## Appendix

A Broader Impact Statement ..... 15
B Glossary of Symbols ..... 16
C General-Policy Definitions ..... 17
C. 1 General-Policy Definitions ..... 17
D Proofs ..... 18
D. 1 Main Theorem ..... 18
D. 2 Inverse Model Implies Density Ratio ..... 20
E Derivation of Density Ratio Objective ..... 22
F Markov State Abstractions and Kinematic Inseparability ..... 23
F. 1 Example MDP ..... 23
F. 2 "Strongly Markov" Implies KI ..... 23
G Implementation Details for Visual Gridworld ..... 25
G. 1 Computing Resources ..... 25
G. 2 Network Architectures ..... 25
G. 3 Hyperparameters ..... 26
H Implementation Details for DeepMind Control ..... 27
H. 1 Computing Resources ..... 27
H. 2 Markov Network Architecture ..... 27
H. 3 Hyperparameters ..... 27
I Additional Representation Visualizations ..... 29
J Gridworld Results for Increased Pretraining Time ..... 30
K DeepMind Control Experiment with RBF-DQN ..... 31
L Investigating the Value of Smoothness in Markov Abstractions ..... 32
M DeepMind Control Ablation Study ..... 33

## A Broader Impact Statement

Our approach enables agents to automatically construct useful, Markov state representations for reinforcement learning from rich observations. Since most reinforcement learning algorithms implicitly assume Markov abstract state representations, and since agents may struggle to learn when that assumption is violated, this work has the potential to benefit a large number of algorithms.
Our training objective is designed for neural networks, which are not guaranteed to converge to a global optimum when trained with stochastic gradient descent. Typically, such training objectives will only approximately satisfy the theoretical conditions they encode. However, this is not a drawback of our method-it applies to any representation learning technique that uses neural networks. Moreover, as neural network optimization techniques improve, our method will converge to a Markov abstraction, whereas other approaches may not. In the meantime, systems in safety-critical domains should ensure that they can cope with non-Markov abstractions without undergoing catastrophic failures.
We have shown experimentally that our method is effective in a variety of domains; however, other problem domains may require additional hyperparameter tuning, which can be expensive. Nevertheless, one benefit of our method is that Markov abstractions can be learned offline, without access to reward information. This means our algorithm could be used, in advance, to learn an abstraction for some problem domain, and then subsequent tasks in that environment (perhaps with different reward functions) could avoid the problem of perceptual abstraction as well as any associated training costs.

## B Glossary of Symbols

Here we provide a glossary of the most commonly-used symbols appearing in the rest of the paper.


## C General-Policy Definitions

## C. 1 General-Policy Definitions

The definitions of $T_{\phi, t}^{\pi}$ and $R_{\phi, t}^{\pi}$ in Section 4 only apply when the policy $\pi$ is a member of $\Pi_{\phi}$. Here we derive more general definitions that are valid for arbitrary policies, not just those in $\Pi_{\phi}$. For the special case where $\pi \in \Pi_{\phi}$, these definitions are equivalent to equations (4) and (5).

## Abstract transition probabilities.

$$
\begin{aligned}
& \operatorname{Pr}\left(z^{\prime} \mid a, z\right) \\
& =\sum_{x^{\prime} \in X} \operatorname{Pr}\left(z^{\prime} \mid x^{\prime}, a, z\right) \operatorname{Pr}\left(x^{\prime} \mid a, z\right) \\
= & \sum_{x^{\prime} \in X: \phi\left(x^{\prime}\right)=z^{\prime}} \operatorname{Pr}\left(x^{\prime} \mid a, z\right) \\
= & \sum_{x^{\prime} \in X: \phi\left(x^{\prime}\right)=z^{\prime}} \sum_{x \in X} \operatorname{Pr}\left(x^{\prime} \mid a, x, z\right) \operatorname{Pr}(x \mid a, z) \\
= & \sum_{x^{\prime} \in X: \phi\left(x^{\prime}\right)=z^{\prime}} \sum_{x \in X} \operatorname{Pr}\left(x^{\prime} \mid a, x, z\right) \frac{\operatorname{Pr}(a \mid x, z) \operatorname{Pr}(x \mid z)}{\sum_{\tilde{x} \in X} \operatorname{Pr}(a \mid \tilde{x}, z) \operatorname{Pr}\left(\tilde{x}_{t} \mid z_{t}\right)} \\
T_{\phi, t}^{\pi}\left(z^{\prime} \mid a, z\right):= & \sum_{x^{\prime} \in X: \phi\left(x^{\prime}\right)=z^{\prime}} \sum_{x \in X: \phi(x)=z} T\left(x^{\prime} \mid a, x\right) \frac{\pi_{t}(a \mid x) B_{\phi, t}^{\pi}(x \mid z)}{\sum_{\tilde{x} \in X} \pi_{t}(a \mid \tilde{x}) B_{\phi, t}^{\pi}(\tilde{x} \mid z)}
\end{aligned}
$$

## Abstract rewards.

$$
\begin{aligned}
& \sum_{r \in R} r \operatorname{Pr}\left(r \mid z^{\prime}, a, z\right) \\
= & \sum_{x^{\prime} \in X} \sum_{x \in X} \sum_{r \in R} r \operatorname{Pr}\left(r, x^{\prime}, x \mid z^{\prime}, a, z\right) \\
= & \sum_{x^{\prime} \in X} \sum_{x \in X} \sum_{r \in R} r \operatorname{Pr}\left(r \mid x^{\prime}, z^{\prime}, a, x, z\right) \operatorname{Pr}\left(x^{\prime}, x \mid z^{\prime}, a, z\right) \\
= & \sum_{x^{\prime} \in X} \sum_{x \in X} R\left(x^{\prime}, a, x\right) \frac{\operatorname{Pr}\left(z^{\prime} \mid x^{\prime}, a, x, z\right) \operatorname{Pr}\left(x^{\prime}, x \mid a, z\right)}{\operatorname{Pr}\left(z^{\prime} \mid a, z\right)} \\
= & \sum_{x^{\prime} \in X: \phi\left(x^{\prime}\right)=z^{\prime}} \sum_{x \in X} R\left(x^{\prime}, a, x\right) \frac{\operatorname{Pr}\left(x^{\prime} \mid a, x, z\right) \operatorname{Pr}(x \mid a, z)}{\operatorname{Pr}\left(z^{\prime} \mid a, z\right)} \\
= & \sum_{x^{\prime} \in X: \phi\left(x^{\prime}\right)=z^{\prime}} \sum_{x \in X} R\left(x^{\prime}, a, x\right) \frac{\operatorname{Pr}\left(x^{\prime} \mid a, x\right)}{\operatorname{Pr}\left(z^{\prime} \mid a, z\right)} \frac{\operatorname{Pr}(a \mid x) \operatorname{Pr}(x \mid z)}{\sum_{\tilde{x} \in X} \operatorname{Pr}(a \mid \tilde{x}) \operatorname{Pr}(\tilde{x} \mid z)} \\
R_{\phi, t}^{\pi}\left(z^{\prime}, a, z\right):= & \sum_{x^{\prime} \in X: \phi\left(x^{\prime}\right)=z^{\prime}} \sum_{x \in X: \phi(x)=z} R\left(x^{\prime}, a, x\right) \frac{T\left(x^{\prime} \mid a, x\right) \pi_{t}(a \mid x) B_{\phi, t}^{\pi}(x \mid z)}{T_{\phi, t}^{\pi}\left(z^{\prime} \mid a, z\right) \sum_{\tilde{x} \in X} \pi_{t}(a \mid \tilde{x}) B_{\phi, t}^{\pi}(\tilde{x} \mid z)}
\end{aligned}
$$

## D Proofs

Here we provide proofs of Theorem 1 and its corollary, which state that the Inverse Model and Density Ratio conditions are sufficient for $\phi$ and $M_{\phi}$ to be Markov. Then, to complement the counterexample from Section 4.2, we also present and prove a second theorem which states that, given the Inverse Model condition, the Density Ratio condition is in fact necessary for a Markov abstraction.

## D. 1 Main Theorem

The proof of Theorem 1 makes use of two lemmas: Lemma D.1, that equal $k$-step and $(k-1)$-step belief distributions imply equal $k$-step and $(k-1)$-step transition models and reward functions, and Lemma D.2, that equal $k$-step and ( $k-1$ )-step belief distributions imply equal $(k+1)$-step and $k$-step belief distributions. Since the lemmas apply for any arbitrary policy, we use the general-policy definitions from Appendix C.

Lemma D.1. Given an MDP $M$, abstraction $\phi$, policy $\pi$, initial state distribution $P_{0}$, and any $k \geq 1$, if $B_{\phi, t}^{\pi(k)}\left(x_{t} \mid z_{t},\left\{a_{t-i}, z_{t-i}\right\}_{i=1}^{k}\right)=B_{\phi, t}^{\pi(k-1)}\left(x_{t} \mid z_{t},\left\{a_{t-i}, z_{t-i}\right\}_{i=1}^{k-1}\right)$, then for all $a_{t} \in A$ and $z_{t+1} \in Z$ :

$$
\begin{aligned}
T_{\phi, t}^{\pi(k)}\left(z_{t+1} \mid\left\{a_{t-i}, z_{t-i}\right\}_{i=0}^{k}\right) & =T_{\phi, t}^{\pi(k-1)}\left(z_{t+1} \mid\left\{a_{t-i}, z_{t-i}\right\}_{i=0}^{k-1}\right) \\
\cap \quad R_{\phi, t}^{\pi(k)}\left(z_{t+1},\left\{a_{t-i}, z_{t-i}\right\}_{i=0}^{k}\right) & =R_{\phi, t}^{\pi(k-1)}\left(z_{t+1},\left\{a_{t-i}, z_{t-i}\right\}_{i=0}^{k-1}\right) .
\end{aligned}
$$

In the proof below, we start with $B_{\phi, t}^{\pi(k)}=B_{\phi, t}^{\pi(k-1)}$, and repeatedly multiply or divide both sides by the same quantity, or take the same summations of both sides, to obtain $T_{\phi, t}^{\pi(k)}=T_{\phi, t}^{\pi(k-1)}$, then apply the same process again, making use of the fact that $T_{\phi, t}^{\pi(k)}=T_{\phi, t}^{\pi(k-1)}$, to obtain $R_{\phi, t}^{\pi(k)}=R_{\phi, t}^{\pi(k-1)}$.

Proof:

$$
B_{\phi, t}^{\pi(k)}\left(x_{t} \mid z_{t},\left\{a_{t-i}, z_{t-i}\right\}_{i=1}^{k}\right)=B_{\phi, t}^{\pi(k-1)}\left(x_{t} \mid z_{t},\left\{a_{t-i}, z_{t-i}\right\}_{i=1}^{k-1}\right) .
$$

Let $a_{t} \in A$ be any action.

$$
\begin{align*}
& \Rightarrow \quad \pi_{t}\left(a_{t} \mid x_{t}\right) B_{\phi, t}^{\pi(k)}\left(x_{t} \mid z_{t},\left\{a_{t-i}, z_{t-i}\right\}_{i=1}^{k}\right)=\pi_{t}\left(a_{t} \mid x_{t}\right) B_{\phi, t}^{\pi(k-1)}\left(x_{t} \mid z_{t},\left\{a_{t-i}, z_{t-i}\right\}_{i=1}^{k-1}\right) \\
& \Rightarrow \quad \frac{\pi_{t}\left(a_{t} \mid x_{t}\right) B_{\phi, t}^{\pi(k)}\left(x_{t} \mid z_{t},\left\{a_{t-i}, z_{t-i}\right\}_{i=1}^{k}\right)}{\sum_{\tilde{x}_{t} \in X} \pi_{t}\left(a_{t} \mid \tilde{x}_{t}\right) B_{\phi, t}^{\pi(k)}\left(\tilde{x}_{t} \mid z_{t},\left\{a_{t-i}, z_{t-i}\right\}_{i=1}^{k}\right)}=\frac{\pi_{t}\left(a_{t} \mid x_{t}\right) B_{\phi, t}^{\pi(k-1)}\left(x_{t} \mid z_{t},\left\{a_{t-i}, z_{t-i}\right\}_{i-1}^{k-1}\right)}{\sum_{\tilde{x}_{t} \in X} \pi_{t}\left(a_{t} \mid \tilde{x}_{t}\right) B_{\phi, t}^{\pi(k-1)}\left(\tilde{x}_{t} \mid z_{t},\left\{a_{t-i}, z_{t-i}\right\}_{i=1}^{k-1}\right)} . \tag{6}
\end{align*}
$$

Let

$$
\begin{equation*}
C_{\phi, t}^{\pi(k)}:=\frac{\pi_{t}\left(a_{t} \mid x_{t}\right) B_{\phi, t}^{\pi(k)}\left(x_{t} \mid z_{t},\left\{a_{t-i}, z_{t-i}\right\}_{i=1}^{k}\right)}{\sum_{\tilde{x}_{t} \in X} \pi_{t}\left(a_{t} \mid \tilde{x}_{t}\right) B_{\phi, t}^{\pi(k)}\left(\tilde{x}_{t} \mid z_{t},\left\{a_{t-i}, z_{t-i}\right\}_{i=1}^{k}\right)} \tag{7}
\end{equation*}
$$

Combining (6) and (7), we obtain:

$$
\begin{array}{rlrl}
C_{\phi, t}^{\pi(k)} & =C_{\phi, t}^{\pi(k-1)} \\
\Rightarrow \quad \sum_{x_{t} \in z_{t}} \sum_{x_{t+1} \in z_{t+1}} T\left(x_{t+1} \mid a_{t}, x_{t}\right) C_{\phi, t}^{\pi(k)} & =\sum_{x_{t} \in z_{t}} \sum_{x_{t+1} \in z_{t+1}} T\left(x_{t+1} \mid a_{t}, x_{t}\right) C_{\phi, t}^{\pi(k-1)} \\
\Leftrightarrow & T_{\phi, t}^{\pi(k)}\left(z_{t+1} \mid\left\{a_{t-i}, z_{t-i}\right\}_{i=0}^{k}\right) & =T_{\phi, t}^{\pi(k-1)}\left(z_{t+1} \mid\left\{a_{t-i}, z_{t-i}\right\}_{i=0}^{k-1}\right) . \tag{9}
\end{array}
$$

Additionally, we can combine (8) and (9) and apply the same approach for rewards:

$$
\begin{array}{rlrl} 
& C_{\phi, t}^{\pi(k)} & =C_{\phi, t}^{\pi(k-1)} \\
\Rightarrow & \frac{C_{\phi, t}^{\pi(k)}}{T_{\phi, t}^{\pi(k)}\left(z_{t+1} \mid\left\{a_{t-i}, z_{t-i}\right\}_{i=0}^{k}\right)} & =\frac{C_{\phi, t}^{\pi(k-1)}}{T_{\phi, t}^{\pi(k-1)}\left(z_{t+1} \mid\left\{a_{t-i}, z_{t-i}\right\}_{i=0}^{k-1}\right)} \\
\Rightarrow \quad & \frac{T\left(x_{t+1} \mid a_{t}, z_{t}\right) C_{\phi, t}^{\pi(k)}}{T_{\phi, t}^{\pi(k)}\left(z_{t+1} \mid\left\{a_{t-i}, z_{t-i}\right\}_{i=0}^{k}\right)} & =\frac{T\left(x_{t+1} \mid a_{t}, z_{t}\right) C_{\phi, t}^{\pi(k-1)}}{T_{\phi, t}^{\pi(k-1)}\left(z_{t+1} \mid\left\{a_{t-i}, z_{t-i}\right\}_{i=0}^{k-1}\right)} \\
R_{\phi, t}^{\pi(k)}\left(z_{t+1},\left\{a_{t-i}, z_{t-i}\right\}_{i=0}^{k}\right) & =R_{\phi, t}^{\pi(k-1)}\left(z_{t+1},\left\{a_{t-i}, z_{t-i}\right\}_{i=0}^{k-1}\right) . \tag{10}
\end{array}
$$

Lemma D.2. Given an MDP $M$, abstraction $\phi$, policy $\pi$, and initial state distribution $P_{0}$, if for all $t \geq k$, $z_{t} \in Z$, and $x_{t} \in X$ such that $\phi\left(x_{t}\right)=z_{t}$, it holds that $B_{\phi, t}^{\pi(k)}\left(x_{t} \mid z_{t},\left\{a_{t-i}, z_{t-i}\right\}_{i=1}^{k}\right)=$ $B_{\phi, t}^{\pi(k-1)}\left(x_{t} \mid z_{t},\left\{a_{t-i}, z_{t-i}\right\}_{i=1}^{k-1}\right)$, then for all $z_{t+1} \in Z, x_{t+1} \in X: \phi\left(x_{t+1}\right)=z_{t+1}$,

$$
B_{\phi, t}^{\pi(k+1)}\left(x_{t+1} \mid z_{t+1},\left\{a_{t-i}, z_{t-i}\right\}_{i=0}^{k}\right)=B_{\phi, t}^{\pi(k)}\left(x_{t+1} \mid z_{t+1},\left\{a_{t-i}, z_{t-i}\right\}_{i=0}^{k-1}\right)
$$

To prove this lemma, we invoke Lemma D. 1 to obtain $T_{\phi, t}^{\pi(k)}=T_{\phi, t}^{\pi(k-1)}$, and then follow the same approach as before, performing operations to both sides until we achieve the desired result.

Proof:
Let $T_{\phi, t}^{\pi(k)}$ be defined via (4) and (3). Applying Lemma D. 1 to the premise gives:

$$
T_{\phi, t}^{\pi(k)}\left(z_{t+1} \mid\left\{a_{t-i}, z_{t-i}\right\}_{i=0}^{k}\right)=T_{\phi, t}^{\pi(k-1)}\left(z_{t+1} \mid\left\{a_{t-i}, z_{t-i}\right\}_{i=0}^{k-1}\right)
$$

Returning to the premise, we have:

$$
\begin{aligned}
& B_{\phi, t}^{\pi(k)}\left(x_{t} \mid z_{t},\left\{a_{t-i}, z_{t-i}\right\}_{i=1}^{k}\right)=B_{\phi, t}^{\pi(k-1)}\left(x_{t} \mid z_{t},\left\{a_{t-i}, z_{t-i}\right\}_{i=1}^{k-1}\right) \\
& \Rightarrow \quad \sum_{\phi, t}^{\pi(k)}\left(x_{t} \mid z_{t},\left\{a_{t-i}, z_{t-i}\right\}_{i=1}^{k}\right) \\
& T_{\phi, t}^{\pi(k)}\left(z_{t+1} \mid\left\{a_{t-i}, z_{t-i}\right\}_{i=0}^{k}\right)=\frac{B_{\phi, t}^{\pi(k-1)}\left(x_{t} \mid z_{t},\left\{a_{t-i}, z_{t-i}\right\}_{i=1}^{k-1}\right)}{T_{\phi, t}^{\pi(k-1)}\left(z_{t+1} \mid\left\{a_{t-i}, z_{t-i}\right\}_{i=0}^{k-1}\right)} \\
& \Rightarrow \quad \sum_{\substack{x_{t} \in X: \\
\phi\left(x_{t}\right)=z_{t}}} \frac{T\left(x_{t+1} \mid a_{t}, x_{t}\right) B_{\phi, t}^{\pi(k)}\left(x_{t} \mid z_{t},\left\{a_{t-i}, z_{t-i}\right\}_{i=1}^{k}\right)}{T_{\phi, t}^{\pi(k)}\left(z_{t+1} \mid\left\{a_{t-i}, z_{t-i}\right\}_{i=0}^{k}\right)}=\sum_{\substack{x_{t} \in X: \\
\phi\left(x_{t}\right)=z_{t}}} \frac{T\left(x_{t+1} \mid a_{t}, x_{t}\right) B_{\phi, t}^{\pi(k-1)}\left(x_{t} \mid z_{t},\left\{a_{t-i}, z_{t-i}\right\}_{i=1}^{k-1}\right)}{T_{\phi, t}^{\pi(k-1)}\left(z_{t+1} \mid\left\{a_{t-i}, z_{t-i}\right\}_{i=0}^{k-1}\right)} \\
& \Rightarrow \quad B_{\phi, t}^{\pi(k+1)}\left(x_{t+1} \mid z_{t+1},\left\{a_{t-i}, z_{t-i}\right\}_{i=0}^{k}\right)=B_{\phi, t}^{\pi(k)}\left(x_{t+1} \mid z_{t+1},\left\{a_{t-i}, z_{t-i}\right\}_{i=0}^{k-1}\right) .
\end{aligned}
$$

We now summarize the proof of the main theorem. We begin by showing that the belief distributions $B_{\phi, t}^{\pi(k)}$ and $B_{\phi, t}^{\pi(k-1)}$ must be equal for $k=1$, and use Lemma D. 1 to prove the base case of the theorem. Then we use Lemma D. 2 to prove that the theorem holds in general.

## Proof of Theorem 1:

Base case. For each $\pi \in \Pi_{\phi}$, let $B_{\phi, t}^{\pi}$ be defined via (2). Then, starting from the Density Ratio condition, for any $z_{t+1}, z_{t} \in Z$, and $x_{t+1} \in X$ such that $\phi\left(x_{t+1}\right)=z_{t+1}$, and any action $a_{t-1} \in A$ :

$$
\begin{array}{rlrl} 
& & \frac{P_{\phi, t}^{\pi}\left(z_{t} \mid z_{t-1}\right)}{P_{\phi, t}^{\pi}\left(z_{t}\right)} & =\frac{P_{t}^{\pi}\left(x_{t} \mid z_{t-1}\right)}{P_{t}^{\pi}\left(x_{t}\right)} \\
\Rightarrow \quad & \frac{P_{\phi, t}^{\pi}\left(z_{t} \mid z_{t-1}\right)}{P_{\phi, t}^{\pi}\left(z_{t}\right)} & =\sum_{x_{t-1} \in X} \frac{P_{t}^{\pi}\left(x_{t} \mid x_{t-1}\right)}{P_{t}^{\pi}\left(x_{t}\right)} B_{\phi, t}^{\pi}\left(x_{t-1} \mid z_{t-1}\right) \\
\Rightarrow \quad \frac{P_{t}^{\pi}\left(x_{t}\right)}{P_{\phi, t}^{\pi}\left(z_{t}\right)} & =\sum_{x_{t-1} \in X} \frac{P_{t}^{\pi}\left(x_{t} \mid x_{t-1}\right) B_{\phi, t}^{\pi}\left(x_{t-1} \mid z_{t-1}\right)}{P_{\phi, t}^{\pi}\left(z_{t} \mid z_{t-1}\right)} \cdot \frac{I_{t}^{\pi}\left(a_{t-1} \mid x_{t}, x_{t-1}\right) \pi_{\phi, t}\left(a_{t-1} \mid z_{t-1}\right)}{I_{\phi, t}\left(a_{t-1} \mid z_{t}, z_{t-1}\right) \pi_{t}\left(a_{t-1} \mid x_{t-1}\right)} \\
\Rightarrow \quad & \frac{\mathbb{1}\left[\phi\left(x_{t}\right)=z_{t}\right] P_{t}^{\pi}\left(x_{t}\right)}{\sum_{\tilde{x}_{t} \in X} P_{t}^{\pi}\left(\tilde{x}_{t}\right)} & =\mathbb{1}\left[\phi\left(x_{t}\right)=z_{t}\right] \sum_{x_{t-1} \in X} \frac{T\left(x_{t} \mid a_{t-1}, x_{t-1}\right) B_{\phi, t}^{\pi}\left(x_{t-1} \mid z_{t-1}\right)}{T_{\phi, t}^{\pi}\left(z_{t} \mid a_{t-1}, z_{t-1}\right)} \\
\Rightarrow \quad B_{\phi, t}\left(x_{t} \mid z_{t}\right) & =\frac{\mathbb{1}\left[\phi\left(x_{t}\right)=z_{t}\right] \sum_{x_{t-1} \in X} T_{\phi}\left(x_{t} \mid a_{t-1}, x_{t-1}\right) B_{\phi, t}^{\pi}\left(x_{t-1} \mid z_{t-1}\right)}{\sum_{\tilde{x}_{t} \in z_{t}} \sum_{\tilde{x}_{t-1} \in z_{t-1}} T\left(\tilde{x}_{t} \mid a_{t-1}, \tilde{x}_{t-1}\right) B_{\phi, t}^{\pi}\left(\tilde{x}_{t-1} \mid z_{t-1}\right)} \\
\Rightarrow \quad B_{\phi, t}^{\pi(0)}\left(x_{t} \mid z_{t}\right) & =B_{\phi, t}^{\pi(1)}\left(x_{t} \mid z_{t}, a_{t-1}, z_{t-1}\right) \tag{11}
\end{array}
$$

Here (11) satisfies the conditions of Lemma D. 1 (with $k=1$ ), therefore, for all $a_{t} \in A$ :

$$
\begin{aligned}
T_{\phi, t}^{\pi(0)}\left(z_{t+1} \mid a_{t}, z_{t}\right) & =T_{\phi, t}^{\pi(1)}\left(z_{t+1} \mid a_{t}, z_{t}, a_{t-1}, z_{t-1}\right) \\
\text { and } \quad R_{\phi, t}^{\pi(0)}\left(z_{t+1}, a_{t}, z_{t}\right) & =R_{\phi, t}^{\pi(1)}\left(z_{t+1}, a_{t}, z_{t}, a_{t-1}, z_{t-1}\right)
\end{aligned}
$$

This proves the theorem for $k=1$.
Induction on $k$. Note that (11) also allows us to apply Lemma D.2. Therefore, by induction on $k$ :

$$
\begin{equation*}
B_{\phi, t}^{\pi(k+1)}\left(x_{t+1} \mid z_{t+1},\left\{a_{t-i}, z_{t-i}\right\}_{i=0}^{k}\right)=B_{\phi, t}^{\pi}\left(x_{t+1} \mid z_{t+1}\right) \quad \forall k \geq 1 \tag{12}
\end{equation*}
$$

When (12) holds, we informally say that the belief distribution is Markov.
Proof of Corollary 1.1: Follows directly from Definition 2 and Lemma D. 1 via induction on $k$.
This first corollary says that a Markov abstraction implies a Markov abstract state representation. The next one says that, if a belief distribution is non-Markov over some horizon $n$, it must also be non-Markov when conditioning on a single additional timestep.

Corollary 1.2. If there exists some $n \geq 1$ such that $B_{\phi, t}^{\pi(n)} \neq B_{\phi, t}^{\pi}$, then $B_{\phi, t}^{\pi(1)} \neq B_{\phi, t}^{\pi}$.
Proof: Suppose such an $n$ exists, and assume for the sake of contradiction that $B_{\phi, t}^{\pi(1)}=B_{\phi, t}^{\pi}$. Then by Lemma D. $2, B_{\phi, t}^{\pi(k)}=B_{\phi, t}^{\pi}$ for all $k \geq 1$. However this is impossible, since we know there exists some $n \geq 1$ such that $B_{\phi, t}^{\pi(n)} \neq B_{\phi, t}^{\pi}$. Therefore $B_{\phi, t}^{\pi(1)} \neq B_{\phi, t}^{\pi}$.

## D. 2 Inverse Model Implies Density Ratio

As discussed in Section 4.2, the Inverse Model condition is not sufficient to ensure a Markov abstraction. In fact, what is missing is precisely the Density Ratio condition. Theorem 1 already states that, given the Inverse Model condition, the Density Ratio condition is sufficient for an abstraction to be Markov over its policy class; the following theorem states that it is also necessary.

Theorem 2. If $\phi: X \rightarrow Z$ is a Markov abstraction of MDP $M=(X, A, R, T, \gamma)$ for any policy in the policy class $\Pi_{\phi}$, and the Inverse Model condition of Theorem 1 holds for every timestep $t$, then the Density Ratio condition also holds for every timestep $t$.

## Proof:

Since $\phi$ is a Markov abstraction, equation (12) holds for any $k \geq 1$. Fixing $k=1$, we obtain:

$$
\begin{align*}
& B_{\phi, t}^{\pi(1)}\left(x^{\prime} \mid z^{\prime}, a, z\right)=B_{\phi, t}^{\pi(0)}\left(x^{\prime} \mid z^{\prime}\right) \\
& \frac{\mathbb{1}\left[\phi\left(x^{\prime}\right)=z^{\prime}\right] \sum_{\tilde{x} \in X} T\left(x^{\prime} \mid a, \tilde{x}\right) B_{\phi, t}^{\pi}(\tilde{x} \mid z)}{\sum_{\tilde{x}^{\prime} \in X: \phi\left(x^{\prime}\right)=z^{\prime}} \sum_{\tilde{x} \in X: \phi(x)=z} T\left(\tilde{x}^{\prime} \mid a, \tilde{x}\right) B_{\phi, t}^{\pi}(\tilde{x} \mid z)}=\frac{\mathbb{1}\left[\phi\left(x^{\prime}\right)=z^{\prime}\right] P_{t}^{\pi}\left(x^{\prime}\right)}{\sum_{\tilde{x})^{\prime} \in X: \phi\left(\tilde{x}^{\prime}\right)=z^{\prime}} P_{t}^{\pi}\left(\tilde{x}^{\prime}\right)} \\
& \frac{\sum_{\tilde{x} \in X} T\left(x^{\prime} \mid a, \tilde{x}\right) B_{\phi, t}^{\pi}(\tilde{x} \mid z)}{T_{\phi, t}^{\pi}\left(z^{\prime} \mid a, z\right)}=\frac{P_{t}^{\pi}\left(x^{\prime}\right)}{P_{\phi, t}^{\pi}\left(z^{\prime}\right)} \\
& \frac{\sum_{\tilde{x} \in X} T\left(x^{\prime} \mid a, \tilde{x}\right) B_{\phi, t}^{\pi}(\tilde{x} \mid z)}{P_{t}^{\pi}\left(x^{\prime}\right)}=\frac{T_{\phi, t}^{\pi}\left(z^{\prime} \mid a, z\right)}{P_{\phi, t}^{\pi}\left(z^{\prime}\right)} \\
& \sum_{\tilde{x} \in X} \frac{I_{t}^{\pi}\left(a \mid x^{\prime}, \tilde{x}\right) P_{t}^{\pi}\left(x^{\prime} \mid \tilde{x}\right) B_{\phi, t}^{\pi}(\tilde{x} \mid z)}{P_{t}^{\pi}\left(x^{\prime}\right) \pi(a \mid \tilde{x})}=\frac{I_{\phi, t}^{\pi}\left(a \mid z^{\prime}, z\right) P_{\phi, t}^{\pi}\left(z^{\prime} \mid z\right)}{P_{\phi, t}^{\pi}\left(z^{\prime}\right) \pi(a \mid z)} \\
& \sum_{\tilde{x} \in X} I_{t}^{\pi}\left(a \mid x^{\prime}, \tilde{x}\right) P_{t}^{\pi}\left(x^{\prime} \mid \tilde{x}\right) B_{\phi, t}^{\pi}(\tilde{x} \mid z)  \tag{13}\\
& P_{t}^{\pi}\left(x^{\prime}\right)=\frac{I_{\phi, t}^{\pi}\left(a \mid z^{\prime}, z\right) P_{\phi, t}^{\pi}\left(z^{\prime} \mid z\right)}{P_{\phi, t}^{\pi}\left(z^{\prime}\right)}
\end{align*}
$$

Here we apply the Inverse Model condition, namely that $I_{t}^{\pi}\left(a \mid x^{\prime}, x\right)=I_{\phi, t}^{\pi}\left(a \mid z^{\prime}, z\right)$ for all $z, z^{\prime} \in Z$; $x, x^{\prime} \in X$, such that $\phi\left(x^{\prime}\right)=z^{\prime}$ and $\phi(x)=z$.

$$
\begin{aligned}
\frac{\sum_{\tilde{x} \in X} P_{t}^{\pi}\left(x^{\prime} \mid \tilde{x}\right) B_{\phi, t}^{\pi}(\tilde{x} \mid z)}{P_{t}^{\pi}\left(x^{\prime}\right)} & =\frac{P_{\phi, t}^{\pi}\left(z^{\prime} \mid z\right)}{P_{\phi, t}^{\pi}\left(z^{\prime}\right)} \\
\frac{P_{t}^{\pi}\left(x^{\prime} \mid z\right)}{P_{t}^{\pi}\left(x^{\prime}\right)} & =\frac{P_{\phi, t}^{\pi}\left(z^{\prime} \mid z\right)}{P_{\phi, t}^{\pi}\left(z^{\prime}\right)} \quad \text { (Density Ratio) }
\end{aligned}
$$

It may appear that Theorems 1 and 2 together imply that the Inverse Model and Density Ratio conditions are necessary and sufficient for an abstraction to be Markov over its policy class; however, this is not quite true. Both conditions, taken together, are sufficient for an abstraction to be Markov, and, given the Inverse Model condition, the Density Ratio condition is necessary. Examining equation (13), we see that, had we instead assumed the Density Ratio condition for Theorem 2 (rather than the Inverse Model condition), we would not recover $I_{t}^{\pi}\left(a \mid x^{\prime}, x\right)=I_{\phi, t}^{\pi}\left(a \mid z^{\prime}, z\right)$, but rather $\sum_{\tilde{x} \in X} I_{t}^{\pi}\left(a \mid x^{\prime}, \tilde{x}\right) B_{\phi, t}^{\pi}(\tilde{x} \mid z)=I_{\phi, t}^{\pi}\left(a \mid z^{\prime}, z\right)$. That is, the Inverse Model condition would only be guaranteed to hold in expectation, but not for arbitrary $x \in X$.

## E Derivation of Density Ratio Objective

Our Density Ratio objective in Section 5 is based on the following derivation, adapted from Tiao (2017).

Suppose we have a dataset consisting of samples $X_{c}=\left\{{x_{c}^{\prime}}^{(i)}\right\}_{i=1}^{n_{c}}$ drawn from conditional distribution $\operatorname{Pr}\left(x^{\prime} \mid x\right)$, and samples $X_{m}=\left\{x_{m}^{\prime}{ }^{(j)}\right\}_{j=1}^{n_{m}}$ drawn from marginal distribution $\operatorname{Pr}\left(x^{\prime}\right)$. We assign label $y=1$ to samples from $X_{c}$ and $y=0$ to samples from $X_{m}$, and our goal is to predict the label associated with each sample. To construct an estimator, we rename the two distributions $p\left(x^{\prime} \mid y=1\right):=\operatorname{Pr}\left(x^{\prime} \mid z\right)$ and $p\left(x^{\prime} \mid y=0\right):=\operatorname{Pr}\left(x^{\prime}\right)$ and rewrite the density ratio $\delta\left(x^{\prime}\right):=\frac{\operatorname{Pr}\left(x^{\prime} \mid x\right)}{\operatorname{Pr}\left(x^{\prime}\right)}$ as follows:

$$
\begin{equation*}
\delta\left(x^{\prime}\right)=\frac{p\left(x^{\prime} \mid y=1\right)}{p\left(x^{\prime} \mid y=0\right)}=\frac{p\left(y=1 \mid x^{\prime}\right) p\left(x^{\prime}\right)}{p(y=1)} \frac{p(y=0)}{p\left(y=0 \mid x^{\prime}\right) p\left(x^{\prime}\right)}=\frac{n_{m}}{\left(n_{m}+n_{c}\right)} \frac{\left(n_{m}+n_{c}\right)}{n_{c}} \frac{p\left(y=1 \mid x^{\prime}\right)}{p\left(y=0 \mid x^{\prime}\right)}=\frac{n_{m}}{n_{c}} \frac{p\left(y=1 \mid x^{\prime}\right)}{1-p\left(y=1 \mid x^{\prime}\right)} \tag{14}
\end{equation*}
$$

When $n_{c}=n_{m}=N$, which is the case for our implementation, the leading fraction can be ignored. To estimate $\delta\left(x^{\prime}\right)$, we can simply train a classifier $g\left(x^{\prime}, x ; \theta_{g}\right)$ to approximate $p\left(y=1 \mid x^{\prime}\right)$ and then substitute $g$ for $p\left(y=1 \mid x^{\prime}\right)$ in (14).
However, we need not estimate $\delta\left(x^{\prime}\right)$ to satisfy the Density Ratio condition; we need only ensure $\delta_{\phi}\left(z^{\prime}\right)=\mathbb{E}_{B_{\phi}}\left[\delta\left(x^{\prime}\right)\right]$. We therefore repeat the derivation for abstract states and obtain $\delta_{\phi}\left(z^{\prime}\right):=$ $\frac{\operatorname{Pr}\left(z^{\prime} \mid z\right)}{\operatorname{Pr}\left(z^{\prime}\right)}=\frac{n_{m}}{n_{c}} \frac{q\left(y=1 \mid z^{\prime}\right)}{1-q\left(y=1 \mid z^{\prime}\right)}$, where $q$ is our renamed distribution, and modify our classifier $g$ to accept abstract states instead of ground states. The labels are the same regardless of whether we use ground or abstract state pairs, so training will cause $g$ to approach $p$ and $q$ simultaneously, thus driving $\delta_{\phi}\left(z^{\prime}\right) \rightarrow \delta\left(x^{\prime}\right)$ and satisfying the Density Ratio condition.

## F Markov State Abstractions and Kinematic Inseparability

As discussed in Section 3, the notion of kinematic inseparability (Misra et al., 2020) is closely related to Markov abstraction. Recall that two states $x_{1}^{\prime}$ and $x_{2}^{\prime}$ are defined to be kinematically inseparable if $\operatorname{Pr}\left(x, a \mid x_{1}^{\prime}\right)=\operatorname{Pr}\left(x, a \mid x_{2}^{\prime}\right)$ and $T\left(x^{\prime \prime} \mid a, x_{1}^{\prime}\right)=T\left(x^{\prime \prime} \mid a, x_{2}^{\prime}\right)$ (which the authors call "backwards" and "forwards" KI, respectively). Misra et al. (2020) define kinematic inseparability abstractions over the set of all possible "roll-in" distributions $u(x, a)$ supported on $X \times A$, and technically, the backwards KI probabilities $\operatorname{Pr}\left(x, a \mid x^{\prime}\right)$ depend on $u$. However, to support choosing a policy class, we can just as easily define $u$ in terms of a policy: $u(x, a):=\pi(a \mid x) P_{t}^{\pi}(x)$. This formulation leads to:

$$
P_{t}^{\pi}\left(x, a \mid x^{\prime}\right):=\frac{T\left(x^{\prime} \mid a, x\right) \pi(a \mid x) P_{t}^{\pi}(x)}{\sum_{\tilde{x} \in X, \tilde{a} \in A} T\left(x^{\prime} \mid \tilde{a}, \tilde{x}\right) \pi(\tilde{a} \mid \tilde{x}) P_{t}^{\pi}(\tilde{x})}
$$

Definition 3. An abstraction $\phi: X \rightarrow Z$ is a kinematic inseparability abstraction of MDP $M=$ ( $X, A, R, T, \gamma$ ) over policy class $\Pi_{C}$, if for all policies $\pi \in \Pi_{C}$, and all $a \in A ; x, x_{1}^{\prime}, x_{2}^{\prime}, x^{\prime \prime} \in X$ such that $\phi\left(x_{1}^{\prime}\right)=\phi\left(x_{2}^{\prime}\right) ; P_{t}^{\pi}\left(x, a \mid x_{1}^{\prime}\right)=P_{t}^{\pi}\left(x, a \mid x_{2}^{\prime}\right)$ and $T\left(x^{\prime \prime} \mid a, x_{1}^{\prime}\right)=T\left(x^{\prime \prime} \mid a, x_{2}^{\prime}\right)$.

Similarly, we can define forward-or backward-KI abstractions where only $T\left(x^{\prime \prime} \mid a, x_{1}^{\prime}\right)=$ $T\left(x^{\prime \prime} \mid a, x_{2}^{\prime}\right)$-or respectively, $P_{t}^{\pi}\left(x, a \mid x_{1}^{\prime}\right)=P_{t}^{\pi}\left(x, a \mid x_{2}^{\prime}\right)$-is guaranteed to hold. A KI abstraction is one that is both forward KI and backward KI.

The KI conditions are slightly stronger conditions than those of Theorem 1, as the following example demonstrates.

## F. 1 Example MDP

The figure below modifies the transition dynamics of the MDP in Section 2, such that the action $a_{1}$ has the same effect everywhere: to transition to either central state, $x_{1}$ or $x_{2}$, with equal probability.


Figure 5: An MDP and a Markov abstraction that is not a KI abstraction.
By the same reasoning as in Section 4.2, the Inverse Model condition holds here, but now, due to the shared transition dynamics of action $a_{1}$, the Density Ratio condition holds as well, for any policy in $\Pi_{\phi}$. We can apply Theorem 1 to see that the abstraction is Markov, or we can simply observe that conditioning the belief distribution $B_{\phi, t}^{\pi}\left(x \mid z_{B}\right)$ on additional history has no effect, since any possible trajectory ending in $z_{B}$ leads to the same 50-50 distribution over ground states $x_{1}$ and $x_{2}$. Either way, $\phi$ is Markov by Definition 2.
This abstraction also happens to satisfy the backwards KI condition, since $P_{t}^{\pi}\left(x, a \mid x_{1}^{\prime}\right)=P_{t}^{\pi}\left(x, a \mid x_{2}^{\prime}\right)$ for any $(x, a)$ pair and any policy. However, clearly $T\left(x^{\prime} \mid a_{0}, x_{1}\right) \neq T\left(x^{\prime} \mid a_{0}, x_{2}\right)$, and therefore the forwards KI condition does not hold and this is not a KI abstraction.
This example shows that the Markov conditions essentially take the stance that because $\pi$ is restricted to the policy class $\Pi_{\phi}$, knowing the difference between $x_{1}$ and $x_{2}$ doesn't help, because $\pi$ must have the same behavior for both ground states. By contrast, the KI conditions take the stance that because $x_{1}$ and $x_{2}$ have different dynamics, the agent may wish to change its behavior based on which state it sees, so it ought to choose an abstraction that does not limit decision making in that respect.

## F. 2 "Strongly Markov" Implies KI

In Section 5, we mentioned that the Density Ratio training objective was stronger than necessary to ensure the corresponding condition of Theorem 1. Instead of encoding the condition $\frac{\operatorname{Pr}\left(x^{\prime} \mid z\right)}{\operatorname{Pr}\left(x^{\prime}\right)}=$ $\frac{\operatorname{Pr}\left(z^{\prime} \mid z\right)}{\operatorname{Pr}\left(z^{\prime}\right)}$, we discussed how the contrastive training procedure actually encodes the stronger condition $\frac{\operatorname{Pr}\left(x^{\prime} \mid x\right)}{\operatorname{Pr}\left(x^{\prime}\right)}=\frac{\operatorname{Pr}\left(z^{\prime} \mid z\right)}{\operatorname{Pr}\left(z^{\prime}\right)}$ that holds for each ground state $x$ individually, rather than just in expectation.

Let us call the latter condition the Strong Density Ratio condition, and call its combination with the Inverse Model condition the Strong Markov conditions.

Clearly the Strong Markov conditions imply the original Markov conditions, and, as the following theorem shows, they also imply the KI conditions.

Theorem 3. If $\phi: X \rightarrow Z$ is an abstraction of $M D P M=(X, A, R, T, \gamma)$ such that for any policy $\pi$ in the policy class $\Pi_{\phi}$, both the Inverse Model condition of Theorem 1 and the Strong Density Ratio condition-i.e. $\frac{P_{t}^{\pi}\left(x^{\prime} \mid x\right)}{P_{t}^{\pi}\left(x^{\prime}\right)}=\frac{P_{\phi, t}^{\pi}\left(z^{\prime} \mid z\right)}{P_{\phi, t}^{\pi}\left(z^{\prime}\right)}$, for all $z, z^{\prime} \in Z ; x, x^{\prime} \in X$ such that $\phi(x)=z$ and $\phi\left(x^{\prime}\right)=z^{\prime}$-hold for every timestep $t$, then $\phi$ is a kinematic inseparability abstraction.

Proof:
Starting from the Inverse Model condition, we have $I_{t}^{\pi}\left(a \mid x^{\prime}, x\right)=I_{\phi, t}^{\pi}\left(a \mid z^{\prime}, z\right)$ for all $z, z^{\prime} \in Z$; $x, x^{\prime} \in X$ such that $\phi(x)=z$ and $\phi\left(x^{\prime}\right)=z^{\prime}$. Independently varying either $x^{\prime} \in \phi^{-1}\left(z^{\prime}\right)$ or $x \in \phi^{-1}(z)$, we obtain the following:

$$
\begin{align*}
& {\left[\text { Vary } x^{\prime}\right] \rightarrow \quad I_{t}^{\pi}\left(a \mid x_{1}^{\prime}, x\right)=I_{t}^{\pi}\left(a \mid x_{2}^{\prime}, x\right)}  \tag{15}\\
& {[\text { Vary } x] \rightarrow}  \tag{16}\\
& I_{t}^{\pi}\left(a \mid x^{\prime}, x_{1}\right)=I_{t}^{\pi}\left(a \mid x^{\prime}, x_{2}\right)
\end{align*}
$$

Similarly, if we start from the Strong Density Ratio condition, we obtain:

$$
\begin{align*}
& {\left[\operatorname{Vary} x^{\prime}\right] \rightarrow \quad \frac{P_{t}^{\pi}\left(x_{1}^{\prime} \mid x\right)}{P_{t}^{\pi}\left(x_{1}^{\prime}\right)}=\frac{P_{t}^{\pi}\left(x_{2}^{\prime} \mid x\right)}{P_{t}^{\pi}\left(x_{2}^{\prime}\right)}}  \tag{17}\\
& {[\operatorname{Vary} x] \rightarrow \frac{P_{t}^{\pi}\left(x^{\prime} \mid x_{1}\right)}{P_{t}^{\pi}\left(x^{\prime}\right)}=\frac{P_{t}^{\pi}\left(x^{\prime} \mid x_{2}\right)}{P_{t}^{\pi}\left(x^{\prime}\right)}} \tag{18}
\end{align*}
$$

If we apply Bayes' theorem to (17), we can cancel terms in the result, and also in (18), to obtain:

$$
\begin{align*}
P_{t}^{\pi}\left(x \mid x_{1}^{\prime}\right) & =P_{t}^{\pi}\left(x \mid x_{2}^{\prime}\right)  \tag{19}\\
\text { and } \quad P_{t}^{\pi}\left(x^{\prime} \mid x_{1}\right) & =P_{t}^{\pi}\left(x^{\prime} \mid x_{2}\right) . \tag{20}
\end{align*}
$$

Combining (15) with (19), we obtain the backwards KI condition:

$$
\begin{align*}
I_{t}^{\pi}\left(a \mid x_{1}^{\prime}, x\right) P_{t}^{\pi}\left(x \mid x_{1}^{\prime}\right) & =I_{t}^{\pi}\left(a \mid x_{2}^{\prime}, x\right) P_{t}^{\pi}\left(x \mid x_{2}^{\prime}\right) \\
P_{t}^{\pi}\left(x, a \mid x_{1}^{\prime}\right) & =P_{t}^{\pi}\left(x, a \mid x_{2}^{\prime}\right) \tag{BackwardsKI}
\end{align*}
$$

Similarly, we can combine (16) with (20) to obtain the forwards KI condition:

$$
\begin{aligned}
I_{t}^{\pi}\left(a \mid x^{\prime}, x_{1}\right) P_{t}^{\pi}\left(x^{\prime} \mid x_{1}\right) & =I_{t}^{\pi}\left(a \mid x^{\prime}, x_{2}\right) P_{t}^{\pi}\left(x^{\prime} \mid x_{2}\right) \\
T\left(x^{\prime} \mid a, x_{1}\right) \pi\left(a \mid x_{1}\right) & =T\left(x^{\prime} \mid a, x_{2}\right) \pi\left(a \mid x_{2}\right) \\
T\left(x^{\prime} \mid a, x_{1}\right) & =T\left(x^{\prime} \mid a, x_{2}\right)
\end{aligned}
$$

(Forwards KI)

Thus, we see that the training objectives in Section 5 encourage learning a kinematic inseparability abstraction in addition to a Markov abstraction. This helps avoid representation collapse by ensuring that we do not group together any states for which a meaningful kinematic distinction can be made.

## G Implementation Details for Visual Gridworld

The visual gridworld is a $6 \times 6$ grid with four discrete actions: up, down, left, and right. Observed states are generated by converting the agent's $(x, y)$ position to a one-hot image representation (see Figure 1, left). The image displays each position in the $6 \times 6$ grid as a $3 \mathrm{px}-\mathrm{by}-3 \mathrm{px}$ patch, inside of which we light up one pixel (in the center) and then smooth it using a truncated Gaussian kernel. This results in an $18 \times 18$ image (where 3 px-by- 3 px grid cells are equidistant), to which we then add per-pixel noise from another truncated Gaussian. During pretraining, there are no rewards or terminal states. During training, for each random seed, a single state is designated to be the goal state, and the agent receives -1 reward per timestep until it reaches the goal state, at which point a new episode begins, with the agent in a random non-goal location.

## G. 1 Computing Resources

To build the figures in the paper, we pretrained 5 different abstractions, and trained 7 different agents, each with 300 seeds. Each 3000 -step pretraining run takes about 1 minute, and each training run takes about 30 seconds, on a 2016 MacBook Pro 2 GHz i5 with no GPU, for a total of about 42 compute hours. We ran these jobs on a computing cluster with comparable processors or better.

## G. 2 Network Architectures

```
FeatureNet(
    (phi): Encoder(
        (0): Reshape(-1, 252)
        (1): Linear(in_features=252, out_features=32, bias=True)
        (2): Tanh()
        (3): Linear(in_features=32, out_features=2, bias=True)
        (4): Tanh()
    )
    (inv_model): InverseNet(
        (0): Linear(in_features=4, out_features=32, bias=True)
        (1): Tanh()
        (2): Linear(in_features=32, out_features=4, bias=True)
    )
    (contr_model): ContrastiveNet(
        (0): Linear(in_features=4, out_features=32, bias=True)
        (1): Tanh()
        (2): Linear(in_features=32, out_features=1, bias=True)
        (3): Sigmoid()
    )
)
QNet(
    (Q): Linear(in_features=2, out_features=32, bias=True)
    (1): ReLU()
    (2): Linear(in_features=32, out_features=4, bias=True)
)
AutoEncoder / PixelPredictor(
    (phi): Encoder(
        (0): Reshape(-1, 252)
        (1): Linear(in_features=252, out_features=32, bias=True)
        (2): Tanh()
        (3): Linear(in_features=32, out_features=2, bias=True)
        (4): Tanh()
    )
    (phi_inverse): Decoder(
        (0): Linear(in_features=2, out_features=32, bias=True)
        (1): Tanh()
        (2): Linear(in_features=32, out_features=252, bias=True)
        (3): Tanh()
        (4): Reshape(-1, 21, 12)
    )
    MSELoss()
)
```


## G. 3 Hyperparameters

We tuned the DQN hyperparameters until it learned effectively with expert features (i.e. ground-truth $(x, y)$ position), then we left the DQN hyperparameters fixed while tuning pretraining hyperparameters. For pretraining, we considered 3000 and 30,000 gradient updates (see Appendix J), and batch sizes within $\{512,1024,2048\}$. We found that the higher batch size was helpful for stabilizing the offline representations. We also did some informal experiments with latent dimensionality above 2 , such as 3 or 10, which produced similar results: representations were still Markov, but harder to interpret. We use 2 dimensions in the paper for ease of visualization. We did not tune the loss coefficients, but we include ablations where either $\alpha$ or $\beta$ is set to zero.

| Hyperparameter | Value |
| :--- | :---: |
| Number of seeds | 300 |
| Optimizer | Adam |
| Learning rate | 0.003 |
| Batch size | 2048 |
| Gradient updates | 3000 |
| Latent dimensions | 2 |
| Number of conditional samples, $n_{c}$ | 1 |
| Number of marginal samples, $n_{m}$ | 1 |
| Loss coefficients |  |
| $\mathcal{L}_{\text {Inverse }}(\alpha)$ | 1.0 |
| $\mathcal{L}_{\text {Contrastive }}(\beta)$ | 1.0 |
| $\mathcal{L}_{\text {Smoothness }}(\eta)$ | 0.0 |
| Table 2: Pretraining hyperparameters |  |


| Hyperparameter | Value |
| :--- | :---: |
| Number of seeds | 300 |
| Number of episodes | 100 |
| Maximum steps per episode | 1000 |
| Optimizer | Adam |
| Learning rate | 0.003 |
| Batch size | 16 |
| Discount factor, $\gamma$ | 0.9 |
| Starting exploration probability, $\epsilon_{0}$ | 1.0 |
| Final exploration probability, $\epsilon$ | 0.05 |
| Epsilon decay period | 2500 |
| Replay buffer size | 10000 |
| Initialization steps | 500 |
| Target network copy period | 50 |
| Table 3: DQN hyperparameters |  |

## H Implementation Details for DeepMind Control

We use the same RAD network architecture and code implementation as Laskin et al. (2020a), which we customized to add our Markov objective. We note that there was a discrepancy between the batch size in their code implementation (128) and what was reported in the original paper (512); we chose the former for our experiments.
The SAC (expert) results used the code implementation from Yarats \& Kostrikov (2020).
The DBC, DeepMDP, CPC, and SAC-AE results are from Zhang et al. (2021), except for Ball_in_Cup, which they did not include in their experimental evaluation. We ran their DBC code independently (with the same settings they used) to produce our Ball_in_Cup results.

## H. 1 Computing Resources

To build the graph in Section 7, we trained 4 agents on 6 domains, with 10 seeds each. Each training run (to 500,000 steps) takes between 24 and 36 hours (depending on the action repeat for that domain), on a machine with two Intel Xeon Gold 5122 vCPUs and shared access to one Nvidia 1080Ti GPU, for a total of approximately 7200 compute hours. We ran these jobs on a computing cluster with comparable hardware or better.

## H. 2 Markov Network Architecture

```
MarkovHead(
    InverseModel(
        (body): Sequential(
            (0): Linear(in_features=100, out_features=1024, bias=True)
            (1): ReLU()
            (2): Linear(in_features=1024, out_features=1024, bias=True)
            (3): ReLU()
        )
        (mean_linear): Linear(in_features=1024, out_features=action_dim, bias
            =True)
        (log_std_linear): Linear(in_features=1024, out_features=action_dim,
            bias=True)
    )
    ContrastiveModel(
            (model): Sequential(
            (0): Linear(in_features=100, out_features=1024, bias=True)
            (1): ReLU()
            (2): Linear(in_features=1024, out_features=1, bias=True)
        )
    )
    BCEWithLogitsLoss()
)
```


## H. 3 Hyperparameters

When tuning our algorithm, we left all RAD hyperparameters fixed except init_steps, which we increased to equal 10 episodes across all domains to provide adequate coverage for pretraining. We compensated for this change by adding (init_steps - 1 K ) catchup learning steps to ensure both methods had the same number of RL updates. This means our method is at a slight disadvantage, since RAD can begin learning from reward information after just 1 K steps, but our method must wait until after the first 10 episodes of uniform random exploration. Otherwise, we only considered changes to the Markov hyperparameters (see Table 5). We set the Markov learning rate equal to the RAD learning rate for each domain (and additionally considered 5e-5 for cheetah only). We tuned the $\mathcal{L}_{\text {Inv }}$ loss coefficient within $\{0.1,1.0,10.0,30.0\}$, and the $\mathcal{L}_{S m o o t h}$ loss coefficient within $\{0$, $10.0,30.0\}$. The other hyperparameters, including the network architecture, we did not change from their initial values.


## I Additional Representation Visualizations

Here we visualize abstraction learning progress for the $6 \times 6$ visual gridworld domain for six random seeds. Each figure below displays selected frames (progressing from left to right) of a different abstraction learning method (top to bottom): $\mathcal{L}_{\text {Markov }} ; \mathcal{L}_{\text {Inv }}$ only; $\mathcal{L}_{\text {Ratio }}$ only; autoencoder; pixel prediction. The networks are initialized identically for each random seed. Color denotes ground-truth $(x, y)$ position, which is not shown to the agent. These visualizations span 30,000 training steps (columns, left to right: after $1,100,200,700,3 \mathrm{~K}, 10 \mathrm{~K}$, and 30 K steps, respectively). In particular, the third column from the right shows the representations after 3000 steps, which we use for the results in the main text. We show additional learning curves for the final representations in Appendix J.


Figure 7

## J Gridworld Results for Increased Pretraining Time

Since some of the representations in Appendix I appeared not to have converged after just 3000 training steps, we investigated whether the subsequent learning performance would improve with more pretraining. We found that increasing the number of pretraining steps from 3000 to 30,000 improves the learning performance of $\phi_{\text {Ratio }}$ and $\phi_{\text {Autoenc }}$ and $\phi_{\text {PixelPred }}$ (see Figure 8 ), with the latter representation now matching the performance of $\phi_{M a r k o v}$.
It is perhaps unsurprising that the pixel prediction model eventually recovers the performance of the Markov abstraction, because the pixel prediction task is a valid way to ensure Markov abstract states. However, as we discuss in Sec. 3.2, the pixel prediction objective is misaligned with the basic goal of state abstraction, since it must effectively throw away no information. It is clear from Figures 7 and 8 that our method is able to reliably learn a Markov representation about ten times faster than pixel prediction, which reflects the fact that the latter is a fundamentally more challenging objective.


Figure 8: Mean episode reward for the visual gridworld navigation task, using representations that were pretrained for 3,000 steps (left) versus 30,000 steps (right). Increased pretraining time improves the performance of $\phi_{\text {Ratio }}$, $\phi_{\text {Autoenc }}$ and $\phi_{\text {Pixel Pred }}$. ( 300 seeds; 5-point moving average; shaded regions denote $95 \%$ confidence intervals.)

## K DeepMind Control Experiment with RBF-DQN

Recently, Asadi et al. (2021) showed how to use radial basis functions for value-function based RL in problems with continuous action spaces. When trained with ground-truth state information, RBF-DQN achieved state-of-the-art performance on several continuous control tasks; however, to our knowledge, the algorithm has not yet been used for image-based domains.
We trained RBF-DQN from stacked image inputs on "Finger, Spin," one of the tasks from Section 7, customizing the authors' PyTorch implementation to add our Markov training objective. We do not use any data augmentation or smoothness loss, and we skip the pretraining phase entirely; we simply add the Markov objective as an auxiliary task during RL. Here we again observe that adding the Markov objective improves learning performance over the visual baseline and approaches the performance of using ground-truth state information (see Figure 9).


Figure 9: Mean episode reward for RBF-DQN on "Finger, Spin" with Markov, visual, and expert features. Adding the Markov objective dramatically improves performance over the visual baseline. (Markov - 5 seeds; Visual - 6 seeds; Expert - 3 seeds; shaded regions denote $95 \%$ confidence intervals).

## L Investigating the Value of Smoothness in Markov Abstractions

Given a Markov abstraction $\phi$, we can always generate another abstraction $\phi^{\prime}$ by adding a procedural reshuffling of the $\phi$ representation's bits. Since the $\phi^{\prime}$ representation contains all the information that was in the original representation, $\phi^{\prime}$ is also a Markov abstraction. However, the new representation may be highly inefficient for learning.
To demonstrate this, we ran an experiment where we optionally relabeled the positions in the $6 \times 6$ gridworld domain, and trained two agents: one using the smooth, true $(x, y)$ positions, and one using the non-smooth, relabeled positions. Although both representations are Markov, and contain exactly the same information, we observe that the agent trained on the non-smooth positions performed significantly worse (see Figure 10).

The loss term $\mathcal{L}_{S m o o t h}$ used in Section 7 penalizes consecutive states that are more than $d_{0}$ away from each other, thereby encouraging representations to have a high degree of smoothness in addition to being Markov. This approach is similar to the temporal coherence loss proposed by Jonschkowski \& Brock (2015).


Figure 10: Mean episode reward for the $6 \times 6$ gridworld navigation task, comparing original $(x, y)$ position with rearranged $(x, y)$ position. ( 300 seeds; 5-point moving average; shaded regions denote $95 \%$ confidence intervals).

## M DeepMind Control Ablation Study

We ran an ablation study to evaluate which aspects of our training objective were most beneficial for the DeepMind Control domains. We considered the RAD implementation and its Markov variant from Section 7, as well as modifications to the Markov objective that removed either the pretraining phase or the smoothness loss, $\mathcal{L}_{\text {Smooth }}$ (see Figure 11).
Overall, the ablations perform slightly worse than the original Markov objective. Both ablations still have better performance than RAD on three of six domains, but are tied or slightly worse on the others. Interestingly, removing pretraining actually results in a slight improvement over Markov on Finger. Removing smoothness tends to degrade performance, although, for Cheetah, it leads to the fastest initial learning phase of any method.

We suspect the results on Cheetah are worse than the RAD baseline because the experiences used to learn the representations do not cover enough of the state space. Learning then slows down as the agent starts to see more states from outside of those used to train its current representation. As we point out in Section 3.6, it can be helpful to incorporate a more sophisticated exploration strategy to overcome these sorts of state coverage issues. Our approach is agnostic to the choice of exploration algorithm, and we see this as an important direction for future work.


Figure 11: Ablation results for DeepMind Control Suite. Each plot shows mean episode reward vs. environment steps. (Markov and RAD - 10 seeds; others - 6 seeds; 5-point moving average; shaded regions denote $90 \%$ confidence intervals).

