM, IMPORT and PEARL also use the one-hot problem ID, during meta-training. We also report the optimal returns achievable with no exploration as "No exploration." Where applicable, all methods use the same architecture. The full architecture and approach details are in Appendix B.3.

We report the average returns achieved by each approach in trials with one exploration and one exploitation episode, averaged over 3 seeds with 1-standard deviation error bars (full details in Appendix B). We evaluate each approach on 100 meta-testing trials, every 2K meta-training trials. In all plots, the training timesteps includes all timesteps from both exploitation and exploration episodes in meta-training trials.

### 6.1. Didactic Experiments

We first evaluate on the grid worlds shown in Figures 5a and 5b. The state consists of the agent’s  $(x, y)$ -position, a one-hot indicator of the object at the agent’s position (none, bus, map, pot, or fridge), a one-hot indicator of the agent’s inventory (none or an ingredient), and the goal. The actions are *move* up, down, left, or right; *ride bus*, which, at a bus, teleports the agent to another bus of the same color; *pick up*, which, at a fridge, fills the agent’s inventory with the fridge’s ingredients; and *drop*, which, at the pot, empties the agent’s inventory into the pot. Episodes consist of 20 timesteps and the agent receives  $-0.1$  reward at each timestep until the goal, described below, is met (details in Appendix B.1; qualitative results in Appendix B.2).

**Targeted exploration.** We first test if these methods can efficiently explore in the presence of distractions in two versions of the benchmark in Figure 5a: *distracting bus* and *map*. In both, there are 4 possible goals (the 4 green locations). During each episode, a goal is randomly sampled. Reaching the goal yields  $+1$  reward and ends the episode. The 4 colored buses each lead to near a different potential green goal location when ridden bus and in different problems




Figure 7. Cooking results: only DREAM achieves optimal reward on training problems (left), and on generalizing to unseen problems (middle). 3D visual navigation results: only DREAM reads the sign and solves the task (right).

$\mu$ , their destinations are set to be 1 of the  $4!$  different permutations. The *distracting bus* version tests if the agent can ignore distractions by including unhelpful gray buses, which are never needed to optimally reach any goal. In different problems, the gray buses lead to different permutations of the gray locations. The *map* version tests if the agent can leverage objects for exploration by including a map that reveals the destinations of the colored buses when touched.

Figure 6 shows the results after 1M steps. DREAM learns to optimally explore and thus receives optimal reward in both versions: In *distracting bus*, it ignores the unhelpful gray buses and learns the destinations of all helpful buses by riding them. In *map*, it learns to leverage informative objects by visiting the map. During exploitation, DREAM immediately reaches the goal by riding the correct colored bus. In contrast, IMPORT and E-RL<sup>2</sup> get stuck in a local optimum, indicative of the coupling problem (Section 4.1), which achieves the same returns as no exploration at all. They do not explore the helpful buses or map and consequently sub-optimally exploit by just walking to the goal. VARIBAD learns slower, likely because it learns a dynamics model, but eventually matches the sub-optimal returns of IMPORT and RL<sup>2</sup> in  $\sim 3M$  steps (not shown).

PEARL achieves sub-optimal returns, even with infinite meta-training (see line for PEARL-UB), as follows. TS explores by sampling a problem ID from its posterior and executing its policy conditioned on this ID. Since for any given problem (bus configuration) and goal, the optimal problem-specific policy rides the one bus leading to the goal, TS does not explore optimally (i.e., explore all the buses or read the map), even with the optimal problem-specific policy and true posterior problem distribution.

Recall that DREAM tries to discard extraneous task-irrelevant information from the problem ID with an information bottleneck that minimizes the mutual information  $I(z; \mu)$  between problem IDs and the encoder  $F_\psi(z | \mu)$ . This makes exploration targeted, since DREAM only ex-

plores to recover information in  $z$ . We hypothesize that the bottleneck only improves exploration in domains with distracting task-irrelevant information to discard from the problem ID. This empirically holds when we ablate the information bottleneck from DREAM, plotted under DREAM (no bottleneck): In *distracting bus*, DREAM without the bottleneck wastes its exploration on the distracting gray unhelpful buses and consequently achieves low returns, as seen in Figure 6 (left). In contrast, *map* and other below domains do not contain any distracting information in the problem ID. Consistent with our hypothesis, DREAM achieves comparable returns with or without the information bottleneck in these domains, as seen in Figure 6 (right) and Figure 7.

**Generalization to new problems.** We test generalization to unseen problems in a cooking benchmark (Figure 5b). The fridges each contain 1 of 4 different (color-coded) ingredients, determined by the problem ID. The fridges’ contents are unobserved until the agent uses the “pick up” action at the fridge. Goals (recipes) specify placing 2 correct ingredients in the pot in the right order. The agent receives positive reward for picking up and placing the correct ingredients, and negative reward for using wrong ingredients. We hold out 1 of the  $4^3 = 64$  problems for meta-testing.

Figure 7 shows the results on training (left) and held-out (middle) problems. Only DREAM achieves near-optimal returns on both. During exploration, it investigates each fridge with the “pick up” action, and then directly retrieves the correct ingredients during exploitation. E-RL<sup>2</sup> gets stuck in a local optimum, only sometimes exploring the fridges. This achieves 3.8x lower returns, only slightly higher than no exploration at all. Here, leveraging the problem ID actually hurts IMPORT compared to E-RL<sup>2</sup>. IMPORT successfully solves the task, given access to the problem ID, but fails without it. As before, VARIBAD learns slowly and TS (PEARL-UB) cannot learn optimal exploration.






Figure 8. 3D Visual Navigation. The agent must read the sign to determine what colored object to go to.

## 6.2. Sparse-Reward 3D Visual Navigation

We conclude with a challenging benchmark testing both sophisticated exploration and scalability to pixel inputs. We modify a benchmark from Kamienny et al. (2020) to increase both the exploration and scalability challenge by including more objects and a visual sign, illustrated in Figure 8. In the 3 different problems, the sign on the right says “blue”, “red” or “green.” The goals specify whether the agent should collect the key or block. The agent receives +1 reward for collecting the correct object (color specified by the sign, shape specified by the goal), -1 reward for the wrong object, and 0 reward otherwise. The agent begins the episode on the far side of the barrier and must walk around the barrier to visually “read” the sign. The agent’s observations are  $80 \times 60$  RGB images and its actions are to rotate left or right, move forward, or end the episode.

DREAM is the only method that learns to read the sign and achieve reward (Figure 7 right). All end-to-end approaches get stuck in local optima due to the chicken-and-egg coupling problem, where they do not learn to read the sign and hence stay away from all the objects, in fear of receiving negative reward. This achieves close to 0 returns, consistent with the results in Kamienny et al. (2020). As before, PEARL-UB cannot learn optimal exploration.

## 7. Conclusion

In summary, this work identifies a chicken-and-egg problem that end-to-end meta-RL approaches suffer from, where learning good exploitation requires already having learned good exploration and vice-versa. This creates challenging local optima, since typically neither exploration nor exploitation is good at the beginning of meta-training. We show that appropriately leveraging simple one-hot problem IDs allows us to break this cyclic dependency with DREAM. Consequently, DREAM has strong empirical performance on meta-RL problems requiring complex exploration, as well as substantial theoretical sample complexity improvements in the tabular setting. Though prior works also leverage

the problem ID and use decoupled objectives that avoid the chicken-and-egg problem, no other existing approaches can recover optimal exploration empirically and theoretically like DREAM.

**Reproducibility.** Our code is publicly available at <https://github.com/ezliu/dream>.

## Acknowledgements

We thank Luisa Zintgraf for her insights about VARIBAD. We also thank Sahaana Suri, Suraj Nair, Minae Kwon, and Ramtin Keramati for their feedback on drafts of this paper.

We thank Arkira Chantaratananond for creating animations of the grid world tasks in the project web page. Icons used in this paper were made by Freepik, ThoseIcons, dDara, Pixel perfect, ThoseIcons, mynamepong, Icongeek26, and Vitaly Gorbachev from [www.flaticon.com](http://www.flaticon.com).

EL is supported by a National Science Foundation Graduate Research Fellowship under Grant No. DGE-1656518. AR is supported by a Google PhD Fellowship and Open Philanthropy Project AI Fellowship. This work was also supported in part by Google.

## References

- Agarwal, R., Liang, C., Schuurmans, D., and Norouzi, M. Learning to generalize from sparse and underspecified rewards. *arXiv preprint arXiv:1902.07198*, 2019.
- Alemi, A. A., Fischer, I., Dillon, J. V., and Murphy, K. Deep variational information bottleneck. *arXiv preprint arXiv:1612.00410*, 2016.
- Andrychowicz, M., Denil, M., Gomez, S., Hoffman, M. W., Pfau, D., Schaul, T., Shillingford, B., and Freitas, N. D. Learning to learn by gradient descent by gradient descent. In *Advances in neural information processing systems*, pp. 3981–3989, 2016.
- Barber, D. and Agakov, F. V. The IM algorithm: a variational approach to information maximization. In *Advances in neural information processing systems*, 2003.
- Bellemare, M., Srinivasan, S., Ostrovski, G., Schaul, T., Saxton, D., and Munos, R. Unifying count-based exploration and intrinsic motivation. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 1471–1479, 2016.
- Bengio, S., Bengio, Y., Cloutier, J., and Gecsei, J. On the optimization of a synaptic learning rule. In *Preprints Conf. Optimality in Artificial and Biological Neural Networks*, volume 2, 1992.
- Bengio, Y., Bengio, S., and Cloutier, J. Learning a synaptic learning rule. In *IJCNN-91-Seattle International Joint*

- Conference on Neural Networks*, volume 2, pp. 969–969, 1991.
- Burda, Y., Edwards, H., Storkey, A., and Klimov, O. Exploration by random network distillation. *arXiv preprint arXiv:1810.12894*, 2018.
- Chevalier-Boisvert, M. Gym-Miniworld environment for openai gym. <https://github.com/maximecb/gym-miniworld>, 2018.
- Dorfman, R. and Tamar, A. Offline meta reinforcement learning. *arXiv preprint arXiv:2008.02598*, 2020.
- Duan, Y., Schulman, J., Chen, X., Bartlett, P. L., Sutskever, I., and Abbeel, P. RL<sup>2</sup>: Fast reinforcement learning via slow reinforcement learning. *arXiv preprint arXiv:1611.02779*, 2016.
- Eysenbach, B., Gupta, A., Ibarz, J., and Levine, S. Diversity is all you need: Learning skills without a reward function. *arXiv preprint arXiv:1802.06070*, 2018.
- Fakoor, R., Chaudhari, P., Soatto, S., and Smola, A. J. Meta-Q-learning. *arXiv preprint arXiv:1910.00125*, 2019.
- Finn, C., Abbeel, P., and Levine, S. Model-agnostic meta-learning for fast adaptation of deep networks. In *International Conference on Machine Learning (ICML)*, 2017.
- Gregor, K., Rezende, D. J., and Wierstra, D. Variational intrinsic control. *arXiv preprint arXiv:1611.07507*, 2016.
- Gupta, A., Mendonca, R., Liu, Y., Abbeel, P., and Levine, S. Meta-reinforcement learning of structured exploration strategies. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 5302–5311, 2018.
- Gurumurthy, S., Kumar, S., and Sycara, K. Mame: Model-agnostic meta-exploration. *arXiv preprint arXiv:1911.04024*, 2019.
- Hiraoka, T., Imagawa, T., Tangkaratt, V., Osa, T., Onishi, T., and Tsuruoka, Y. Meta-model-based meta-policy optimization. *arXiv preprint arXiv:2006.02608*, 2020.
- Hochreiter, S., Younger, A. S., and Conwell, P. R. Learning to learn using gradient descent. In *International Conference on Artificial Neural Networks (ICANN)*, pp. 87–94, 2001.
- Houthoofd, R., Chen, X., Duan, Y., Schulman, J., Turck, F. D., and Abbeel, P. Vime: Variational information maximizing exploration. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 1109–1117, 2016.
- Houthoofd, R., Chen, Y., Isola, P., Stadie, B., Wolski, F., Ho, O. J., and Abbeel, P. Evolved policy gradients. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 5400–5409, 2018.
- Humphrik, J., Galashov, A., Hasenclever, L., Ortega, P. A., Teh, Y. W., and Heess, N. Meta reinforcement learning as task inference. *arXiv preprint arXiv:1905.06424*, 2019.
- Kaelbling, L. P., Littman, M. L., and Cassandra, A. R. Planning and acting in partially observable stochastic domains. *Artificial intelligence*, 101(1):99–134, 1998.
- Kamienny, P., Pirota, M., Lazaric, A., Lavril, T., Usunier, N., and Denoyer, L. Learning adaptive exploration strategies in dynamic environments through informed policy regularization. *arXiv preprint arXiv:2005.02934*, 2020.
- Kapturowski, S., Ostrovski, G., Quan, J., Munos, R., and Dabney, W. Recurrent experience replay in distributed reinforcement learning. In *International Conference on Learning Representations (ICLR)*, 2019.
- Kingma, D. and Ba, J. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, 2015.
- Leibfried, F., Pascual-Diaz, S., and Grau-Moya, J. A unified bellman optimality principle combining reward maximization and empowerment. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 7869–7880, 2019.
- Liu, E. Z., Hashemi, M., Swersky, K., Ranganathan, P., and Ahn, J. An imitation learning approach for cache replacement. *arXiv preprint arXiv:2006.16239*, 2020a.
- Liu, E. Z., Keramati, R., Seshadri, S., Guu, K., Pasupat, P., Brunskill, E., and Liang, P. Learning abstract models for strategic exploration and fast reward transfer. *arXiv preprint arXiv:2007.05896*, 2020b.
- Mendonca, R., Gupta, A., Kravev, R., Abbeel, P., Levine, S., and Finn, C. Guided meta-policy search. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 9653–9664, 2019.
- Mishra, N., Rohaninejad, M., Chen, X., and Abbeel, P. A simple neural attentive meta-learner. *arXiv preprint arXiv:1707.03141*, 2017.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540): 529–533, 2015.
- Nagabandi, A., Clavera, I., Liu, S., Fearing, R. S., Abbeel, P., Levine, S., and Finn, C. Learning to adapt in dynamic, real-world environments through meta-reinforcement learning. *arXiv preprint arXiv:1803.11347*, 2018.

- Naik, D. K. and Mammone, R. J. Meta-neural networks that learn by learning. In *[Proceedings 1992] IJCNN International Joint Conference on Neural Networks*, volume 1, pp. 437–442, 1992.
- Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., and Lerer, A. Automatic differentiation in pytorch, 2017.
- Pathak, D., Agrawal, P., Efros, A. A., and Darrell, T. Curiosity-driven exploration by self-supervised prediction. In *Computer Vision and Pattern Recognition (CVPR)*, pp. 16–17, 2017.
- Rakelly, K., Zhou, A., Quillen, D., Finn, C., and Levine, S. Efficient off-policy meta-reinforcement learning via probabilistic context variables. *arXiv preprint arXiv:1903.08254*, 2019.
- Rothfuss, J., Lee, D., Clavera, I., Asfour, T., and Abbeel, P. Promp: Proximal meta-policy search. *arXiv preprint arXiv:1810.06784*, 2018.
- Russo, D., Roy, B. V., Kazerouni, A., Osband, I., and Wen, Z. A tutorial on thompson sampling. *arXiv preprint arXiv:1707.02038*, 2017.
- Sæmundsson, S., Hofmann, K., and Deisenroth, M. P. Meta reinforcement learning with latent variable gaussian processes. *arXiv preprint arXiv:1803.07551*, 2018.
- Santoro, A., Bartunov, S., Botvinick, M., Wierstra, D., and Lillicrap, T. One-shot learning with memory-augmented neural networks. *arXiv preprint arXiv:1605.06065*, 2016.
- Schmidhuber, J. *Evolutionary principles in self-referential learning, or on learning how to learn: the meta-meta-... hook*. PhD thesis, Technische Universität München, 1987.
- Stadie, B., Yang, G., Houthoofd, R., Chen, P., Duan, Y., Wu, Y., Abbeel, P., and Sutskever, I. The importance of sampling in meta-reinforcement learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 9280–9290, 2018.
- Thompson, W. R. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3):285–294, 1933.
- Thrun, S. and Pratt, L. *Learning to learn*. Springer Science & Business Media Springer Science & Business Media, 2012.
- van der Maaten, L. and Hinton, G. Visualizing data using t-SNE. *Journal of machine learning research*, 9(0):2579–2605, 2008.
- van Hasselt, H., Guez, A., and Silver, D. Deep reinforcement learning with double Q-learning. In *Association for the Advancement of Artificial Intelligence (AAAI)*, volume 16, pp. 2094–2100, 2016.
- Wang, J. X., Kurth-Nelson, Z., Tirumala, D., Soyer, H., Leibo, J. Z., Munos, R., Blundell, C., Kumaran, D., and Botvinick, M. Learning to reinforcement learn. *arXiv preprint arXiv:1611.05763*, 2016a.
- Wang, Z., Schaul, T., Hessel, M., Hasselt, H. V., Lanctot, M., and Freitas, N. D. Dueling network architectures for deep reinforcement learning. In *International Conference on Machine Learning (ICML)*, 2016b.
- Warde-Farley, D., de Wiele, T. V., Kulkarni, T., Ionescu, C., Hansen, S., and Mnih, V. Unsupervised control through non-parametric discriminative rewards. *arXiv preprint arXiv:1811.11359*, 2018.
- Yang, Y., Caluwaerts, K., Iscen, A., Tan, J., and Finn, C. Norml: No-reward meta learning. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, pp. 323–331, 2019.
- Yu, T., Quillen, D., He, Z., Julian, R., Hausman, K., Finn, C., and Levine, S. Meta-world: A benchmark and evaluation for multi-task and meta reinforcement learning. *arXiv preprint arXiv:1910.10897*, 2019.
- Zhang, J., Wang, J., Hu, H., Chen, Y., Fan, C., and Zhang, C. Learn to effectively explore in context-based meta-RL. *arXiv preprint arXiv:2006.08170*, 2020.
- Zhou, A., Jang, E., Kappler, D., Herzog, A., Khansari, M., Wohlhart, P., Bai, Y., Kalakrishnan, M., Levine, S., and Finn, C. Watch, try, learn: Meta-learning from demonstrations and reward. *arXiv preprint arXiv:1906.03352*, 2019a.
- Zhou, W., Pinto, L., and Gupta, A. Environment probing interaction policies. *arXiv preprint arXiv:1907.11740*, 2019b.
- Zintgraf, L., Shiarlis, K., Igl, M., Schulze, S., Gal, Y., Hofmann, K., and Whiteson, S. Varibad: A very good method for bayes-adaptive deep RL via meta-learning. *arXiv preprint arXiv:1910.08348*, 2019.