
Action-Sufficient State Representation Learning for Control with Structural Constraints

Biwei Huang^{*1} Chaochao Lu^{*23} Liu Leqi¹ José Miguel Hernández-Lobato² Clark Glymour¹
Bernhard Schölkopf³ Kun Zhang⁴¹

Abstract

Perceived signals in real-world scenarios are usually high-dimensional and noisy, and finding and using their representation that contains essential and sufficient information required by downstream decision-making tasks will help improve computational efficiency and generalization ability in the tasks. In this paper, we focus on partially observable environments and propose to learn a minimal set of state representations that capture sufficient information for decision-making, termed *Action-Sufficient state Representations* (ASRs). We build a generative environment model for the structural relationships among variables in the system and present a principled way to characterize ASRs based on structural constraints and the goal of maximizing cumulative reward in policy learning. We then develop a structured sequential Variational Auto-Encoder to estimate the environment model and extract ASRs. Our empirical results on CarRacing and VizDoom demonstrate a clear advantage of learning and using ASRs for policy learning. Moreover, the estimated environment model and ASRs allow learning behaviors from imagined outcomes in the compact latent space to improve sample efficiency.

1. Introduction

State-of-the-art reinforcement learning (RL) algorithms leveraging deep neural networks are usually data hungry and lack interpretability. For example, to attain expert-level performance on tasks such as chess or Atari games, deep RL systems usually require many orders of magnitude more

training data than human experts (Tsividis et al., 2017). One of the reasons is that our perceived signals in real-world scenarios, e.g., images, are usually high-dimensional and may contain much irrelevant information for decision-making of the task at hand. This makes it difficult and expensive for an agent to directly learn optimal policies from raw observational data. Fortunately, the underlying states that directly guide decision-making could be much lower-dimensional (Schölkopf, 2019; Bengio, 2019). One example is that when crossing the street, our decision on when to cross relies on the traffic lights. The useful state of traffic lights (e.g., its color) can be represented by a single binary variable, while the perceived image is high-dimensional. It is essential to extract and exploit such lower-dimensional states to improve the efficiency and interpretability of the decision-making process.

Recently, representation learning algorithms have been designed to learn abstract features from high-dimensional and noisy observations. Exploiting the abstract representations, instead of the raw data, has been shown to perform subsequent decision-making more efficiently (Lesort et al., 2018). Representative methods along this line include deep Kalman filters (Krishnan et al., 2015), deep variational Bayes filters (Karl et al., 2016), world models (Ha & Schmidhuber, 2018), PlaNet (Hafner et al., 2018), DeepMDP (Gelada et al., 2019), stochastic latent actor-critic (Lee et al., 2019), SimPLe (Kaiser et al., 2019), Bisimulation-based methods (Zhang et al., 2021), Dreamer (Hafner et al., 2019; 2020), and others (Srinivas et al., 2020; Shu et al., 2020). Moreover, if we can properly model and estimate the underlying transition dynamics, then we can perform model-based RL or planning, which can effectively reduce interactions with the environment (Ha & Schmidhuber, 2018; Hafner et al., 2018; 2019; 2020).

Despite the effectiveness of the above approaches to learning abstract features, current approaches usually fail to take into account whether the extracted state representations are sufficient and necessary for downstream policy learning. State representations that contain insufficient information may lead to sub-optimal policies, while those with redundant information may require more samples and more complex models for training. We address this problem by modeling the generative process and selection procedure induced

^{*}Equal contribution ¹Carnegie Mellon University ²University of Cambridge ³Max Planck Institute for Intelligent Systems, Tübingen ⁴Mohamed bin Zayed University of Artificial Intelligence. Correspondence to: Kun Zhang <kunz1@cmu.edu>.

by reward maximization; by considering a generative environment model involving observed states, state-transition dynamics, and rewards, and explicitly characterizing structural relationships among variables in the RL system, we propose a principled approach to learning minimal sufficient state representations. We show that only the state dimensions that have direct or indirect edges to the reward variable are essential and should be considered for decision making. Furthermore, they can be learned by maximizing their ability to predict the action, given that the cumulative reward is included in the prediction model, while at the same time achieving their minimality w.r.t. the mutual information with observations as well as their dimensionality. The contributions of this paper are summarized as follows:

- We construct a generative environment model, which includes the observation function, transition dynamics, and reward function, and explicitly characterizes structural relationships among variables in the RL system.
- We characterize a minimal sufficient set of state representations, termed Action-Sufficient state Representations (ASRs), for the downstream policy learning by making use of structural constraints and the goal of maximizing cumulative reward in policy learning.
- In light of the characterization, we develop Structured Sequential Variational Auto-Encoder (SS-VAE), which explicitly encodes structural relationships among variables, for reliable identification of ASRs.
- Accordingly, policy learning can be done separately from representation learning, and the policy function only relies on a set of low-dimensional state representations, which improve both model and sample efficiency. Moreover, the estimated environment model and ASRs allow learning behaviors from imagined outcomes in the compact latent space, which effectively reduce possibly risky explorations.

2. Environment Model with Structural Constraints

In order to characterize a set of minimal sufficient state representations for downstream policy learning, we first formulate a generative environment model in partially observable Markov decision process (POMDP), and then show how to explicitly embed structural constraints over variables in the RL system and leverage them.

Suppose we have sequences of observations $\{\langle o_t, a_t, r_t \rangle\}_{t=1}^T$, where $o_t \in \mathcal{O}$ denotes perceived signals at time t , such as high-dimensional images, with \mathcal{O} being the observation space, $a_t \in \mathcal{A}$ is the performed action with \mathcal{A} being the action space, and $r_t \in \mathcal{R}$ represents the reward variable with \mathcal{R} being the reward space. We denote the underlying states, which are latent, by $\vec{s}_t \in \mathcal{S}$, with \mathcal{S} being the state space. We describe the generating process of

the environment model as follows:

$$\begin{cases} o_t = f(\vec{s}_t, e_t), \\ r_t = g(\vec{s}_{t-1}, a_{t-1}, \epsilon_t), \\ \vec{s}_t = h(\vec{s}_{t-1}, a_{t-1}, \eta_t), \end{cases} \quad (1)$$

where f , g , and h represent the observation function, reward function, and transition dynamics, respectively, and e_t , ϵ_t , and η_t are corresponding independent and identically distributed (i.i.d.) random noises. The latent states \vec{s}_t form an MDP: given \vec{s}_{t-1} and a_{t-1} , \vec{s}_t are independent of states and actions before $t-1$. Moreover, the action a_{t-1} directly influences latent states \vec{s}_t , instead of perceived signals o_t , and the reward is determined by the latent states (and the action) as well. The perceived signals o_t are generated from the underlying states \vec{s}_t , contaminated by random noise e_t . We also consider noise ϵ_t in the reward function to capture unobserved factors that may affect the reward, as well as measurement noise.

It is commonplace that the action variable a_{t-1} may not influence every dimension of \vec{s}_t , and the reward r_t may not be influenced by every dimension of \vec{s}_{t-1} as well, and furthermore there are structural relationships among different dimensions of \vec{s}_t . Figure 1 gives an illustrative graphical representation, where $s_{3,t-1}$ influences $s_{2,t}$, a_{t-1} does not have an edge to $s_{3,t}$, and among the states, only $s_{2,t-1}$ and $s_{3,t-1}$ have edges to r_t . We use $R_t = \sum_{\tau=t}^{\infty} \gamma^{\tau-t} r_\tau$ to denote the discounted cumulative reward starting from time t , where $\gamma \in [0, 1]$ is the discounted factor that determines how much immediate rewards are favored over more distant rewards.

To reflect such constraints, we explicitly encode the graph structure over variables, including the structure over different dimensions of \vec{s} and the structures from a_{t-1} to \vec{s}_t , \vec{s}_{t-1} to r_t , and \vec{s}_t to o_t . Accordingly, we re-formulate (1) as follows:

$$\begin{cases} o_t = f(D_{\vec{s} \rightarrow o} \odot \vec{s}_t, e_t), \\ r_t = g(D_{\vec{s} \rightarrow r} \odot \vec{s}_{t-1}, D_{a \rightarrow r} \odot a_{t-1}, \epsilon_t), \\ s_{i,t} = h_i(D_{\vec{s}(\cdot, i)} \odot \vec{s}_{t-1}, D_{a \rightarrow \vec{s}(\cdot, i)} \odot a_{t-1}, \eta_{i,t}), \end{cases} \quad (2)$$

for $i = 1, \dots, d$, where $\vec{s}_t = (s_{1,t}, \dots, s_{d,t})^\top$, \odot denotes element-wise product, and $D_{(\cdot)}$ are binary matrices indicating the graph structure over variables. Specifically, $D_{\vec{s} \rightarrow o} \in \{0, 1\}^{d \times 1}$ represents the graph structure from d -dimensional \vec{s}_t to o_t , $D_{\vec{s} \rightarrow r} \in \{0, 1\}^{d \times 1}$ the structure from \vec{s}_{t-1} to the reward variable r_t , $D_{a \rightarrow r} \in \{0, 1\}$ the structure from the action variable a_{t-1} to the reward variable r_t , $D_{\vec{s}} \in \{0, 1\}^{d \times d}$ denotes the graph structure from d -dimensional \vec{s}_{t-1} to d -dimensional \vec{s}_t and $D_{\vec{s}(\cdot, i)}$ is its i -th column, and $D_{a \rightarrow \vec{s}} \in \{0, 1\}^{1 \times d}$ corresponds to the graph structure from a_{t-1} to \vec{s}_t with $D_{a \rightarrow \vec{s}(\cdot, i)}$ representing its i -th column. For example, $D_{\vec{s}(j, i)} = 0$ means that there is no edge from $s_{j,t-1}$ to $s_{i,t}$. Here, we assume that the environment model, as well as the structural constraints, is invariant across time instance t .




Figure 1: A graphical illustration of the generative environment model. Grey nodes denote observed variables and white nodes represent unobserved variables. Here, a_{t-1} does not have an edge to $s_{3,t}$, and only $s_{2,t-1}$ and $s_{3,t-1}$ have edges to r_t , and moreover, we take into account the structural relationships among different dimensions of latent states \vec{s}_t . The solid lines represent causal relations, while the dashed lines represent predictions of a_t from \vec{s}_t , which is not causal, and moreover, dashed lines mean that the relations may not exist and they may differ under different policies.

2.1. Minimal Sufficient State Representations

Given observational sequences $\{(o_t, a_t, r_t)\}_{t=1}^T$, we aim to learn minimal sufficient state representations for the downstream policy learning. In the following, we first characterize the state dimensions that are indispensable for policy learning, when the environment model, including structural relationships, is given. Then we provide criteria to achieve sufficiency and minimality of the estimated state representations, when only $\{(o_t, a_t, r_t)\}_{t=1}^T$, but not the environment model, is given.

Finding minimal sufficient state dimensions with a given environment model. RL agents learn to choose appropriate actions according to the current state vector \vec{s}_t to maximize the future cumulative reward, in which some dimensions may be redundant for policy learning. Then how can we identify a minimal subset of state dimensions that are sufficient to choose optimal actions? Below, we first give the definition of *Action-Sufficient state Representations (ASRs)* according to the graph structure. We further show in Proposition 1 that ASRs are minimal sufficient for policy learning, and they can be characterized by leveraging the (conditional) independence/dependence relations among the quantities, under the Markov condition and faithfulness assumption (Pearl, 2000; Spirtes et al., 1993).

Definition 1 (Action-Sufficient State Representations (ASRs)). *Given the graphical representation corresponding to the environment model, such as the representation in*

Figure 1, we define recursively ASRs that affect the future reward as: (1) $s_{i,t} \in \vec{s}_t^{\text{ASR}}$ has an edge to the reward in the next time-step r_{t+1} , or (2) $s_{i,t} \in \vec{s}_t^{\text{ASR}}$ has an edge to another state dimension in the next time-step $s_{j,t+1}$, such that the same component at time t is in ASRs, i.e., $s_{j,t} \in \vec{s}_t^{\text{ASR}}$.

Proposition 1. *Under the assumption that the graphical representation, corresponding to the environment model, is Markov and faithful to the measured data, $\vec{s}_t^{\text{ASR}} \subseteq \vec{s}_t$ are a minimal subset of state dimensions that are sufficient for policy learning, and $s_{i,t} \in \vec{s}_t^{\text{ASR}}$ if and only if $s_{i,t} \not\perp\!\!\!\perp R_{t+1} | a_{t-1:t}, \vec{s}_{t-1}^{\text{ASR}}$.*

Proposition 1 can be shown based on the global Markov condition and the faithfulness assumption, which connects d-separation¹ to conditional independence/dependence relations. A proof is given in Appendix. According to the above proposition, it is easy to see that for the graph given in Figure 1, we have $\vec{s}_t^{\text{ASR}} = (s_{2,t}, s_{3,t})^\top$. That is, we only need $(s_{2,t}, s_{3,t})^\top$, instead of \vec{s}_t , for the downstream policy learning.

Minimal sufficient state representation learning from observed sequences. In practice, we usually do not have access to the latent states or the environment model, but instead only the observed sequences $\{(o_t, a_t, r_t)\}_{t=1}^T$. Then how can we learn the ASRs from the raw high-dimensional inputs such as images? We denote by \vec{s}_t the estimated whole latent state representations and $\tilde{\vec{s}}_t^{\text{ASR}} \subseteq \tilde{\vec{s}}_t$ the estimated minimal sufficient state representations for policy learning.

As discussed above, ASRs and R_{t+1} are dependent given $a_{t-1:t}$ and $\tilde{\vec{s}}_{t-1}^{\text{ASR}}$, while other state dimensions are independent of R_{t+1} , so we can learn the ASRs by maximizing²

$$I(\tilde{\vec{s}}_t^{\text{ASR}}; R_{t+1} | a_{t-1:t}, \tilde{\vec{s}}_{t-1}^{\text{ASR}}) - I(\tilde{\vec{s}}_t^{\text{C}}; R_{t+1} | a_{t-1:t}, \tilde{\vec{s}}_{t-1}^{\text{ASR}}), \quad (3)$$

where $\tilde{\vec{s}}_t^{\text{C}} = \tilde{\vec{s}}_t \setminus \tilde{\vec{s}}_t^{\text{ASR}}$, and $I(\cdot)$ denotes mutual information. Such regularization is used to achieve minimal sufficient state representations for policy learning; that is, only $\tilde{\vec{s}}_{t-1}^{\text{ASR}}$ are useful for policy learning, while $\tilde{\vec{s}}_{t-1}^{\text{C}}$ are not. Furthermore, the mutual information can be represented as a form of conditional entropy, e.g., $I(\tilde{\vec{s}}_t^{\text{ASR}}; R_{t+1} | a_{t-1:t}, \tilde{\vec{s}}_{t-1}^{\text{ASR}}) = H(\tilde{\vec{s}}_t^{\text{ASR}} | a_{t-1:t}, \tilde{\vec{s}}_{t-1}^{\text{ASR}}) - H(\tilde{\vec{s}}_t^{\text{ASR}} | R_{t+1}, a_{t-1:t}, \tilde{\vec{s}}_{t-1}^{\text{ASR}})$, where $H(\cdot)$ denotes the con-

¹A path p is said to be d-separated by a set of nodes Z if and only if (1) p contains a chain $i \rightarrow m \rightarrow j$ or a fork $i \leftarrow m \rightarrow j$ such that the middle node m is in Z , or (2) p contains a collider $i \rightarrow m \leftarrow j$ such that the middle node m is not in Z and such that no descendant of m is in Z .

²Note that, when observational sequences are generated from a random policy (i.e., $\vec{s}_t \perp\!\!\!\perp a_t$), learning the ASRs can be also implemented in a simpler way, which is given in Appendix.

ditional entropy, with

$$\begin{aligned} & H(\tilde{s}^{\text{ASR}}|a_{t-1:t}, \tilde{s}_{t-1}^{\text{ASR}}) \\ &= -\mathbb{E}_{q_\phi} \mathbb{E}_{p_{\alpha_1}} \{\log p_{\alpha_1}(\tilde{s}^{\text{ASR}}|a_{t-1:t}, \tilde{s}_{t-1}^{\text{ASR}})\} \\ &= -\mathbb{E}_{q_\phi} \mathbb{E}_{p_{\alpha_1}} \{\log p_{\alpha_1}(\tilde{s}^{\text{ASR}}|a_{t-1:t}, \tilde{D}^{\text{ASR}} \odot \tilde{s}_{t-1}^{\text{ASR}})\}, \end{aligned}$$

and

$$\begin{aligned} & H(\tilde{s}^{\text{ASR}}|R_{t+1}, a_{t-1:t}, \tilde{s}_{t-1}^{\text{ASR}}) \\ &= -\mathbb{E}_{q_\phi} \mathbb{E}_{p_{\alpha_2}} \{\log p_{\alpha_2}(\tilde{s}^{\text{ASR}}|R_{t+1}, a_{t-1:t}, \tilde{D}^{\text{ASR}} \odot \tilde{s}_{t-1}^{\text{ASR}})\}, \end{aligned}$$

where p_{α_i} , for $i = 1, 2$, denotes the probabilistic predictive model of \tilde{s}^{ASR} with parameters α_i , and $q_\phi(\tilde{s}_t|\tilde{s}_{t-1}, \mathbf{y}_{1:t}, a_{1:t-1})$ is the probabilistic inference model of \tilde{s}_t with parameters ϕ and $\mathbf{y}_t = (o_t^\top, r_t^\top)$, and $\tilde{D}^{\text{ASR}} \in \{0, 1\}^{\tilde{d} \times 1}$ is a binary vector indicating which dimensions of \tilde{s}_t are in \tilde{s}_t^{ASR} , so $\tilde{D}^{\text{ASR}} \odot \tilde{s}_t$ gives ASRs \tilde{s}_t^{ASR} . Similarly, we can also represent $I(\tilde{s}^{\text{C}}; R_{t+1}|a_{t-1:t}, \tilde{s}_{t-1}^{\text{ASR}})$ in such a way.

Although with the above regularization, we can achieve ASRs theoretically, in practice, we further add another regularization to achieve minimality of the representation by minimizing conditional mutual information between observed high-dimensional signals \mathbf{y}_t and the ASR \tilde{s}_t^{ASR} at time t given data at previous time instances, similar to that in information bottleneck (Tishby et al., 1999), and meanwhile minimizing the dimensionality of ASRs with sparsity constraints:

$$\lambda_1 \sum_{t=2}^T I(\mathbf{y}_t; \tilde{s}_t^{\text{ASR}}|\mathbf{y}_{1:t-1}, a_{1:t-1}, \tilde{s}_{t-1}^{\text{ASR}}) + \lambda_2 \|\tilde{D}^{\text{ASR}}\|_1,$$

where the conditional mutual information can be upper bound by a KL-divergence:

$$I(\mathbf{y}_t; \tilde{s}_t^{\text{ASR}}|\mathbf{y}_{1:t-1}, a_{1:t-1}, \tilde{s}_{t-1}^{\text{ASR}}) \leq \mathbb{E}\{\text{KL}(q_{\phi'}\|p_\gamma)\}, \quad (4)$$

with $q_{\phi'} \equiv q_{\phi'}(\tilde{s}_t^{\text{ASR}}|\tilde{s}_{t-1}^{\text{ASR}}, \mathbf{y}_{1:t}, a_{1:t-1})$ and $p_\gamma \equiv p_\gamma(\tilde{s}_t^{\text{ASR}}|\tilde{s}_{t-1}^{\text{ASR}}, a_{t-1}; D_{\tilde{s}}, D_{a \rightarrow \tilde{s}})$ being the transition dynamics of \tilde{s}_t with parameters γ , and the expectation is over $p(\tilde{s}_{t-1}, \mathbf{y}_{1:t}, a_{1:t-1})$.

Furthermore, Proposition 1 shows that given the (estimated) environment model, only those state dimensions that have a directed path to the reward variable are the ASRs. In our learning procedure, we also take into account the relationship between the learned states \tilde{s}_t and the reward, and leverage such structural constraints for learning the ASRs. Denote by $\tilde{D}^{\text{ASR}} \in \{0, 1\}^{\tilde{d} \times 1}$ a binary vector indicating whether the corresponding state dimension in \tilde{s}_t has a directed path to the reward variable. Consequently, we enforce the similarity between \tilde{D}^{ASR} and \tilde{D}^{ASR} by adding an L_1 norm on $\tilde{D}^{\text{ASR}} - \tilde{D}^{\text{ASR}}$. Therefore, the ASRs can be

learned by maximizing the following function:

$$\begin{aligned} \mathcal{L}^{\text{min \& suff}} = & \lambda_3 \underbrace{\sum_{t=1}^T \{I(\tilde{s}^{\text{ASR}}; R_{t+1}|a_{t-1:t}, \tilde{s}_{t-1}^{\text{ASR}}) - I(\tilde{s}^{\text{C}}; R_{t+1}|a_{t-1:t}, \tilde{s}_{t-1}^{\text{ASR}})\}}_{\text{Sufficiency \& Minimality}} \\ & - \lambda_4 \|\tilde{D}^{\text{ASR}} - \tilde{D}^{\text{ASR}}\|_1 - \lambda_1 \underbrace{\sum_{t=1}^T \mathbb{E}\{\text{KL}(q_\phi\|p_\gamma)\}}_{\text{Further restrictions of minimality}} - \lambda_2 \|\tilde{D}^{\text{ASR}}\|_1, \end{aligned} \quad (5)$$

where λ 's are regularization terms, and note that \tilde{D}^{ASR} can be directly derived from the estimated structural matrices $D_{a \rightarrow r}$ and $D_{\tilde{s}(\cdot, i)}$. The constraint in Eq. 5 provides a principled way to achieve minimal sufficient state representations. Notice that it is just part of the objective function to maximize, and it will be involved in the complete objective function in Fig. 2 to learn the whole environment model.

Remarks. By explicitly involving structural constraints, we achieve minimal sufficient state representations from the view of generative process underlying the RL problem and the selection procedure induced by reward maximization, which enjoys the following advantages. 1) The structural information provides an interpretable and intuitive picture of the generating process. 2) Accordingly, it also provides an interpretable and intuitive way to characterize a minimal sufficient set of state representations for policy learning, which removes unrelated information. 3) There is no information loss when representation learning and policy learning are done separately, which is computationally more efficient. 4) The generative environment model is fixed, independent of the behavior policy that is performed. Furthermore, based on the estimated environment model and ASRs, it is flexible to use a wide range of policy learning methods, and one can also perform model-based RL, which effectively reduce possibly risky explorations.

3. Structured Sequential VAE for the Estimation of ASRs

In this section, we give estimation procedures for the environment model and ASRs, as well as the identifiability guarantee in linear cases.

Identifiability in Linear-Gaussian Cases. Below, we first show the identifiability guarantee in the linear case, as a special case of Eq. (2):

$$\begin{cases} o_t = D_{\tilde{s} \rightarrow o}^\top \tilde{s}_t + e_t, \\ r_{t+1} = D_{\tilde{s} \rightarrow r}^\top \tilde{s}_t + D_{a \rightarrow r}^\top a_t + \epsilon_{t+1}, \\ \tilde{s}_t = D_{\tilde{s}}^\top \tilde{s}_{t-1} + D_{a \rightarrow \tilde{s}}^\top a_{t-1} + \eta_t. \end{cases} \quad (6)$$

In the linear case, $D_{\tilde{s} \rightarrow o}$, $D_{\tilde{s} \rightarrow r}$, $D_{a \rightarrow r}$, $D_{\tilde{s}}$, and $D_{a \rightarrow \tilde{s}}$ are linear coefficients, indicating corresponding graph structures and also the strength. Denote the covariance matrices of e_t and ϵ_t by Σ_e and Σ_ϵ , respectively. Further let

$\ddot{D}_{\tilde{s}^*o} := (D_{\tilde{s}^*o}^\top, D_{\tilde{s}^*r}^\top)^\top$. The following proposition shows that the environment model in the linear case is identifiable up to some orthogonal transformation on certain coefficient matrices from observed data $\{o_t, a_t, r_t\}_{t=1}^T$.

Proposition 2 (Identifiability). *Suppose the perceived signal o_t , the reward r_t , and the latent states \tilde{s}_t follow a linear environment model. If assumptions A1~A4 (given in Appendix D) hold and with the second-order statistics of the observed data $\{o_t, a_t, r_t\}_{t=1}^T$, the noise variances Σ_e and Σ_ϵ , D_{a^*r} , $\ddot{D}_{\tilde{s}^*o}^\top D_{\tilde{s}^*s}^k D_{a^*s}^\top$ (with $k \geq 0$), and $\ddot{D}_{\tilde{s}^*o}^\top \ddot{D}_{\tilde{s}^*o}$ are uniquely identified.*

This proposition shows that in the linear case, with the second-order statistics of the observed data, we can identify the parameters up to orthogonal transformations. In particular, suppose the linear environment model with parameters $(D_{\tilde{s}^*o}, D_{\tilde{s}^*r}, D_{a^*r}, D_{\tilde{s}^*s}, D_{a^*s}, \Sigma_e, \Sigma_\epsilon)$ and that with $(\tilde{D}_{\tilde{s}^*o}, \tilde{D}_{\tilde{s}^*r}, \tilde{D}_{a^*r}, \tilde{D}_{\tilde{s}^*s}, \tilde{D}_{a^*s}, \tilde{\Sigma}_e, \tilde{\Sigma}_\epsilon)$ are observationally equivalent. Then we have $\tilde{D}_{\tilde{s}^*o} = U \ddot{D}_{\tilde{s}^*o}$, $\tilde{D}_{a^*r} = D_{a^*r}$, $\tilde{D}_{\tilde{s}^*s} = U^\top D_{\tilde{s}^*s} U$, $\tilde{D}_{a^*s} = D_{a^*s} U$, $\tilde{\Sigma}_e = \Sigma_e$, and $\tilde{\Sigma}_\epsilon = \Sigma_\epsilon$, where U is an orthogonal matrix.

General Nonlinear Cases. To handle general nonlinear cases with the generative process given in Eq. (2), we develop a Structured Sequential VAE (SS-VAE) to learn the model (including structural constraints) and infer latent state representations \tilde{s}_t and ASRs \tilde{s}_t^{ASR} , with the input $\{o_t, a_t, r_t\}_{t=1}^T$. Specifically, the latent state dimensions are organized with structures, captured by $D_{\tilde{s}^*}$, to achieve conditional independence. The structural relationships over perceived signals, latent states, the action variable, and the reward variable are also embedded as free parameters (i.e., $D_{\tilde{s}^*o}, D_{\tilde{s}^*r}, D_{a^*r}, D_{a^*s}$) into SS-VAE. Moreover, we aim to learn state representations \tilde{s}_t and ASRs \tilde{s}_t^{ASR} that satisfy the following properties: (i) \tilde{s}_t should capture sufficient information of observations o_t , r_t , and a_t , that is, it should be enough to enable reconstruction. (ii) The state representations should allow for accurate predictions of the next state and also the next observation. (iii) The transition dynamics should follow an MDP. (iv) \tilde{s}_t^{ASR} are minimal sufficient state representations for the downstream policy learning.

Let $\mathbf{y}_{1:T} = \{(o_t^\top, r_t^\top)^\top\}_{t=1}^T$. To achieve the above properties, we maximize the objective function shown in Fig. 2, which contains the reconstruction error at each time instance, the one-step prediction error of observations, the KL divergence to constrain the latent space, and moreover, the MDP restrictions on transition dynamics, the sufficiency and minimality guarantee of state representations for policy learning, as well as sparsity constraints on the graph structure. We denote by p_θ the generative model with parameters θ and structural constraints $D_{(\cdot)}$, q_ϕ the inference model with parameters ϕ , p_γ the transition dynamics, and p_{α_i} the predictive model of ASRs. Each factor in p_γ , q_ϕ , and p_{α_i} is modeled with a mixture of Gaussians (MoGs), to

approximate a wide class of continuous distributions.

Below are the details of each component in the above objective function:

- **Reconstruction and prediction components:** These two parts are commonly used in sequential VAE. They aim to minimize the reconstruction error and prediction error of the perceived signal o_t and the reward r_t .
- **Transition component:** To achieve the property that state representations satisfy an MDP, we explicitly model the transition dynamics: $\log p_\gamma(\tilde{s}_t | \tilde{s}_{t-1}, a_{t-1}; D_{\tilde{s}^*}, D_{a^*s})$. In particular, $\tilde{s}_t | \tilde{s}_{t-1}$ is modelled with a mixture of Gaussians: $\sum_{k=1}^K \pi_k \mathcal{N}(\boldsymbol{\mu}_k(\tilde{s}_{t-1}, a_{t-1}), \Sigma_k(\tilde{s}_{t-1}, a_{t-1}))$, where K is the number of mixtures, $\boldsymbol{\mu}_k(\cdot)$ and $\Sigma_k(\cdot)$ are given by multi-layer perceptrons (MLP) with inputs \tilde{s}_{t-1} and a_{t-1} , parameters γ , and structural constraints $D_{\tilde{s}^*}$ and D_{a^*s} . This explicit constraint on state dynamics is essential for establishing a Markov chain in latent space and for learning a representation for long-term predictions. Note that unlike in traditional VAE (Kingma & Welling, 2013), we do not assume that different dimensions in \tilde{s}_t are marginally independent, but model their structural relationships explicitly to achieve conditional independence.
- **KL-divergence constraint:** The KL divergence is used to constrain the state space with multiple purposes: (1) It is used in the lower bound of $\log P(\mathbf{y}_{1:T})$ to achieve conditional disentanglement between $q_\phi(\tilde{s}_{i,t} | \cdot)$ and $q_\phi(\tilde{s}_{j,t} | \cdot)$ for $i \neq j$, (2) and also to achieve further restrictions of minimality of ASRs.
- **Sufficiency & minimality constraints:** We achieve minimal sufficient state representations for the downstream policy learning by leveraging the conditional mutual information between \tilde{s}_t^{ASR} and R_{t+1} , and structural constraints. For details, please refer to Section 2.1.
- **Sparsity constraints:** According to the edge-minimality property (Zhang & Spirtes, 2011), we additionally put sparsity constraints on structural matrices to achieve better identifiability. In particular, we use L_1 norm of the structural matrices as regularizers in the objective function to achieve sparsity of the solution.

Figure 3 gives the diagram of the neural network architecture in model training. We use SS-VAE to learn the environment model and ASRs. Specifically, the encoder, which is used to learn the inference model $q_\phi(\tilde{s}_t | \tilde{s}_{t-1}, \mathbf{y}_{1:t}, a_{1:t-1})$, includes a Long Short-Term Memory (LSTM (Hochreiter & Schmidhuber, 1997)) to encode the sequential information with output h_t and a Mixture Density Network (MDN (Bishop, 1994)) to output the parameters of MoGs. At each time instance, the input $\langle o_{t+1}, r_{t+1}, a_t \rangle$ is projected to the encoder and a sample of \tilde{s}_{t+1} is inferred from q_ϕ as output. The generated sample further acts as an input to the decoder, together with a_{t+1} and structural matrices $D_{\tilde{s}^*o}$, $D_{\tilde{s}^*r}$, and

$$\begin{aligned}
 & \mathcal{L}(\mathbf{y}_{1:T}; (\theta, \phi, \gamma, \alpha, D_{(\cdot)})) \\
 &= \sum_{t=1}^{T-2} \mathbb{E}_{q_\phi} \left\{ \underbrace{\log p_\theta(o_t | \tilde{s}_t; D_{\tilde{s} \rightarrow o}) + \log p_\theta(r_{t+1} | \tilde{s}_t, a_t; D_{\tilde{s} \rightarrow r}, D_{a \rightarrow r})}_{\text{Reconstruction}} + \underbrace{\log p_\theta(o_{t+1} | \tilde{s}_t) + \log p_\theta(r_{t+2} | \tilde{s}_t, a_{t+1})}_{\text{Prediction}} \right\} \\
 &+ \lambda_3 \sum_{t=1}^T \underbrace{\left\{ I(\tilde{s}^{\text{ASR}}; R_{t+1} | a_{t-1:t}, \tilde{s}_{t-1}^{\text{ASR}}) - I(\tilde{s}^{\text{C}}, R_{t+1} | a_{t-1:t}, \tilde{s}_{t-1}^{\text{ASR}}) \right\}}_{\text{Sufficiency \& Minimality}} \\
 &- \lambda_1 \sum_{t=1}^T \mathbb{E} \left\{ \text{KL}(q_{\phi'}(\tilde{s}_t^{\text{ASR}} | \tilde{s}_{t-1}^{\text{ASR}}, \mathbf{y}_{1:t}, a_{1:t-1}) \| p_\gamma(\tilde{s}_t^{\text{ASR}} | \tilde{s}_{t-1}^{\text{ASR}}, a_{t-1}; D_{\tilde{s}}, D_{a \rightarrow \tilde{s}})) \right\} - \lambda_2 \|\tilde{D}_{\text{ASR}}\|_1 \\
 &\underbrace{- (\lambda_5 \|D_{\tilde{s} \rightarrow o}\|_1 + \lambda_6 \|D_{\tilde{s} \rightarrow r}\|_1 + \lambda_7 \|D_{\tilde{s}}\|_1 + \lambda_8 \|D_{a \rightarrow \tilde{s}}\|_1 + \lambda_4 \|\tilde{D}_{\text{ASR}} - \hat{D}_{\text{ASR}}\|_1)}_{\text{Conditional disentanglement \& Further restrictions of minimality}}, \\
 &\underbrace{\hspace{10em}}_{\text{Sparsity}},
 \end{aligned}$$

Figure 2: Our objective function.

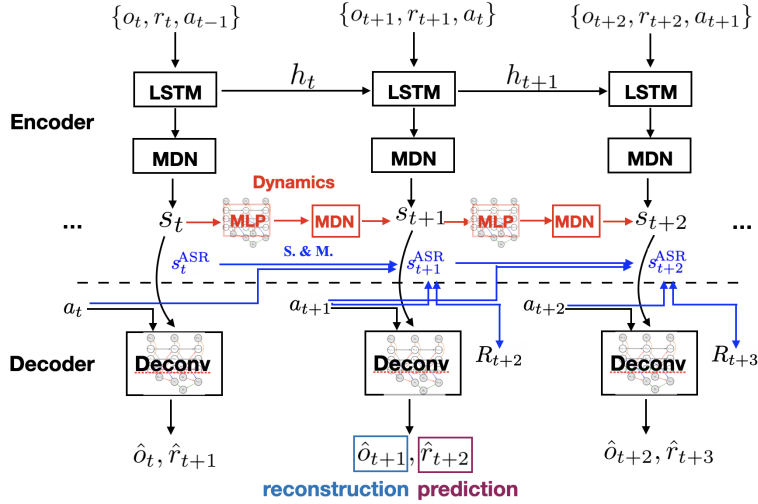


Figure 3: Diagram of neural network architecture to learn state representations. The corresponding structural constraints are involved in “Deconv” and “MLP”, and “S.&M.” represents the regularization part for minimal sufficient state representation learning.

$D_{a \rightarrow r}$. Then the decoder outputs \hat{o}_{t+1} and \hat{r}_{t+2} . Moreover, the state dynamics which satisfies a Markov process and is embedded with structural constraints $D_{\tilde{s}}$ and $D_{a \rightarrow \tilde{s}}$, is modeled with an MLP and MDN, marked with red in Figure 3. The part for minimal sufficient representations (denoted by S.&M.) uses MLP and is marked with blue. During training, we approximate the expectation in \mathcal{L} by sampling and then jointly learn all parameters by maximizing \mathcal{L} using stochastic gradient descent.

4. Policy Learning with ASRs

After estimating the generative environment model, we are ready to learn the optimal policy, where the policy function only depends on low-dimensional ASRs, instead of high-dimensional images. The entire procedure roughly contains the following three parts: (1) data collection with a random or sub-optimal policy, (2) environment model estimation

(with details in Section 3), and (3) policy learning with ASRs. Notably, the generative environment model is fixed, regardless of the behavior policy that is used to generate the data, and after learning the environment model, as well as the inference model for ASRs, our framework is flexible for both model-free and model-based policy learning.

Model-Free Policy Learning. For model-free policy learning, we make use of the learned environment model to infer ASRs \tilde{s}_t^{ASR} from past observed sequences $\{o_{\leq t}, r_{\leq t}, a_{\leq t-1}\}$ and then predict the action with the estimated low-dimensional ASRs. Our method is flexible to use a wide range of model-free methods; for example, one may use deep Q-learning for discrete actions (Mnih et al., 2015) and deep deterministic policy gradient (DDPG) for continuous actions (Lillicrap et al., 2015). Algorithm 1 in Appendix G gives the detailed procedure of model-free policy learning with ASRs in partially observable environ-

ments.

Model-Based Policy Learning. The downside of model-free RL algorithms is that they are usually data hungry, requiring very large amounts of interactions. On the contrary, model-based RL algorithms enjoy much better sample efficiency. Hence, we make use of the learned generative environment model, including the transition dynamics, observation function, and reward function, for model-based policy optimization. Based on the generative environment model, one can learn behaviors from imagined outcomes to increase sample-efficiency and mitigate heavy and possibly risky interactions with the environment. We present the procedure of the classic Dyna algorithm (Sutton, 1990; Sutton & Barto, 2018) with ASRs in Algorithm 2 in Appendix G.

5. Experiments

To evaluate the proposed approach, we conducted experiments on both CarRacing environment (Klimov, 2016) with an illustration in Figure 4 and VizDoom (Kempka et al., 2016) environment with an illustration in Figure 5, following the setup in the world model (Ha & Schmidhuber, 2018) for a fair comparison. It is known that CarRacing is very challenging—the recent world model (Ha & Schmidhuber, 2018) is the first known solution to achieve the score required to solve the task. Without stated otherwise, all results were averaged across five random seeds, with standard deviation shown in the shaded area.

5.1. CarRacing Experiment

CarRacing is a continuous control task with three continuous actions: steering left/right, acceleration, and brake. Reward is -0.1 every frame and $+1000/N$ for every track tile visited, where N is the total number of tiles in track. It is obvious that the CarRacing environment is partially observable: by just looking at the current frame, although we can tell the position of the car, we know neither its direction nor velocity that are essential for controlling the car. For a fair comparison, we followed a similar setting as in Ha & Schmidhuber (2018). Specifically, we collected a dataset of $10k$ random rollouts of the environment, and each runs with a random policy until failure. The dimensionality of latent states \tilde{s}_t was set to $\tilde{d} = 32$, determined by hyperparameter tuning.

Analysis of ASRs. To demonstrate the structures over observed frames, latent states, actions, and rewards, we visualized the learned $D_{\tilde{s} \rightarrow o}$, $D_{\tilde{s} \rightarrow r}$, $D_{\tilde{s}}$, and $D_{a \rightarrow \tilde{s}}$, as shown in Figure 6. Intuitively, we can see that $D_{\tilde{s} \rightarrow r}$ and $D_{a \rightarrow \tilde{s}}$ have many values close to zero, meaning that the reward is only influenced by a small number of state dimensions, and not many state dimensions are influenced by the action. Furthermore, from $D_{\tilde{s}}$, we found that there are influences from $\tilde{s}_{i,t}$ to $\tilde{s}_{i,t+1}$ (diagonal values) for most state dimensions, which




Figure 4: An illustration of Car Racing environment.




Figure 5: An illustration of VizDoom take cover scenario.

is reasonable because we want to learn an MDP over the underlying states, while the connections across states (off-diagonal values) are much sparser. Compared to the original 32-dim latent states, ASRs have only 21 dimensions. Below, we empirically showed that the low-dimensional ASRs significantly improve the policy learning performance in terms of both efficiency and efficacy.

Comparison Between Model-Free and Model-Based ASRs. We applied both model-free (DDPG) (Lillicrap et al., 2015) and model-based (Dyna and Prioritized Sweeping) algorithms (Sutton, 1990) to ASRs (with 21-dims). As shown in Figure 7, interestingly, by taking advantage of the learned generative model, model-based ASRs is superior to model-free ASRs at a faster rate, which demonstrates the effectiveness of the learned model. It also shows that with the estimated environment model and ASRs, we can learn behaviors from imagined outcomes to improve sample-efficiency.

Comparison with VRL, SLAC, PlaNet, DBC, and Dreamer. We also compared the proposed framework of policy learning with ASRs (with 21-dims) with 1) the same learning strategy but with vanilla representation learning (VRL, implemented without the components for minimal sufficient state representations as in Eq. (5)), 2) SLAC (Lee et al., 2019), 3) PlaNet (Hafner et al., 2018), 4) DBC (Zhang et al., 2021), and 5) Dreamer (Hafner et al., 2019). For a fair comparison, the latent dimensions of VRL, PlaNet, SLAC, DBC and Dreamer are set to 21 as well, and we require all of them to have the model capacity similar to ours (i.e., similar model architectures). From Figure 7, we can see that our methods, both model-free and model-based, obviously outperform others. It is worth noting that the huge performance difference between ASRs and VRL shows that the components for minimal sufficient state representations play a pivotal role in our objective.

Comparison with World Models. In light of the fact that world models (Ha & Schmidhuber, 2018) achieved good performance in CarRacing, we further compared our method (with 21-dim ASRs) with the world model. For a fair comparison, following Ha & Schmidhuber (2018), we also used the Covariance-Matrix Adaptation Evolution




Figure 6: Visualization of estimated structural matrices $D_{\bar{s} \rightarrow o}$, $D_{\bar{s} \rightarrow r}$, $D_{a \rightarrow \bar{s}}$, and $D_{\bar{s}}$ in Car Racing.




Figure 7: Cumulative rewards of model-based ASRs, model-free ASRs, VRL, SLAC, PlaNet, Dreamer and DBC evaluated on CarRacing.




Figure 8: Fitness Value of ASRs compared to world models evaluated on Car-Racing, including mean score, max score, and the best average score.




Figure 9: Ablation study of latent dynamics prediction (LDP) evaluated on Car Racing with model-free ASR.

Strategy (CMA-ES) (Hansen, 2016) with a population of 64 agents to optimize the parameters of the controller. In addition, following a similar setting as in Ha & Schmidhuber (2018) (where the agent’s fitness value is defined as the average cumulative reward of the 16 random rollouts), we show the fitness values of the best performer (max) and the population (mean) at each generation (Figure 8). We also took the best performing agent at the end of every 25 generations and tested it over 1024 random rollout scenarios to record the average (best avg score). It is obvious that our method (denoted by *ASR-**) has a more efficient and also efficacy training process. The best average score of ASRs is 65 higher than that of world models.

Comparison with Dreamer and DBC with Background Distraction. We further compared ASRs (with 21-dims) with Dreamer and DBC when there are natural video distractors in CarRacing; we chose Dreamer and DBC, because their performance are relatively better than other comparisons when there are no distractors. Specifically, we followed Zhang et al. (2021) to incorporate natural video from the Kinetics dataset (Kay et al., 2017) as background in CarRacing. Similarly, for a fair comparison, we require all of them to have the same latent dimensions and have the similar model capacity. As shown in Table 1, we can see that our method outperforms both Dreamer and DBC.

Ablation Study. We further performed ablation studies on latent dynamics prediction; that is, we compared with the case when the transition dynamics in Fig. 2 is not explicitly

modeled, but is replaced with a standard normal distribution. Figure 9 shows that by explicitly modelling the transition dynamics (denoted by *with LDP*), the cumulative reward has an obvious improvement over the one without modelling the transition dynamics (denoted by *without LDP*).

5.2. VizDoom Experiment

We also applied the proposed method to VizDoom *take over* scenario (Kempka et al., 2016), which is a discrete control problem with two actions: move left and move right. Reward is +1 at each time step while alive, and the cumulative reward is defined to be the number of time steps the agent manages to stay alive during an episode.

Considering that in the *take over* scenario the action space is discrete, we applied the widely used DQN (Mnih et al., 2013) on ASRs for policy learning. In addition to the comparisons with VRL (as in CarRacing) and DQN on raw observations, we further compared with another common approach to POMDPs: DRQN (Hausknecht & Stone, 2015). As shown in Figure 10, DQN on ASRs achieve a much better performance than all other comparisons, and in particular, DQN on ASRs outperforms DRQN on observations by around 400 on average in terms of cumulative reward. Similarly, we applied model-based (Dyna) algorithms (Sutton, 1990) to ASRs (with 21-dims). As shown in Figure 10, we can draw the same conclusion that by taking advantage of the learned generative model, model-based ASRs is superior to model-free ASRs at a faster rate. We also applied ASRs to world models, where Figure 11 shows that our method with

Model	Cumulative Rewards
Dreamer	621±124.5
DBC	803±112.5
Model-free ASRs	938±87.2
Model-based ASRs	954±98.6

Table 1: Comparisons with Dreamer and DBC in CarRacing with natural video distractors, after 2000 training episodes, with standard error.




Figure 10: Comparing ASRs and SOTA methods evaluated on VizDoom.




Figure 11: Fitness value of ASRs (with CMA-ES) compared to world models evaluated on VizDoom.

ASRs (denoted by *ASR*-*) achieves a better performance.

6. Related Work

In the past few years, a number of approaches have been proposed to learn low-dimensional Markovian representations, which capture the variation in the environment generated by the agent’s actions, without direct supervision (Lesort et al., 2018; Krishnan et al., 2015; Karl et al., 2016; Ha & Schmidhuber, 2018; Watter et al., 2015; Zhang et al., 2018; Kulkarni et al., 2016; Mahadevan & Maggioni, 2007; Gelada et al., 2019; Gregor et al., 2018; Ghosh et al., 2019; Zhang et al., 2021). Common strategies for such state representation learning include reconstructing the observation, learning a forward model, or learning an inverse model. Furthermore, prior knowledge, such as temporal continuity (Wiskott & Sejnowski, 2002), can be added to constrain the state space.

Recently, much attention has been paid to world models, which try to learn an abstract representation of both spatial and temporal aspects of the high-dimensional input sequences (Watter et al., 2015; Ebert et al., 2017; Ha & Schmidhuber, 2018; Hafner et al., 2018; Zhang et al., 2019b; Gelada et al., 2019; Kaiser et al., 2019; Hafner et al., 2019; 2020). Based on the learned world model, agents can perform model-based RL or planning. Our proposed method is also in the class of world models, which models the generative environment model, and additionally, encodes structural constraints and achieves the sufficiency and minimality of the estimated state representations from the view of generative and selection process. In contrast, Shu et al. (2020) makes use of contrastive loss, as an alternative of reconstruction loss; however, it only focuses on the transition dynamics and also fails to ensure the sufficiency and minimality. Another line of approaches of state representation learning is based on predictive state representations (PSRs) (Littman & Sutton, 2002; Singh et al., 2004). A recent approach generalizes PSRs to nonlinear predictive models, by exploiting the coarsest partition of histories into classes that are maximally predictive of the future (Zhang et al., 2019a). Moreover, bisimulation-based methods have also attracted

much attention (Castro, 2020; Zhang et al., 2021).

On the other hand, our work is also related to Bayesian network learning and causal discovery (Spirites et al., 1993; Pearl, 2000; Huang* et al., 2020). For example, Strehl et al. (2007) considers factorized-state MDP with structures being modeled with dynamic Bayesian network or decision trees. Incorporating such structure information has shown benefits in several machine learning tasks (Zhang* et al., 2020; Huang et al., 2019), and in this paper, we show its advantages in POMDPs.

7. Conclusions and Future Work

In this paper, we develop a principled framework to characterize a minimal set of state representations that suffice for policy learning, by making use of structural constraints and the goal of maximizing cumulative reward in policy learning. Accordingly, we propose SS-VAE to reliably extract such a set of state representations from raw observations. The estimated environment model and ASRs allow learning behaviors from imagined outcomes in the compact latent space, which effectively reduce sample complexity and possibly risky interactions with the environment. The proposed approach shows promising results on complex environments—CarRacing and Vizdoom. The future work along this direction include investigating identifiability conditions in general nonlinear cases and extending the approach to cover heterogeneous environments, where the generating processes may change over time or across domains.

Acknowledgement

BH would like to acknowledge the support of Apple Scholarship. KZ would like to acknowledge the support by the National Institutes of Health (NIH) under Contract R01HL159805, by the NSF-Convergence Accelerator Track-D award #2134901, and by a grant from Apple.

References

- Bengio, Y. The consciousness prior. *arXiv preprint arXiv:1709.08568*, 2019.
- Bishop, C. M. Mixture density networks. In *Technical Report NCRG/4288, Aston University, Birmingham, UK*, 1994.
- Castro, P. S. Scalable methods for computing state similarity in deterministic Markov decision processes. *AAAI*, 2020.
- Ebert, F., Finn, C., Lee, A. X., and Levine, S. Self-supervised visual planning with temporal skip connections. *ArXiv Preprint ArXiv:1710.05268*, 2017.
- Gelada, C., Kumar, S., Buckman, J., Nachum, O., and Bellemare, M. G. Deepmdp: Learning continuous latent space models for representation learning. In *International Conference on Machine Learning (ICML)*, 2019.
- Ghosh, D., Gupta, A., and Levine, S. Learning actionable representations with goal conditioned policies. *ICLR*, 2019.
- Gregor, K., Papamakarios, G., Besse, F., Buesing, L., and Weber, T. Temporal difference variational auto-encoder. *arXiv preprint arXiv:1806.03107*, 2018.
- Ha, D. and Schmidhuber, J. World models. In *Advances in Neural Information Processing Systems*, 2018.
- Hafner, D., Lillicrap, T., Fischer, I., Villegas, R., Ha, D., Lee, H., and Davidson, J. Learning latent dynamics for planning from pixels. *arXiv preprint arXiv:1811.04551*, 2018.
- Hafner, D., Lillicrap, T., Ba, J., and Norouzi, M. Dream to control: Learning behaviors by latent imagination. *arXiv preprint arXiv:1912.01603*, 2019.
- Hafner, D., Lillicrap, T., Norouzi, M., and Ba, J. Mastering atari with discrete world models. *arXiv preprint arXiv:2010.02193*, 2020.
- Hansen, N. The CMA evolution strategy: A tutorial. *arXiv preprint arXiv:1604.00772*, 2016.
- Hausknecht, M. and Stone, P. Deep recurrent q-learning for partially observable mdps. In *2015 AAAI Fall Symposium Series*, 2015.
- Hochreiter, S. and Schmidhuber, J. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- Huang, B., Zhang, K., Gong, M., and Glymour, C. Causal discovery and forecasting in nonstationary environments with state-space models. In *International Conference on Machine Learning (ICML)*, 2019.
- Huang*, B., Zhang*, K., Zhang, J., Ramsey, J., Sanchez-Romero, R., Glymour, C., and Schölkopf, B. Causal discovery from heterogeneous/nonstationary data. *JMLR*, 21(89), 2020.
- Kaiser, L., Babaeizadeh, M., Milos, P., Osinski, B., Campbell, R. H., Czechowski, K., . . . , and Michalewski, H. Model-based reinforcement learning for Atari. *arXiv preprint arXiv:1903.00374*, 2019.
- Karl, M., Soelch, M., Bayer, J., and van der Smagt, P. Deep variational bayes filters: Unsupervised learning of state space models from raw data. *arXiv preprint arXiv:1605.06432*, 2016.
- Kay, W., Carreira, J., Simonyan, K., Zhang, B., Hillier, C., Vijayanarasimhan, S., Viola, F., Green, T., Back, T., Natsev, P., et al. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*, 2017.
- Kempka, M., Wydmuch, M., Runc, G., Toczek, J., and Jaśkowski, W. ViZDoom: A Doom-based AI research platform for visual reinforcement learning. In *IEEE Conference on Computational Intelligence and Games*, pp. 341–348, Santorini, Greece, Sep 2016. IEEE. URL <http://arxiv.org/abs/1605.02097>. The best paper award.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Klimov, O. Carracing-v0. <http://gym.openai.com/>, 2016.
- Krishnan, R., Shalit, U., and Sontag, D. Deep kalman filters. *arXiv preprint arXiv:1511.05121*, 2015.
- Kulkarni, T. D., Saedi, A., Gautam, S., and Gershman, S. J. Deep successor reinforcement learning. *arXiv preprint arXiv:1606.02396*, 2016.
- Lee, A. X., Nagabandi, A., Abbeel, P., and Levine, S. Stochastic latent actor-critic: Deep reinforcement learning with a latent variable model. *arXiv preprint arXiv:1907.00953*, 2019.
- Lesort, T., Díaz-Rodríguez, N., Goudou, J. F., and Filliat, D. State representation learning for control: An overview. *Neural Networks*, 108:379–392, 2018.
- Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., and Wierstra, D. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.
- Littman, M. L. and Sutton, R. S. Predictive representations of state. In *In Advances in neural information processing systems*, pp. 1555–1561, 2002.

- Mahadevan, S. and Maggioni, M. Proto-value functions: A laplacian framework for learning representation and control in markov decision processes. *Journal of Machine Learning Research (JMLR)*, 8:2169–2231, 2007.
- Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., and Riedmiller, M. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., Petersen, S., Beattie, C., Sadik, A., Antonoglou, I., King, H., Kumaran, D., Wierstra, D., Legg, S., and Hassabis, D. Human-level control through deep reinforcement learning. *Nature*, 518(7540): 529–533, 2015.
- Pearl, J. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, Cambridge, 2000.
- Schölkopf, B. Causality for machine learning. *arXiv preprint arXiv:1911.10500*, 2019.
- Shu, R., Nguyen, T., Chow, Y., Pham, T., Than, K., Ghavamzadeh, M., Ermon, S., and Bui, H. Predictive coding for locally-linear control. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*. PMLR, 2020.
- Singh, S., James, M. R., and Rudary, M. R. Predictive state representations: A new theory for modeling dynamical systems. In *In Proceedings of the 20th conference on Uncertainty in artificial intelligence*, pp. 512–519, 2004.
- Spirtes, P., Glymour, C., and Scheines, R. *Causation, Prediction, and Search*. Springer-Verlag Lectures in Statistics, 1993.
- Srinivas, A., Laskin, M., and Abbeel, P. Curl: Contrastive unsupervised representations for reinforcement learning. *ICML*, 2020.
- Strehl, A. L., Diuk, C., and Littman, M. L. Efficient structure learning in factored-state mdps. In *AAAI*, 2007.
- Sutton, R. S. Integrated architectures for learning, planning, and reacting based on approximating dynamic programming. In *Machine learning proceedings 1990*, pp. 216–224. Elsevier, 1990.
- Sutton, R. S. and Barto, A. G. *Reinforcement learning: An introduction*. MIT press, 2018.
- Tishby, N., Pereira, F. C., and Bialek, W. The information bottleneck method. *The 37th annual Allerton Conference on Communication, Control, and Computing*, 1999.
- Tsividis, P. A., Pouncy, T., Xu, J. L., Tenenbaum, J. B., and Gershman, S. J. Human learning in atari. In *AAAI Spring Symposium Series*, 2017.
- Watter, M., Springenberg, J., Boedecker, J., and Riedmiller, M. Embed to control: A locally linear latent dynamics model for control from raw images. *NeurIPS*, 2015.
- Wiskott, L. and Sejnowski, T. J. Slow feature analysis: Un-supervised learning of invariances. *Neural Computation*, 14(4):715–770, 2002.
- Zhang, A., Lipton, Z. C., Pineda, L., Azizzadenesheli, K., Anandkumar, A., Itti, L., Pineau, J., and Furlanello, T. Learning causal state representations of partially observable environments. *arXiv preprint arXiv:1906.10437*, 2019a.
- Zhang, A., McAllister, R., Calandra, R., Gal, Y., and Levine, S. Learning invariant representations for reinforcement learning without reconstruction. *ICLR*, 2021.
- Zhang, J. and Spirtes, P. Intervention, determinism, and the causal minimality condition. *Synthese*, 182(3):335–347, 2011.
- Zhang, K. and Hyvärinen, A. A general linear non-gaussian state-space model: Identifiability, identification, and applications. In *Asian Conference on Machine Learning*, pp. 113–128, 2011.
- Zhang*, K., Gong*, M., Stojanov, P., Huang, B., Liu, Q., and Glymour, C. Domain adaptation as a problem of inference on graphical models. In *NeurIPS*, 2020.
- Zhang, M., Vikram, S., Smith, L., Abbeel, P., Johnson, M. J., and Levine, S. Solar: deep structured representations for model-based reinforcement learning. *arXiv preprint arXiv:1808.09105*, 2018.
- Zhang, M., Vikram, S., Smith, L., Abbeel, P., Johnson, M., and Levine, S. Self-supervised visual planning with temporal skip connections. *ICML*, 2019b.

Appendices for “Action-Sufficient State Representation Learning for Control with Structural Constraints”

A. Proof of Proposition 1

We first give the definitions of the Markov condition and the faithfulness assumption, which will be used in the proof.

Definition 2 (Global Markov Condition (Spirtes et al., 1993; Pearl, 2000)). *The distribution p over a set of variables \mathbf{V} satisfies the global Markov property on graph G if for any partition (A, B, C) such that if B d -separates A from C , then $p(A, C|B) = p(A|B)p(C|B)$.*

Definition 3 (Faithfulness Assumption (Spirtes et al., 1993; Pearl, 2000)). *There are no independencies between variables that are not entailed by the Markov Condition.*

Below, we give the proof of Proposition 1.

Proof. The proof contains the following three steps.

- In step 1, we show that a state dimension $s_{i,t}$ is in ASRs, that is, it has a directed path to $r_{t+\tau}$ and the path does not go through any action variable, if and only if $s_{i,t} \not\perp\!\!\!\perp R_{t+1} | a_{t-1:t}, \vec{s}_{t-1}$.
- In step 2, we show that for $s_{i,t}$ with $s_{i,t} \not\perp\!\!\!\perp R_{t+1} | a_{t-1:t}, \vec{s}_{t-1}$, if and only if $s_{i,t} \not\perp\!\!\!\perp R_{t+1} | a_{t-1:t}, \vec{s}_{t-1}^{\text{ASR}}$.
- In step 3, we show that ASRs \vec{s}_t^{ASR} are minimal sufficient for policy learning.

Step 1: We first show that if a state dimension $s_{i,t}$ is in ASRs, then $s_{i,t} \not\perp\!\!\!\perp R_{t+1} | a_{t-1:t}, \vec{s}_{t-1}$.

We prove it by contradiction. Suppose that $s_{i,t}$ is independent of R_{t+1} given $a_{t-1:t}$ and \vec{s}_{t-1} . Then according to the faithfulness assumption, we can see from the graph that $s_{i,t}$ does not have a directed path to $r_{t+\tau}$, which contradicts to the assumption, because, otherwise, $a_{t-1:t}$ and \vec{s}_{t-1} cannot break the paths between $s_{i,t}$ and R_{t+1} which leads to the dependence.

We next show that if $s_{i,t} \not\perp\!\!\!\perp R_{t+1} | a_{t-1:t}, \vec{s}_{t-1}$, then $s_{i,t} \in \vec{s}_t^{\text{ASR}}$.

Similarly, by contradiction suppose that $s_{i,t}$ does not have a directed path to $r_{t+\tau}$. From the graph, it is easy to see that $a_{t-1:t}$ and \vec{s}_{t-1} must break the path between $s_{i,t}$ and R_{t+1} . According to the Markov assumption, $s_{i,t}$ is independent of R_{t+1} given $a_{t-1:t}$ and \vec{s}_{t-1} , which contradicts to the assumption. Since we have a contradiction, it must be that $s_{i,t}$ has a directed path to $r_{t+\tau}$.

Step 2: In step 1, we have shown that $s_{i,t} \not\perp\!\!\!\perp R_{t+1} | a_{t-1:t}, \vec{s}_{t-1}$, if and only if it has a directed path to $r_{t+\tau}$. From the graph, it is easy to see that for those state dimensions which have a directed path to $r_{t+\tau}$, $a_{t-1:t}$ and $\vec{s}_{t-1}^{\text{ASR}}$ cannot break the path between $s_{i,t}$ and R_{t+1} . Moreover, for those state dimensions which do not have a directed path to $r_{t+\tau}$, $a_{t-1:t}$ and $\vec{s}_{t-1}^{\text{ASR}}$ are enough to break the path between $s_{i,t}$ and R_{t+1} .

Therefore, for $s_{i,t}$ with $s_{i,t} \not\perp\!\!\!\perp R_{t+1} | a_{t-1:t}, \vec{s}_{t-1}$, if and only if $s_{i,t} \not\perp\!\!\!\perp R_{t+1} | a_{t-1:t}, \vec{s}_{t-1}^{\text{ASR}}$.

Step 3: In the previous steps, it has been shown that if a state dimension $s_{i,t}$ is in ASRs, then $s_{i,t} \not\perp\!\!\!\perp R_{t+1} | a_{t-1:t}, \vec{s}_{t-1}^{\text{ASR}}$, and if a state dimension $s_{i,t}$ is not in ASRs, then $s_{i,t} \perp\!\!\!\perp R_{t+1} | a_{t-1:t}, \vec{s}_{t-1}^{\text{ASR}}$. This implies that \vec{s}_t^{ASR} are minimal sufficient for policy learning to maximize the future reward.

□

B. Learning the ASRs under a Random Policy

When collecting the data, that are used to learn the environment model and ASRs, with random actions, it is apparent to observe that actions a_t and ASRs \vec{s}_t^{ASR} are dependent only conditioning on the cumulative reward and previous state—this is a type of dependence relationship induced by selection on the effect (reward). We can then learn the ASRs by maximizing

$$I(\vec{s}_t^{\text{ASR}}; a_t | R_{t+1}, \vec{s}_{t-1}^{\text{ASR}}) - I(\vec{s}_t^{\text{C}}; a_t | R_{t+1}, \vec{s}_{t-1}^{\text{ASR}}), \quad (7)$$

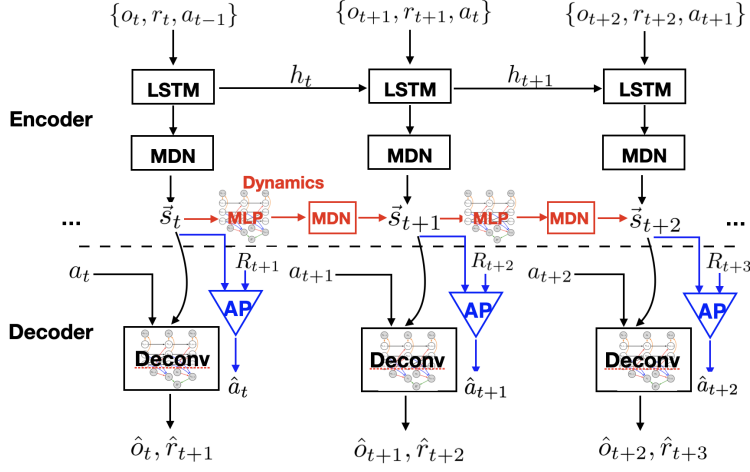


Figure 12: Diagram of neural network architecture to learn state representations. The corresponding structural constraints are involved in “Deconv” and “MLP”, and “AP” represents the action prediction part for sufficient state representation learning.

where $\tilde{s}^C = \tilde{s} \setminus \tilde{s}^{\text{ASR}}$, and I denotes mutual information. Since $I(\tilde{s}_t^{\text{ASR}}; a_t | R_{t+1}, \tilde{s}_{t-1}^{\text{ASR}}) = H(a_t | R_{t+1}, \tilde{s}_{t-1}^{\text{ASR}}) - H(a_t | \tilde{s}_{t-1:t}^{\text{ASR}}, R_{t+1})$, where $H(\cdot)$ denotes the conditional entropy, we can estimate ASRs by maximizing $H(a_t | R_{t+1}, \tilde{s}_{t-1}^{\text{ASR}}) - H(a_t | \tilde{s}_{t-1:t}^{\text{ASR}}, R_{t+1})$, with

$$H(a_t | \tilde{s}_{t-1:t}^{\text{ASR}}, R_{t+1}) = -\mathbb{E}_{q_{\phi, \alpha_1}} \{ \log p_{\alpha_1}(a_t | \tilde{s}_{t-1:t}^{\text{ASR}}, R_{t+1}) \} = -\mathbb{E}_{q_{\phi, \alpha_1}} \{ \log p_{\alpha_1}(a_t | \tilde{D}^{\text{ASR}} \odot \tilde{s}_{t-1:t}, R_{t+1}) \},$$

and

$$H(a_t | R_{t+1}, \tilde{s}_{t-1}^{\text{ASR}}) = -\mathbb{E}_{q_{\phi, \alpha_2}} \{ \log p_{\alpha_2}(a_t | R_{t+1}, \tilde{s}_{t-1}^{\text{ASR}}) \} = -\mathbb{E}_{q_{\phi, \alpha_2}} \{ \log p_{\alpha_2}(a_t | \tilde{D}^{\text{ASR}} \odot \tilde{s}_{t-1}^{\text{ASR}}, R_{t+1}) \},$$

where p_{α_i} , for $i = 1, 2$, denotes the probabilistic predictive model of a_t with parameters α_i , q_{ϕ, α_i} is the joint distribution over \tilde{s}_t and a_t with $q_{\phi, \alpha_i} = q_{\phi} p_{\alpha_i}$, and $q_{\phi}(\tilde{s}_t | \tilde{s}_{t-1}, \mathbf{y}_{1:t}, a_{1:t-1})$ is the probabilistic inference model of \tilde{s}_t with parameters ϕ and $\mathbf{y}_t = (o_t^T, r_t^T)$, and $\tilde{D}^{\text{ASR}} \in \{0, 1\}^{d \times 1}$ is a binary vector indicating which dimensions of \tilde{s}_t are in \tilde{s}_t^{ASR} , so $\tilde{D}^{\text{ASR}} \odot \tilde{s}_t$ gives ASRs \tilde{s}_t^{ASR} . Similarly, we can also represent $I(\tilde{s}_t^C; a_t | R_{t+1}, \tilde{s}_{t-1}^{\text{ASR}})$ in the same way. Accordingly, the “sufficiency & Minimality” term in the objective function as shown in Figure 2 should be replaced by (7), and the corresponding diagram of neural network architecture is as shown in Figure 12.

C. Further Regularization of Minimality

In this section, we give the detailed derivation on the further regularization of minimality of state representations given in Section 2.1, which is similar to that in the information bottleneck. We achieve it by minimizing conditional mutual information between observed high-dimensional signals \mathbf{y}_t , where $\mathbf{y}_t = \{o_t^T, r_t^T\}$, and the ASR \tilde{s}_t^{ASR} at time t given data at previous time instances, and meanwhile minimizing the dimensionality of ASRs with sparsity constraints:

$$\lambda_1 \sum_{t=2}^T I(\mathbf{y}_t; \tilde{s}_t^{\text{ASR}} | \mathbf{y}_{1:t-1}, a_{1:t-1}, \tilde{s}_{t-1}) + \lambda_2 \|\tilde{D}^{\text{ASR}}\|_1.$$

Note that in the above conditional mutual information, we need to conditional on the previous states \tilde{s}_{t-1} , instead of $\tilde{s}_{t-1}^{\text{ASR}}$, which two give different conditional mutual information. It can be shown by contradiction. Suppose $I(\mathbf{y}_t; \tilde{s}_t^{\text{ASR}} | \mathbf{y}_{1:t-1}, a_{1:t-1}, \tilde{s}_{t-1}) = I(\mathbf{y}_t; \tilde{s}_t^{\text{ASR}} | \mathbf{y}_{1:t-1}, a_{1:t-1}, \tilde{s}_{t-1}^{\text{ASR}})$, and denote $\tilde{s}^C = \tilde{s} \setminus \tilde{s}^{\text{ASR}}$. Then the equivalence implies that \tilde{s}_{t-1}^C is independent of o_t (where $o_t \in \mathbf{y}_t$) given $\{\mathbf{y}_{1:t-1}, a_{1:t-1}, \tilde{s}_{t-1}^{\text{ASR}}\}$. It is obviously violated for the example given in Figure 1, where $\tilde{s}^C = s_1$ and $\tilde{s}^{\text{ASR}} = \{s_2, s_3\}$, and $s_{1,t-1}$ is dependent on o_t given $\{\mathbf{y}_{1:t-1}, a_{1:t-1}, s_{2,t-1}, s_{3,t-1}\}$. Hence, conditioning on \tilde{s}_{t-1} and $\tilde{s}_{t-1}^{\text{ASR}}$ give different conditional mutual information. Therefore, in the above conditional mutual information, we need to condition on the previous states \tilde{s}_{t-1} .

Moreover, the conditional mutual information $I(\mathbf{y}_t; \tilde{\mathbf{s}}_t^{\text{ASR}} | \mathbf{y}_{1:t-1}, a_{1:t-1}, \tilde{\mathbf{s}}_{t-1})$ can be upper bound by a KL-divergence, and below we denote $\{\mathbf{y}_{1:t-1}, a_{1:t-1}, \tilde{\mathbf{s}}_{t-1}\}$ by \mathbf{z}_t :

$$\begin{aligned}
 & I(\mathbf{y}_t; \tilde{\mathbf{s}}_t^{\text{ASR}} | \mathbf{y}_{1:t-1}, a_{1:t-1}, \tilde{\mathbf{s}}_{t-1}) \\
 & \equiv I(\mathbf{y}_t; \tilde{\mathbf{s}}_t^{\text{ASR}} | \mathbf{z}_t) \\
 & \equiv \mathbb{E}_{p(\mathbf{y}_t, \tilde{\mathbf{s}}_t^{\text{ASR}}, \mathbf{z}_t)} \left\{ \log \frac{q_\phi(\tilde{\mathbf{s}}_t^{\text{ASR}} | \mathbf{y}_t, \mathbf{z}_t)}{p(\tilde{\mathbf{s}}_t^{\text{ASR}} | \mathbf{z}_t)} \right\} \\
 & = \mathbb{E}_{p(\mathbf{y}_t, \tilde{\mathbf{s}}_t^{\text{ASR}}, \mathbf{z}_t)} \left\{ \log \frac{q_\phi(\tilde{\mathbf{s}}_t^{\text{ASR}} | \mathbf{y}_t, \mathbf{z}_t) p_\gamma(\tilde{\mathbf{s}}_t^{\text{ASR}} | \tilde{\mathbf{s}}_{t-1}, a_{t-1})}{p(\tilde{\mathbf{s}}_t^{\text{ASR}} | \mathbf{z}_t) p_\gamma(\tilde{\mathbf{s}}_t^{\text{ASR}} | \tilde{\mathbf{s}}_{t-1}, a_{t-1})} \right\} \\
 & = \mathbb{E}_{p(\mathbf{y}_t, \tilde{\mathbf{s}}_t^{\text{ASR}}, \mathbf{z}_t)} \left\{ \log \frac{q_\phi(\tilde{\mathbf{s}}_t^{\text{ASR}} | \mathbf{y}_t, \mathbf{z}_t)}{p_\gamma(\tilde{\mathbf{s}}_t^{\text{ASR}} | \tilde{\mathbf{s}}_{t-1}, a_{t-1})} \right\} - \mathbb{E}_{p(\mathbf{z}_t)} \left\{ \text{KL}(p(\tilde{\mathbf{s}}_t^{\text{ASR}} | \mathbf{z}_t) \| p_\gamma(\tilde{\mathbf{s}}_t^{\text{ASR}} | \tilde{\mathbf{s}}_{t-1}, a_{t-1})) \right\} \\
 & \leq \mathbb{E}_{p(\mathbf{y}_t, \mathbf{z}_t)} \left[\text{KL}(q_\phi(\tilde{\mathbf{s}}_t^{\text{ASR}} | \mathbf{y}_t, \mathbf{z}_t) \| p_\gamma(\tilde{\mathbf{s}}_t^{\text{ASR}} | \tilde{\mathbf{s}}_{t-1}, a_{t-1})) \right] \\
 & \equiv \mathbb{E}_{p(\tilde{\mathbf{s}}_{t-1}, \mathbf{y}_{1:t}, a_{1:t-1})} \left[\text{KL}(q_\phi(\tilde{\mathbf{s}}_t^{\text{ASR}} | \tilde{\mathbf{s}}_{t-1}, \mathbf{y}_{1:t}, a_{1:t-1}) \| p_\gamma(\tilde{\mathbf{s}}_t^{\text{ASR}} | \tilde{\mathbf{s}}_{t-1}, a_{t-1}; D_{\tilde{\mathbf{s}}}, D_{a^* \tilde{\mathbf{s}}})) \right]
 \end{aligned}$$

with p_γ being the transition dynamics of $\tilde{\mathbf{s}}_t$ with parameters γ .

D. Assumptions of Proposition 2

To show the identifiability of the model in the linear case, we make the following assumptions:

- A1. $d_o + d_r \geq d_s$, where $|o_t| = d_o$, $|r_t| = d_r$, and $|s_t| = d_s$.
- A2. $(D_{\tilde{\mathbf{s}}^* o}^\top, D_{\tilde{\mathbf{s}}^* r}^\top)$ is full column rank and $D_{\tilde{\mathbf{s}}}^\top$ is full rank.
- A3. The control signal a_t is i.i.d. and the state $\tilde{\mathbf{s}}_t$ is stationary.
- A4. The process noise has a unit variance, i.e., $\text{var}(\eta_t) = I$.

E. Proof of Proposition 2

Proof. The proof of the linear case without control signals has been shown in Zhang & Hyvarinen (2011). Below, we give the identifiability proof in the linear-Gaussian case with control signals:

$$\begin{cases} o_t = D_{\tilde{\mathbf{s}}^* o}^\top \tilde{\mathbf{s}}_t + e_t, \\ r_{t+1} = D_{\tilde{\mathbf{s}}^* r}^\top \tilde{\mathbf{s}}_t + D_{a^* r}^\top a_t + \epsilon_{t+1}, \\ \tilde{\mathbf{s}}_t = D_{\tilde{\mathbf{s}}}^\top \tilde{\mathbf{s}}_{t-1} + D_{a^* \tilde{\mathbf{s}}}^\top a_{t-1} + \eta_t. \end{cases} \quad (8)$$

Let $\mathbf{y}_{t+1} = [o_t^\top, r_{t+1}^\top]^\top$, $\ddot{D}_{\tilde{\mathbf{s}}^* o} = [D_{\tilde{\mathbf{s}}^* o}^\top, D_{\tilde{\mathbf{s}}^* r}^\top]^\top$, $\ddot{D}_{a^* r} = [\mathbf{0}^\top, D_{a^* r}^\top]^\top$, and $\ddot{e}_t = [e_t^\top, \epsilon_{t+1}^\top]^\top$. Then the above equation can be represented as:

$$\begin{cases} \mathbf{y}_t = \ddot{D}_{\tilde{\mathbf{s}}^* o}^\top \tilde{\mathbf{s}}_t + \ddot{D}_{a^* r}^\top a_t + \ddot{e}_t, \\ \tilde{\mathbf{s}}_t = D_{\tilde{\mathbf{s}}}^\top \tilde{\mathbf{s}}_{t-1} + D_{a^* \tilde{\mathbf{s}}}^\top a_{t-1} + \eta_t. \end{cases} \quad (9)$$

Because the dynamic system is linear and Gaussian, we make use of the second-order statistics of the observed data to show the identifiability. We first consider the cross-covariance between \mathbf{y}_{t+k} and a_t :

$$\begin{cases} \text{Cov}(\mathbf{y}_{t+k}, a_t) = \ddot{D}_{\tilde{\mathbf{s}}^* o}^\top D_{\tilde{\mathbf{s}}}^{k-1} D_{a^* \tilde{\mathbf{s}}}^\top \cdot \text{Var}(a_t), & \text{if } k > 0, \\ \text{Cov}(\mathbf{y}_{t+k}, a_t) = \ddot{D}_{a^* r}^\top \cdot \text{Var}(a_t), & \text{if } k = 0. \end{cases} \quad (10)$$

Thus, from the cross-covariance between \mathbf{y}_{t+k} and a_t , we can identify $\ddot{D}_{\tilde{\mathbf{s}}^* o}^\top D_{a^* \tilde{\mathbf{s}}}^\top$, $\ddot{D}_{a^* r}$, and $\ddot{D}_{\tilde{\mathbf{s}}^* o}^\top D_{\tilde{\mathbf{s}}}^k D_{a^* \tilde{\mathbf{s}}}^\top$ for $k > 0$.

Next, we consider the auto-covariance function of $\tilde{\mathbf{s}}$. Define the auto-covariance function of $\tilde{\mathbf{s}}$ at lag k as $\mathbf{R}_{\tilde{\mathbf{s}}}(k) = \mathbb{E}[\tilde{\mathbf{s}}_t \tilde{\mathbf{s}}_{t+k}^\top]$, and similarly for $\mathbf{R}_{\mathbf{y}}(k)$. Clearly, $\mathbf{R}_{\tilde{\mathbf{s}}}(-k) = \mathbf{R}_{\tilde{\mathbf{s}}}(k)^\top$ and $\mathbf{R}_{\mathbf{y}}(-k) = \mathbf{R}_{\mathbf{y}}(k)^\top$. Then we have

$$\begin{cases} \mathbf{R}_{\tilde{\mathbf{s}}}(k) = \mathbf{R}_{\tilde{\mathbf{s}}}(k-1) \cdot D_{\tilde{\mathbf{s}}}, & \text{if } k > 0, \\ \mathbf{R}_{\tilde{\mathbf{s}}}(k) = \mathbf{R}_{\tilde{\mathbf{s}}}^\top(1) \cdot D_{\tilde{\mathbf{s}}} + D_{a^* \tilde{\mathbf{s}}}^\top \text{Var}(a_{t-1}) D_{a^* \tilde{\mathbf{s}}} + I, & \text{if } k = 0. \end{cases} \quad (11)$$

Below, we first consider the case where $d_o + d_r = d_s$. Let $\tilde{\mathbf{y}}_t = \ddot{D}_{\tilde{\mathbf{s}}^* o}^\top \tilde{\mathbf{s}}_t$, so $\mathbf{y}_t = \tilde{\mathbf{y}}_t + \ddot{D}_{a^* r}^\top a_{t-1} + \ddot{e}_t$ and $\mathbf{R}_{\tilde{\mathbf{y}}}(k) = \ddot{D}_{\tilde{\mathbf{s}}^* o}^\top \mathbf{R}_{\tilde{\mathbf{s}}}(k) \ddot{D}_{\tilde{\mathbf{s}}^* o}$. $\mathbf{R}_{\tilde{\mathbf{y}}}(k)$ satisfies the recursive property:

$$\begin{cases} \mathbf{R}_{\tilde{\mathbf{y}}}(k) = \mathbf{R}_{\tilde{\mathbf{y}}}(k-1) \cdot \Omega^\top, & \text{if } k > 0, \\ \mathbf{R}_{\tilde{\mathbf{y}}}(k) = \mathbf{R}_{\tilde{\mathbf{y}}}^\top(1) \cdot \Omega^\top + \ddot{D}_{\tilde{\mathbf{s}}^* o}^\top (D_{a^* \tilde{\mathbf{s}}}^\top \text{Var}(a_{t-1}) D_{a^* \tilde{\mathbf{s}}} + I) \ddot{D}_{\tilde{\mathbf{s}}^* o}, & \text{if } k = 0, \end{cases} \quad (12)$$

where $\Omega = \ddot{D}_{\bar{s} \rightarrow o}^\top D_{\bar{s}} \ddot{D}_{\bar{s} \rightarrow o}^{-1}$.

Denote $S_k = \ddot{D}_{\bar{s} \rightarrow o}^\top D_{\bar{s}}^{k-1} D_{a \rightarrow \bar{s}}^\top \cdot \text{Var}(a_t)$. Then we can derive the recursive property for $\mathbf{R}_y(k)$:

$$\begin{cases} \mathbf{R}_y(k) = \mathbf{R}_y(k-1) \cdot \Omega^\top - \ddot{D}_{a \rightarrow r}^\top S_{k-1}^\top \Omega^\top + \ddot{D}_{a \rightarrow r}^\top S_k^\top, & \text{if } k > 1, \\ \mathbf{R}_y(k) = \mathbf{R}_y(k-1) \cdot \Omega^\top - \ddot{D}_{a \rightarrow r}^\top \text{Var}^\top(a_t) \ddot{D}_{a \rightarrow r} \Omega^\top - \Sigma_e \Omega^\top + \ddot{D}_{a \rightarrow r}^\top S_k^\top, & \text{if } k = 1, \\ \mathbf{R}_y(k) = \mathbf{R}_y^\top(1) \cdot \Omega^\top + (\ddot{D}_{a \rightarrow r}^\top \text{Var}(a_t) \ddot{D}_{a \rightarrow r} + \Sigma_e) \\ \quad + \ddot{D}_{\bar{s} \rightarrow o}^\top (D_{a \rightarrow \bar{s}}^\top \text{Var}(a_t) D_{a \rightarrow \bar{s}} + I) \ddot{D}_{\bar{s} \rightarrow o}, & \text{if } k = 0. \end{cases}$$

When $k = 2$, we have

$$\mathbf{R}_y(2) = \mathbf{R}_y(1) \cdot \Omega^\top - \ddot{D}_{a \rightarrow r}^\top S_1^\top \Omega^\top + \ddot{D}_{a \rightarrow r}^\top S_2^\top.$$

The above equation can be re-organized as

$$(\mathbf{R}_y(2) - \ddot{D}_{a \rightarrow r}^\top \cdot S_2^\top) = (\mathbf{R}_y(1) - \ddot{D}_{a \rightarrow r}^\top \cdot S_1^\top) \cdot \Omega^\top.$$

Because $\ddot{D}_{a \rightarrow r}$ and S_k are identifiable, and suppose $(\mathbf{R}_y(1) - \ddot{D}_{a \rightarrow r}^\top \cdot S_1^\top)$ is invertible, $\Omega = \ddot{D}_{\bar{s} \rightarrow o}^\top D_{\bar{s}} \ddot{D}_{\bar{s} \rightarrow o}^{-1}$ is identifiable.

We further consider $\mathbf{R}_y(0)$ and $\mathbf{R}_y(1)$ and write down the following form:

$$\begin{aligned} & \begin{bmatrix} \mathbf{R}_y(0) - \ddot{D}_{\bar{s} \rightarrow o}^\top (D_{a \rightarrow \bar{s}}^\top \text{Var}(a_{t-1}) D_{a \rightarrow \bar{s}} + I) \ddot{D}_{\bar{s} \rightarrow o} \\ \mathbf{R}_y(1) \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{R}_y^\top(1) \\ \mathbf{R}_y(0) \end{bmatrix} \cdot \Omega^\top + \begin{bmatrix} \ddot{D}_{a \rightarrow r}^\top \text{Var}(a_t) \ddot{D}_{a \rightarrow r} \\ -\ddot{D}_{a \rightarrow r}^\top \text{Var}^\top(a_t) \ddot{D}_{a \rightarrow r} \Omega^\top + \ddot{D}_{a \rightarrow r}^\top S_1^\top \end{bmatrix} + \Sigma_e \begin{bmatrix} I \\ -\Omega^\top \end{bmatrix}. \end{aligned}$$

From the above two equations we can then identify Σ_e and $\ddot{D}_{\bar{s} \rightarrow o}^\top (D_{a \rightarrow \bar{s}}^\top \text{Var}(a_{t-1}) D_{a \rightarrow \bar{s}} + I) \ddot{D}_{\bar{s} \rightarrow o}$, and because $\ddot{D}_{\bar{s} \rightarrow o}^\top D_{a \rightarrow \bar{s}}^\top$ is identifiable, $\ddot{D}_{\bar{s} \rightarrow o}^\top \ddot{D}_{\bar{s} \rightarrow o}$ is identifiable.

In summary, we have shown the identifiability of $\ddot{D}_{a \rightarrow r}$, $\ddot{D}_{\bar{s} \rightarrow o}^\top D_{a \rightarrow \bar{s}}^\top$, $\ddot{D}_{\bar{s} \rightarrow o}^\top D_{\bar{s}}^k D_{a \rightarrow \bar{s}}^\top$, $\ddot{D}_{\bar{s} \rightarrow o}^\top \ddot{D}_{\bar{s} \rightarrow o}$, and Σ_e . Furthermore, $\ddot{D}_{\bar{s} \rightarrow o}$, $D_{\bar{s}}$, and $D_{a \rightarrow \bar{s}}$ are identified up to some orthogonal transformations. That is, suppose the model in Eq. 8 with parameters $(D_{\bar{s} \rightarrow o}, D_{\bar{s} \rightarrow r}, D_{a \rightarrow r}, D_{\bar{s}}, D_{a \rightarrow \bar{s}}, \Sigma_e, \Sigma_\epsilon)$ and that with $(\tilde{D}_{\bar{s} \rightarrow o}, \tilde{D}_{\bar{s} \rightarrow r}, \tilde{D}_{a \rightarrow r}, \tilde{D}_{\bar{s}}, \tilde{D}_{a \rightarrow \bar{s}}, \tilde{\Sigma}_e, \tilde{\Sigma}_\epsilon)$ are observationally equivalent, we then have $\tilde{D}_{\bar{s} \rightarrow o} = U \ddot{D}_{\bar{s} \rightarrow o}$, $\tilde{D}_{a \rightarrow r} = D_{a \rightarrow r}$, $\tilde{D}_{\bar{s}} = U^\top D_{\bar{s}} U$, $\tilde{D}_{a \rightarrow \bar{s}} = D_{a \rightarrow \bar{s}} U$, $\tilde{\Sigma}_e = \Sigma_e$, and $\tilde{\Sigma}_\epsilon = \Sigma_\epsilon$, where U is an orthogonal matrix.

Next, we extend the above results to the case where $d_o + d_r > d_s$. Let $\ddot{D}_{\bar{s} \rightarrow o(i, \cdot)}$ be the i -th row of $\ddot{D}_{\bar{s} \rightarrow o}$. Recall that $\ddot{D}_{\bar{s} \rightarrow o}^\top$ is of full column rank. Then for any i , one can show that there always exist $d_s - 1$ rows of $\ddot{D}_{\bar{s} \rightarrow o}$, such that they, together with $\ddot{D}_{\bar{s} \rightarrow o(i, \cdot)}$, form a $d_s \times d_s$ full-rank matrix, denoted by $\ddot{\ddot{D}}_{\bar{s} \rightarrow o(i, \cdot)}$. Then from the observed data corresponding to $\ddot{\ddot{D}}_{\bar{s} \rightarrow o(i, \cdot)}$, $\ddot{\ddot{D}}_{\bar{s} \rightarrow o(i, \cdot)}$ is determined up to orthogonal transformations. Thus, $\ddot{D}_{\bar{s} \rightarrow o}$ is identified up to orthogonal transformations. Similarly, $D_{a \rightarrow r}$, $D_{\bar{s}}$, and $D_{a \rightarrow \bar{s}}$ are identified up to orthogonal transformations. Furthermore, $\text{Cov}(\ddot{D}_{\bar{s} \rightarrow o}^\top \tilde{s}_t + D_{a \rightarrow r}^\top a_t)$ is determined by $\ddot{D}_{\bar{s} \rightarrow o}$, $\ddot{D}_{a \rightarrow r}$, $D_{\bar{s}}$, and $D_{a \rightarrow \bar{s}}$. Because $\text{Cov}(y_t) = \text{Cov}(\ddot{D}_{\bar{s} \rightarrow o}^\top \tilde{s}_t + D_{a \rightarrow r}^\top a_t) + \Sigma_{\tilde{e}}$, $\Sigma_{\tilde{e}}$ is identifiable.

One may further add sparsity constraints on $D_{\bar{s} \rightarrow o}$, $D_{\bar{s} \rightarrow r}$, $D_{\bar{s}}$, and $D_{a \rightarrow \bar{s}}$, to select more sparse structures among the equivalent ones. For example, one may add sparsity constraints on the rows of $D_{\bar{s} \rightarrow o}$. Note this corresponds to the mask on the elements of \tilde{s}_t in Eq. 2; if the full row is 0, then the corresponding dimension of \tilde{s}_t is not selected. □

F. More Estimation Details for General Nonlinear Models

The generative model p_θ can be further factorized as follows:

$$\begin{aligned} & \log p_\theta(\mathbf{y}_{1:T} | \tilde{\mathbf{s}}_{1:T}, a_{1:T-1}; D_{\bar{s} \rightarrow o}, D_{\bar{s} \rightarrow r}, D_{a \rightarrow r}) \\ &= \log p_\theta(o_{1:T} | \tilde{\mathbf{s}}_{1:T}; D_{\bar{s} \rightarrow o}) + \log p_\theta(r_{1:T} | \tilde{\mathbf{s}}_{1:T}, a_{1:T-1}; D_{\bar{s} \rightarrow r}, D_{a \rightarrow r}) \\ &= \sum_{t=1}^T \log p_\theta(o_t | \tilde{s}_t; D_{\bar{s} \rightarrow o}) + \log p_\theta(r_t | \tilde{s}_{t-1}, a_{t-1}; D_{\bar{s} \rightarrow r}, D_{a \rightarrow r}), \end{aligned} \quad (13)$$

where both $p_\theta(o_t | \tilde{s}_t; D_{\bar{s} \rightarrow o})$ and $p_\theta(r_t | \tilde{s}_{t-1}, a_{t-1}; D_{\bar{s} \rightarrow r}, D_{a \rightarrow r})$ are modelled by mixture of Gaussians, with $D_{\bar{s} \rightarrow o}$ indicating the existence of edges from \tilde{s}_t to o_t and $D_{\bar{s} \rightarrow r}$ indicating the existence of edges from \tilde{s}_{t-1} to r_t .

The inference model $q_\phi(\tilde{\mathbf{s}}_{1:T}|\mathbf{y}_{1:T}, a_{1:T-1})$ is factorized as

$$\begin{aligned} & \log q_\phi(\tilde{\mathbf{s}}_{1:T}|\mathbf{y}_{1:T}, a_{1:T-1}) \\ &= \log q_\phi(\tilde{\mathbf{s}}_1|\mathbf{y}_1, a_0) + \sum_{t=2}^T \log q_\phi(\tilde{\mathbf{s}}_t|\tilde{\mathbf{s}}_{t-1}, \mathbf{y}_{1:t}, a_{1:t-1}), \end{aligned}$$

where both $q_\phi(\tilde{\mathbf{s}}_1|\mathbf{y}_1, a_0)$ and $q_\phi(\tilde{\mathbf{s}}_t|\tilde{\mathbf{s}}_{t-1}, \mathbf{y}_{1:t}, a_{1:t-1})$ are modelled with mixture of Gaussians.

The transition dynamics p_γ is factorized as

$$\log p_\gamma(\tilde{\mathbf{s}}_{1:T}|a_{1:T-1}; D_{\tilde{\mathbf{s}}(\cdot, i)}, D_{a \rightarrow \tilde{\mathbf{s}}(\cdot, i)}) = \sum_{t=1}^T \log p_\gamma(\tilde{\mathbf{s}}_t|\tilde{\mathbf{s}}_{t-1}, a_{t-1}; D_{\tilde{\mathbf{s}}(\cdot, i)}, D_{a \rightarrow \tilde{\mathbf{s}}(\cdot, i)}), \quad (14)$$

with $\tilde{\mathbf{s}}_t|\tilde{\mathbf{s}}_{t-1}$ modelled with mixture of Gaussians.

Thus, the KL divergence can be represented as follows:

$$\begin{aligned} & \text{KL}(q_\phi(\tilde{\mathbf{s}}_{1:T}|\mathbf{y}_{1:T}, a_{1:T-1})\|p_\gamma(\tilde{\mathbf{s}}_{1:T})) \\ &= \text{KL}(q_\phi(\tilde{\mathbf{s}}_1|\mathbf{y}_1, a_0)\|p_\gamma(\tilde{\mathbf{s}}_1)) + \sum_{t=2}^T \mathbb{E}_{q_\phi} [\text{KL}(q_\phi(\tilde{\mathbf{s}}_t|\tilde{\mathbf{s}}_{t-1}, \mathbf{y}_{1:t}, a_{1:t-1})\|p_\gamma(\tilde{\mathbf{s}}_t|\tilde{\mathbf{s}}_{t-1}))]. \end{aligned} \quad (15)$$

In practice, KL divergence with mixture of Gaussians is hard to implement, so instead, we used the following objective function:

$$\begin{aligned} & \text{KL}(q_\phi(\tilde{\mathbf{s}}_1|\mathbf{y}_1, a_0)\|p_{\gamma'}(\tilde{\mathbf{s}}_1)) + \sum_{t=2}^T \mathbb{E}_{q_\phi} [\text{KL}(q_\phi(\tilde{\mathbf{s}}_t|\tilde{\mathbf{s}}_{t-1}, \mathbf{y}_{1:t}, a_{1:t-1})\|p_{\gamma'}(\tilde{\mathbf{s}}_t|\tilde{\mathbf{s}}_{t-1}))] \\ & + \lambda \sum_{t=1}^T \log p_\gamma(\tilde{\mathbf{s}}_t|\tilde{\mathbf{s}}_{t-1}, a_{t-1}; D_{\tilde{\mathbf{s}}(\cdot, i)}, D_{a \rightarrow \tilde{\mathbf{s}}(\cdot, i)}) \end{aligned} \quad (16)$$

where $p_{\gamma'}$ is a standard multivariate Gaussian $\mathcal{N}(\vec{0}, I_d)$.

G. More Details for Policy Learning with ASRs

Algorithm 1 gives the procedure of model-free policy learning with ASRs in partially observable environments. Specifically, it starts from model initialization (line 1) and data collection with a random policy (line 2). Then it updates the environment model and identifies the set of ASRs with the collected data (line 3), after which, the main procedure of policy optimization follows. In particular, because we do not directly observe the states $\tilde{\mathbf{s}}_t$, on lines 8 and 12, we infer $q_\phi(\tilde{\mathbf{s}}_{t+1}^{\text{ASR}}|o_{\leq t+1}, r_{\leq t+1}, a_{\leq t})$ and sample $\tilde{\mathbf{s}}_{t+1}^{\text{ASR}}$ from the posterior. The sampled ASRs are then stored in the buffer (line 13). Furthermore, we randomly sample a minibatch of N transitions to optimize the policy (lines 14 and 15). One may perform various RL algorithms on the ASRs, such as deep deterministic policy gradient (DDPG (Lillicrap et al., 2015)) or Q-learning (Mnih et al., 2015).

Algorithm 2 presents the procedure of the classic model-based Dyna algorithm with ASRs. Lines 17-22 make use of the learned environment model to predict the next step, including $\tilde{\mathbf{s}}_{t+1}^{\text{ASR}}$ and r_{t+1} , and update the Q function n times. Specifically, in our implementation, the hyper-parameter n is 20. Based on the learned model, the agent learns behaviors from imagined outcomes in the compact latent space, which helps to increase sample efficiency.

H. Additional Experiments and Details

H.1. CarRacing Experiment

CarRacing is a continuous control task with three continuous actions: steering left/right, acceleration, and brake. Reward is -0.1 every frame and $+1000/N$ for every track tile visited, where N is the total number of tiles in track. It is obvious that the CarRacing environment is partially observable: by just looking at the current frame, although we can tell the position of the car, we know neither its direction nor velocity that are essential for controlling the car.

For a fair comparison, we followed a similar setting as in Ha & Schmidhuber (2018). Specifically, we collected a dataset of $10k$ random rollouts of the environment, and each runs with random policy until failure, for model estimation. The dimensionality of latent states $\tilde{\mathbf{s}}_t$ was set to $\tilde{d} = 32$, and regularization parameters was set to $\lambda_1 = 1, \lambda_2 = 1, \lambda_3 = 1, \lambda_4 = 1, \lambda_5 = 1, \lambda_6 = 6, \lambda_7 = 10, \lambda_8 = 0.1$, which are determined by hyperparameter tuning.

Algorithm 1 Model-Free Policy Learning with ASRs in Partially Observable Environments

- 1: Randomly initialize neural networks and initialize replay buffer \mathcal{B} .
 - 2: Apply random control signals and record multiple rollouts.
 - 3: Estimate the model given in (2) with the recorded data (according to Section 3).
 - 4: Identify indices of ASRs according to the learned graph structure and the criteria in Prop. 1.
 - 5: **for** episode = 1, ..., M **do**
 - 6: Initialize a random process \mathcal{N} for action exploration.
 - 7: Receive initial observations o_1 and r_1 .
 - 8: Infer the posterior $q_\phi(\vec{s}_1^{\text{ASR}}|o_1, r_1)$ and sample \vec{s}_1^{ASR} .
 - 9: **for** t = 1, ..., T **do**
 - 10: Select action $a_t = \pi(\vec{s}_t^{\text{ASR}}) + \mathcal{N}_t$ according to the current policy and exploration noise.
 - 11: Execute action a_t and receive reward r_{t+1} and observation o_{t+1} .
 - 12: Infer the posterior $q_\phi(\vec{s}_{t+1}^{\text{ASR}}|o_{\leq t+1}, r_{\leq t+1}, a_{\leq t})$ and sample $\vec{s}_{t+1}^{\text{ASR}}$.
 - 13: Store transition $(\vec{s}_t^{\text{ASR}}, a_t, r_{t+1}, \vec{s}_{t+1}^{\text{ASR}})$ in \mathcal{B} .
 - 14: Sample a random minibatch of N transitions $(\vec{s}_i^{\text{ASR}}, a_i, r_{i+1}, \vec{s}_{i+1}^{\text{ASR}})$ from \mathcal{B} .
 - 15: Update network parameters using a specified RL algorithm (e.g., DQN or DDPG).
 - 16: **end for**
 - 17: **end for**
-

Algorithm 2 Model-Based Policy Learning with ASRs in Partially Observable Environments

- 1: Randomly initialize neural networks and initialize replay buffer \mathcal{B} .
 - 2: Apply random control signals and record multiple rollouts.
 - 3: Estimate the model given in (2) with the recorded data (according to Section 3).
 - 4: Identify indices of ASRs according to the learned graph structure and the criteria in Prop. 1.
 - 5: **for** episode = 1, ..., M **do**
 - 6: Initialize a random process \mathcal{N} for action exploration.
 - 7: Receive initial observations o_1 and r_1 .
 - 8: Infer the posterior $q_\phi(\vec{s}_1^{\text{ASR}}|o_1, r_1)$ and sample \vec{s}_1^{ASR} .
 - 9: **for** t = 1, ..., T **do**
 - 10: Select action $a_t = \pi(\vec{s}_t^{\text{ASR}}) + \mathcal{N}_t$ according to the current policy and exploration noise.
 - 11: Execute action a_t and receive reward r_{t+1} and observation o_{t+1} .
 - 12: Infer the posterior $q_\phi(\vec{s}_{t+1}^{\text{ASR}}|o_{\leq t+1}, r_{\leq t+1}, a_{\leq t})$ and sample $\vec{s}_{t+1}^{\text{ASR}}$.
 - 13: Store transition $(\vec{s}_t^{\text{ASR}}, \vec{s}_t, a_t, r_{t+1}, \vec{s}_{t+1}^{\text{ASR}}, \vec{s}_{t+1}, o_{t+1})$ in \mathcal{B} .
 - 14: Sample a random minibatch of N transitions $(\vec{s}_i^{\text{ASR}}, a_i, r_{i+1}, \vec{s}_{i+1}^{\text{ASR}})$ from \mathcal{B} .
 - 15: Update network parameters using a specified RL algorithm (e.g., DQN or DDPG).
 - 16: Update the model given in (2) with the recorded data from \mathcal{B} (according to Section 3).
 - 17: **for** p = 1, ..., n **do**
 - 18: Sample a random minibatch of pairs of (\vec{s}_t, a_t) from \mathcal{B} .
 - 19: Predict $(\vec{s}_{t+1}^{\text{ASR}}, r_{t+1})$ according to the model given in (2).
 - 20: Update network parameters using a specified RL algorithm (e.g., DQN or DDPG).
 - 21: **end for**
 - 22: **end for**
 - 23: **end for**
-

Without sparsity constraints. Figure 13 gives the estimated structural matrices $D_{\tilde{s} \rightarrow o}$, $D_{\tilde{s} \rightarrow r}$, $D_{a \rightarrow \tilde{s}}$, and $D_{\tilde{s}}$ in CarRacing, without the explicit sparsity constraints, where the connections are very dense.

Difference between our SS-VAE and Planet, Dreamer. Both our method and Planet (Hafner et al., 2018) and Dreamer (Hafner et al., 2019) are world model-based methods. The differences are mainly in two aspects: (1) our method explicitly considers the structural relationships among variables in the RL system, and (2) it guarantees minimal sufficient state representations for policy learning. Previous approaches usually fail to take into account whether the extracted state representations are sufficient and necessary for downstream policy learning. Moreover, as for the component of recurrent networks, SS-VAE uses LSTM that only contains the stochastic part, while PlaNet and Dreamer use RSSM that contains both deterministic and stochastic components.




Figure 13: Visualization of estimated structural matrices $D_{\tilde{s} \rightarrow o}$, $D_{\tilde{s} \rightarrow r}$, $D_{a \rightarrow \tilde{s}}$, and $D_{\tilde{s}}$ in Car Racing, without the explicit sparsity constraints.

H.2. VizDoom Experiment

We also applied the proposed method to VizDoom (Kempka et al., 2016). VizDoom provides many scenarios and we chose the *take cover* scenario. Unlike CarRacing, *take cover* is a discrete control problem with two actions: move left and move right. Reward is +1 at each time step while alive, and the cumulative reward is defined to be the number of time steps the agent manages to stay alive during an episode. Therefore, in order to survive as long as possible, the agent has to learn how to avoid fireballs shot by monsters from the other side of the room. In this task, *solving* is defined as attaining the average survival time of greater than 750 time steps over 100 consecutive episodes, each running for a maximum of 2100 time steps.

Following a similar setting as in Ha & Schmidhuber (2018), we collected a dataset of 10k random rollouts of the environment, and each runs with random policy until failure. The dimensionality of latent state \tilde{s}_t is set to $\tilde{d} = 32$. We also set $\lambda_1 = 1$, $\lambda_2 = 1$, $\lambda_3 = 1$, $\lambda_4 = 1$, $\lambda_5 = 1$, $\lambda_6 = 6$, $\lambda_7 = 10$, $\lambda_8 = 0.1$. By tuning thresholds, we finally reported all the results on the 21-dim ASRs, which achieved the best results in all the experiments.

Analysis of ASRs. Similar to the analysis in CarRacing, we also visualized the learned $D_{\tilde{s} \rightarrow o}$, $D_{\tilde{s} \rightarrow r}$, $D_{\tilde{s}}$, and $D_{a \rightarrow \tilde{s}}$ in VizDoom, as shown in Figure 14. Intuitively, we can see that $D_{\tilde{s} \rightarrow r}$ and $D_{a \rightarrow \tilde{s}}$ have many values close to zero, meaning that the reward is only influenced by a small number of state dimensions, and not many state dimensions are influenced by the action. Furthermore, from $D_{\tilde{s}}$, we found that the connections across states are sparse.




Figure 14: Visualization of estimated structural matrices $D_{\tilde{s} \rightarrow o}$, $D_{\tilde{s} \rightarrow r}$, $D_{a \rightarrow \tilde{s}}$, and $D_{\tilde{s}}$ in VizDoom.

I. Detailed Model Architectures

In the car racing experiment, the original screen images were resized to $64 \times 64 \times 3$ pixels. The encoder consists of three components: a preprocessor, an LSTM, and an MDN. The preprocessor architecture is presented in Figure 15, which takes as input the images, actions and rewards, and its output acts as the input to LSTM. We used 256 hidden units in the LSTM and used a five-component Gaussian mixture in the MDN. The decoder also consists of three components: a current observation reconstructor (Figure 16), a next observation predictor (Figure 17), and a reward predictor (Figure 18). The architecture of the transition/dynamics is shown in Figure 19, and its output is also modelled by an MDN with a five-component Gaussian mixture. In the VizDoom experiment, we used the same image size and the same architectures except that the LSTM has 512 hidden units and the action has one dimension. It is worth emphasising that we applied weight normalization to all the parameters of the architectures above except for the structural matrices $D_{(\cdot)}$.

In DDPG, both actor network and critic network are modelled by two fully connected layers of size 300 with ReLU and batch normalisation. Similarly, in DQN (Mnih et al., 2013) on both ASRs and SSSs, the Q network is also modelled by two fully connected layers of size 300 with ReLU and batch normalisation. However, in DQN on observations, it is modelled by three convolutional layers (i.e., $\text{relu conv } 32 \times 8 \times 8 \rightarrow \text{relu conv } 64 \times 4 \times 4 \rightarrow \text{relu conv } 64 \times 3 \times 3$) followed by two additional fully connected layers of size 64. In DRQN (Hausknecht & Stone, 2015) on observations, we used the same architecture as in DQN on observations but padded an extra LSTM layer with 256 hidden units as the final layer.




Figure 15: Network architecture of preprocessor.




Figure 16: Network architecture of observation reconstruction.

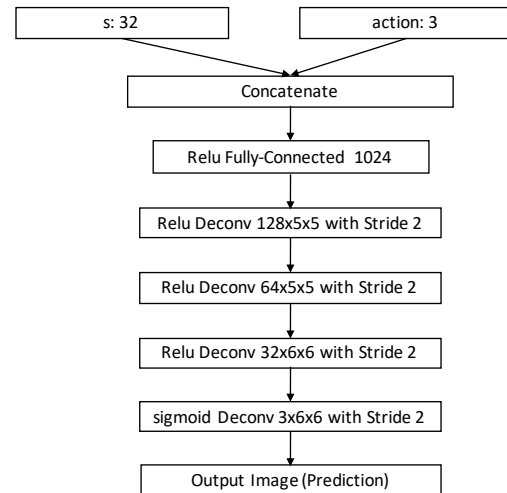


Figure 17: Network architecture of observation prediction.

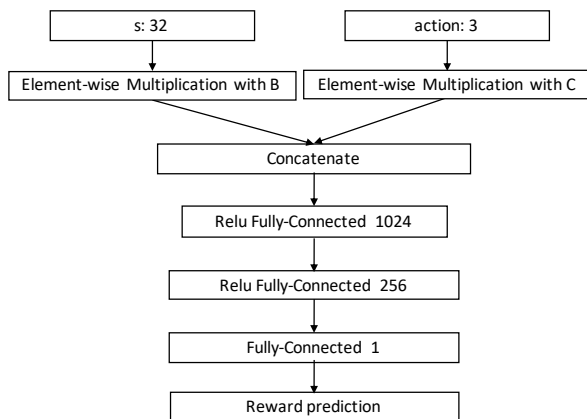


Figure 18: Network architecture of reward.

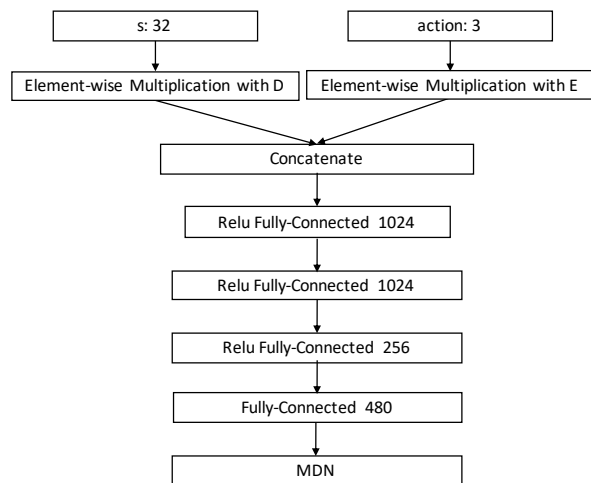


Figure 19: Network architecture of transition/dynamics.