On Learning Informative Trajectory Embeddings for Imitation, Classification and Regression

Zichang Ge* Singapore Management University Singapore zichang.ge.2023@phdcs.smu.edu.sg

> Arunesh Sinha Rutgers University New Brunswick, NJ, USA arunesh.sinha@rutgers.edu

ABSTRACT

In real-world sequential decision making tasks like autonomous driving, robotics, and healthcare, learning from observed stateaction trajectories is critical for tasks like imitation, classification, and clustering. For example, self-driving cars must replicate human driving behaviors, while robots and healthcare systems benefit from modeling decision sequences, whether or not they come from expert data. Existing trajectory encoding methods often focus on specific tasks or rely on reward signals, limiting their ability to generalize across domains and tasks.

Inspired by the success of embedding models like CLIP and BERT in static domains, we propose a novel method for embedding state-action trajectories into a latent space that captures the skills and competencies in the dynamic underlying decision-making processes. This method operates without the need for reward labels, enabling better generalization across diverse domains and tasks. Our contributions are threefold: (1) We introduce a trajectory embedding approach that captures multiple abilities from state-action data. (2) The learned embeddings exhibit strong representational power across downstream tasks, including imitation, classification, clustering, and regression. (3) The embeddings demonstrate unique properties, such as controlling agent behaviors in IQ-Learn and an additive structure in the latent space. Experimental results confirm that our method outperforms traditional approaches, offering more flexible and powerful trajectory representations for various applications. Our code is available at https://github.com/Erasmo1015/vte.

KEYWORDS

Representation Learning; Sequential Decision Making

ACM Reference Format:

Zichang Ge^{*}, Changyu Chen^{*}, Arunesh Sinha, and Pradeep Varakantham. 2025. On Learning Informative Trajectory Embeddings for Imitation, Classification and Regression. In Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Detroit, Michigan, USA, May 19 – 23, 2025, IFAAMAS, 9 pages.

This work is licensed under a Creative Commons Attribution International 4.0 License. Changyu Chen* Singapore Management University Singapore cychen.2020@phdcs.smu.edu.sg

Pradeep Varakantham Singapore Management University Singapore pradeepv@smu.edu.sg

1 INTRODUCTION

Learning from state-action trajectories is a key requirement in sequential decision-making tasks, driving applications such as imitation, classification, regression, and clustering. For instance, autonomous vehicles need to mimic human driving behavior in scenarios like lane merging, while in robotics, trajectory learning is essential for replicating complex manipulation tasks. Traditional representation learning methods, though successful in static fields like computer vision (e.g., CLIP [40]) and natural language processing (e.g., BERT [13]), often struggle to generalize to these sequential settings. In dynamic environments where trajectories unfold over time and reward signals may not always be present, the challenges of learning effective representations become more pronounced. This limitation raises a crucial question:

How can we learn informative trajectory embeddings that capture the dynamic decision-making processes driving these trajectories?

Although prior work, such as goal-conditioned learning [2] has explored state-action trajectory representation, these methods rely on external labels such as goals or rewards, limiting their applicability across diverse domains. Most other works of representation learning in MDP [8, 10, 35, 36, 38, 47, 49, 50] focus on state representation learning, losing information of action sequences. A related topic is the work on skill (or options) extraction [24] from trajectories; however, skills or options capture information only about sub-trajectories. We find that a naive average of skills found in a trajectory does not provide informative embedding of the trajectory and thus cannot reach the return of the demonstrations (see ablation experiments in Section 5.6).

In this work, we propose a novel approach that learns the informative embeddings of state-action trajectories. Our approach has two stages. First, we use a skill extractor designed for sequential decision making inspired by Jiang et al. [29]. We leverage Hierarchical State Space Models (HSSM) to extract a probability distribution of multiple possible skills from the trajectory. Next, this skill distribution is input into a shallow transformer and trained with a Variational autoencoder (VAE) [31] loss. This setup outputs a latent ability vector encapsulating the trajectory's ability level. The process resembles a VAE, where the trajectory passes through a bottleneck, retaining key information for the decoder to reconstruct the trajectory.

Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Y. Vorobeychik, S. Das, A. Nowé (eds.), May 19 – 23, 2025, Detroit, Michigan, USA. © 2025 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

^{*}Equal Contribution.

Our approach provides several key advantages over previous methods. First, our approach does not require external labels, such as rewards or goal conditions, which are typically necessary in works like Zeng et al. [51]. Second, our method encapsulates the entire trajectory's information without any extra labels such as task [24] or rewards [2]. Third, our method can effectively learn from a dataset of trajectories generated by diverse policies, extracting latent representations that distinguish the trajectories generated by different policies in the latent space. Last but not least, our latent ability vectors enable a variety of downstream tasks including *conditional imitation* learning that recovers the diverse policies that generated the dataset of trajectories, classification of trajectories, and regression tasks to predict the return.

Our experimental results highlight several important characteristics of the latent ability vector. Our experiments demonstrate that previous baselines struggle to generate meaningful representations from the dataset of trajectories generated by diverse policies, resulting in poor recovery of the diverse policies that generated these trajectories. We show similar comparison for classification, clustering, and regression. Lastly, we demonstrate its strong representational power by showing that perturbing different dimensions of the trajectory embedding vector leads to distinct behavioral changes in agents.

To summarize, our main contributions include:

- Unsupervised trajectory encoding. We present a method that effectively extracts informative embeddings of state-action trajectories without any reward/goal labels. Our approach hinges on two key design choices: (1) learning trajectory embeddings through skill abstraction (representation of sub-trajectories), and (2) leveraging a transformer to capture the temporal nature of the skill sequence.
- Diverse downstream tasks. Our learned ability vector demonstrates strong representation across tasks like data generation (via imitation learning with ability embedding), classification, clustering, and regression.
- **Disentangled representation.** Our ability embedding shows interesting properties such as different dimensions of ability vectors controlling different behaviors of the agent in the conditional imitation learned policy, and an intuitive distance structure in the trajectory embedding space.

2 RELATED WORK

Representation Learning in MDP Representation learning in Markov Decision Processes (MDPs) is centered on extracting meaningful features from unlabelled trajectories to enhance performance in downstream tasks. Prior works primarily focus on either learning state representations [8, 10, 34–36, 38, 47, 49, 50] or constructing world models [14, 22, 23, 27, 37, 41]. Recent advancements, such as GCPC [51], expand the focus from sub-trajectory representations to embeddings in trajectory space. This work leverages sequence models like GPT and BERT to encode trajectories, with the resulting trajectory representations utilized to improve subsequent policy learning in offline reinforcement learning (RL) settings. However, these approaches still depend on reward (goal) labels, which differentiates them from our proposed method (which does not require reward/goal labels). We develop a continuous state-action trajectory embedding in an unsupervised setting. This approach allows us to learn compact information representations that holistically capture the inherent policy behavior patterns present in the trajectories, in a broader context without the reliance on labeled data.

Trajectory Embedding There are several existing methods that embed trajectories. Grover et al. [20] proposed learning the trajectory embeddings using constrastive learning; however, their approach relies on labeled trajectories, whereas ours does not. Tangkaratt et al. [43] and related methods target learning the expert policies from the diverse-quality demonstrations but do not involve learning trajectory-level representations. Other methods focus on encoding the state trajectory rather than encoding state-action trajectory. Hawke et al. [25] obtains the embeddings through optical flows and some other sources to improve the model performance in autonomous driving. Gavenski et al. [17] utilizes the path signatures to automatically encode the constraints.

Imitation Learning Behavioral cloning (BC) is a basic offline imitation learning method that replicates expert actions without leveraging dynamics information. Advanced methods include GAIL [26] and its variants [4, 15, 32], which optimize policies using a GANlike [19] framework to learn from expert demonstrations in an online setting. In addition, offline imitation learning works ValueDICE [33] and its variants [3, 9, 28] focus on dynamics-aware approaches to minimize KL-divergence and integrate SAC updates.

Imitation Learning (IL) and Inverse Reinforcement Learning (IRL) techniques often assume optimal demonstrations [1, 11, 18]. The assumption often fails in many real-world scenarios. Kaiser et al. [30] analyze five sources of suboptimality, including unnecessary or incorrect actions and limited demonstration scenarios.

Thus, many works [5–7] have utilized sub-optimal demonstrations to learn an optimal policy. Some [7] utilizes ranking information among trajectories as a supervision signal, though this can be costly and error-prone. Other approaches include pre-labeling demonstrations as expert or non-expert [45], using crowd-sourced data with confidence scores [48], and bootstraps from sub-optimal demonstrations to synthesize optimality-parameterized data [12].

However, a significant gap remains in imitating behaviors across different levels of optimality, or "ability levels," whether sub-optimal or optimal. When dealing with mixtures of demonstrations of varying quality, the underlying behavior patterns that generated these demonstrations are often overlooked. Few works like Behavior Transformer (BeT) [42] learn a multi-modal policy, allowing it to reconstruct different modes of behaviors. But the policy is parameterized by mixture of Gaussian rather than Gaussian prior, without an informative representation to represent the trajectory. To address this, our approach reconstructs behavior patterns at different ability levels by learning continuous trajectory embeddings in an unsupervised manner. These embeddings, which capture the ability levels, facilitate various downstream tasks, including imitation of diverse policies, trajectory classification, and disentangled representation of behaviors.

3 PRELIMINARIES

Problem setting We consider environments represented as a Markov decision process (MDP), which is defined by a tuple



Figure 1: Illustration of VTE Framework. For the encoder, by exploiting the pretrained SE-Logit from Section 4.1, we extract the skill variable $z_{1:T}$ and the boundary variable $m_{1:T}$. These are then passed through separate MLPs, mapping $z_{1:T}$ and $m_{1:T}$ to $e_{1:T}^{z}$ and $e_{1:T}^{m}$, respectively, which are of equal size. At each time step, we concatenate these embeddings, resulting in $e_{1:T}^{z,m}$, where $e_{i}^{z,m} = \text{Concat}(e_{i}^{z}, e_{i}^{m})$. $e_{1:T}^{z,m}$ is then fed into a transformer to compute the posterior $q_{\phi}(e|\tau)$. For the decoder, the action a_{t} is predicted from the state x_{t} , conditioned on the trajectory embedding.

 $(X, \mathcal{A}, p_0, P, r, \gamma)$. X, \mathcal{A} denote state and action spaces, p_0 and $P(\mathbf{x}' | \mathbf{x}, \mathbf{a})$ represent the initial state distribution and the dynamics. The reward function is $r(\mathbf{x}, \mathbf{a}) \in \mathbb{R}$, and $\gamma \in (0, 1)$ is the discount factor. Π denotes the set of all stationary stochastic policies that map states in X to actions in \mathcal{A} .

We assume access to an offline dataset, $\mathcal{D} = \{\tau_i\}_{i=1}^N$, where each trajectory consists of a sequence of states $\mathbf{x} \in X$ and actions $\mathbf{a} \in \mathcal{A}$: $\tau_i = \{(\mathbf{x}_0, \mathbf{a}_0), (\mathbf{x}_1, \mathbf{a}_1), (\mathbf{x}_2, \mathbf{a}_2), ...\}$. The trajectories are assumed to be generated by a policy that is conditional on an ability level, \mathbf{e} :

$$\boldsymbol{x}_0 \sim p_0, \boldsymbol{e} \sim p(\boldsymbol{e}), \boldsymbol{a}_t \sim \pi(\boldsymbol{a}_t | \boldsymbol{x}_t, \boldsymbol{e}), \boldsymbol{x}_{t+1} \sim P(\boldsymbol{x}_{t+1} | \boldsymbol{a}_t, \boldsymbol{x}_t) \quad (1)$$

e is unobserved and it determines the specific policy $\pi(\cdot|\mathbf{x}, e)$ used to generate the trajectory. Our goal is to learn this latent variable, e_{τ} , also referred to as trajectory embedding for any given trajectory, τ . We utilize an offline dataset \mathcal{D} of trajectories to train our mechanisms. We further show that such informative representations of the trajectory enable various downstream tasks, including policy imitation $\pi(\cdot|\mathbf{x}, e)$, trajectory classification, and regression on the trajectory's return.

Variational Autoencoders [31] Our approach that learns the latent representation of the trajectory is inspired by the well known Variational AutoEncoder (VAE) framework. VAE has an encoder-decoder architecture, where the encoder (a neural network with weights ϕ) learns a probability distribution on a latent space given an input data point. During training, the decoder (a neural network with weights θ) is used to reconstruct the input given the latent space encoding. The learning happens by maximizing the evidence lower bound (ELBO) of the intractable log-likelihood log $p(\tau)$:

$$\log p(\tau) \ge -D_{KL}(q_{\phi}(\boldsymbol{e}|\tau) \| \boldsymbol{p}(\boldsymbol{e})) + \mathbb{E}_{q_{\phi}(\boldsymbol{e}|\tau)}[\log p_{\theta}(\tau|\boldsymbol{e})], \quad (2)$$

where \boldsymbol{e} is a sample in the latent space from the approximate posterior distribution $q_{\phi}(\boldsymbol{e}|\tau)$. In the first term, the prior p is chosen as standard Normal distribution. The second term corresponds to

a reconstruction of an observed sample generated with the likelihood $p(\tau|\boldsymbol{e})$. A common choice for the approximate posterior qis a Gaussian distribution, $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, where $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ are outputs of the encoder q_{ϕ} network (as described in [31]). Once trained, we can draw samples in the latent space and the decoder can generate samples in the space of observations.

Skill extraction via compression [29] Skill, or options learning, derives higher-level abstractions from state-action sequences, which can help compress the entire sequence. The Learning Options via Compression (LOVE) approach [29] has proven effective by modeling state-action sequences as a generative process that relies on specific latent (unobserved) variables at each time step.

In this process, the latent skill variable $z \in \mathbb{Z}$ represents the skill used at a given time step, where \mathbb{Z} is the set of possible skills. Another latent variable, $m \in \{0, 1\}$ indicates whether a new skill starts (1) or the current skill continues (0). At each step, m and z influence the hidden state s, which in turn affects the observed state x. Over time, z is influenced by the current m and the previous skill, while m depends on the previous state. For a detailed explanation of the graphical model, we refer readers to the LOVE paper.

Mathematically, the generative process for the action $a_{1:T}$ conditional on the observation $x_{1:T}$ is:

$$p(z_{1:T}, s_{1:T}, m_{1:T}, a_{1:T} | x_{1:T}) =$$

$$\prod_{t=1}^{T} p(a_t | s_t) p(m_t | s_{t-1}) p(s_t | x_t, z_t) p(z_t | x_t, z_{t-1}, m_{t-1})$$
(3)

The skill learning in LOVE is achieved by maximizing the likelihood of the sequences while penalizing the description length of the skills. Due to intractability of the above likelihood, the authors introduce a variational distribution:

$$q_{\phi} (\mathbf{z}_{1:T}, \mathbf{s}_{1:T}, \mathbf{m}_{1:T} \mid \mathbf{x}_{1:T}, \mathbf{a}_{1:T}) =$$

$$\prod_{t=1}^{T} q_{\phi} (\mathbf{m}_{t} \mid \mathbf{x}_{1:t}) q_{\phi} (\mathbf{z}_{t} \mid \mathbf{z}_{t-1}, \mathbf{m}_{t}, \mathbf{x}_{1:T}, \mathbf{a}_{1:T}) q_{\phi} (\mathbf{s}_{t} \mid \mathbf{z}_{t}, \mathbf{x}_{t})$$
(4)

Overall, this yields a model with 3 learned components: 1) A state abstraction posterior $q_{\phi}(s_t|z_t, x_t)$; 2) A termination policy $q_{\phi}(m_t|x_{1:t})$ that decides if the previous skill ends; 3) A skill posterior $q_{\phi}(z_t|z_{t-1}, m_t, x_{1:T}, a_{1:T})$ to determine the current skill z_t .

LOVE penalizes the description length of the skills to improve the quality of the learned skills. The description length of the skills is measured by:

INFOCOST
$$(\phi; p_z) = -\mathbb{E}_{\substack{\tau_1:T, \\ \boldsymbol{m}_1:T, \\ \boldsymbol{z}_1:T}} \left[\sum_{t=1}^T \log p_z(z_t) \boldsymbol{m}_t \right].$$

Combining INFOCOST with maximal likelihood objective, the authors propose to solve the following optimization problem:

$$\min_{\phi, p_z} \operatorname{InfoCost}(\phi; p_z) \quad \text{s.t. } \mathcal{L}_{\text{ELBO}}(\phi) \le C,$$
(5)

where $\mathcal{L}_{\text{ELBO}}(\phi)$ is the negated evidence lower bound of the likelihood defined by Eq. (3) (detailed description of $\mathcal{L}_{\text{ELBO}}(\phi)$ can be found in Appendix B of Jiang et al. [29]). Once solved, one can infer the skill variables $\mathbf{z}_{1:T}$ and boundary variables $\mathbf{m}_{1:T}$ given the trajectory of observations $\mathbf{x}_{1:T}$. $\mathbf{z}_{1:T}$, alone with $\mathbf{m}_{1:T}$, provides all the information about learned skills.

Imitation Learning Imitation learning (IL) aims to learn a policy for performing a task based solely on expert demonstrations, which consist of state-action trajectories without any reinforcement signals [26]. IQ-Learn [16] has been proposed as an efficient and robust imitation learning algorithm that learns a single Q-function, implicitly capturing both reward and policy. The theoretical framework of IQ-Learn builds on the classic inverse reinforcement learning (IRL) objective:

$$\max_{r \in \mathcal{R}} \min_{\pi \in \Pi} L(\pi, r) = \mathbb{E}_{\rho_E}[r(s, a)] - \mathbb{E}_{\rho_\pi}[r(s, a)] - H(\pi) - \psi(r),$$

where ρ_E and ρ_{π} denote the occupancy measures of the expert policy and the learned policy, respectively, *r* represents a learnable reward function, $H(\pi)$ refers to the entropy of policy π , and ψ is a convex reward regularizer.

The authors showed that this objective can be achieved by only maximizing the Q function in the following objective:

$$\mathcal{J}(\pi, Q) = \mathbb{E}_{\rho_E} \left[f \left(Q - \gamma \mathbb{E}_{s' \sim P(\cdot | s, a)} V^{\pi} \left(s' \right) \right) \right]$$
$$- (1 - \gamma) \mathbb{E}_{s_0 \sim \rho_0} \left[V^{\pi} \left(s_0 \right) \right],$$
$$V^{\pi}(s) = \mathbb{E}_{a \sim \pi(\cdot | s)} \left[Q(s, a) - \log \pi(a | s) \right],$$

where f is a concave function associated with the choice of the reward regularizer ψ , and ρ_0 represents the initial state distribution. For a fixed Q, the soft actor-critic (SAC) [21] update: $\max_{\pi} \mathbb{E}_{s \sim \beta, a \sim \pi(\cdot|s)} [Q(s, a) - \log \pi(a|s)]$, brings π closer to π_Q . Here, β refers to the distribution of previously sampled states. Although imitation learning algorithms such as IQ-Learn are typically used to learn a policy from the expert demonstrations, in Section 4, we discuss how these methods can, in principle, learn from demonstrations of varying quality, including non-expert data. Empirically, we show that IQ-Learn can be applied to a mixture of expert and non-expert data using the trajectory embedding inferred by our approach.

4 METHODOLOGY

Our approach has two stages, as illustrated in Figure 1: Section 4.1 covers skill extraction and Section 4.2 explains transformer usage and VAE style learning.

4.1 Learning Skills via Compression

Our main idea is based on the assumption that skills (which are temporal abstractions of the trajectory) naturally combine detailed, step-by-step information and capture the patterns within segments of the trajectory. This suggests that learning an embedding for the entire trajectory by focusing on the skill space (i.e., the space of these higher-level skills) would be easier than trying to learn it directly from the raw state-action space. To achieve this, we utilize the skill learning technique called LOVE, as introduced in the work by Jiang et al. [29].

By solving the optimization problem of Eq. (5), we can readily acquire the latent variable values, $z_{1:T}$, skill change variable values, $m_{1:T}$, the learned termination policy, $q_{\phi}(m_t|\mathbf{x}_{1:t})$, and skill posterior, $q_{\phi}(\mathbf{z}_t|\mathbf{z}_{t-1}, \mathbf{m}_t, \mathbf{x}_{1:T}, \mathbf{a}_{1:T})$ (refer to Eq. (4)). For brevity, we denote this process by $\mathbf{z}_{1:T}, \mathbf{m}_{1:T} = \text{SE}_{\phi}(\mathbf{x}_{1:T}, \mathbf{a}_{1:T})$, where SE_{ϕ} is named skill extractor.

However, we noticed that the skills we sampled were not providing enough useful information to recover the trajectories. To address this and make the most of the skill knowledge without losing any detail, we capture the full information about the distributions of z_t (the skill) and m_t (the indicator for whether a new skill starts). In the LOVE model, z_t is sampled from a categorical distribution, and m_t is sampled from a Bernoulli distribution.

Instead of just working with these samples, we aim to capture more detail by using the logit vectors that describe the underlying distributions. Specifically, we represent z_t as a vector in \mathbb{R}^l (real-valued space of length l) and m_t as a vector in \mathbb{R}^2 . This approach, with a slight abuse of notation, allows us to work with the full distributions rather than just the samples. We denote this process as $z_{1:T}$, $m_{1:T}$ = SE-Logit_{ϕ}($x_{1:T}$, $a_{1:T}$), meaning the logits are derived from the input state and action sequences.

4.2 Variational Trajectory Encoding

One naive way of producing trajectory embeddings is to apply the mean pooling operation on $z_{1:T}$:

$$\boldsymbol{e} = \frac{1}{T} \sum_{i=1}^{T} \boldsymbol{z}_i. \tag{6}$$

However, we find this embedding cannot capture information at the trajectory level as it only contains information about the skills that are present in the trajectory. Thus, we propose to learn a more informative embedding by combining not just skills but also the ordering of those skills in the trajectory. We employ a VAE to compute this embedding. Starting from Eq. (2), we show how to construct the required terms in the VAE loss. We first describe how to obtain the encoding part of VAE, i.e., computing $q_{\phi}(\boldsymbol{e}|\tau)$, which is represented using a normal distribution, $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$.

- Obtain $\boldsymbol{z}_{1:T}, \boldsymbol{m}_{1:T} = \text{SE-Logit}_{\phi}(\boldsymbol{x}_{1:T}, \boldsymbol{a}_{1:T})$
- Perform $e_{1:T}^{z} = MLP(z_{1:T})$ and $e_{1:T}^{m} = MLP(m_{1:T})$. The multilayer perceptron (MLP) converts *z* and *m* to the same size.
- Concatenate for each time step, obtaining $e_{1:T}^{z,m}$, where $e_i^{z,m} = \text{Concat}(e_i^z, e_i^m)$
- Process e^{z,m}_{1:T} via a shallow transformer, and obtain a mean pooling of its output.
- Obtain mean *μ* and standard deviation Σ by mapping the transformer's output via a fully connected neural layer.

Next, we describe the decoding part and how to optimize the reconstruction loss term $\mathbb{E}_{q_{\phi}(\boldsymbol{e}|\tau)}[\log p_{\theta}(\tau|\boldsymbol{e})]$. As a trajectory is a sequence of states and actions, we obtain the following result:

PROPOSITION 1. For any given environment, $\arg \max_{\theta} \log p_{\theta}(\tau | \boldsymbol{e}) = \arg \max_{\theta} \sum_{t} \log p_{\theta}(\boldsymbol{a}_t | \boldsymbol{x}_t, \boldsymbol{e}).$

PROOF. The proof readily follows from the fact that $p_{\theta}(\tau|\boldsymbol{e}) = p(\boldsymbol{x}_0)p_{\theta}(\boldsymbol{a}_0|\boldsymbol{x}_0, \boldsymbol{e})p(\boldsymbol{x}_1|\boldsymbol{x}_0, \boldsymbol{a}_0)p_{\theta}(\boldsymbol{a}_1|\boldsymbol{x}_1, \boldsymbol{e})\dots$ Applying the log, the product becomes a sum and only the $p_{\theta}(\boldsymbol{a}_t|\boldsymbol{x}_t, \boldsymbol{e})$ terms depend on θ .

With this result above, optimizing the reconstruction loss is exactly the same as behavioral cloning [39]. Thus, we use a neural network to parameterize the policy $p_{\theta}(a_t|x_t, e)$, and optimize $\sum_t \mathbb{E}_{q_{\phi}(e|\tau)}[\log p_{\theta}(a_t|x_t, e)]$ for both ϕ and θ . This term is optimized for ϕ using the standard reparameterization trick from VAEs.

We refer to our approach as Variational Trajectory Encoding (VTE). Compared with the embedding from the mean pooling of the skill embeddings in Eq. (6), our approach shows better performance (see Section 5.6). Furthermore, we analyze the differences between trajectory embeddings produced by mean pooling and our method in the Appendix^{*}.

4.3 Trajectory Embedding for Downstream Tasks

In this section, we utilize the trajectory embedding obtained from VTE to address different downstream tasks: (1) Imitating trajectories from different expertise level policies (e.g., Expert, Good, Bad); (2) Classifying expertise level of the trajectory; and (3) Predicting return from a trajectory, without a reward label.

4.3.1 Imitating Trajectories of Varying Abilities. In this task, we consider an offline dataset \mathcal{D} consisting of trajectories collected from a mixture of expert and non-expert policies. Recall the trajectory generation process described by Eq. (1), where \mathbf{e} captures the intrinsic *ability level* of the policy that generated this trajectory. We define this ability level more concretely (in Section 5.1) as a range of returns of the policy that generated this trajectory. Our goal is to learn these policies $\pi(\cdot | \mathbf{x}, \mathbf{e})$ that generate the trajectories.

To learn $\pi(\cdot | \mathbf{x}, \mathbf{e})$, we utilize the recent IQ-Learn framework [16] and modify it to a conditional version. The IQ-Learn approach

maintains an actor π and critic Q neural network. We make the conditional version by introducing conditions into the actor and critic neural network. We have conditional actor, $\pi(\cdot|\mathbf{x}, \mathbf{e})$, and critic, $Q(\mathbf{x}, \mathbf{a}|\mathbf{e})$. Details are provided in Algorithm 1 in the Appendix.

4.3.2 Classification and Regression. For classification task, we train a classifier based on multi-layer perceptron (MLP) to predict the ability level with the same dataset used in imitation learning task (see Section 4.3.1), i.e. we learn a mapping $f_{\text{CLS}} : \mathbb{R}^d \to \{1, ..., M\}$ from the trajectory embedding to its ability level, where *d* is the dimension of the trajectory embedding and *M* is the number of ability levels. For regression task, we train an MLP $f_{\text{REG}} : \mathbb{R}^d \to \mathbb{R}$ to predict the return of the trajectory.

Note that the labels (ability level and return) are inaccessible to our trajectory encoding algorithm; we include them here solely to evaluate embedding quality. The classifier achieves high accuracy only if our method learns high-quality trajectory embeddings, but it will fail when the embeddings are of low quality. In the worst scenario, all embeddings are identical, and the classifier has to make random guesses. The same rationale applies to the regression task.

5 EXPERIMENTS

In this section, we answer the following questions through experiments: (1) Is the learned trajectory embedding well-structured? (2) How effective is the trajectory embedding on downstream tasks? (3) Does the trajectory embedding exhibit specific properties? (4) How does a simple mean polling of skill embedding perform?

We answer question (1) in Section 5.2 and question (2) in Sections 5.3 and 5.4. Question (3) is extensively analyzed in Section 5.5, and finally an ablation experiment is done in Section 5.6 for answering question (4).

5.1 Experiment Setup

All experiments were run on NVIDIA Quadro RTX 6000 GPUs, CUDA 11.0 with Python version 3.7.12 in Pytorch 1.8.0. Hyperparameter settings are in the Appendix.

Ability Level As mentioned in Section 4.3.1, our downstream tasks are trained on a dataset with varying ability levels. We define the ability level of a policy by the average return of its collected trajectories. For instance, a low-ability policy generates trajectories with returns being (400 ± 100), while an expert-ability policy generates trajectories with returns being (2000 ± 100).

Dataset Due to lack of a public dataset for our tasks, we generated the dataset ourselves. We selected three environments from MuJoCo [44], Hopper, Walker2D, and Half-Cheetah. For each environment, we trained an RL agent using Soft Actor Critic (SAC) to the expert level, and saved checkpoints throughout the training. We then took three checkpoints to generate trajectories, corresponding to *low, medium*, and *expert* ability levels, with 300 trajectories generated per ability level. The return information is provided in Table 1. **Baselines** GCPC [51] is the closest work to ours, which utilizes the encoder-decoder transformer to encode trajectories. However, it is designed for offline RL and thus relies on rewards in the dataset. We evaluate two versions of GCPC: the original, which includes rewards, referred to as **GCPC**, and an adapted version that removes rewards from the input, referred to as **GCPC-NR** (No Reward). We also modify our VTE framework by replacing the skill extractor

^{*}Please refer to the Appendix available at https://arxiv.org/abs/2501.09327.



Figure 2: tSNE Clustering Analysis

 Table 1: Dataset returns across different ability levels for each

 environment

Environment	Low	Medium	Expert
Hopper	409.0 ±4.9	885.3±49.2	3206.8±18.1
Walker2D	3064.2 ± 43.3	4359.6 ± 51.1	5912.2 ± 25.3
Half Cheetah	2402.9 ±21.2	$\textbf{4197.6}{\pm}54.0$	$\textbf{6321.9}{\pm}61.3$

SE-Logit with an MLP, referred to as **VTE-MLP** (see Section 4). Additionally, we introduce a strong baseline where the downstream task has access to the ability levels. For the imitation learning task, we consider this baseline an upper bound, which we call **Known-Abl** (Known Ability).

Evaluation Metrics For classification, we use accuracy as the performance metric. For other downstream tasks, we measure performance by calculating the relative L2 norm of the difference between the learned and dataset returns. This is done by averaging the L2 norm difference for each learned return and its corresponding target return, and then normalizing by the target return.

5.2 VTE Generates Well-Structured Embeddings

We perform clustering analysis using tSNE [46] on the latent trajectory embeddings to verify whether different ability levels can be distinguished. As shown in Figure 2, three distinct clusters emerge in each environment, corresponding to the low, medium, and expert ability levels, confirming that our method successfully learns the latent vectors that separate ability levels. We also provide the clustering results using Principal Component Analysis (PCA) in the appendix.

5.3 VTE Enables Ability-Conditioned Imitating

We assess the imitation performance of our method across three environments: Hopper, Walker2D, and Half-Cheetah, each with three skill levels: low, medium, and expert. We extract the latent trajectory embeddings using our approach and then learn a policy conditioned on these vectors using conditional IQ-Learn. The objective is to match the target return of that trajectory as observed in data, rather than learning the optimal policy.

Table 2 summarizes the results against the baselines using the relative L2 norm loss between the average evaluation return of the policy conditioned on a trajectory embedding and the return of that trajectory in the dataset, where a lower value is better. We have bolded the best results and underlined the second-best results for each environment, with all results expressed as percentages. Our method consistently outperforms both baselines in the imitation task, achieving performance comparable to the upper bound set by Known-Abl. Despite the fact that Known-Abl and GCPC utilize additional reward information, our method surpasses GCPC across all environments and demonstrates lower error rates than Known-Abl in the Hopper environment. This demonstrates that our trajectory embedding significantly enhances imitation learning when the dataset consists of mixed ability levels.

Figure 3 also presents the evaluation of returns on Hopper. For each ability level, our method facilitates learning conditioned on the latent trajectory embeddings, achieving results close to the upper-bound of the Known-Abl. In contrast, the VTE-MLP shows instability during training.

5.4 VTE Facilitates Trajectory Classification and Regression

We present the classification results in Table 3 and the regression of rewards in Table 4. It is evident that VTE-MLP and VTE achieves 100% classification accuracy within 80 epochs. For the regression task, the different methods appear to achieve similar results.

5.5 Property of Trajectory Embedding

Perturbed Conditions. Figure 4 visually compares behaviors in the Walker2D and Hopper environments, highlighting the distinct behaviors generated when specific dimensions of the trajectory embeddings are perturbed in opposing directions. Perturbing different dimensions produces varied behaviors. The figure highlights results for two dimensions in each environment. The left and right columns display results from the Walker2D and Hopper environments. In

	Environment	Known-Abl	VTE-MLP	GCPC	GCPC-NR	VTE	
	Hopper	<u>9.2</u> ±16.7	24.6 ± 20.5	21.3 ±25.9	27.6±22.5	0.9 ±2.8	
	Walker2D	7.7 ±14.8	45.3 ± 23.1	27.8 ± 18.3	46.2 ± 23.4	13.1±20.3	
	Half Cheetah	1.3 ±1.1	$19.1 {\pm} 7.0$	$\textbf{28.7}{\pm}20.0$	$47.3{\pm}24.9$	<u>1.5</u> ±1.2	
1500 Sp 1000 500 00 100k	200k 3	3000 52000 2000 2000 2000	Mediu 100k	1111 200k	4000 Spr 2000 300k 00	100	Expert
Trainir	ig step		Training	step		100	Fraining step
		— Known-Abl	- VTE-MLP -	GCPC – GCPC	-NR – VTE		

Table 2: Relative L2 norm difference between the learned Returns and the target Returns.

Figure 3: Evaluation curve of returns on Hopper for different ability levels.



Figure 4: Overall visual comparison of change in behavior in Walker2D environment, presented in the left column, and Hopper environment presented in the right column, when a dimension of the trajectory embedding is changed. Each of the four boxes is a perturbation result on one of the 10 vector dimensions of the trajectory embedding. Inside each box there are three different rows showing a sequence of frames. First Row: The value of the dimension is decreased. Second Row: The value is not perturbed as a control group. Third Row: The value of the dimension is increased (see more description in text).

Table 3: Classification Accurac	y on Ability Levels(%)	•
---------------------------------	------------------------	---

Environment	GCPC-NR	VTE-MLP	VTE
Hopper	34.2	100.0	100.0
Walker2D	34.2	100.0	100.0
Half Cheetah	32.9	100.0	100.0

each box, we present the results of perturbing one dimension: the middle row represents the control group without perturbation, while the first and third rows show the effects of decreasing and increasing the value of that dimension, respectively.

Table 4: Relative Regression Error on Rewards(%).

Environment	GCPC-NR	VTE-MLP	VTE
Hopper	3.8	2.7	3.1
Walker2D	0.7	0.8	0.7
Half Cheetah	0.9	0.7	0.8

In the top-left box, the Walker2D agent leans forward during walking when the value decreases and falls backward when it increases. In the bottom-left box, the agent walks leisurely with a decreased value but sprints with larger body movements when the



Figure 5: Heatmap of Wasserstein distance (see definition in text) between distribution of trajectory embeddings for different ability levels.

Table 5: Range of the evaluation returns on Walker2D shown across ability levels for the returns seen in dataset, and for the conditional policies learned from embeddings of mean skill pooling and our VTE approach.

Method	Low	Medium	Expert
In Dataset	3064.2 ±43.3	4359.6 ±51.1	5912.2 ±25.3
Skill Pooling	1058.3 ± 752.3	776.0±438.5	796.9±561.7
VTE	2045.9 ± 1101.9	3782.9 ± 1043.5	5053.0 ± 1834.4

value is increased. In the top-right box, the Hopper agent bends its leg with a reduced value but keeps it rigid when the value is increased. In the bottom-right box, the agent eagerly hops forward with a decreased value but remains still with an increased value.

These results clearly demonstrate that perturbing a single dimension in the latent trajectory embedding produces contrasting behaviors, illustrating the high degree of disentanglement in our learned representations.

Measure Distance Between Polices. We compute Wasserstein distances between policies at different ability levels, as shown in Figure 5. The Wasserstein distance quantifies the minimum "cost" to transform one probability distribution into another. For each environment, we collect 10 trajectory embeddings to form an empirical distribution, representing the policy's ability level. We then calculate the Wasserstein distance between the distributions from different ability levels. For each policy, we also compare distributions within the same corresponding ability level to assess intra-level variability. The results demonstrate that distances within the same ability level are relatively small, while most distances between different ability levels are significantly larger. This indicates that the learned trajectory embeddings effectively distinguish between policies of different ability levels.

5.6 Ablation Experiment

As mentioned in Section 4.2, just a mean pooling of skills does not produce desired results. Here we present a result that shows the variation in returns (from imitating) based on mean pooled skill embedding and our VTE embeddings. The result in Table 5 on WALKER2D show that the mean pooled skill embedding shows less variation in returns (across ability levels), thereby not able to learn the returns of low, medium, and expert ability levels.

6 CONCLUSION

In conclusion, this work introduces a novel unsupervised approach for encoding state-action trajectories into informative embeddings without the need for external reward or goal labels. The method leverages hierarchical skill abstraction and a transformer and VAEbased architecture to capture the temporal dynamics of trajectory skills. The resulting informative trajectory embedding demonstrates strong representation capabilities across various downstream tasks, including imitation learning, classification, clustering, and regression. Moreover, the disentangled nature of the learned embedding allows for intuitive control of agent behaviors and meaningful differentiation in the trajectory embedding space.

7 ACKNOWLEDGEMENT

This research/project is supported by the National Research Foundation Singapore and DSO National Laboratories under the AI Singapore Programme (AISG Award No: AISG2-RP-2020-017) and the grant W911NF-24-1-0038 from the US Army Research Office.

REFERENCES

- Pieter Abbeel and Andrew Y Ng. 2004. Apprenticeship learning via inverse reinforcement learning. In Proceedings of the twenty-first international conference on Machine learning. 1.
- [2] Anurag Ajay, Yilun Du, Abhi Gupta, Joshua Tenenbaum, Tommi Jaakkola, and Pulkit Agrawal. 2022. Is conditional generative modeling all you need for decisionmaking? arXiv preprint arXiv:2211.15657 (2022).
- [3] Oleg Arenz and Gerhard Neumann. 2020. Non-adversarial imitation learning and its connections to adversarial methods. arXiv preprint arXiv:2008.03525 (2020).
- [4] Nir Baram, Oron Anschel, and Shie Mannor. 2016. Model-based adversarial imitation learning. arXiv preprint arXiv:1612.02179 (2016).
- [5] Daniel Brown, Wonjoon Goo, Prabhat Nagarajan, and Scott Niekum. 2019. Extrapolating beyond suboptimal demonstrations via inverse reinforcement learning from observations. In *International conference on machine learning*. PMLR, 783– 792.
- [6] Daniel S Brown, Wonjoon Goo, and Scott Niekum. 2020. Better-thandemonstrator imitation learning via automatically-ranked demonstrations. In

Conference on robot learning. PMLR, 330-359.

- [7] Benjamin Burchfiel, Carlo Tomasi, and Ronald Parr. 2016. Distance minimization for reward learning from scored trajectories. In Proceedings of the AAAI Conference on Artificial Intelligence, Vol. 30.
- [8] Micah Carroll, Orr Paradise, Jessy Lin, Raluca Georgescu, Mingfei Sun, David Bignell, Stephanie Milani, Katja Hofmann, Matthew Hausknecht, Anca Dragan, et al. 2022. Uni [mask]: Unified inference in sequential decision problems. Advances in neural information processing systems 35 (2022), 35365–35378.
- [9] Alex J Chan and Mihaela van der Schaar. 2021. Scalable bayesian inverse reinforcement learning. arXiv preprint arXiv:2102.06483 (2021).
- [10] Lili Chen, Kevin Lu, Aravind Rajeswaran, Kimin Lee, Aditya Grover, Misha Laskin, Pieter Abbeel, Aravind Srinivas, and Igor Mordatch. 2021. Decision transformer: Reinforcement learning via sequence modeling. Advances in neural information processing systems 34 (2021), 15084–15097.
- [11] Letian Chen, Rohan Paleja, Muyleng Ghuy, and Matthew Gombolay. 2020. Joint goal and strategy inference across heterogeneous demonstrators via reward network distillation. In Proceedings of the 2020 ACM/IEEE international conference on human-robot interaction. 659–668.
- [12] Letian Chen, Rohan Paleja, and Matthew Gombolay. 2021. Learning from suboptimal demonstration via self-supervised reward regression. In *Conference on robot learning*. PMLR, 1262–1277.
- [13] Jacob Devlin. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805 (2018).
- [14] Yiming Ding, Ignasi Clavera, and Pieter Abbeel. 2020. Mutual information maximization for robust plannable representations. arXiv preprint arXiv:2005.08114 (2020).
- [15] Justin Fu, Katie Luo, and Sergey Levine. 2017. Learning robust rewards with adversarial inverse reinforcement learning. arXiv preprint arXiv:1710.11248 (2017).
- [16] Divyansh Garg, Shuvam Chakraborty, Chris Cundy, Jiaming Song, and Stefano Ermon. 2021. Iq-learn: Inverse soft-q learning for imitation. Advances in Neural Information Processing Systems 34 (2021), 4028–4039.
- [17] Nathan Gavenski, Juarez Monteiro, Felipe Meneguzzi, Michael Luck, and Odinaldo Rodrigues. 2024. Explorative imitation learning: A path signature approach for continuous environments. In ECAI 2024. IOS Press, 1551–1558.
- [18] Matthew Gombolay, Reed Jensen, Jessica Stigile, Sung-Hyun Son, and Julie Shah. 2016. Apprenticeship scheduling: Learning to schedule from human experts. AAAI Press/international joint conferences on artificial intelligence.
- [19] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. Advances in neural information processing systems 27 (2014).
- [20] Aditya Grover, Maruan Al-Shedivat, Jayesh Gupta, Yuri Burda, and Harrison Edwards. 2018. Learning policy representations in multiagent systems. In International conference on machine learning. PMLR, 1802–1811.
- [21] Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. 2018. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International conference on machine learning*. PMLR, 1861– 1870.
- [22] Danijar Hafner, Timothy Lillicrap, Mohammad Norouzi, and Jimmy Ba. 2020. Mastering atari with discrete world models. arXiv preprint arXiv:2010.02193 (2020).
- [23] Nicklas Hansen, Yixin Lin, Hao Su, Xiaolong Wang, Vikash Kumar, and Aravind Rajeswaran. 2022. Modem: Accelerating visual model-based reinforcement learning with demonstrations. arXiv preprint arXiv:2212.05698 (2022).
- [24] Karol Hausman, Jost Tobias Springenberg, Ziyu Wang, Nicolas Heess, and Martin Riedmiller. 2018. Learning an embedding space for transferable robot skills. In International Conference on Learning Representations.
- [25] Jeffrey Hawke, Richard Shen, Corina Gurau, Siddharth Sharma, Daniele Reda, Nikolay Nikolov, Przemysław Mazur, Sean Micklethwaite, Nicolas Griffiths, Amar Shah, et al. 2020. Urban driving with conditional imitation learning. In 2020 IEEE International Conference on Robotics and Automation (ICRA). IEEE, 251–257.
- [26] Jonathan Ho and Stefano Ermon. 2016. Generative adversarial imitation learning. Advances in neural information processing systems 29 (2016), 4565–4573.
- [27] Michael Janner, Qiyang Li, and Sergey Levine. 2021. Offline reinforcement learning as one big sequence modeling problem. Advances in neural information processing systems 34 (2021), 1273-1286.
- [28] Daniel Jarrett, Ioana Bica, and Mihaela van der Schaar. 2020. Strictly batch imitation learning by energy-based distribution matching. Advances in Neural Information Processing Systems 33 (2020), 7354–7365.

- [29] Yiding Jiang, Evan Liu, Benjamin Eysenbach, J Zico Kolter, and Chelsea Finn. 2022. Learning options via compression. Advances in Neural Information Processing Systems 35 (2022), 21184–21199.
- [30] Michael Kaiser, Holger Friedrich, and Rudiger Dillmann. 1995. Obtaining good performance from a bad teacher. In *Programming by Demonstration vs. Learning* from Examples Workshop at ML, Vol. 95. Citeseer.
- [31] Diederik P Kingma. 2013. Auto-encoding variational bayes. arXiv preprint arXiv:1312.6114 (2013).
 [32] Ilya Kostrikov, Kumar Krishna Agrawal, Debidatta Dwibedi, Sergey Levine, and
- [32] Ilya Kostrikov, Kumar Krishna Agrawal, Debidatta Dwibedi, Sergey Levine, and Jonathan Tompson. 2018. Discriminator-Actor-Critic: Addressing Sample Inefficiency and Reward Bias in Adversarial Imitation Learning. In International Conference on Learning Representations.
- [33] Ilya Kostrikov, Ofir Nachum, and Jonathan Tompson. 2019. Imitation Learning via Off-Policy Distribution Matching. In International Conference on Learning Representations.
- [34] Michael Laskin, Aravind Srinivas, and Pieter Abbeel. 2020. Curl: Contrastive unsupervised representations for reinforcement learning. In *International conference* on machine learning. PMLR, 5639–5650.
- [35] Fangchen Liu, Hao Liu, Aditya Grover, and Pieter Abbeel. 2022. Masked autoencoding for scalable and generalizable decision making. Advances in Neural Information Processing Systems 35 (2022), 12608–12618.
- [36] Suraj Nair, Aravind Rajeswaran, Vikash Kumar, Chelsea Finn, and Abhinav Gupta. 2022. R3m: A universal visual representation for robot manipulation. arXiv preprint arXiv:2203.12601 (2022).
- [37] Tung D Nguyen, Rui Shu, Tuan Pham, Hung Bui, and Stefano Ermon. 2021. Temporal predictive coding for model-based planning in latent space. In *International Conference on Machine Learning*. PMLR, 8130–8139.
- [38] Simone Parisi, Aravind Rajeswaran, Senthil Purushwalkam, and Abhinav Gupta. 2022. The unsurprising effectiveness of pre-trained vision models for control. In international conference on machine learning. PMLR, 17359–17371.
- [39] Dean A Pomerleau. 1988. Alvinn: An autonomous land vehicle in a neural network. Advances in neural information processing systems 1 (1988).
- [40] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In International conference on machine learning. PMLR, 8748–8763.
- [41] Younggyo Seo, Danijar Hafner, Hao Liu, Fangchen Liu, Stephen James, Kimin Lee, and Pieter Abbeel. 2023. Masked world models for visual control. In *Conference* on Robot Learning. PMLR, 1332–1344.
- [42] Nur Muhammad Shafiullah, Zichen Cui, Ariuntuya Arty Altanzaya, and Lerrel Pinto. 2022. Behavior transformers: Cloning k modes with one stone. Advances in neural information processing systems 35 (2022), 22955–22968.
- [43] Voot Tangkaratt, Bo Han, Mohammad Emtiyaz Khan, and Masashi Sugiyama. 2020. Variational imitation learning with diverse-quality demonstrations. In International Conference on Machine Learning. PMLR, 9407–9417.
- [44] Emanuel Todorov, Tom Erez, and Yuval Tassa. 2012. Mujoco: A physics engine for model-based control. In 2012 IEEE/RSJ international conference on intelligent robots and systems. IEEE, 5026–5033.
- [45] Michal Valko, Mohammad Ghavamzadeh, and Alessandro Lazaric. 2013. Semisupervised apprenticeship learning. In European workshop on reinforcement learning. PMLR, 131–142.
- [46] Laurens Van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-SNE. Journal of machine learning research 9, 11 (2008).
- [47] Philipp Wu, Arjun Majumdar, Kevin Stone, Yixin Lin, Igor Mordatch, Pieter Abbeel, and Aravind Rajeswaran. 2023. Masked trajectory models for prediction, representation, and control. In *International Conference on Machine Learning*. PMLR, 37607–37623.
- [48] Yueh-Hua Wu, Nontawat Charoenphakdee, Han Bao, Voot Tangkaratt, and Masashi Sugiyama. 2019. Imitation learning from imperfect demonstration. In International Conference on Machine Learning. PMLR, 6818–6827.
- [49] Tete Xiao, Ilija Radosavovic, Trevor Darrell, and Jitendra Malik. 2022. Masked visual pre-training for motor control. arXiv preprint arXiv:2203.06173 (2022).
- [50] Mengjiao Yang and Ofir Nachum. 2021. Representation matters: Offline pretraining for sequential decision making. In *International Conference on Machine Learning*. PMLR, 11784–11794.
- [51] Zilai Zeng, Ce Zhang, Shijie Wang, and Chen Sun. 2024. Goal-conditioned predictive coding for offline reinforcement learning. Advances in Neural Information Processing Systems 36 (2024).