
EMI: Exploration with Mutual Information

Hyungseok Kim^{*1,2} Jaekyeom Kim^{*1,2} Yeonwoo Jeong^{1,2} Sergey Levine³ Hyun Oh Song^{1,2}

Abstract

Reinforcement learning algorithms struggle when the reward signal is very sparse. In these cases, naive random exploration methods essentially rely on a random walk to stumble onto a rewarding state. Recent works utilize intrinsic motivation to guide the exploration via generative models, predictive forward models, or discriminative modeling of novelty. We propose EMI, which is an exploration method that constructs embedding representation of states and actions that does not rely on generative decoding of the full observation but extracts predictive signals that can be used to guide exploration based on forward prediction in the representation space. Our experiments show competitive results on challenging locomotion tasks with continuous control and on image-based exploration tasks with discrete actions on Atari. The source code is available at <https://github.com/snu-mlab/EMI>.

1. Introduction

The central task in reinforcement learning is to learn policies that would maximize the total reward received from interacting with the unknown environment. Although recent methods have been demonstrated to solve a range of complex tasks (Mnih et al., 2015; Schulman et al., 2015; 2017), the success of these methods hinges on whether the agent constantly receives the intermediate reward feedback or not. In case of challenging environments with sparse reward signals, these methods struggle to obtain meaningful policies unless the agent luckily stumbles into the rewarding or predefined goal states.

To this end, prior works on exploration generally utilize some kind of intrinsic motivation mechanism to provide a

^{*}Equal contribution ¹Seoul National University, Department of Computer Science and Engineering ²Neural Processing Research Center ³UC Berkeley, Department of Electrical Engineering and Computer Sciences. Correspondence to: Hyun Oh Song <hyunoh@snu.ac.kr>.




Figure 1: Visualization of a sample trajectory in our learned embedding space.

measure of novelty. These measures can be based on density estimation via generative models (Bellemare et al., 2016; Fu et al., 2017; Oh et al., 2015), predictive forward models (Stadie et al., 2015; Houthoofd et al., 2016), or discriminative methods that aim to approximate novelty (Pathak et al., 2017). Methods based on predictive forward models and generative models must model the distribution over state observations, which can make them difficult to scale to complex, high-dimensional observation spaces.

Our aim in this work is to devise a method for exploration that does not require a direct generation of high-dimensional state observations, while still retaining the benefits of being able to measure novelty based on the forward prediction. If exploration is performed by seeking out states that maximize surprise, the problem, in essence, is in measuring surprise, which requires a representation where functionally similar states are close together, and functionally distinct states are far apart.

In this paper, we propose to learn compact representations for both the states (ϕ) and actions (ψ) simultaneously satisfying the following criteria: First, given the representations of state and the corresponding next state, the uncertainty of the representation of the corresponding action should be minimal. Second, given the representations of the state and the corresponding action, the uncertainty of the representation of the corresponding next state should also be minimal. Third, the action embedding representation (ψ) should seamlessly support both continuous and discrete actions. Finally, we impose a linear dynamics model in the representation space which can also explain the rare irreducible error under the dynamics model. Given the representation, we guide the exploration by measuring surprise based on forward

prediction and a relative increase in diversity in the embedding representation space. Figure 1 illustrates an example visualization of our learned state embedding representations (ϕ) and sample trajectories in the representation space in Montezuma’s Revenge.

We present two main technical contributions that make this into a practical exploration method. First, we describe how compact state and action representations can be constructed via variational divergence estimation of mutual information without relying on generative decoding of full observations (Nowozin et al., 2016). Second, we show that imposing linear topology on the learned embedding representation space (such that the transitions are linear), thereby offloading most of the modeling burden onto the embedding function itself, provides an essential informative measure of surprise when visiting novel states.

For the experiments, we show that we can use our representations on a range of complex image-based tasks and robotic locomotion tasks with continuous actions. We report significantly improved results compared to a number of recent intrinsic motivation based exploration methods (Fu et al., 2017; Pathak et al., 2017) on several challenging Atari tasks and robotic locomotion tasks with sparse rewards.

2. Related works

Our work is related to the following strands of active research:

Unsupervised representation learning via mutual information estimation Recent literature on unsupervised representation learning generally focuses on extracting latent representations maximizing an approximate lower bound on the mutual information between the code and the data. In the context of generative adversarial networks (Goodfellow et al., 2014), Chen et al. (2016); Belghazi et al. (2018) aim at maximizing the approximation of mutual information between the latent code and the raw data. Belghazi et al. (2018) estimates the mutual information with neural networks via Donsker & Varadhan (1983) estimation to learn better generative models. Hjelm et al. (2018) builds on the idea and trains a decoder-free encoding representation maximizing the mutual information between the input image and the representation. Furthermore, the method uses f -divergence (Nowozin et al., 2016) estimation of Jensen-Shannon divergence rather than the KL divergence to estimate the mutual information for better numerical stability. Oord et al. (2018) estimates mutual information via an autoregressive model and makes predictions on local patches in an image. Thomas et al. (2017) aims to learn the representations that maximize the causal relationship between the distributed policies and the representation of changes in the state. Nachum et al. (2018) connects mutual information estimators to represen-

tation learning in hierarchical RL.

Exploration with intrinsic motivation Prior works on exploration mostly employ intrinsic motivation to estimate the measure of novelty or surprisal to guide the exploration. Mohamed & Rezende (2015) introduced the connection between mutual information estimation and empowerment for intrinsic motivation. Bellemare et al. (2016); Ostrovski et al. (2017) utilize density estimation via CTS (Bellemare et al., 2014) generative model and PixelCNN (van den Oord et al., 2016) and derive pseudo-counts as the intrinsic motivation. Fu et al. (2017) avoids building explicit density models by training K-exemplar models that distinguish a state from all other observed states. Some methods train predictive forward models (Stadie et al., 2015; Houthoofd et al., 2016; Oh et al., 2015) and estimate the prediction error as the intrinsic motivation. Oh et al. (2015) employs generative decoding of the full observation via recursive autoencoders and thus can be challenging to scale for high dimensional observations. VIME (Houthoofd et al., 2016) approximates the environment dynamics, uses the information gain of the learned dynamics model as intrinsic rewards, and showed encouraging results on robotic locomotion problems. However, the method needs to update the dynamics model per each observation and is unlikely to be scalable for complex tasks with high dimensional states such as Atari games.

RND (Burda et al., 2018) trains a network to predict the output of a fixed randomly initialized target network and uses the prediction error as the intrinsic reward but the method does not report the results on continuous control tasks. ICM (Pathak et al., 2017) transforms the high dimensional states to feature space and imposes cross entropy and Euclidean loss so the action and the feature of the next state are predictable. However, ICM does not utilize mutual information like VIME to directly measure the uncertainty and is limited to discrete actions. Our method (EMI) is also reminiscent of (Kohonen & Somervuo, 1998) in the sense that we seek to construct a decoder-free latent space from the high dimensional observation data with a topology in the latent space. In contrast to the prior works on exploration, we seek to construct the representation under linear topology and does not require decoding the full observation but seek to encode the essential predictive signal that can be used for guiding the exploration.

3. Preliminaries

We consider a Markov decision process defined by the tuple $(\mathcal{S}, \mathcal{A}, P, r, \gamma)$, where \mathcal{S} is the set of states, \mathcal{A} is the set of actions, $P : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}_+$ is the environment transition distribution, $r : \mathcal{S} \rightarrow \mathbb{R}$ is the reward function, and $\gamma \in (0, 1)$ is the discount factor. Let π denote a stochastic policy over actions given states. Denote $\mathbb{P}_0 : \mathcal{S} \rightarrow \mathbb{R}_+$ as the distribution of initial state s_0 . The discounted sum of

expected rewards under the policy π is defined by

$$\eta(\pi) = \mathbb{E}_\tau \left[\sum_{t=0}^{\infty} \gamma^t r(s_t) \right],$$

where $\tau = (s_0, a_0, \dots, a_{T-1}, s_T)$ denotes the trajectory, $s_0 \sim \mathbb{P}_0(s_0)$, $a_t \sim \pi(a_t | s_t)$, and $s_{t+1} \sim P(s_{t+1} | s_t, a_t)$. The objective in policy based reinforcement learning is to search over the space of parameterized policies (*i.e.* neural network) $\pi_\theta(a | s)$ in order to maximize $\eta(\pi_\theta)$.

Also, denote $\mathbb{P}_{SAS'}^\pi$ as the joint probability distribution of singleton experience tuples (s, a, s') starting from $s_0 \sim \mathbb{P}_0(s_0)$ and following the policy π . Furthermore, define $\mathbb{P}_A^\pi = \int_{S \times S'} d\mathbb{P}_{SAS'}^\pi$ as the marginal distribution of actions, $\mathbb{P}_{S'}^\pi = \int_A d\mathbb{P}_{SAS'}^\pi$ as the marginal distribution of states and the corresponding next states, $\mathbb{P}_{S \times A}^\pi = \int_{S'} d\mathbb{P}_{SAS'}^\pi$ as the marginal distribution of the next states, and $\mathbb{P}_{SA}^\pi = \int_{S'} d\mathbb{P}_{SAS'}^\pi$ as the marginal distribution of states and the actions following the policy π .

4. Methods

Our goal is to construct the embedding representation of the observation and action (discrete or continuous) for complex dynamical systems that does not rely on generative decoding of the full observation, but still provides a useful predictive signal that can be used for exploration. This requires a representation where functionally similar states are close together, and functionally distinct states are far apart. We approach this objective from the standpoint of maximizing mutual information under several criteria.

4.1. Mutual information maximizing state and action embedding representations

In this subsection, we introduce the desiderata for our objective and discuss the variational divergence lower bound for efficient computation of the objective. We denote the embedding function of states $\phi_\alpha : \mathcal{S} \rightarrow \mathbb{R}^d$ and actions $\psi_\beta : \mathcal{A} \rightarrow \mathbb{R}^d$ with parameters α and β (*i.e.* neural networks) respectively. We seek to learn the embedding function of states (ϕ_α) and actions (ψ_β) satisfying the following two criteria:

1. Given the embedding representation of states and the actions $[\phi_\alpha(s); \psi_\beta(a)]$, the uncertainty of the embedding representation of the corresponding next states $\phi_\alpha(s')$ should be minimal and vice versa.
2. Given the embedding representation of states and the corresponding next states $[\phi_\alpha(s); \phi_\alpha(s')]$, the uncertainty of the embedding representation of the corresponding actions $\psi_\beta(a)$ should also be minimal and vice versa.

Intuitively, the first criterion translates to maximizing the mutual information between $[\phi_\alpha(s); \psi_\beta(a)]$, and $\phi_\alpha(s')$

which we define as $\mathcal{I}_S(\alpha, \beta)$ in Equation (1). And the second criterion translates to maximizing the mutual information between $[\phi_\alpha(s); \phi_\alpha(s')]$ and $\psi_\beta(a)$ defined as $\mathcal{I}_A(\alpha, \beta)$ in Equation (2).

$$\begin{aligned} \text{maximize}_{\alpha, \beta} \mathcal{I}_S(\alpha, \beta) &:= \mathcal{I}([\phi_\alpha(s); \psi_\beta(a)]; \phi_\alpha(s')) \\ &= \mathcal{D}_{\text{KL}}(\mathbb{P}_{SAS'}^\pi \| \mathbb{P}_{SA}^\pi \otimes \mathbb{P}_{S'}^\pi) \end{aligned} \quad (1)$$

$$\begin{aligned} \text{maximize}_{\alpha, \beta} \mathcal{I}_A(\alpha, \beta) &:= \mathcal{I}([\phi_\alpha(s); \phi_\alpha(s')]; \psi_\beta(a)) \\ &= \mathcal{D}_{\text{KL}}(\mathbb{P}_{SAS'}^\pi \| \mathbb{P}_{S'}^\pi \otimes \mathbb{P}_A^\pi) \end{aligned} \quad (2)$$

Mutual information is not bounded from above and maximizing mutual information is notoriously difficult to compute in high dimensional settings. Motivated by (Hjelm et al., 2018; Belghazi et al., 2018), we compute the variational divergence lower bound of mutual information (Nowozin et al., 2016). Concretely, variational divergence (f -divergence) representation is a tight estimator for the mutual information of two random variables X and Z , derived as in Equation (3).

$$\begin{aligned} \mathcal{I}(X; Z) &= \mathcal{D}_{\text{KL}}(\mathbb{P}_{XZ} \| \mathbb{P}_X \otimes \mathbb{P}_Z) \\ &\geq \sup_{\omega \in \Omega} \mathbb{E}_{\mathbb{P}_{XZ}} T_\omega(x, z) - \log \mathbb{E}_{\mathbb{P}_X \otimes \mathbb{P}_Z} \exp(T_\omega(x, z)), \end{aligned} \quad (3)$$

where $T_\omega : \mathcal{X} \times \mathcal{Z} \rightarrow \mathbb{R}$ is a differentiable transform with parameter ω . Furthermore, for better numerical stability, we utilize a different measure between the joint and marginals than the KL-divergence. In particular, we employ Jensen-Shannon divergence (JSD) (Hjelm et al., 2018) which is bounded both from below and above by 0 and $\log(4)$ ¹.

Theorem 1. *The lower bound of mutual information using Jensen-Shannon divergence is*

$$\begin{aligned} \mathcal{I}^{(\text{JSD})}(X; Z) &\geq \sup_{\omega \in \Omega} \mathbb{E}_{\mathbb{P}_{XZ}} [-\text{sp}(-T_\omega(x, z))] \\ &\quad - \mathbb{E}_{\mathbb{P}_X \otimes \mathbb{P}_Z} [\text{sp}(T_\omega(x, z))] + \log(4) \end{aligned}$$

Proof.

$$\begin{aligned} \mathcal{I}^{(\text{JSD})}(X; Z) &= \mathcal{D}_{\text{JSD}}(\mathbb{P}_{XZ} \| \mathbb{P}_X \otimes \mathbb{P}_Z) \\ &\geq \sup_{\omega \in \Omega} \mathbb{E}_{\mathbb{P}_{XZ}} [S_\omega(x, z)] - \mathbb{E}_{\mathbb{P}_X \otimes \mathbb{P}_Z} [\text{JSD}^*(S_\omega(x, z))] \\ &= \sup_{\omega \in \Omega} \mathbb{E}_{\mathbb{P}_{XZ}} [-\text{sp}(-T_\omega(x, z))] \\ &\quad - \mathbb{E}_{\mathbb{P}_X \otimes \mathbb{P}_Z} [\text{sp}(T_\omega(x, z))] + \log(4), \end{aligned}$$

where the inequality in the second line holds from the definition of f -divergence (Nowozin et al., 2016). In the third line, we substituted $S_\omega(x, z) = \log(2) - \log(1 + \exp(-T_\omega(x, z)))$ and Fenchel conjugate of Jensen-Shannon divergence, $\text{JSD}^*(t) = -\log(2 - \exp(t))$. \square

¹In (Nowozin et al., 2016), the authors derive the lower bound of $D_{\text{JSD}} = D_{\text{KL}}(P||M) + D_{\text{KL}}(Q||M)$, instead of $D_{\text{JSD}} = \frac{1}{2}(D_{\text{KL}}(P||M) + D_{\text{KL}}(Q||M))$, where $M = \frac{1}{2}(P + Q)$.




Figure 2: Computational architecture for estimating $\mathcal{I}_S^{(\text{JSD})}$ and $\mathcal{I}_A^{(\text{JSD})}$ for image-based observations.

From Theorem 1, we have,

$$\begin{aligned} & \underset{\alpha, \beta}{\text{maximize}} \mathcal{I}_S^{(\text{JSD})}(\alpha, \beta) \\ & \geq \underset{\alpha, \beta}{\text{maximize}} \sup_{\omega_S \in \Omega_S} \mathbb{E}_{\mathbb{P}_{SAS'}^\pi} [-\text{sp}(-T_{\omega_S}(\phi_\alpha(s), \psi_\beta(a), \phi_\alpha(s'))))] \\ & - \mathbb{E}_{\mathbb{P}_{SA}^\pi \otimes \mathbb{P}_{S'}^\pi} [\text{sp}(T_{\omega_S}(\phi_\alpha(s), \psi_\beta(a), \phi_\alpha(s')))] + \log 4, \end{aligned} \quad (4)$$

$$\begin{aligned} & \underset{\alpha, \beta}{\text{maximize}} \mathcal{I}_A^{(\text{JSD})}(\alpha, \beta) \\ & \geq \underset{\alpha, \beta}{\text{maximize}} \sup_{\omega_A \in \Omega_A} \mathbb{E}_{\mathbb{P}_{SAS'}^\pi} [-\text{sp}(-T_{\omega_A}(\phi_\alpha(s), \psi_\beta(a), \phi_\alpha(s'))))] \\ & - \mathbb{E}_{\mathbb{P}_{S\tilde{S}'}^\pi \otimes \mathbb{P}_{\tilde{A}}^\pi} [\text{sp}(T_{\omega_A}(\phi_\alpha(s), \psi_\beta(\tilde{a}), \phi_\alpha(s')))] + \log 4, \end{aligned} \quad (5)$$

where $\text{sp}(z) = \log(1 + \exp(z))$. The expectations in Equation (4) and Equation (5) are approximated using the empirical samples trajectories τ . Note, the samples $s' \sim \mathbb{P}_{S'}^\pi$, and $\tilde{a} \sim \mathbb{P}_{\tilde{A}}^\pi$ from the marginals are obtained by dropping (s, a) and (s, s') in samples (s, a, s') and (s, \tilde{a}, s') from $\mathbb{P}_{SAS'}^\pi$. Figure 2 illustrates the computational architecture for estimating the lower bounds on \mathcal{I}_S and \mathcal{I}_A .

4.2. Embedding the linear dynamics model with the error model

Since the embedding representation space is learned, it is natural to impose a topology on it (Kohonen, 1983). In EMI, we impose a simple and convenient topology where transitions are linear since this spares us from having to also represent a complex dynamical model. This allows us to offload most of the modeling burden onto the embedding function itself, which in turn provides us with a useful and informative measure of surprise when visiting novel states. Once the embedding representations are learned, this linear dynamics model allows us to measure *surprise* in terms of the residual error under the model or measure *diversity* in terms of the similarity in the embedding space. Section 4.3 discusses the intrinsic reward computation procedure in more detail.

Concretely, we seek to learn the representation of states $\phi(s)$ and the actions $\psi(a)$ such that the representation of the corresponding next state $\phi(s')$ follow linear dynamics *i.e.* $\phi(s') = \phi(s) + \psi(a)$. Intuitively, we would like the nonlinear aspects of the dynamics to be offloaded to the neural networks $\phi(\cdot), \psi(\cdot)$ so that in the \mathbb{R}^d embedding space, the dynamics become linear. Regardless of the expressivity of the neural networks, however, there always exists irreducible error under the linear dynamic model. For example, the state transition which leads the agent from one room to another in Atari environments (*i.e.* Venture, Montezuma's revenge, *etc.*) or the transition leading the agent in the same position under certain actions (*i.e.* Agent bumping into a wall when navigating a maze) would be extremely challenging to explain under the linear dynamics model.

To this end, we introduce the error model $S_\gamma : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^d$, which is another neural network taking the state and action as input, estimating the irreducible error under the linear model. Motivated by the work of Candès et al. (2011), we seek to minimize Frobenius norm of the error term so that the error term contributes on sparingly unexplainable occasions. Equation (6) shows the embedding learning problem under linear dynamics with modeled errors.

$$\begin{aligned} & \underset{\alpha, \beta, \gamma}{\text{minimize}} \underbrace{\|S_\gamma\|_{2,0}}_{\text{error minimization}} \\ & \text{subject to } \underbrace{\Phi'_\alpha = \Phi_\alpha + \Psi_\beta + S_\gamma}_{\text{embedding linear dynamics}}, \end{aligned} \quad (6)$$

where we used the matrix notation for compactness. $\Phi_\alpha, \Psi_\beta, S_\gamma$ denotes the matrices of respective embedding representations stacked columns wise. Relaxing the matrix $\|\cdot\|_{2,0}$ norm with Frobenius norm, Equation (7) shows our final learning objective.

$$\begin{aligned} \underset{\alpha, \beta, \gamma}{\text{minimize}} \quad & \|\Phi'_\alpha - (\Phi_\alpha + \Psi_\beta + S_\gamma)\|_F^2 \\ & + \lambda_{\text{error}} \|S_\gamma\|_F^2 + \lambda_{\text{info}} \mathcal{L}_{\text{info}}, \end{aligned} \quad (7)$$

where $\mathcal{L}_{\text{info}}$ denotes the following mutual information term.

$$\begin{aligned} \mathcal{L}_{\text{info}} = \quad & \inf_{\omega_S \in \Omega_S} \mathbb{E}_{\mathbb{P}_{SAS'}^\pi} \text{sp}(-T_{\omega_S}(\phi_\alpha(s), \psi_\beta(a), \phi_\alpha(s'))) \\ & + \mathbb{E}_{\mathbb{P}_{SA}^\pi \otimes \mathbb{P}_{S'}^\pi} \text{sp}(T_{\omega_S}(\phi_\alpha(s), \psi_\beta(a), \phi_\alpha(\tilde{s}'))) \\ & + \inf_{\omega_A \in \Omega_A} \mathbb{E}_{\mathbb{P}_{SAS'}^\pi} \text{sp}(-T_{\omega_A}(\phi_\alpha(s), \psi_\beta(a), \phi_\alpha(s'))) \\ & + \mathbb{E}_{\mathbb{P}_{SA}^\pi \otimes \mathbb{P}_A^\pi} \text{sp}(T_{\omega_A}(\phi_\alpha(s), \psi_\beta(\tilde{a}), \phi_\alpha(s'))) \end{aligned}$$

$\lambda_{\text{error}}, \lambda_{\text{info}}$ are hyperparameters which control the relative contributions of the linear dynamics error and the mutual information term. In practice, for image-based experiments, we found the optimization process to be more stable when we further regularize the distribution of action embedding representation to follow a predefined prior distribution. Concretely, we regularize the action embedding distribution to follow a standard normal distribution via $\mathcal{D}_{\text{KL}}(\mathbb{P}_{\psi}^\pi \parallel \mathcal{N}(0, I))$ similar to VAEs (Kingma & Welling, 2013). Intuitively, this has the effect of grounding the distribution of action embedding representation (and consequently the state embedding representation) across different iterations of the learning process.

Note, regularizing the distribution of *state* instead of *action* embeddings renders the optimization process much more unstable. This is because the distribution of states are much more likely to be skewed than the distribution of actions, especially during the initial stage of optimization, so the Gaussian approximation becomes much less accurate in contrast to the distribution of actions. In Section 5.5, we compare the state and action embeddings as regularization targets in terms of the quality of the learned embedding functions.

4.3. Intrinsic reward augmentation

We consider a formulation based on the prediction error under the linear dynamics model as shown in Equation (8). This formulation incorporates the error term and makes sure we differentiate the irreducible error that does not contribute as the novelty.

$$r_e(s_t, a_t, s'_t) = \|\phi(s_t) + \psi(a_t) + S(s_t, a_t) - \phi(s'_t)\|^2 \quad (8)$$

Algorithm 1 shows the complete procedure in detail. The choice of different intrinsic reward formulation and the computation of $\mathcal{L}_{\text{info}}$ are fully described in supplementary Section 2 and 3.

Algorithm 1 Exploration with mutual information state and action embeddings (EMI)

initialize $\alpha, \beta, \gamma, \omega_A, \omega_S$
for $i = 1, \dots, \text{MAXITER}$ **do**
 Collect samples $\{(s_t, a_t, s'_t)\}_{t=1}^n$ with policy π_θ
 Compute prediction error intrinsic rewards $\{r_e(s_t, a_t, s'_t)\}_{t=1}^n$ following Equation (8)
for $j = 1, \dots, \text{OPTITER}$ **do**
for $k = 1, \dots, \lfloor \frac{n}{m} \rfloor$ **do**
 Sample a minibatch $\{(s_{t_l}, a_{t_l}, s'_{t_l})\}_{l=1}^m$
 Update $\alpha, \beta, \gamma, \omega_A, \omega_S$ using the Adam update rule to minimize Equation (7)
end for
end for
 Augment the intrinsic rewards with environment reward r_{env} as $r = r_{\text{env}} + \eta r_e$ and update the policy network π_θ using any RL method
end for

5. Experiments

We compare the experimental performance of EMI to recent prior works on both low-dimensional locomotion tasks with continuous control from rllab benchmark (Duan et al., 2016) and the complex vision-based tasks with discrete control from the Arcade Learning Environment (Bellemare et al., 2013). For the locomotion tasks, we chose SwimmerGather and SparseHalfCheetah environments for direct comparison against the prior work of (Fu et al., 2017). SwimmerGather is a hierarchical task where a two-link robot needs to reach green pellets, which give positive rewards, instead of red pellets, which give negative rewards. SparseHalfCheetah is a challenging locomotion task where a cheetah-like robot does not receive any rewards until it moves 5 units in one direction.

For vision-based tasks, we selected Freeway, Frostbite, Venture, Montezuma’s Revenge, Gravitar, and Solaris for comparison with recent prior works (Pathak et al., 2017; Fu et al., 2017; Burda et al., 2018). These six Atari environments feature very sparse reward feedback and often contain many moving distractor objects which can be challenging for the methods that rely on explicit decoding of the full observations (Oh et al., 2015). Table 1 shows the overall performance of EMI compared to the baseline methods in all tasks.

5.1. Implementation Details

We compare all exploration methods using the same RL procedure, in order to provide a fair comparison. Specifically, we use TRPO (Schulman et al., 2015), a policy gradient method that can be applied to both continuous and discrete action spaces. Although the absolute performance on each task depends strongly on the choice of RL algorithm, com-

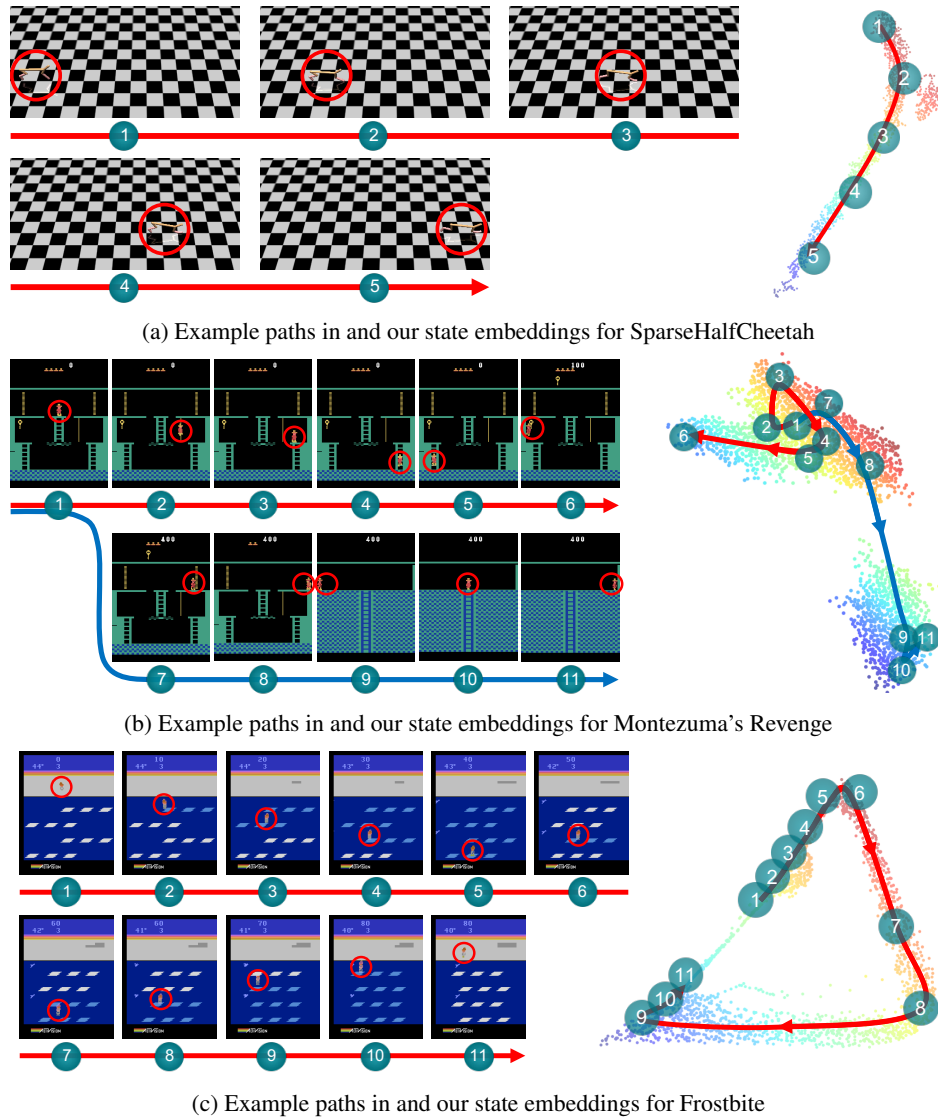


Figure 3: Example sample paths in our learned embedding representations. Note the embedding dimensionality d is 2, and thus we did not use any dimensionality reduction techniques.

paring the different methods with the same RL procedure allows us to control for this source of variability. Also, we observed TRPO is less sensitive to changes in hyperparameters than A3C (see Mnih et al. (2016)) making the comparisons easier.

In the locomotion experiments, we use a 2-layer fully connected neural network as the policy network. In the Atari experiments, we use a 2-layer convolutional neural network followed by a single layer fully connected neural network. We convert the 84×84 input RGB frames to grayscale images and resize them to 52×52 images following the practice in Tang et al. (2017). The embedding dimensionality is set to $d = 2$ and intrinsic reward coefficient is set to $\eta = 0.001$ in all of the environments. We use Adam (Kingma & Ba, 2015) optimizer to train embedding net-

works. Please refer to supplementary Section 1 for more details.

5.2. Locomotion tasks with continuous control

We compare EMI with TRPO (Schulman et al., 2015), EX2 (Fu et al., 2017), ICM (Pathak et al., 2017) and RND (Burda et al., 2018) on two challenging locomotion environments: SwimmerGather and SparseHalfCheetah. Figures 4a and 4b shows that EMI outperforms all baseline methods on both tasks. Figure 3a visualizes the scatter plot of the learned state embeddings and an example trajectory for the SparseHalfCheetah experiment. The figure shows that the learned representation successfully preserves the similarity in observation space.

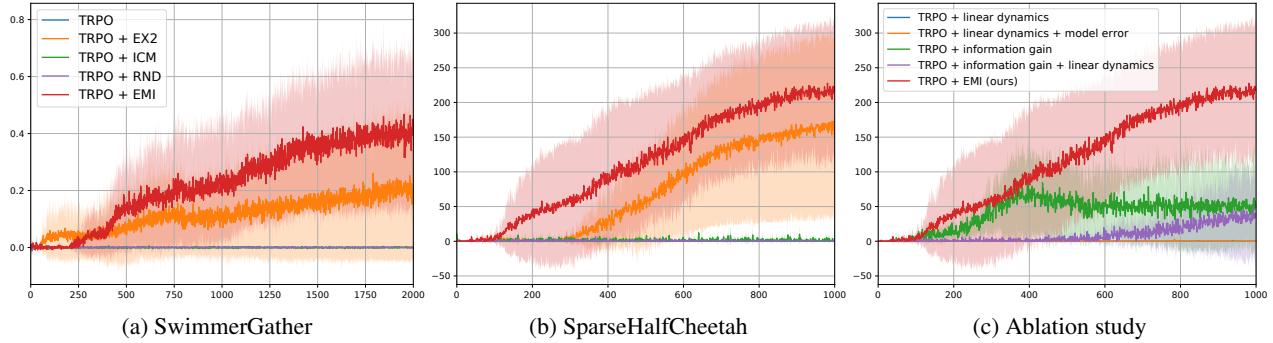


Figure 4: (a), (b): Performance of EMI on locomotion tasks with sparse rewards compared to the baseline methods. The solid lines show the mean reward (y-axis) of 5 different seeds at each iteration (x-axis) and the shaded area represents one standard deviation from the mean. (c): Ablation result on SparseHalfCheetah. Each iteration represents 50K time steps for SwimmerGather and 5K time steps for SparseHalfCheetah.

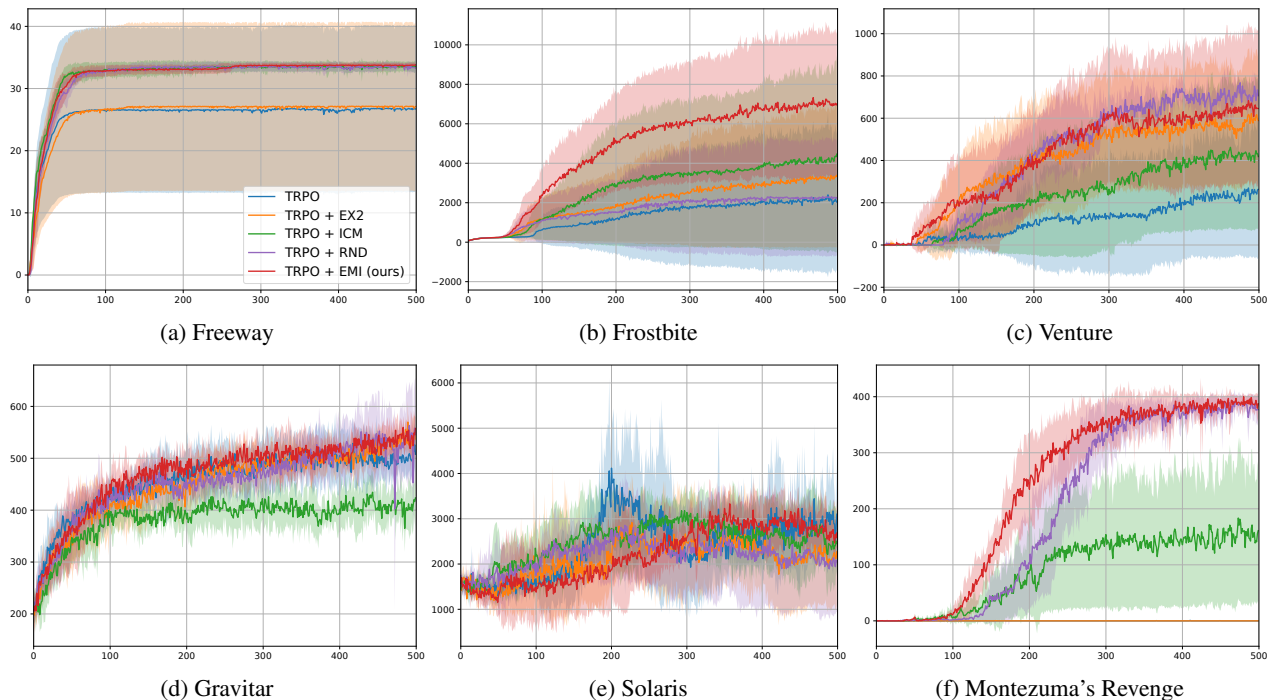


Figure 5: Performance of EMI on sparse reward Atari environments compared to the baseline methods. The solid lines show the mean reward (y-axis) of 5 different seeds at each iteration (x-axis). Each iteration represents 100K time steps.

5.3. Vision-based tasks with discrete control

For vision-based exploration tasks, our results in Figure 5 show that EMI significantly outperforms the TRPO, EX2, ICM baselines on Frostbite and Montezuma's Revenge, and show competitive performance against RND. Figures 3b and 3c illustrate our learned state embeddings ϕ . Since our embedding dimensionality is set to $d = 2$, we directly visualize the scatter plot of the embedding representation in 2D. Figure 3b shows that the embedding space naturally separates state samples into two clusters each of which

corresponds to different rooms in Montezuma's revenge. Figure 3c shows smooth sample transitions along the embedding space in Frostbite where functionally similar states are close together and distinct states are far apart.

5.4. Ablation study

We perform an ablation study showing the effect of removing each term in the objective in Equation (7) on SparseHalfCheetah. First, removing the information gain term collapses the embedding space and the agent fails to get

	EMI	EX2	ICM	RND	AE-SimHash	VIME	TRPO
SwimmerGather	0.438	0.200	0	0	0.258	0.196	0
SparseHalfCheetah	218.1	153.7	1.4	3.4	0.5	98.0	0
Freeway	33.8	27.1	33.6	33.3	33.5	-	26.7
Frostbite	7002	3387	4465	2227	5214	-	2034
Venture	646	589	418	707	445	-	263
Gravitar	558	550	424	546	482	-	508
Solaris	2688	2276	2453	2051	4467	-	3101
Montezuma	387	0	161	377	75	-	0

Table 1: Mean reward comparison of baseline methods. We compare EMI with EX2 (Fu et al., 2017), ICM (Pathak et al., 2017), RND (Burda et al., 2018), AE-SimHash (Tang et al., 2017), VIME (Houthoofd et al., 2016), and TRPO (Schulman et al., 2015). The EMI, EX2, ICM, RND, and TRPO columns show the mean reward of 5 different seeds consistent with the settings in Figure 4 and Figure 5. The AE-SimHash and VIME columns show the results from the original papers. All methods in the table are implemented based on TRPO policy. The results of MuJoCo experiments are reported at 5M and 100M time steps respectively. The results of Atari experiments are reported at 50M time steps.

any rewards as shown in Figure 4c. Also, we observed that adding the model error term (Purple versus Red in the figure) shows drastic performance improvement. We observed that modeling the linear dynamics error helps stabilize the embedding learning process during training. Please refer to supplementary Section 4, 5, and 6 for further analyses.

5.5. Regularization of embedding distributions

In order to visually examine the learned embedding representations, we designed a simple image-based 2D environment which we call *BoxImage*.




Figure 6: Example observations from BoxImage. White agent moves inside the black box.

In BoxImage, the agent exists at a position with real-valued coordinates and moves by performing actions in a confined 2D space. Then the agent receives the top-down view of the environment as image states (examples shown in Figure 6). For the implementation details, please refer to supplementary Section 7.

When the state embedding function $\phi : \mathbb{R}^{52 \times 52} \rightarrow \mathbb{R}^2$, and the action embedding function $\psi : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ are trained with the regularization on the *action* embedding distribution with $D_{\text{KL}}(\mathbb{P}_{\psi}^{\pi} \parallel \mathcal{N}(0, I))$, the learned embedding representations successfully represent the distributions of the agent’s 2D positions and actions, as shown in Figure 7. On the other hand, employing the regularization on the *state* embedding distribution with $D_{\text{KL}}(\mathbb{P}_{\phi}^{\pi} \parallel \mathcal{N}(0, I))$ results in severe degradation in the embedding quality, mainly due to the skewness of the state sample distribution.

6. Conclusion

We presented EMI, a practical exploration method that does not rely on the direct generation of high dimensional observations and instead extracts the predictive signal that can




Figure 7: (Left) Agent’s actual 2D positions and actions at the top and bottom respectively. (Center) Learned state and action embeddings when the *action* embedding is regularized. (Right) Learned state and action embeddings when the *state* embedding is regularized.

be used for exploration within a compact representation space. Our results on challenging robotic locomotion tasks with continuous actions and high dimensional image-based games with sparse rewards show that our approach transfers to a wide range of tasks. As future work, we would like to explore utilizing the learned linear dynamic model for optimal planning in the embedding representation space. In particular, we would like to investigate how an optimal trajectory from a state to a given goal in the embedding space under the linear representation topology translates to the optimal trajectory in the observation space under complex dynamical systems.

Acknowledgements

This work was partially supported by Samsung Advanced Institute of Technology and Institute for Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No.2019-0-01367, BabyMind). Hyun Oh Song is the corresponding author.

References

- Belghazi, I., Rajeswar, S., Baratin, A., Hjelm, R. D., and Courville, A. Mutual information neural estimation. In *International Conference on Machine Learning*, volume 2018, 2018.
- Bellemare, M., Veness, J., and Talvitie, E. Skip context tree switching. In *International Conference on Machine Learning*, pp. 1458–1466, 2014.
- Bellemare, M., Srinivasan, S., Ostrovski, G., Schaul, T., Saxton, D., and Munos, R. Unifying count-based exploration and intrinsic motivation. In *Advances in Neural Information Processing Systems*, pp. 1471–1479, 2016.
- Bellemare, M. G., Naddaf, Y., Veness, J., and Bowling, M. The arcade learning environment: An evaluation platform for general agents. *Journal of Artificial Intelligence Research*, 47:253–279, 2013.
- Burda, Y., Edwards, H., Storkey, A., and Klimov, O. Exploration by random network distillation. *arXiv preprint arXiv:1810.12894*, 2018.
- Candès, E. J., Li, X., Ma, Y., and Wright, J. Robust principal component analysis? *Journal of the ACM (JACM)*, 58(3): 11, 2011.
- Chen, X., Duan, Y., Houthoofd, R., Schulman, J., Sutskever, I., and Abbeel, P. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. In *Advances in neural information processing systems*, pp. 2172–2180, 2016.
- Donsker, M. D. and Varadhan, S. S. Asymptotic evaluation of certain markov process expectations for large time. iv. *Communications on Pure and Applied Mathematics*, 36(2):183–212, 1983.
- Duan, Y., Chen, X., Houthoofd, R., Schulman, J., and Abbeel, P. Benchmarking deep reinforcement learning for continuous control. In *International Conference on Machine Learning*, pp. 1329–1338, 2016.
- Fu, J., Co-Reyes, J., and Levine, S. Ex2: Exploration with exemplar models for deep reinforcement learning. In *Advances in Neural Information Processing Systems*, pp. 2577–2587, 2017.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial nets. In *Advances in neural information processing systems*, pp. 2672–2680, 2014.
- Hjelm, R. D., Fedorov, A., Lavoie-Marchildon, S., Grewal, K., Trischler, A., and Bengio, Y. Learning deep representations by mutual information estimation and maximization. *arXiv preprint arXiv:1808.06670*, 2018.
- Houthoofd, R., Chen, X., Duan, Y., Schulman, J., De Turck, F., and Abbeel, P. Vime: Variational information maximizing exploration. In *Advances in Neural Information Processing Systems*, pp. 1109–1117, 2016.
- Kingma, D. P. and Ba, J. L. Adam: A method for stochastic optimization. In *Proceedings of the 3rd International Conference on Learning Representations (ICLR)*, 2015.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Kohonen, T. Representation of information in spatial maps which are produced by self-organization. In *Synergetics of the Brain*, pp. 264–273. Springer, 1983.
- Kohonen, T. and Somervuo, P. Self-organizing maps of symbol strings. *Neurocomputing*, 21(1-3):19–30, 1998.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540): 529, 2015.
- Mnih, V., Badia, A. P., Mirza, M., Graves, A., Lillicrap, T., Harley, T., Silver, D., and Kavukcuoglu, K. Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*, pp. 1928–1937, 2016.
- Mohamed, S. and Rezende, D. J. Variational information maximisation for intrinsically motivated reinforcement learning. In *Advances in neural information processing systems*, pp. 2125–2133, 2015.
- Nachum, O., Gu, S., Lee, H., and Levine, S. Near-optimal representation learning for hierarchical reinforcement learning. *arXiv preprint arXiv:1810.01257*, 2018.
- Nowozin, S., Cseke, B., and Tomioka, R. f-gan: Training generative neural samplers using variational divergence minimization. In *Advances in Neural Information Processing Systems*, pp. 271–279, 2016.
- Oh, J., Guo, X., Lee, H., Lewis, R. L., and Singh, S. Action-conditional video prediction using deep networks in atari games. In *Advances in neural information processing systems*, pp. 2863–2871, 2015.
- Oord, A. v. d., Li, Y., and Vinyals, O. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- Ostrovski, G., Bellemare, M. G., Oord, A. v. d., and Munos, R. Count-based exploration with neural density models. *arXiv preprint arXiv:1703.01310*, 2017.

- Pathak, D., Agrawal, P., Efros, A. A., and Darrell, T. Curiosity-driven exploration by self-supervised prediction. In *International Conference on Machine Learning*, volume 2017, 2017.
- Schulman, J., Levine, S., Abbeel, P., Jordan, M., and Moritz, P. Trust region policy optimization. In *International Conference on Machine Learning*, volume 2015, 2015.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- Stadie, B. C., Levine, S., and Abbeel, P. Incentivizing exploration in reinforcement learning with deep predictive models. *arXiv preprint arXiv:1507.00814*, 2015.
- Tang, H., Houthoofd, R., Foote, D., Stooke, A., Chen, X., Duan, Y., Schulman, J., DeTurck, F., and Abbeel, P. # exploration: A study of count-based exploration for deep reinforcement learning. In *Advances in Neural Information Processing Systems*, pp. 2753–2762, 2017.
- Thomas, V., Pondard, J., Bengio, E., Sarfati, M., Beaudoin, P., Meurs, M.-J., Pineau, J., Precup, D., and Bengio, Y. Independently controllable factors. *arXiv preprint arXiv:1708.01289*, 2017.
- van den Oord, A., Kalchbrenner, N., Espeholt, L., Vinyals, O., Graves, A., et al. Conditional image generation with pixelcnn decoders. In *Advances in Neural Information Processing Systems*, pp. 4790–4798, 2016.