## A Omitted Derivations

## A. 1 Jacobian of $\hat{\theta}$ with respect to $\psi$

We begin with an observation: first-order optimality conditions for $\hat{\theta}$ necessitate that

$$
\begin{equation*}
\left.\frac{\partial \hat{\eta}}{\partial \theta}\right|_{\hat{\theta}}=0 \tag{8}
\end{equation*}
$$

Then, the implicit function theorem tells us that for sufficiently small $\Delta \psi$, there exists $\Delta \theta$ as a function of $\Delta \psi$ such that

$$
\begin{equation*}
\left.\frac{\partial \hat{\eta}}{\partial \theta}\right|_{\hat{\theta}+\Delta \theta, \psi+\Delta \psi}=0 \tag{9}
\end{equation*}
$$

To first order, we have

$$
\begin{equation*}
\left.\left.\frac{\partial \hat{\eta}}{\partial \theta}\right|_{\hat{\theta}+\Delta \theta, \psi+\Delta \psi} \approx \underbrace{\frac{\partial \hat{\eta}}{\partial \theta}}_{0}\right|_{\hat{\theta}}+\left.\frac{\partial^{2} \hat{\eta}}{\partial \theta \partial \theta^{\top}}\right|_{\hat{\theta}} \Delta \theta+\left.\frac{\partial^{2} \hat{\eta}}{\partial \theta \partial \psi^{\top}}\right|_{\hat{\theta}} \Delta \psi \tag{10}
\end{equation*}
$$

Thus, solving for $\Delta \theta$ as a function of $\Delta \psi$ and taking the limit as $\Delta \psi \rightarrow 0$, we obtain

$$
\begin{equation*}
\frac{\partial \hat{\theta}}{\partial \psi^{\top}}=-\left.\left(\left.\frac{\partial^{2} \hat{\eta}}{\partial \theta \partial \theta^{\top}}\right|_{\hat{\theta}}\right)^{-1} \frac{\partial^{2} \hat{\eta}}{\partial \theta \partial \psi^{\top}}\right|_{\hat{\theta}} \tag{11}
\end{equation*}
$$

## A. 2 Policy Hessian

Fix dynamics $T$, and let $\pi_{\theta}(\tau)$ denote the probability density of trajectory $\tau$ under policy $\pi_{\theta}$. Then we have

$$
\begin{equation*}
\frac{\partial \log \pi_{\theta}(\tau)}{\partial \theta}=\frac{\frac{\partial \pi_{\theta}(\tau)}{\partial \theta}}{\pi_{\theta}(\tau)} \quad \text { i.e. } \quad \frac{\partial \pi_{\theta}(\tau)}{\partial \theta}=\pi_{\theta}(\tau) \frac{\partial \log \pi_{\theta}(\tau)}{\partial \theta} \tag{12}
\end{equation*}
$$

Thus we get the basic (REINFORCE) policy gradient

$$
\begin{equation*}
\frac{\partial \eta}{\partial \theta}=\frac{\partial}{\partial \theta} \int \pi_{\theta}(\tau) R_{\psi}(\tau) \mathrm{d} \tau=\int \frac{\partial \pi_{\theta}(\tau)}{\partial \theta} R_{\psi}(\tau) \mathrm{d} \tau=\underset{\tau \sim \pi_{\theta}, T}{\mathbb{E}}\left[\frac{\partial \log \pi_{\theta}(\tau)}{\partial \theta} R_{\psi}(\tau)\right] \tag{13}
\end{equation*}
$$

Differentiating our earlier expression for $\frac{\partial \pi_{\theta}(\tau)}{\partial \theta}$ once more, and then reusing that same expression again, we have

$$
\begin{align*}
\frac{\partial^{2} \pi_{\theta}(\tau)}{\partial \theta \partial \theta^{\top}} & =\frac{\partial \pi_{\theta}(\tau)}{\partial \theta} \frac{\partial \log \pi_{\theta}(\tau)}{\partial \theta^{\top}}+\pi_{\theta}(\tau) \frac{\partial^{2} \log \pi_{\theta}(\tau)}{\partial \theta \partial \theta^{\top}}  \tag{14}\\
& =\pi_{\theta}(\tau)\left(\frac{\partial \log \pi_{\theta}(\tau)}{\partial \theta} \frac{\partial \log \pi_{\theta}(\tau)}{\partial \theta^{\top}}+\frac{\partial^{2} \log \pi_{\theta}(\tau)}{\partial \theta \partial \theta^{\top}}\right) \tag{15}
\end{align*}
$$

Thus

$$
\begin{align*}
\frac{\partial^{2} \eta}{\partial \theta \partial \theta^{\top}} & =\frac{\partial^{2}}{\partial \theta \partial \theta^{\top}} \int \pi_{\theta}(\tau) R_{\psi}(\tau) \mathrm{d} \tau  \tag{16}\\
& =\int \frac{\partial^{2} \pi_{\theta}(\tau)}{\partial \theta \partial \theta^{\top}} R_{\psi}(\tau) \mathrm{d} \tau  \tag{17}\\
& =\int \pi_{\theta}(\tau)\left(\frac{\partial \log \pi_{\theta}(\tau)}{\partial \theta} \frac{\partial \log \pi_{\theta}(\tau)}{\partial \theta^{\top}}+\frac{\partial^{2} \log \pi_{\theta}(\tau)}{\partial \theta \partial \theta^{\top}}\right) R_{\psi}(\tau) \mathrm{d} \tau  \tag{18}\\
& =\underset{\tau \sim \pi_{\theta}, T}{\mathbb{E}}\left[\left(\frac{\partial \log \pi_{\theta}(\tau)}{\partial \theta} \frac{\partial \log \pi_{\theta}(\tau)}{\partial \theta^{\top}}+\frac{\partial^{2} \log \pi_{\theta}(\tau)}{\partial \theta \partial \theta^{\top}}\right) R_{\psi}(\tau)\right] \tag{19}
\end{align*}
$$

## B Implementation detail

The section discusses how to compute $A x$ using standard automatic differentiation packages. We first define the following function:

$$
\begin{equation*}
\eta_{h}\left(\theta_{1}, \theta_{2}, \theta_{3}, \theta, \widehat{T}_{\phi}, \psi\right)=\underset{\pi_{\theta}, \widehat{T}_{\phi}}{\mathbb{E}}\left[\left(\log \pi_{\theta_{1}}\left(a_{t} \mid s_{t}\right) \log \pi_{\theta_{2}}\left(a_{t} \mid s_{t}\right)+\log \pi_{\theta_{3}}\left(a_{t} \mid s_{t}\right)\right) R_{\psi}(\tau)\right] \tag{20}
\end{equation*}
$$

where $\theta_{1}, \theta_{2}, \theta_{3}$ are parameter copies of $\theta$. We then use Hessian-vector product to avoid directly computing the second derivatives. Specifically, we compute the two parts in Eq. 7) respectively by first differentiating $\eta_{h}$ w.r.t $\theta_{1}$ and $\theta_{2}^{\top}$

$$
\begin{equation*}
g_{1}=\frac{\partial}{\partial \theta_{1}}\left(\frac{\partial \eta_{h}}{\partial \theta_{2}^{\top}} \cdot x\right)=\frac{\partial \log \pi_{\theta_{1}}\left(a_{t} \mid s_{t}\right)}{\partial \theta_{1}}\left(\frac{\partial \log \pi_{\theta_{2}}\left(a_{t} \mid s_{t}\right)}{\partial \theta_{2}^{\top}} \cdot x\right) R_{\psi}(\tau) \tag{21}
\end{equation*}
$$

and then differentiate $\eta_{h}$ w.r.t $\theta_{3}$ for twice

$$
\begin{equation*}
g_{2}=\frac{\partial}{\partial \theta_{3}}\left(\frac{\partial \eta_{h}}{\partial \theta_{3}^{\top}} \cdot x\right)=\frac{\partial}{\partial \theta_{3}}\left(\frac{\partial \log \pi_{\theta_{3}}\left(a_{t} \mid s_{t}\right)}{\partial \theta_{3}^{\top}} \cdot x\right) R_{\psi}(\tau), \tag{22}
\end{equation*}
$$

and thus we have $A x=g_{1}+g_{2}$.

## C Hyper-parameters

We experimented with the following task settings: Hopper -2 D with $x$ velocity and $z$ height from $\Psi=[-2,2] \times[1.2,2.0]$, Walker-2D with $x$ velocity and $z$ height from $\Psi=[-2,2] \times[1.0,1.8]$, Ant-2D with $x$ velocity and $y$ velocity from $\Psi=[-3,3] \times[-3,3]$, Ant -3 D with $x$ velocity, $y$ velocity and $z$ height from $\Psi=[-3,3] \times[-3,3] \times[0.4,0.6]$, Cheetah-Highdim with $\Psi=[-1,1]^{18}$. We also list the coefficient of the parameterized reward functions in Table 1

Table 1: Coefficient in parameterized reward functions

|  | Hopper2D | Walker2D | Ant2D | Ant3D |
| :---: | :---: | :---: | :---: | :---: |
| $c_{1}$ | 1 | 1 | 1 | 1 |
| $c_{2}$ | 0 | 0 | 1 | 1 |
| $c_{3}$ | 5 | 5 | 0 | 30 |

The hyper-parameters of MAML and PEARL are mostly taken directly from the supplied implementation of [13] and [38]. We run MAML for 500 training iterations: for each iteration, MAML uses a meta-batch size of 40 (the number of tasks sampled at each iteration) and a batch size of 20 (the number of rollouts used to compute the policy gradient updates). Overall, MAML requires 80 million samples during meta training. For PEARL, we first collect a batch of training tasks (150) by uniformly sampling from $\Psi$. We run PEARL for 500 training iterations: for each iteration, PEARL randomly sample 5 tasks and collects 1000 samples for each task from both prior (400) and posterior (600) of the context variables; for each gradient update, PEARL uses a meta-batch size of 10 and optimizes the parameters of actor, critic and context encoder by 4000 steps of gradient descent. Overall, PEARL requires 2.5 million samples during meta training.
For AdMRL, we first do zero-shot adaptation for each task by 40 virtual steps ( $n_{\text {zeroshot }}=40$ ). We then perform SLBO [25] by interleaving data collection, dynamical model fitting and policy updates, where we use 3 outer iterations $\left(n_{\text {slbo }}=3\right)$ and 20 inner iterations $\left(n_{\text {inner }}=20\right)$. Algorithm 2 shows the pseudo code of the virtual training procedure. For each inner iteration, we update model for 100 steps ( $n_{\text {model }}=100$ ), and update policy for 20 steps ( $n_{\text {policy }}=20$ ), each with 10000 virtual samples $\left(n_{\text {trpo }}=10000\right)$. For the first task, we use $n_{\text {slbo }}=10$ (for Hopper2D, Walker2D) or $n_{\text {slbo }}=20$ (for Ant2D, Ant3D, Cheetah-Highdim). For all tasks, we sweep the learning rate $\alpha$ in $\{1,2,4,8,16,32\}$ and we use $\alpha=2$ for Hopper2D, $\alpha=8$ for Walker2D, $\alpha=4$ for Ant2D and Ant3D, $\alpha=16$ for Cheetah-Highdim. To compute the gradient w.r.t the task parameters, we do 200 iterations of conjugate gradient descent.

## D Examples of high-dimensional tasks

Figure 7 shows some trajectories in the high-dimensional task Cheetah-Highdim.

```
Algorithm 2 Virtual Training in AdMRL
    procedure VirtualTraining \((\theta:\) policy, \(\phi:\) model, \(\psi:\) task, \(\mathcal{D}:\) data, \(n:\) virtual steps)
        for \(n\) iterations do
            Optimize virtual dynamics \(\widehat{T}_{\phi}\) over \(\phi\) with data sampled from \(\mathcal{D}\) by \(n_{\text {model }}\) steps
            for \(n_{\text {policy }}\) iterations do
                    \(\mathcal{D}^{\prime} \leftarrow\left\{\right.\) collect \(n_{\text {trpo }}\) samples from the learned dynamics \(\left.\widehat{T}_{\phi}\right\}\)
                    Optimize \(\pi_{\theta}\) by running TRPO on \(\mathcal{D}^{\prime}\)
```



Figure 7: The high-dimensional tasks are surprisingly often semantically meaningful. Policies learned in these tasks can have diverse behaviors, such as front flip (top row), back flip (middle row), jumping (bottom row), etc.

