
Predictive Coding for Locally-Linear Control

Rui Shu^{*1} Tung Nguyen^{*23} Yinlam Chow⁴ Tuan Pham²³ Khoat Than²³ Mohammad Ghavamzadeh⁵
Stefano Ermon¹ Hung Bui²

Abstract

High-dimensional observations and unknown dynamics are major challenges when applying optimal control to many real-world decision making tasks. The Learning Controllable Embedding (LCE) framework addresses these challenges by embedding the observations into a lower dimensional latent space, estimating the latent dynamics, and then performing control directly in the latent space. To ensure the learned latent dynamics are predictive of next-observations, all existing LCE approaches decode back into the observation space and explicitly perform next-observation prediction—a challenging high-dimensional task that furthermore introduces a large number of nuisance parameters (i.e., the decoder) which are discarded during control. In this paper, we propose a novel *information-theoretic* LCE approach and show theoretically that explicit next-observation prediction can be replaced with predictive coding. We then use predictive coding to develop a *decoder-free* LCE model whose latent dynamics are amenable to locally-linear control. Extensive experiments on benchmark tasks show that our model reliably learns a controllable latent space that leads to superior performance when compared with state-of-the-art LCE baselines.

1. Introduction

With the rapid growth of systems equipped with powerful sensing devices, it is important to develop algorithms that are capable of controlling systems from high-dimensional raw sensory inputs (e.g., pixel input). However, scaling stochastic optimal control and reinforcement learning (RL) methods to high-dimensional unknown environments re-

mains an open challenge. To tackle this problem, a common approach is to employ various heuristics to embed the high-dimensional observations into a lower-dimensional latent space (Finn et al., 2016; Kurutach et al., 2018; Kaiser et al., 2019). The class of Learning Controllable Embedding (LCE) algorithms (Watter et al., 2015; Banijamali et al., 2018; Hafner et al., 2018; Zhang et al., 2019; Levine et al., 2020) further supplies the latent space with a latent dynamics model to enable planning directly in latent space.

Our present work focuses on this class of LCE algorithms and takes a critical look at the prevailing heuristic used to learn the controllable latent space: next-observation prediction. To ensure that the learned embedding and latent dynamics are predictive of future observations, existing LCE algorithms introduce a decoder during training and explicitly perform next-observation prediction by decoding the predicted latent states back into the observation space. Despite its empirical success (Watter et al., 2015; Banijamali et al., 2018; Zhang et al., 2019; Levine et al., 2020), this approach suffers from two critical drawbacks that motivate the search for better alternatives: (i) it requires the model to handle the challenging task of high-dimensional prediction; (ii) it does so in a parameter-inefficient manner—requiring the use of a decoder that is discarded during control.

To address these concerns, we propose a novel information-theoretic LCE approach for learning a controllable latent space. Our contributions are as follows.

1. We characterize the quality of the learned embedding through the lens of *predictive suboptimality* and show that predictive coding (van den Oord et al., 2018) is sufficient for minimizing predictive suboptimality.
2. Based on predictive coding, we propose a simpler and parameter-efficient model that jointly learns a controllable latent space and latent dynamics specifically amenable to locally-linear controllers.
3. We conduct detailed analyses and empirically characterize how model ablation impacts the learned latent space and control performance.
4. Finally, we show that our method out-performs state-of-the-art LCE algorithms on several benchmark tasks,

^{*}Equal contribution ¹Stanford University ²VinAI ³Hanoi University of Science and Technology ⁴Google Research ⁵Facebook AI Research. Correspondence to: Rui Shu <ruishu@stanford.edu>, Tung Nguyen <v.tungnd13@vinai.io>.

demonstrating predictive coding as a superior alternative to next-observation prediction when learning controllable embeddings.

2. Background

We are interested in controlling non-linear dynamical systems of the form $s_{t+1} = f_S(s_t, u_t) + w$, over the horizon T . In this definition, $s_t \in \mathcal{S} \subseteq \mathbb{R}^{n_s}$ and $u_t \in \mathcal{U} \subseteq \mathbb{R}^{n_u}$ are the state and action of the system at time step $t \in \{0, \dots, T-1\}$, w is the Gaussian system noise, and f_S is the smooth non-linear system dynamics. We are particularly interested in the scenario in which we only have access to the high-dimensional observation $x_t \in \mathcal{X} \subseteq \mathbb{R}^{n_x}$ of each state s_t ($n_x \gg n_s$). This scenario has application in many real-world problems, such as visual-servoing (Espiau et al., 1992), in which we only observe high-dimensional images of the environment and not its underlying state. We further assume that the high-dimensional observations x have been selected such that for any arbitrary control sequence $U = \{u_t\}_{t=0}^{T-1}$, the observation sequence $\{x_t\}_{t=0}^T$ is generated by a stationary Markov process, i.e., $x_{t+1} \sim p(\cdot | x_t, u_t)$, $\forall t \in \{0, \dots, T-1\}$.¹

A common approach to control the non-linear dynamical system described above is to solve the following stochastic optimal control (SOC) problem (Shapiro et al., 2009) that minimizes the expected cumulative cost

$$\min_U L(U, p, c, x_0) := \mathbb{E} \left[\sum_{t=0}^{T-1} c(x_t, u_t) \mid p, x_0 \right], \quad (\text{SOC1})$$

where $c : \mathcal{X} \times \mathcal{U} \rightarrow \mathbb{R}_{\geq 0}$ is the immediate cost function and x_0 is the observation at the initial state s_0 . Throughout the paper, we assume that all immediate cost functions are bounded by $c_{\max} > 0$ and Lipschitz with constant $c_{\text{lip}} > 0$. One form of the immediate cost function that is particularly common in goal tracking problems is $c(x, u) = \|x - x_{\text{goal}}\|^2$, where x_{goal} is the observation at the goal state.

The application of SOC to high-dimensional environments, however, faces several challenges. Since the observations x are high-dimensional and the dynamics in the observation space $p(\cdot | x_t, u_t)$ is unknown, solving (SOC1) is often intractable as it requires solving two difficult problems: high-dimensional dynamics estimation and high-dimensional optimal control. To address these issues, the Learning Controllable Embedding (LCE) framework proposes to learn a low-dimensional latent (embedding) space $\mathcal{Z} \subseteq \mathbb{R}^{n_z}$ ($n_z \ll n_x$) and a latent state dynamics, and then perform optimal control in the latent space instead. This framework includes algorithms such as E2C (Watter et al., 2015), RCE (Banijamali et al., 2018), SOLAR (Zhang et al., 2019), and

¹One method to enable this Markovian assumption is by buffering observations (Mnih et al., 2013) for a number of time steps.

PCC (Levine et al., 2020). By learning a stochastic encoder $E : \mathcal{X} \rightarrow \mathbb{P}(\mathcal{Z})$ and latent dynamics $F : \mathcal{Z} \times \mathcal{U} \rightarrow \mathbb{P}(\mathcal{Z})$, LCE algorithms defines a new SOC in the latent space,

$$\min_U \mathbb{E} \left[L(U, F, \bar{c}, z_0) \mid E(x_0) \right], \quad (\text{SOC2})$$

where z_0 is sampled from the distribution $E(x_0)$, i.e., $z_0 \sim E(z_0 | x_0)$, and $\bar{c} : \mathcal{Z} \times \mathcal{U} \rightarrow \mathbb{R}_{\geq 0}$ is the latent cost function. By solving the much lower-dimensional (SOC2), the resulting optimal control U_2^* is then applied as a feasible solution to (SOC1) and incurs a suboptimality that depends on the choice of the encoder E and latent dynamics F .²

Although Levine et al. (2020) provided an initial theoretical characterization of this SOC suboptimality, the selection of E and F ultimately remains heuristically-driven in all previous works. These heuristics vary across different studies (Levine et al., 2020; Banijamali et al., 2018; Watter et al., 2015; Zhang et al., 2019; Hafner et al., 2018), but the primary approach employed by the existing LCE algorithms is explicit next-observation prediction. By introducing a decoder $D : \mathcal{Z} \rightarrow \mathbb{P}(\mathcal{X})$, the composition $D \circ F \circ E$ is cast as an action-conditional latent variable model; the advances in latent variable modeling (Kingma & Welling, 2013; Rezende et al., 2014; Burda et al., 2015; Johnson et al., 2016; Sohn et al., 2015) are then leveraged to train E , F , and D to perform explicit next-observation prediction by maximizing a lower bound on the log-likelihood $\ln \int D(x_{t+1} | z_{t+1}) F(z_{t+1} | z_t, u_t) E(z_t | x_t) dz_{t:t+1}$, over the dataset whose trajectories are drawn from $p(x_t, u_t, x_{t+1})$.

Next-observation prediction offers a natural way to learn a non-degenerate choice of E and F , and enjoys the merit of being empirically successful. However, it requires the introduction of a decoder D as nuisance parameter that only serves the auxiliary role of training the encoder E and latent dynamics F . The focus of our paper is whether E and F can be successfully selected via a decoder-free heuristic.

3. Information-Theoretic LCE

Existing methods instantiate (SOC2) by learning the encoder E and latent dynamics model F in conjunction with an auxiliary decoder D to explicitly perform next-observation prediction. The auxiliary decoder ensures that the learned representation *can* be used for next-observation prediction, and is discarded after the encoder and latent dynamics model are learned. Not only is this a parameter-inefficient procedure for learning (E, F) , this approach also learns (E, F) by explicitly performing the challenging high-dimensional next-observation prediction. In this section, we

²This suboptimality also depends on \bar{c} , but we assume \bar{c} to be simple, e.g., $\bar{c}(z, u) = \|z - z_{\text{goal}}\|^2$, where $z_{\text{goal}} = E(x_{\text{goal}})$. E thus subsumes the responsibility of defining a latent space that is compatible with \bar{c} . See Appendix A.1 for further justification.




Figure 1. Two high-level approaches to learn an E and F to instantiate (SOC2). One way is to explicitly introduce a decoder D and do next-observation prediction (left), whereas our method uses F as a variational device to train E via predictive coding (right).

propose an information-theoretic approach that can learn (E, F) without decoding and next-observation prediction.

3.1. Predictive Suboptimality of a Representation

Our approach exploits the observation that the sole purpose of the decoder is to ensure that the learned representation is good for next-observation prediction. In other words, the decoder is used to characterize the suboptimality of next-observation prediction when the prediction model is forced to rely on the learned representation. We refer to this concept as *predictive suboptimality* of the learned representation and formally define it as follows.

Definition 1. Let $p(x_{t+1}, x_t, u_t)$ denote the data distribution. Given an encoder $E : \mathcal{X} \rightarrow \mathcal{Z}$,³ let $q(x_{t+1} | x_t, u_t)$ denote the prediction model

$$q(x_{t+1} | x_t, u_t) \propto \psi_1(x_{t+1})\psi_2(E(x_{t+1}), E(x_t), u_t),$$

where ψ_1 and ψ_2 are expressive non-negative functions. We define the predictive suboptimality $\ell_{pred}^*(E)$ of a representation induced by E as the best-case prediction loss

$$\min_q \mathbb{E}_{p(x_{t+1}, x_t, u_t)} D_{KL} [p(x_{t+1} | x_t, u_t) || q(x_{t+1} | x_t, u_t)].$$

Importantly, the function ψ_2 should measure the compatibility of the triplet (x_{t+1}, x_t, u_t) —but is only allowed to do so via the representations $E(x_{t+1})$ and $E(x_t)$. Thus, the behavior of the representation bottleneck plays a critical role in modulating the expressivity of the model q . If E is invertible, then q is a powerful prediction model; if E is a constant, then q can do no better than marginal density estimation of $p(x_{t+1})$.

While it is possible to minimize the predictive suboptimality of E by introducing the latent dynamics model F and decoder D , and then performing next-observation prediction via $D \circ F \circ E$, our key insight is that predictive suboptimality can be bounded by the following mutual information gap (see Appendix A.2 for proof).

³For simplicity, we assume that the encoder E considered here is deterministic.

Lemma 1. Let X_{t+1} , X_t , and U_t be the random variables associated with the data distribution $p(x_{t+1}, x_t, u_t)$. The predictive suboptimality $\ell_{pred}^*(E)$ is upper bounded by the mutual information gap

$$I(X_{t+1}; X_t, U_t) - I(E(X_{t+1}); E(X_t), U_t).$$

Since $I(X_{t+1}; X_t, U_t)$ is a constant and upper bounds $I(E(X_{t+1}); E(X_t), U_t)$ by the data processing inequality, this means we can minimize the predictive suboptimality of E simply by maximizing the mutual information between the future latent state $E(X_{t+1})$ and the current latent state and action pair $(E(X_t), U_t)$ —a form of predictive coding. We denote this mutual information $\ell_{MI}(E)$ as a function of E . To maximize this quantity, we can then leverage the recent advances in variational mutual information approximation (van den Oord et al., 2018; Poole et al., 2019; Belghazi et al., 2018; Nguyen et al., 2010; Hjelm et al., 2018) to train the encoder in a decoder-free fashion.

3.2. Consistency in Prediction of the Next Latent State

A notable consequence of introducing the encoder E is that it can be paired with a latent cost function \bar{c} to define an alternative cost function in the observation space,

$$c_E(x, u) := \mathbb{E} [\bar{c}(z, u) | E(x)],$$

where z is sampled from $E(x)$.⁴ This is particularly useful for high-dimensional SOC problems, where it is difficult to prescribe a meaningful cost function *a priori* in the observation space. For example, for goal tracking problems using visuosensory inputs, prescribing the cost function to be $c(x, u) = \|x - x_{goal}\|^2$ suffers from the uninformative nature of the 2-norm in high-dimensional pixel space (Beyer et al., 1999). In the absence of a prescribed c , a natural proxy for the unknown cost function is to replace it with c_E and consider the new SOC problem,

$$\min_U L(U, p, c_E, x_0). \quad (\text{SOC1-E})$$

⁴In Sections 3.2 and 3.3, we consider the general case of the stochastic encoder in order to extend the analysis in (Levine et al., 2020). This analysis readily carries over to the limiting case when E becomes deterministic.

Assuming (SOC1-E) to be the de facto SOC problem of interest, we wish to learn an F such that the optimal control U_2^* in (SOC2) approximately solves (SOC1-E). One such consideration for the latent dynamics model would be to set F as the true latent dynamics induced by (p, E) , and we refer to such F as the one that is *consistent* with (p, E) .

Our main contribution in this section is to justify—from a control perspective—why selecting a consistent F with respect to (p, E) minimizes the suboptimality incurred from using (SOC2) as an approximation to (SOC1-E). The following lemma (see Appendix A.3 for proof) provides the suboptimality performance gap between the solutions of (SOC2) and (SOC1-E).

Lemma 2. *For any given encoder E and latent dynamics F , let U_{1-E}^* be the solution to (SOC1-E) and U_2^* be a solution to (SOC2). Then, we have the following performance bound between the costs of the control signals U_{1-E}^* and U_2^* :*

$$L(U_{1-E}^*, p, c_E, x_0) \geq L(U_2^*, p, c_E, x_0) - 2\lambda_C \cdot \sqrt{2R_C(E, F)}, \quad (1)$$

where $R_C(E, F) = \mathbb{E}_{p(x_{t+1}, x_t, u_t)}[D_{KL}(E(z_{t+1}|x_{t+1}) || (F \circ E)(z_{t+1}|x_t, u_t))]$ and $\lambda_C = T^2 c_{\max} \bar{U}$.

In Eq. 1, the expectation is over the state-action stationary distribution of the policy used to generate the training samples (uniformly random policy in this work), and \bar{U} is the Lebesgue measure of \mathcal{U} .⁵ Moreover, $E(z_{t+1}|x_{t+1})$ and $(F \circ E)(z_{t+1}|x_t, u_t) = \int F(z_{t+1}|z_t, u_t) E(z_t|x_t) dz_t$ are the probability over the next latent state z_{t+1} . Based on Figure 1, we therefore interpret $R_C(E, F)$ as the measure of discrepancy between the dynamics $x_t \rightarrow x_{t+1} \rightarrow z_{t+1}$ induced by (p, E) versus the latent dynamics model $x_t \rightarrow z_t \rightarrow z_{t+1}$ induced by (E, F) , which we term the *consistency* regularizer. We note that while our resulting bound is similar to Lemma 2 in Levine et al. (2020), there are two key differences. First, our analysis makes explicit the assumption that the cost function c is not prescribed and thus replaced in practice with the proxy cost function c_E based on the heuristically-learned encoder. Second, by making this assumption explicit, our bound is based on samples from the environment dynamics p instead of the next-observation prediction model dynamics \hat{p} as required in Levine et al. (2020).

By restricting the stochastic encoder E to be a distribution with fixed entropy (e.g., by fixing the variance if E is conditional Gaussian), the minimization of the consistency regularizer corresponds to maximizing the log-likelihood of F for predicting z_{t+1} , given (z_t, u_t) , under the dynamics induced by (p, E) . This correspondence holds even in the limiting case of E being deterministic (e.g., fixing the variance to an arbitrarily small value). In other words,

⁵In the case when sampling policy is non-uniform and has no measure-zero set, $1/\bar{U}$ is its minimum measure.

for (SOC2) to approximate (SOC1-E) well, we select F to be a good predictor of the true latent dynamics.

3.3. Suboptimality in Locally-Linear Control

In Section 3.2, we derived the suboptimality of using (SOC2) as a surrogate control objective for (SOC1-E), and showed that the suboptimality depends on the consistency of latent dynamics model F with respect to the true latent dynamics induced by (p, E) .

We now shift our attention to the optimization of (SOC2) itself. Similar to previous works (Watter et al., 2015; Banijamali et al., 2018; Zhang et al., 2019; Levine et al., 2020), we shall specifically consider the class of locally-linear control (LLC) algorithms, e.g., iLQR (Li & Todorov, 2004), for solving (SOC2). The main idea in LLC algorithms is to compute an optimal action sequence by linearizing the dynamics around some nominal trajectory. This procedure implicitly assumes that the latent dynamics F has low curvature, so that local linearization via first-order Taylor expansion yields to a good linear approximation over a sufficiently large radius. As a result, the *curvature* of F will play an important role in the optimizability of (SOC2) via LLC algorithms.

Levine et al. (2020) analyzed the suboptimality incurred from applying LLC algorithms to (SOC2) as a function of the curvature of F . For self-containedness, we summarize their analysis as follows. We shall assume F to be a conditional Gaussian model with a mean prediction function $f_{\mathcal{Z}}(z, u)$. The curvature of $f_{\mathcal{Z}}$ can then be measured via

$$R_{\text{LLC}}(F) = \mathbb{E}_{x, u, \eta} [\|f_{\mathcal{Z}}(z + \eta_z, u + \eta_u) - f_{\mathcal{Z}}(z, u) - (\nabla_z f_{\mathcal{Z}}(z, u) \cdot \eta_z + \nabla_u f_{\mathcal{Z}}(z, u) \cdot \eta_u)\|_2^2 | E].$$

where $\eta = (\eta_z, \eta_u)^\top \sim \mathcal{N}(0, \delta^2 I)$, $\delta > 0$ is a tunable parameter that characterizes the radius of latent state-action space in which the latent dynamics model should have low curvature. Let U_{LLC}^* be a LLC solution to (SOC2). Suppose the nominal latent state-action trajectory $\{(z_t, u_t)\}_{t=0}^{T-1}$ satisfies the condition: $(z_t, u_t) \sim \mathcal{N}((z_{2,t}^*, u_{2,t}^*), \delta^2 I)$, where $\{(z_{2,t}^*, u_{2,t}^*)\}_{t=0}^{T-1}$ is the optimal trajectory of (SOC2). Using Eq. 29 of Levine et al. (2020), one can show that with probability $1 - \eta$, the LLC solution of (SOC2) has the following suboptimality performance gap when compared with the optimal cost of this problem using the solution U_2^* ,

$$L(U_2^*, F, \bar{c}, z_0) \geq L(U_{\text{LLC}}^*, F, \bar{c}, z_0) - 2\lambda_{\text{LLC}} \cdot \sqrt{R_{\text{LLC}}(F)},$$

where

$$\lambda_{\text{LLC}} = T^2 c_{\max} c_{\text{lip}} (1 + \sqrt{2 \log(2T/\eta)}) \sqrt{\bar{U} \bar{X}} / 2,$$

and \bar{X} is the Lebesgue measure with respect to \mathcal{X} . We therefore additionally constrain F to have low curvature so that it is amenable to the application of LLC algorithms.

4. Predictive Coding, Consistency, Curvature

Based on the analysis in Section 3, we identify three desiderata for guiding the selection of the encoder E and latent dynamics model F . We summarize them as follows: (i) *predictive coding* minimizes the predictive suboptimality of the encoder E ; (ii) *consistency* of the latent dynamics model F with respect to (p, E) enables planning directly in the latent space; and (iii) *low-curvature* enables planning in latent space specifically using locally-linear controllers. We refer to these heuristics collectively as Predictive Coding-Consistency-Curvature (PC3). PC3 can be thought of as an information-theoretic extension of the Prediction-Consistency-Curvature (PCC) framework described by Levine et al. (2020)—differing primarily in the replacement of explicit next-observation prediction with predictive coding in the latent space.

In this section, we highlight some of the key design choices involved when instantiating PC3 in practice. In particular, we shall show how to leverage the CPC variational mutual information bound in a parameter-efficient manner and how to enforce the consistency of F with respect to (p, E) without destabilizing training.

4.1. Enforcing Predictive Codes

To estimate the mutual information $\ell_{\text{MI}}(E)$, we employ contrastive predictive coding (CPC) proposed by van den Oord et al. (2018). We perform CPC by introducing a critic $f : \mathcal{Z} \times \mathcal{X} \times \mathcal{U} \rightarrow \mathbb{R}$ to construct the lower bound

$$\begin{aligned} & I(E(X_{t+1}); E(X_t), U_t) \\ & \geq \mathbb{E} \frac{1}{K} \sum_i \ln \frac{\exp f(E(x_{t+1}^{(i)}), E(x_t^{(i)}), u_t^{(i)})}{\frac{1}{K} \sum_j \exp f(E(x_{t+1}^{(i)}), E(x_t^{(j)}), u_t^{(j)})}, \end{aligned} \quad (2)$$

where the expectation is over K i.i.d. samples of (x_{t+1}, x_t, u_t) . Notice that the current latent state-action pair $(E(x_t), u_t)$ is specifically designated as the source of negative samples and used for the contrastive prediction of the next latent state $E(x_{t+1})$. We then tie the critic f to our latent dynamics model F ,

$$\exp f(z_{t+1}, z_t, u_t) := F(z_{t+1} | z_t, u_t).$$

This particular design of the critic has two desirable properties. First, it exploits parameter-sharing to circumvent the instantiation of an auxiliary critic f . Second, it takes advantage of the property that an optimal critic for the lower bound in Eq. (2) is the true latent dynamics (Poole et al., 2019; Ma & Collins, 2018)—which we wish F to approximate. The resulting CPC objective is thus

$$\mathbb{E} \frac{1}{K} \sum_i \ln \frac{F(E(x_{t+1}^{(i)}) | E(x_t^{(i)}), u_t^{(i)})}{\frac{1}{K} \sum_j F(E(x_{t+1}^{(i)}) | E(x_t^{(j)}), u_t^{(j)})},$$

which we denote as $\ell_{\text{cpc}}(E, F)$.

4.2. Enforcing Consistency

Since the true latent dynamics is an optimal critic for the CPC bound, it is tempting to believe that optimizing (E, F) to maximize $\ell_{\text{cpc}}(E, F)$ should be sufficient to encourage the learning of a latent dynamics model F that is consistent with the true latent dynamics induced by (p, E) .

In this section, we show that it is easy to construct a simple counterexample illustrating the non-uniqueness of the true latent dynamics as an optimal critic—and that F may learn to be arbitrarily *inconsistent* with (p, E) while still maximizing $\ell_{\text{cpc}}(E, F)$ under a fixed choice of E . Our simple counterexample proceeds as follows: let E be the identity function, let $\mathcal{X} = \mathcal{U} = \mathbb{R}$, and let $p(x_{t+1}, x_t, u_t)$ be a uniform distribution over the tuples $(1, 1, 1)$ and $(-1, -1, -1)$. Let $F(z_{t+1} | z_t, u_t)$ be a conditional Gaussian distribution with learnable variance $\sigma^2 > 0$ and mean function

$$\mu(z_t, u_t) = \text{sign}(z_t) \cdot \eta,$$

where $\eta > 0$ is a learnable parameter. By symmetry, the bound $\ell_{\text{cpc}}(E, F)$ where $K = 2$ becomes

$$\ln \frac{\exp(-(\eta - 1)^2/\sigma^2)}{\exp(-(\eta - 1)^2/\sigma^2) + \exp(-(\eta + 1)^2/\sigma^2)} + \ln 2.$$

In the denominator, the first term arises from the positive sample (e.g., $(1, 1, 1)$) whereas the second term arises from the negative sample (e.g., $(1, -1, -1)$). One way to maximize this bound would be to set $\eta = 1$ and let $\sigma \rightarrow 0$. Correspondingly, F would approach the true latent dynamics and precisely predict how (z_t, u_t) transitions to z_{t+1} . However, an alternative procedure for maximizing this bound is to fix σ to any positive constant and let $\eta \rightarrow \infty$. In this scenario, F becomes an arbitrarily poor predictor of the underlying latent dynamics.

This counterexample highlights a simple but important characteristic of the CPC bound. In contrast to direct maximum likelihood training of $F(z_{t+1} | z_t, u_t)$ using samples of (z_{t+1}, z_t, u_t) from the true latent dynamics, the contrastive predictive training of the latent dynamics model simply ensures that $F(z_{t+1} | z_t, u_t)$ assigns a *relatively* much higher value to the positive samples than to the negative samples. The fact that the CPC bound may be maximized without learning a consistent dynamics model F may be why previous work by Nachum et al. (2018) using CPC for representation learning in model-free RL chose not to perform model-based latent space control despite also learning an F as a variational artifact from their CPC bound.

Since our goal is to use F in (SOC2) for optimal control, it is critical that we ensure the latent dynamics model F indeed approximates the true latent dynamics. We therefore additionally train F via the maximum likelihood objective

$$\ell_{\text{cons}}(E, F) = \mathbb{E}_{p(x_{t+1}, x_t, u_t)} \ln F(E(x_{t+1}) | E(x_t), u_t).$$

However, naively optimizing (E, F) to maximize both ℓ_{cpc} and ℓ_{cons} is unstable; whereas ℓ_{cpc} is geometry-invariant, ℓ_{cons} is sensitive to non-volume preserving transformations of the latent space (Rezende & Mohamed, 2015; Dinh et al., 2016) and can increase arbitrarily simply by collapsing the latent space. To resolve this issue, we add Gaussian noise $\epsilon \sim \mathcal{N}(0, \sigma^2 I)$ with fixed variance to the next-state encoding $E(x_{t+1})$. Doing so yields the noise-perturbed objectives $\ell_{\text{cpc}+\epsilon}$ and $\ell_{\text{cons}+\epsilon}$. The introduction of noise has two notable effects. First, it imposes an upper bound on the achievable log-likelihood

$$\ell_{\text{cons}+\epsilon}(E, F) \leq -\frac{n_z}{2} \ln 2\pi e \sigma^2$$

based on the entropy of the Gaussian noise. Second, $\ell_{\text{cpc}+\epsilon}$ is now a lower bound to the mutual information between $(E(X_t), U_t)$ and the noise-perturbed $E(X_{t+1}) + \epsilon$,

$$\begin{aligned} & I(E(X_{t+1}) + \epsilon; E(X_t), U_t) \\ & \geq \mathbb{E} \frac{1}{K} \sum_i \ln \frac{F(E(x_{t+1}^{(i)}) + \epsilon^{(i)} | E(x_t^{(i)}), u_t^{(i)})}{\frac{1}{K} \sum_j F(E(x_{t+1}^{(j)}) + \epsilon^{(j)} | E(x_t^{(j)}), u_t^{(j)})}. \end{aligned}$$

Since the noise variance σ^2 is fixed, $\ell_{\text{cpc}+\epsilon}$ can only be maximized by expanding the latent space. By tuning the noise variance σ^2 as a hyperparameter, we can balance the latent space retraction encouraged by $\ell_{\text{cons}+\epsilon}$ with the latent space expansion encouraged by $\ell_{\text{cpc}+\epsilon}$ and thus stabilize the learning of the latent space. For notational simplicity, we shall treat all subsequent mentions of ℓ_{cpc} and ℓ_{cons} to mean their respective noise-perturbed variants, except in the specific ablation conditions where noise is explicitly removed (e.g., the ‘‘w/o ϵ ’’ condition in our experiments). At this point, we wish to note that, given the introduction of ℓ_{cons} , a natural question is whether we should in fact parameter-tie the critic in ℓ_{cpc} to the latent dynamics model in ℓ_{cons} . We demonstrate the value of doing so in Appendix B.5.

4.3. Enforcing Low Curvature

We measure the curvature of F by computing the first-order Taylor expansion error incurred when evaluating at $\bar{z} = z + \eta_z$ and $\bar{u} = u + \eta_u$,

$$\begin{aligned} \ell_{\text{curv}}(F) &= \mathbb{E}_{\eta \sim \mathcal{N}(0, \delta I)} [\|f_{\mathcal{Z}}(\bar{z}, \bar{u}) - (\nabla_z f_{\mathcal{Z}}(\bar{z}, \bar{u})\eta_z \\ & \quad + \nabla_u f_{\mathcal{Z}}(\bar{z}, \bar{u})\eta_u) - f_{\mathcal{Z}}(z, u)\|_2^2]. \end{aligned}$$

Levine et al. (2020) further proposes an amortized version of this objective to accelerate training when the latent dimensionality n_z is large. However, since n_z is relatively small in our benchmark tasks, our initial experimentation suggests amortization to have little wall-clock time impact on these tasks. Our overall objective is thus

$$\max_{E, F} \lambda_1 \ell_{\text{cpc}}(E, F) + \lambda_2 \ell_{\text{cons}}(E, F) - \lambda_3 \ell_{\text{curv}}(F),$$

which maximizes the CPC bound and consistency, while minimizing curvature.

5. Experiments

In this section, we report a thorough ablation study on various components of PC3, as well as compare the performance of our proposed model⁶ with two state-of-the-art LCE baselines: PCC (Levine et al., 2020) and SOLAR (Zhang et al., 2019).⁷ The experiments are based on four image-based control benchmark domains: Planar System, Inverted Pendulum,⁸ Cartpole, and 3-Link Manipulator.

Data generation procedure: In PCC and PC3, each sample is a triplet (x_t, u_t, x_{t+1}) where we (1) sample an underlying state s_t and generate its observation x_t , (2) sample an action u_t , and (3) obtain the next state s_{t+1} from the true dynamics and generate its observation x_{t+1} . In SOLAR, each training sample is an episode $\{x_1, u_1, x_2, \dots, x_T, u_T, x_{T+1}\}$, where T is the control horizon. For fair comparison, we allow SOLAR to collect data uniformly in the state space. See Appendix B.4 for SOLAR performance under its original data generation procedure.

Evaluation metric: We evaluate PC3 and the baselines in terms of control performance. For PC3 and PCC, we apply iLQR algorithm in the latent space with a quadratic cost, $c(z_t, u_t) = (z_t - z_{\text{goal}})^\top Q (z_t - z_{\text{goal}}) + u_t^\top R u_t$, where z_t and z_{goal} are the encoded vectors of the current and goal observation, and $Q = \alpha \cdot I_{n_z}$, $R = \beta \cdot I_{n_u}$. For SOLAR, we use their original local-inference-and-control algorithm.⁹ We report the percentage of time spent in the goal region in the underlying system (Levine et al., 2020).

5.1. Ablation Study

We characterize PC3 by ablating ℓ_{cons} , ℓ_{curv} , and the noise ϵ added to z_{t+1} . For each setting, we report the latent map size,¹⁰ ℓ_{cpc} , ℓ_{cons} , ℓ_{curv} , and the control performance. These statistics are averaged over 10 different models. All settings are run on Pendulum (Balance and Swing Up).

Consistency: In Table 1, we can see that when ℓ_{cons} is omitted, the control performance drops. As discussed in Section 3.2, the latent dynamics model F performs poorly when not explicitly optimized for consistency. This is further demonstrated in Table 2, where we take a pretrained PC3 model, freeze the encoder, and retrain F to maximize either ℓ_{cpc} or ℓ_{cons} . Despite both retrained models achieving similar ℓ_{cpc} scores, it is easy to see that training via ℓ_{cpc} results in much worse latent dynamics in terms of prediction.

⁶<https://github.com/VinAIRResearch/PC3-pytorch>

⁷E2C and RCE, two closely related baselines, are not included, since they are often inferior to PCC (Levine et al., 2020).

⁸Pendulum has two separate tasks: Balance and Swing Up

⁹<https://github.com/sharadmv/parasol>

¹⁰We add the loss $\|\frac{1}{N} \sum_{i=1}^N z_i\|_2^2$ to center the latent map at the origin, then report $\frac{1}{N} \sum_{i=1}^N \|z_i\|_2^2$ as the latent map size.




Figure 2. Inverted pendulum representations. From left to right: PC3, w/o $(\ell_{\text{cons}}, \epsilon)$, w/o (ℓ_{cons}) , w/o ϵ , w/o ℓ_{curv}

Table 1. Ablation analysis. Percentage of steps spent in goal state. From top to bottom: full-model PC3, excluding consistency and latent noise, excluding consistency, excluding latent noise, excluding curvature. For each setting we report the latent map scale, CPC, consistency and curvature loss, and control performance on balance and swing.

Setting	Latent map size	ℓ_{cpc}	ℓ_{cons}	ℓ_{curv}	Balance	Swing Up
PC3	16.2	4.58	2.13	0.03	99.12 \pm 0.66	58.4 \pm 3.53
w/o $(\ell_{\text{cons}}, \epsilon)$	10.47	5.07	-4.13	0.001	34.55 \pm 3.69	17.83 \pm 2.9
w/o ℓ_{cons}	101.52	5.03	-4.87	0.0025	31.08 \pm 3.57	7.46 \pm 1.32
w/o ϵ	0.04	3.27	20.83	0.0009	65.2 \pm 1.11	0 \pm 0
w/o ℓ_{curv}	66.1	4.8	2.34	0.56	96.89 \pm 0.97	21.69 \pm 2.73

Table 2. We took a pretrained PC3 model, froze the encoder E , and then retrained only the latent dynamics model F either without ℓ_{cons} (first row) or without ℓ_{cpc} (second row). Note that we continue to use ℓ_{curv} and add ϵ noise in both settings.

Setting	Latent map size	ℓ_{cpc}	ℓ_{cons}	ℓ_{curv}	Balance	Swing Up
Retrain F w/o ℓ_{cons}	16.2	4.57	-21.93	0.02	46.77 \pm 3.66	18.06 \pm 1.87
Retrain F w/o ℓ_{cpc}	16.2	4.6	2.17	0.03	90.85 \pm 2.33	50.11 \pm 3.74

Noise: The control performance also decreases when we do not add noise to z_{t+1} . This is because the model will collapse the latent space to inflate ℓ_{cons} as shown in Table 1, leading to a degenerate solution. Adding noise to z_{t+1} prevents the map from collapsing; since the noise variance is fixed, ℓ_{cpc} is only maximized by pushing points apart. Indeed, Table 1 shows that when noise is added but ℓ_{cons} is removed, the latent map expands aggressively.

Curvature: Finally, as previously observed in (Levine et al., 2020), imposing low curvature is an important component if we want to use locally-linear control algorithms such as iLQR. Without the curvature loss, the Taylor approximation when running iLQR might not be accurate, leading to poor control performance. The right-most map in Figure 2 shows a depiction of this setting, in which the map has very sharp regions, requiring the transition function to have high-curvature to move in these regions.

5.2. Control Performance Comparison

Experimental pipeline: For each control domain, we run 10 different subtasks (different initial and/or goal states), and report the average performance among these subtasks. For PCC and PC3, we train 10 different models, and each of them will perform all 10 subtasks (which means a total

of $10 \times 10 = 100$ subtasks), and we additionally report the performance of the best model. SOLAR training procedure depends on the specific subtask (i.e., initial and goal state), and since we cannot train 100 different models due to huge computation cost, we train only 1 model for each subtask. All subtasks are shared for three methods.

Result: Table 3 shows that our proposed PC3 model significantly outperforms the baselines by comparing the means and standard error of means on the different control tasks.¹¹ PCC and SOLAR often fail at difficult tasks such as Swing Up and 3-Link. Moreover, SOLAR training procedure depends on the specific task, which makes them unsuitable to be reused for different tasks in the same environment. Figure 3 demonstrates some (randomly selected) latent maps of Planar and Inverted Pendulum domains learned by PCC and PC3. In general, PC3 produces more interpretable latent representation for Pendulum, due to the fact that next observation prediction is too conservative and may force a model to care about things that do not matter to downstream tasks. Finally, in terms of computation, PC3 enjoys huge improvements over the baselines, with $1.85\times$ faster than PCC and $52.8\times$ faster than SOLAR.

¹¹Due to huge computation cost of SOLAR, we lower control horizon for Balance, Swing Up, Cartpole and 3-Link, compared to what was used in the PCC paper.




Figure 3. Top: Planar latent representations; Bottom: Inverted Pendulum latent representations. Left three: PCC, right three: PC3.

Table 3. Percentage steps in goal state for the average model (all) and top 1 model. Since SOLAR is task-specific, it does not have top 1.

Task	PC3 (all)	PCC (all)	SOLAR (all)	PC3 (top 1)	PCC (top 1)
Planar	74.35 ± 0.76	56.6 ± 3.15	71.25 ± 1.51	75.5 ± 0.32	75.5 ± 0.32
Balance	99.12 ± 0.66	91.9 ± 1.72	55.2 ± 1.1	100 ± 0	100 ± 0
Swing Up	58.4 ± 3.53	26.41 ± 2.64	37.78 ± 1.44	84 ± 0	66.9 ± 3.8
Cartpole	96.26 ± 0.95	94.44 ± 1.34	74.8 ± 14.95	97.8 ± 1.4	97.8 ± 1.4
3-link	42.4 ± 3.23	14.17 ± 2.2	6 ± 3.79	78 ± 1.04	45.8 ± 6.4

6. Related Work

LCE Approaches. In contrast to existing LCE methods (Watter et al., 2015; Banijamali et al., 2018; Levine et al., 2020; Zhang et al., 2019; Hafner et al., 2018), our main contribution in PC3 is the development of an information-theoretic approach for minimizing the predictive suboptimality of the encoder and circumventing the need to perform explicit next-observation prediction. In particular, PC3 can be seen as a natural information-theoretic extension of PCC (Levine et al., 2020), which itself extended and improved upon E2C (Watter et al., 2015) and RCE (Banijamali et al., 2018). Compared to SOLAR (Zhang et al., 2019), PC3 (as well as PCC, RCE, and E2C) decouples the representation learning and latent dynamics estimation from control—once the encoder and latent dynamics have been learned for a particular environment, it can be used to solve many SOC problems within the same environment. In contrast, SOLAR is an online algorithm that interleaves model learning with policy optimization. Furthermore, the latent model in SOLAR is restricted to be globally linear, which can potentially impact the control performance.

Information-Theoretic Approaches. Several works have previously explored information-theoretic approaches for representation learning in the reinforcement learning context (Nachum et al., 2018; Anand et al., 2019; Lu et al., 2019). However, these works do not test the quality of their learned representations for the purposes of model-based planning in the latent space, opting instead to leverage the representations for model-free RL. This is particularly no-

table in the case of Nachum et al. (2018), who explicitly learned both an encoder E and latent dynamics model F . As we showed in Section 4.2, maximizing the CPC bound alone may not be sufficient for ensuring that F is a good predictor of the latent dynamics induced by (p, E) . Thus, the resulting (E, F) from predictive coding alone may be unsuitable for multi-step latent planning, as we demonstrate in our ablation analysis in Section 5.1.

7. Conclusion

In this work, we propose a novel information-theoretic Learning Controllable Embedding approach for handling high-dimensional stochastic optimal control. Our approach challenges the necessity of the next-observation prediction in existing LCE algorithms. We show theoretically that predictive coding is a valid alternative to next-observation prediction for learning a representation that minimizes predictive suboptimality. To instantiate information-theoretic LCE, we develop the Predictive Coding-Consistency-Curvature (PC3) model and show that PC3 is a simpler, yet more effective method than existing next-observation prediction-based LCE approaches. We also provide a thorough study on various components of the PC3 objective via ablation analysis to assist the adoption of predictive coding in future LCE research. A natural follow-up would be to study the efficacy of predictive coding when used in conjunction with other techniques in the LCE literature (e.g. latent overshooting) as well as with other controllers beyond the class of locally-linear controllers considered in our present work.

References

- Anand, A., Racah, E., Ozair, S., Bengio, Y., Côté, M.-A., and Hjelm, R. D. Unsupervised state representation learning in atari. In *Advances in Neural Information Processing Systems*, pp. 8766–8779, 2019.
- Banijamali, E., Shu, R., Ghavamzadeh, M., Bui, H., and Ghodsi, A. Robust locally-linear controllable embedding. In *Proceedings of the Twenty First International Conference on Artificial Intelligence and Statistics*, pp. 1751–1759, 2018.
- Belghazi, M., Baratin, A., Rajeshwar, S., Ozair, S., Bengio, Y., Courville, A., and Hjelm, D. Mutual information neural estimation. In *International Conference on Machine Learning*, pp. 531–540, 2018.
- Beyer, K., Goldstein, J., Ramakrishnan, R., and Shaft, U. When is “nearest neighbor” meaningful? In *International conference on database theory*, pp. 217–235. Springer, 1999.
- Burda, Y., Grosse, R., and Salakhutdinov, R. Importance weighted autoencoders. *arXiv preprint arXiv:1509.00519*, 2015.
- Dinh, L., Sohl-Dickstein, J., and Bengio, S. Density estimation using real nvp. *preprint arXiv:1605.08803*, 2016.
- Espiau, B., Chaumette, F., and Rives, P. A new approach to visual servoing in robotics. *IEEE Transactions on Robotics and Automation*, 8(3):313–326, 1992.
- Finn, C., Tan, X., Duan, Y., Darrell, T., Levine, S., and Abbeel, P. Deep spatial autoencoders for visuomotor learning. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 512–519. IEEE, 2016.
- 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.
- Hjelm, R., Fedorov, A., Lavoie-Marchildon, S., Grewal, K., Bachman, P., Trischler, A., and Bengio, Y. Learning deep representations by mutual information estimation and maximization. *preprint arXiv:1808.06670*, 2018.
- Johnson, M., Duvenaud, D., Wiltchko, A., Adams, R., and Datta, S. Composing graphical models with neural networks for structured representations and fast inference. In *Advances in neural information processing systems*, pp. 2946–2954, 2016.
- Kaiser, L., Babaeizadeh, M., Milos, P., Osinski, B., Campbell, R. H., Czechowski, K., Erhan, D., Finn, C., Koza-kowski, P., Levine, S., et al. Model-based reinforcement learning for atari. *preprint arXiv:1903.00374*, 2019.
- Kingma, D. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Kingma, D. and Welling, M. Auto-encoding variational bayes. *preprint arXiv:1312.6114*, 2013.
- Koller, D. and Friedman, N. *Probabilistic graphical models: principles and techniques*. MIT press, 2009.
- Kurutach, T., Tamar, A., Yang, G., Russell, S. J., and Abbeel, P. Learning plannable representations with causal infogan. In *Advances in Neural Information Processing Systems*, pp. 8733–8744, 2018.
- Levine, N., Chow, Y., Shu, R., Li, A., Ghavamzadeh, M., and Bui, H. Prediction, consistency, curvature: Representation learning for locally-linear control. In *International Conference on Learning Representations*, 2020. URL https://openreview.net/forum?id=BJxG_0EtDS.
- Li, W. and Todorov, E. Iterative linear quadratic regulator design for nonlinear biological movement systems. In *ICINCO (1)*, pp. 222–229, 2004.
- Lu, X., Tiomkin, S., and Abbeel, P. Predictive coding for boosting deep reinforcement learning with sparse rewards. *arXiv preprint arXiv:1912.13414*, 2019.
- Ma, Z. and Collins, M. Noise contrastive estimation and negative sampling for conditional models: Consistency and statistical efficiency. *preprint arXiv:1809.01812*, 2018.
- Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., and Riedmiller, M. Playing Atari with deep reinforcement learning. *Preprint arXiv:1312.5602*, 2013.
- Nachum, O., Gu, S., Lee, H., and Levine, S. Near-optimal representation learning for hierarchical reinforcement learning. *preprint arXiv:1810.01257*, 2018.
- Nguyen, X., Wainwright, M. J., and Jordan, M. I. Estimating divergence functionals and the likelihood ratio by convex risk minimization. *IEEE Transactions on Information Theory*, 56(11):5847–5861, 2010.
- Petrik, M., Ghavamzadeh, M., and Chow, Y. Safe policy improvement by minimizing robust baseline regret. In *Advances in Neural Information Processing Systems*, pp. 2298–2306, 2016.
- Poole, B., Ozair, S., van den Oord, A., Alemi, A., and Tucker, G. On variational bounds of mutual information. *preprint arXiv:1905.06922*, 2019.
- Rezende, D. and Mohamed, S. Variational inference with normalizing flows. *preprint arXiv:1505.05770*, 2015.

- Rezende, D., Mohamed, S., and Wierstra, D. Stochastic backpropagation and approximate inference in deep generative models. *preprint arXiv:1401.4082*, 2014.
- Shapiro, A., Dentcheva, D., and Ruszczyński, A. *Lectures on stochastic programming: modeling and theory*. SIAM, 2009.
- Sohn, K., Lee, H., and Yan, X. Learning structured output representation using deep conditional generative models. In *Advances in neural information processing systems*, pp. 3483–3491, 2015.
- van den Oord, A., Li, Y., and Vinyals, O. Representation learning with contrastive predictive coding. *preprint arXiv:1807.03748*, 2018.
- Watter, M., Springenberg, J., Boedecker, J., and Riedmiller, M. Embed to control: A locally linear latent dynamics model for control from raw images. In *Advances in neural information processing systems*, pp. 2746–2754, 2015.
- Zhang, M., Vikram, S., Smith, L., Abbeel, P., Johnson, M., and Levine, S. Solar: Deep structured latent representations for model-based reinforcement learning. In *Proceedings of the 36th International Conference on Machine Learning*, 2019.