

# Flow to Control: Offline Reinforcement Learning with Lossless Primitive Discovery

Yiqin Yang<sup>1\*</sup>, Hao Hu<sup>2\*</sup>, Wenzhe Li<sup>2\*</sup>, Siyuan Li<sup>3</sup>  
Jun Yang<sup>1</sup>, Qianchuan Zhao<sup>1†</sup>, Chongjie Zhang<sup>2†</sup>

<sup>1</sup>Department of Automation, Tsinghua University

<sup>2</sup>Institute for Interdisciplinary Information Sciences, Tsinghua University

<sup>3</sup>Harbin Institute of Technology

{yangyiqi19,huh22,lwz21}@mails.tsinghua.edu.cn

## Abstract

Offline reinforcement learning (RL) enables the agent to effectively learn from logged data, which significantly extends the applicability of RL algorithms in real-world scenarios where exploration can be expensive or unsafe. Previous works have shown that extracting primitive skills from the recurring and temporally extended structures in the logged data yields better learning. However, these methods suffer greatly when the primitives have limited representation ability to recover the original policy space, especially in offline settings. In this paper, we give a quantitative characterization of the performance of offline hierarchical learning and highlight the importance of learning lossless primitives. To this end, we propose to use a flow-based structure as the representation for low-level policies. This allows us to represent the behaviors in the dataset faithfully while keeping the expression ability to recover the whole policy space. We show that such lossless primitives can drastically improve the performance of hierarchical policies. The experimental results and extensive ablation studies on the standard D4RL benchmark show that our method has a good representation ability for policies and achieves superior performance in most tasks.

## Introduction

Online reinforcement learning has succeeded dramatically in various domains, including strategy games (Ye et al. 2020), recommendation systems (Swaminathan and Joachims 2015), and continuous control (Lillicrap et al. 2015). However, it requires extensive interaction with the environment to learn through trial and error. In many real-world problems, like robot learning and autonomous driving, access to an interactive environment can be severely limited due to safety concerns or huge costs (Shalev-Shwartz, Shammah, and Shashua 2016; Singh et al. 2020b). However, in these applications, a large amount of pre-collected data is usually available, including human demonstrations and data from hand-engineered policies. This makes offline RL (Levine et al. 2020; Fu et al. 2020) an appealing approach for effectively learning from such previously logged data.

\*These authors contributed equally.

†Equal advising.

Leveraging the temporally extended and recurring structure has long been an effective paradigm to solve complex and diverse tasks (Dietterich et al. 1998; Sutton, Precup, and Singh 1999; Kulkarni et al. 2016) by facilitating exploration (Pertsch, Lee, and Lim 2020). For example, a kitchen robot may need to finish many sub-tasks to cook a dish: slide cabinet door, move kettle, and open microwave (Gupta et al. 2019). However, such structure in the agent’s logged behavior is less appreciated in the offline setting. Especially, there is little theoretical justification for extracting low-level skills and learning a high-level policy with offline datasets, and it is unclear how offline-learned skills affect the overall performance. This naturally leads to the following question:

*Is there provable benefit from extracting temporally extended primitive skills in the offline setting, and how to extract useful skills to maximize such benefit?*

In this paper, we give an affirmative answer to the first question by analyzing the performance of hierarchical offline learning in linear MDPs. We show that there are provable benefits in learning a hierarchical structure when the low-level skills are well-learned, in the sense that such skills can faithfully represent the logged behavior and be expressive to recover the original policy space. As for the second question, we find that current skill-based offline RL methods are significantly compromised by the limited representation ability of the learned low-level skills to recover the original policy space. Such loss of representation capacity can greatly harm some powerful offline algorithms’ optimization, rendering skill-based methods less impractical in general usage. To recover the policy space in a lossless way and learn more effective skills, we propose to learn an *invertible* function to map latent embeddings to temporally extended actions. Since the mapping is invertible, the RL agent retains full control over the original policy space while faithfully representing the original dataset’s behavior. We name our method Offline Lossless Primitive Discovery (LPD) and evaluate LPD on the standard D4RL benchmark. Experimental results indicate that LPD achieves superior performance on challenging tasks. Extensive ablation studies demonstrate the strong representation ability of the LPD. To the best of our knowledge, our work is the first study analyzing the representation error in offline hierarchical RL.

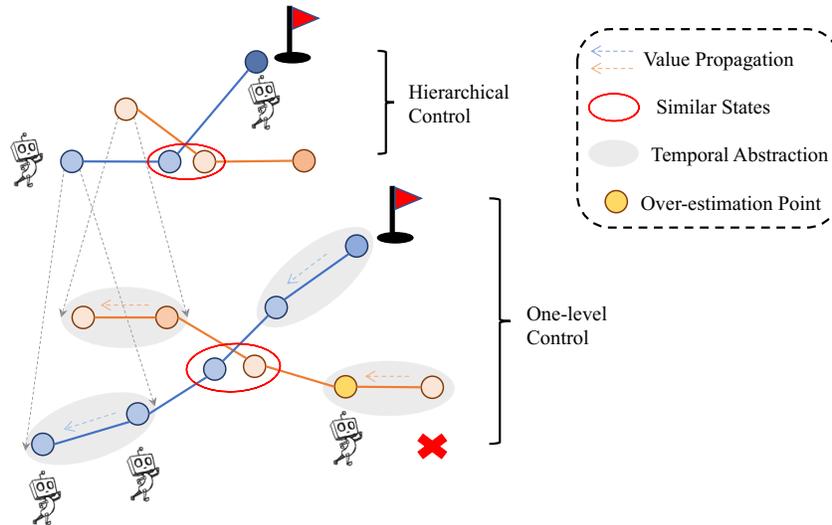


Figure 1: The general description of offline hierarchical learning. The agent makes a decision for every  $c$  step and then follows the behavior-cloned primitive skills for  $c$  steps. The temporally extended primitive skills help the agent to reduce the estimation error and make the right decision by reducing the planning horizon.

## Related Work

**Offline RL.** Current offline RL methods mainly address the extrapolation error issue in value estimation, and they can be roughly divided into policy constraint (Yang et al. 2021; Peng et al. 2019; Nair et al. 2020; Ma et al. 2021), pessimistic value estimation (Kumar et al. 2020; Ma, Jayaraman, and Bastani 2021), and model-based methods (Yu et al. 2021; Kidambi et al. 2020; Hu et al. 2022). Policy constraint methods aim to keep the policy close to the behavior under a probabilistic distance. Pessimistic value estimation methods enforce a regularization constraint on the critic loss to penalize over-generalization. Model-based methods attempt to learn a model from offline data with minimal modification to the policy learning (Argenson and Dulac-Arnold 2020). Differently, we focus on the offline dataset’s temporal structure and explore how to improve these modern offline methods by extracting temporally extended primitive behaviors.

**Hierarchical learning.** Learning primitive skills is closely related to hierarchical models (Dietterich et al. 1998; Sutton, Precup, and Singh 1999; Kulkarni et al. 2016): the low-level policy is the primitive policy extracted from the offline datasets, while the high-level policy is trained using modern offline RL. This structure is similar to online skill discovery works (Sharma et al. 2019; Eysenbach et al. 2018; Shankar and Gupta 2020; Singh et al. 2020a), which use unsupervised objectives to discover skills and use the discovered skills for planning (Sharma et al. 2019), few-shot imitation learning (Nam et al. 2022), or online RL (Nachum et al. 2018). In contrast to these works designed for online settings, we focus on settings where a large dataset of diverse behaviors is provided, but access to the environment is restricted. There are some variants of the above works in offline RL. For example, OPAL (Ajay et al. 2020) explores hierarchical offline learning by

adopting VAE-based primitive policies and achieves sound performance. In the following sections, we briefly name the hierarchical offline RL method incorporated with OPAL as x+OPAL (e.g., IQL+OPAL).

## Preliminaries

We consider infinite-horizon discounted Markov Decision Processes (MDPs), defined by the tuple  $(\mathcal{S}, \mathcal{A}, \mathcal{P}, r, \gamma)$ , where  $\mathcal{S}$  is a state space,  $\mathcal{A}$  is an action space,  $\gamma \in [0, 1)$  is the discount factor and  $\mathcal{P} : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$ ,  $r : \mathcal{S} \times \mathcal{A} \rightarrow [0, r_{\max}]$  are the transition function and reward function, respectively. We also assume a fixed distribution  $\mu_0 \in \Delta(\mathcal{S})$  as the initial state distribution. The goal of an RL agent is to learn a policy  $\pi : \mathcal{S} \rightarrow \Delta(\mathcal{A})$  under dataset  $\mathcal{D}$ , which maximizes the expectation of a discounted cumulative reward:  $J(\pi) = \mathbb{E}_{\mu_0, \pi} [\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t)]$ . We assume the dataset is generated by an unknown behavior policy  $\beta(a|s, z)$ , with prior  $z \sim Z$ . Note that this assumption on data collection is implicit and we do not assume an additional structure of the dataset. We define the dataset for learning primitives as  $\mathcal{D}_{\text{low}} \doteq \{\tau_i = (s_t^i, a_t^i)_{t=0}^{c-1}\}_{i=1}^N$ , where  $\tau_i$  is the sub-trajectory and  $c$  is the skill length. We also create a dataset for high-level policy learning as  $\mathcal{D}_{\text{hi}} = \{(s_0^i, z_i, \sum_{t=0}^{c-1} \gamma^t r_t^i, s_c^i)\}_{i=1}^N$ , where  $z_i$  are learned labels for each sub-trajectory. When learning from the low-level dataset, we consider a finite candidate function class  $\Pi_{\theta}$ . Further, we assume  $\beta \in \Pi_{\theta}$ , i.e., the true low-level policy is realizable in the function class  $\Pi_{\theta}$ .

To make things more concrete, we consider the *linear MDP* (Yang and Wang 2019; Jin et al. 2020) as follows, where the transition kernel and expected reward function are linear with respect to a feature map.

**Definition 1** (Linear MDP). We say an episodic MDP  $(\mathcal{S}, \mathcal{A}, \mathcal{P}, r, \gamma)$  is a linear MDP with known feature map

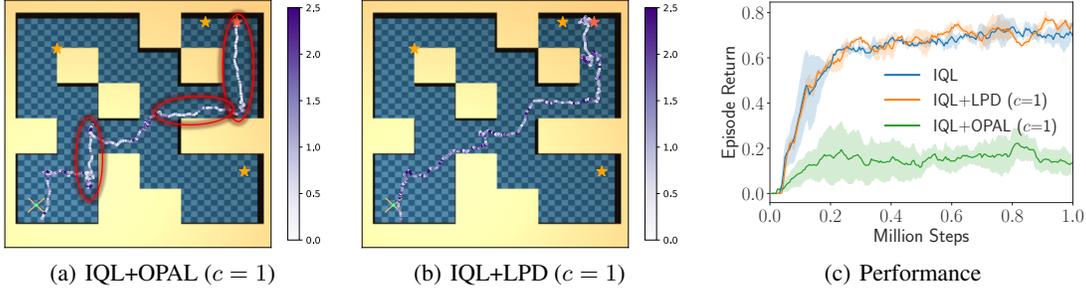


Figure 2: Impact of limited representation on performance. (a) and (b) visualization of the similarity  $\epsilon$  between the agent's decision  $a$  and datasets  $\mathcal{D}$  in the antmaze medium task, which is calculated by  $\epsilon = \min_{\hat{a} \in \mathcal{C}} \|a - \hat{a}\|_1$ , where  $\mathcal{C} = \{(\hat{s}, \hat{a}) \mid \|s - \hat{s}\|_1 \leq 20, (\hat{s}, \hat{a}) \in \mathcal{D}\}$  is the similar state set to the agent. Darker colors correspond to the lower similarity. We find the decoded actions of OPAL are limited to the datasets (red circle), which is opposed to the LPD. (c) Performance on the antmaze-medium-diverse-v0 task with skill length  $c = 1$ .

$\Psi : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}^d, \Phi : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^d$  if there exist an unknown vector  $\omega \in \mathbb{R}^d$  such that

$$\mathcal{P}(s' | s, a) = \Psi(s, a, s')^\top \omega, \mathbb{E}[r(s, a)] = \Phi(s, a)^\top \omega \quad (1)$$

for all  $(s, a, s') \in \mathcal{S} \times \mathcal{A} \times \mathcal{S}$ . And we assume  $\|\Phi(s, a)\|_\infty \leq 1, \|\Psi(s, a, s')\|_\infty \leq 1$  for all  $(s, a, s') \in \mathcal{S} \times \mathcal{A} \times \mathcal{S}$  and  $\|\omega\|_2 \leq \sqrt{d}$ .

We also consider the every- $c$ -step hyper-MDP, where we act every  $c$  step to determine the next primitive  $z$  and keep the reward and the dynamics the same. Then we have the following proposition (See Appendix for proof).

**Proposition 1.** *The every- $c$ -step hyper-MDP induced by the linear MDP  $\mathcal{M}$  and the primitive policy  $\beta(\cdot | s, z)$  is also a linear MDP.*

For the theoretical analysis, we use the *pessimistic value iteration* (PEVI) for high-level policy learning in the hyper-MDP. Please see Algorithm in Appendix for more details. The learned primitive skills  $\pi_\theta(a | s, z)$  and the policy learned from PEVI together comprise a hierarchical policy  $\hat{\pi}_\theta$ . Similarly, We use  $\hat{\pi}_\beta$  to denote the composition of the high-level policy  $\hat{\pi}$  from PEVI and the ground-truth low-level policy  $\beta$ .  $\pi_\beta^*$  is the optimal policy in the hyper-MDP with  $\beta$  as primitive skills and  $\pi^*$  is the optimal policy in  $\mathcal{M}$ . We consider the suboptimality as the evaluation metric, which is defined as the performance difference between the optimal policy  $\pi^*$  and the given policy  $\pi$ . Formally, we have  $\text{SubOpt}(\pi) = J(\pi^*) - J(\pi)$ .

## Theoretical Analysis

In this section, we analyze the benefit of learning a hierarchical structure in the context of offline RL and linear MDPs. To derive a performance bound and examine how a primitive learning method affects the overall performance, we consider the following decomposition:

$$\begin{aligned} \text{SubOpt}(\hat{\pi}_\theta) &= \underbrace{J(\hat{\pi}_\beta) - J(\hat{\pi}_\theta)}_{\text{Primitive Error}} + \\ &\underbrace{J(\pi_\beta^*) - J(\hat{\pi}_\beta)}_{\text{Offline Error}} + \underbrace{J(\pi^*) - J(\pi_\beta^*)}_{\text{Representation Error}}. \end{aligned}$$

The primitive error comes from the generalization error of a learned low-level policy  $\pi_\theta$  and the ground-truth low-level policy  $\beta$ . The offline error comes from learning in a high-level offline dataset. The representation error comes from the limited representation ability of the hierarchical structure. In the following, we upper bound the three different kinds of error, respectively.

The primitive error comes from limited samples of the low-level dataset, which makes the learned low-level policy different from the original primitive. This can be seen as a standard machine learning problem and can be characterized as follows.

**Lemma 1.** *With high probability  $1 - \delta$ ,*

$$J(\hat{\pi}_\beta) - J(\hat{\pi}_\theta) \leq \frac{\gamma c(c+1)r_{\max}}{(1-\gamma)(1-\gamma^c)} \varepsilon_\theta, \quad (2)$$

where  $\varepsilon_\theta = \sqrt{\frac{\ln |\Pi_\theta| / \delta}{N}}$ .  $|\Pi_\theta|$  is the size of the policy class and  $N$  is the number of training samples.

On the contrary, the offline error comes from limited samples of the high-level dataset. With proper pessimism (Jin et al. 2020), we have the following offline-learning bound in the hyper-MDP.

**Lemma 2 (Informal).** *Suppose there exists a finite concentration coefficient  $c^\dagger$  w.r.t. the optimal policy, then with probability  $1 - \delta$ , the policy  $\hat{\pi}$  generated by PEVI satisfies*

$$J(\pi_\beta^*) - J(\hat{\pi}_\beta) \leq \frac{2Cr_{\max}}{(1-\gamma)(1-\gamma^c)} \sqrt{\frac{c^\dagger d^3 \zeta}{N}}, \quad \forall s \in \mathcal{S},$$

where  $C$  is a constant,  $d$  is the dimension of the linear MDP,  $N$  is the size of the dataset and  $\zeta = \log(4dN/(1-\gamma)\delta)$ . Please refer to the Appendix for the Formal description.

Finally, we bound the representation error by quantifying how many of the possible low-level policies can be recovered by different choices of  $z$ . Following Nachum et al. (2018), we make use of an auxiliary inverse primitive model  $\Omega(s, a)$ , which aims to predict which primitive  $z$  will cause  $\beta$  to yield an action  $\tilde{a} = \beta(s, z)$  that induces a next state distribution  $P(s' | s, \tilde{a})$  similar to  $P(s' | s, a)$ . Formally, we have

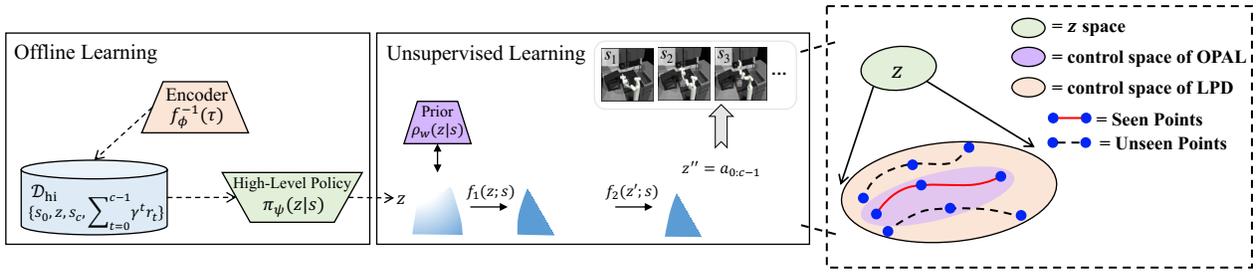


Figure 3: The framework of LPD. We first train the invertible primitive policies  $f_\psi$  by the unsupervised learning. Next, we use  $f_\psi^{-1}$  to create Dataset  $\mathcal{D}_{\text{hi}}$  and further train high-level policy  $\pi_\phi$ . LPD-based primitive policies have adequate control over the original policy space due to their strong representation ability. On the contrary, the control space of the OPAL-based primitives is limited to the vicinity of the dataset (*seen points*).

**Lemma 3.** Consider a mapping  $\Omega : S \times \Pi \rightarrow Z$  and let  $\tilde{\beta} = \beta(s_t, \Omega(s, \pi))$ . We define  $\varepsilon_k : S \times \Pi \rightarrow \mathbb{R}$  for  $k \in [1, c]$  as

$$\varepsilon_k(s_t, \pi) = \text{D}_{\text{TV}}(P_\pi(s_{t+k}|s_{t+k-1}) || P_{\tilde{\beta}}(s_{t+k}|s_{t+k-1})),$$

If

$$\max_{\pi \in \Pi_\theta, k \in [1, c]} \mathbb{E}_\pi[\varepsilon_k(s, \pi)] \leq \varepsilon_\Omega,$$

Then we have

$$J(\pi^*) - J(\pi_\beta^*) \leq \frac{\gamma c(c+1)r_{\max}}{(1-\gamma)(1-\gamma^c)} \varepsilon_\Omega. \quad (3)$$

Together, we have the following theorem.

**Theorem 1.** Under the condition in Lemma 1, 2 and 3, the suboptimality of a policy learned in the hyper-MDP with Algorithm ?? satisfies

$$\text{SubOpt}(\hat{\pi}_\theta) \leq \frac{2Cr_{\max}}{(1-\gamma)(1-\gamma^c)} \sqrt{\frac{c^\dagger d^3 \zeta}{N}} + \frac{\gamma c(c+1)r_{\max}}{(1-\gamma)(1-\gamma^c)} (\varepsilon_\Omega + \varepsilon_\theta), \quad (4)$$

with high probability  $1 - 2\delta$ .

The proofs for all the above lemmas and theorems are provided in Appendix. From Theorem 1, we can see that there is a provable benefit in leveraging the temporal structure in the offline dataset, since a large skill length  $c$  effectively reduces the offline error by a factor of  $\frac{1-\gamma}{1-\gamma^c}$ . However, a large skill length also incurs large representation errors and primitive errors, which forms a trade-off with different choice of  $c$ . We can see that it is crucial to choose a skill learning method that can faithfully represent the original behavior policy while keeping the policy representation ability. It is also important to choose a suitable skill length  $c$  to balance the two errors. Theorem 1 gives a qualitative guidance in choosing the proper length. When the dataset has a bad coverage  $c^\dagger$  or a large dimension  $d$ , we may use a larger skill length. The general description of offline hierarchical learning is shown in Figure 1.

## Method

Following the above theoretical analysis, we find that current skill-based offline RL methods are significantly clipped by the limited representation ability of low-level learned skills to recover the original policy space. Thus, we propose to learn a lossless primitive discovery, which helps the RL agent retain full control over the original policy space while faithfully representing the original dataset’s behavior.

### Lossless Primitive Discovery

We would like to extract a continuous space of temporally-extended primitives  $\pi_\theta(a|s, z)$  from  $\mathcal{D}_{\text{low}}$  which we can later use as an action space for learning downstream tasks with offline RL. A common practice (Ajay et al. 2020) adopts the following objective for learning  $\pi_\theta$ , which maximizes the evidence lower bound (ELBO):

$$\max_{\theta, \psi, w} J(\theta, \phi, w) = \mathbb{E}_{\tau \sim \mathcal{D}_{\text{low}}, z \sim q_\psi(z|\tau)} \left[ \sum_{t=0}^{c-1} \log \pi_\theta(a_t|s_t, z) - \text{D}_{\text{KL}}(q_\psi(z|\tau) || \rho_w(z|s_0)) \right]$$

where  $q_\psi(z|\tau)$  denotes the encoder, which encodes the trajectory  $\tau$  of state-action pairs into distribution in latent space;  $\pi_\theta(a|s, z)$  denotes the decoder, which maximizes the conditional log-likelihood of actions in  $\tau$  given the state and the latent vector;  $\rho_w(z|s_0)$  denotes the prior, which predicts the encoded distribution of the sub-trajectory  $\tau$  from its initial state. The KL-constraint enforces the distribution  $q_\psi(z|\tau)$  to be close to just predicting the latent variable  $z$  given the initial state of this sub-trajectory. The conditioning on the initial state regularizes the distribution  $q_\psi(z|\tau)$  to be consistent over the entire sub-trajectory.

Following the analysis of Theorem 1, we want to minimize the representation error in addition to the primitive error in the low-level policy learning. However, simple VAE-based models can lead to enormous representation errors, as verified in the experiment. To solve this issue, we propose to learn an invertible mapping between skill space and original policy space  $f_\psi : \mathcal{Z} \times \mathcal{S} \rightarrow \mathcal{A}^c$  to replace  $q_\psi(z|\tau)$ , which guarantees our primitive policy is lossless over the original policy space, and the overall objective can be written as

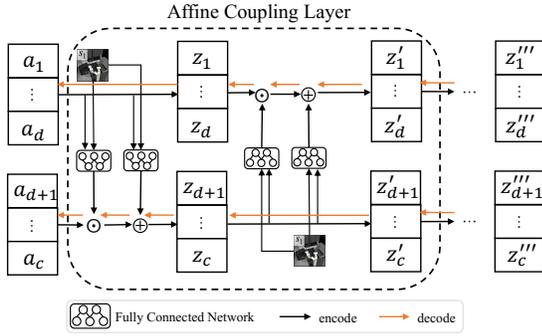


Figure 4: The framework of the flow-based structure. The decoding is  $f_\psi$  and the encoding is  $f_\psi^{-1}$ . The whole process is invertible and lossless.

$$\begin{aligned} \min_{\psi, w} J(\psi, w) &= \widehat{\mathbb{E}}_{\tau \sim \mathcal{D}_{\text{low}}} \left[ -\log p_z(f_\psi^{-1}(a_{0:c-1}; s_0)) \right. \\ &\quad \left. - \log \left| \det \left( \frac{\partial p_z(f_\psi^{-1}(a_{0:c-1}; s_0))}{\partial a_{0:c-1}} \right) \right| \right] \\ \text{s.t. } \widehat{\mathbb{E}}_{\tau \sim \mathcal{D}_{\text{low}}} \left[ \text{D}_{\text{KL}}(f_\psi^{-1}(a_{0:c-1}; s_0) \parallel \rho_w(z|s_0)) \right] &\leq \epsilon_{\text{KL}} \end{aligned}$$

where  $f_\psi$  are parameterized by  $\psi$  and  $f_\psi^{-1} : \mathcal{A}^c \times \mathcal{S} \rightarrow \mathcal{Z}$  is the partial inverse of  $f_\psi$ . The  $a_{0:c-1} = f_\psi(z; s_0)$  denotes the action sequence generated by  $f_\psi$  when given the latent variable  $z$  and state  $s_0$ , and vice versa.  $\det \left( \frac{\partial f(x)}{\partial x^T} \right)$  represents the Jacobian determinant. Based on the works on normalizing flow (Dinh, Sohl-Dickstein, and Bengio 2016; Dinh, Krueger, and Bengio 2014), we can easily represent the invertible function with a neural network and deal with the Jacobian term neatly. Note that our design is general and not limited to *flow*-based structures. Other designs to represent such an invertible  $f$  for primitive skills can be an interesting avenue for future work. The framework of flow-based structure is shown in Figure 4. Please refer to Appendix for the implementation details.

## High-Level Offline Policy Learning

As for the high-level offline policy optimization, we would like to use the learned invertible mapping  $f_\psi(z; s_0)$  and prior  $\rho(z|s_0)$  to improve offline RL for downstream tasks. Similar to Ajay et al. (2020), we relabel the dataset  $\mathcal{D}$  in terms of temporally extended transitions using  $f_\psi^{-1}(a_{0:c-1}; s_0)$ . Specifically, we create a dataset  $\mathcal{D}_{\text{hi}} = \{(s_0^i, z_i, \sum_{t=0}^{c-1} \gamma^t r_t^i, s_c^i)\}_{i=1}^N$  to learn high-level offline policy, where  $z_i \sim f_\psi^{-1}(a_{0:c-1}; s_0)$ . Next, we adopt the strong offline RL algorithm IQL (Kostrikov, Nair, and Levine 2021) to learn  $\pi_\phi(z|s)$ :

---

## Algorithm 1: IQL+LPD algorithm

---

**Input:** Offline dataset  $\mathcal{D}$ .

**Parameter:**  $\phi, \varphi, \vartheta, \psi, w$ .

**Output:** policy  $\pi_\phi$ .

- 1: Learn  $f_\psi$  and  $\rho_w$  with  $\mathcal{D}_{\text{low}}$ .
  - 2: Create dataset  $\mathcal{D}_{\text{hi}}^r$  using the trained mapping  $f_\psi^{-1}$ .
  - 3: **while**  $t = 1$  **to**  $T$  **do**
  - 4:   Train value function  $Q_\vartheta, V_\varphi$  and policy  $\pi_\phi$ .
  - 5:   Update target networks:
  - 6:      $\varphi' \leftarrow (1 - \alpha)\varphi' + \alpha\varphi$ .
  - 7:      $\vartheta' \leftarrow (1 - \alpha)\vartheta' + \alpha\vartheta$ .
  - 8: **end while**
- 

$$J_V(\varphi) = \widehat{\mathbb{E}}_{\tau \sim \mathcal{D}_{\text{hi}}} [L_2^\lambda(Q_{\vartheta'}(s_0, z) - V_\varphi(s_0))]$$

$$J_Q(\vartheta) = \widehat{\mathbb{E}}_{\tau \sim \mathcal{D}_{\text{hi}}} \left[ \left( \sum_{t=0}^{c-1} \gamma^t r_t + \gamma V_\varphi(s_c) - Q_\vartheta(s_0, z) \right)^2 \right]$$

$$J_\pi(\psi) = \widehat{\mathbb{E}}_{\mathcal{D}_{\text{hi}}} [\exp(\beta(Q_\vartheta(s_0, z) - V_\varphi(s_0))) \log \pi_\phi(z|s_0)]$$

where  $L_2^\lambda(u) = |\lambda - \mathbb{I}(u < 0)|u^2$  indicates the expectile regression (Dabney et al. 2018).  $\beta$  is an inverse temperature to control the distribution constraint while maximizing the  $Q$ -values (Yang et al. 2021; Peng et al. 2019). Our complete method is described in Algorithm 1. The graphic overview of our method is shown in Figure 3.

## Experiments

In this section, we aim to address the following questions: (1) How does the limited representation issue affect the agent’s performance, and how does LPD improve on this problem? (2) How does LPD perform compared to strong baseline offline methods and other skill-based offline methods? (3) How does the choice of skill length  $c$  and other factors affect the overall performance of LPD? To answer these questions, we conduct extensive experiments on D4RL tasks. Each experiment result is averaged over five random seeds with a standard deviation.

## Task Description

We evaluate our method on a suite of standard and challenging offline tasks (e.g., D4RL (Fu et al. 2020)) including Franka kitchen, Antmaze, and Adroit. Specifically, the Kitchen tasks involve a Franka robot manipulating multiple objects either in an undirected manner or partially task-directed manner. The task is to use the datasets to arrange objects in the desired configuration, with only a sparse 0-1 completion reward for every object that attains the target configuration. Antmaze requires composing parts of sub-optimal and undirected trajectories to solve a specific point-to-point navigation problem. Adroit tasks require controlling a 24-DoF robotic hand to imitate human behavior. We adopt Kitchen-v0, Antmaze-v0, and Adroit-v0 in our experiments.

Type	Env	IQL+LPD	IQL	CQL	OAMPI	TD3+BC	EMAQ
partial	kitchen	<b>74.9±1.1</b> ↑	46.3	49.8	35.0±3.3	7.5±1.3	<b>74.6±0.6</b>
mixed	kitchen	<b>69.2±1.9</b> ↑	51.0	51.0	47.5±4.1	1.5±0.2	<b>70.8±2.3</b>
complete	kitchen	<b>75.0±0.7</b> ↑	62.5	43.8	10.0±1.9	23.5±2.5	36.9±3.7
fixed	Antmaze-umaze	<b>93.0±1.3</b> ↑	87.5	74.0	64.3±4.6	78.6±4.4	<b>91.0±4.6</b>
play	Antmaze-medium	<b>74.7±2.2</b> ↑	<b>71.2</b>	10.6	0.0±0.0	33.6±2.2	0.0±0.0
play	Antmaze-large	<b>56.2±3.6</b> ↑	39.6	0.2	0.3±0.1	21.4±3.3	0.0±0.0
diverse	Antmaze-umaze	81.6±2.0 ↑	62.2	84.0	60.7±3.9	71.4±4.6	<b>94.0±2.4</b>
diverse	Antmaze-medium	<b>83.7±1.6</b> ↑	70.0	3.0	0.0±0.0	34.7±2.5	0.0±0.0
diverse	Antmaze-large	<b>52.8±1.1</b> ↑	47.5	0.0	0.0±0.0	25.9±2.7	0.0±0.0
human	door	<b>15.1±2.5</b> ↑	4.3	9.9	2.8±0.1	0.0±0.0	-
human	hammer	3.3±0.7 ↑	1.4	<b>4.4</b>	3.9±0.2	0.9±0.1	-
human	pen	63.1±1.6	<b>71.5</b>	37.5	54.6±4.6	39.0±3.6	-
cloned	door	<b>8.1±1.0</b> ↑	1.6	0.4	0.4±0.1	0.0±0.0	0.2±0.3
cloned	hammer	2.1±0.2	2.1	2.1	2.1±0.1	0.3±0.1	1.0±0.7
cloned	pen	<b>65.8±2.7</b> ↑	37.3	39.2	60.0±5.2	25.1±1.9	27.9±3.7

Table 1: Performance of IQL+LPD with six baselines on D4RL tasks with the normalized score metric averaged over five random seeds. The ↑ denotes that LPD achieves a performance improvement over IQL. The results of IQL, CQL, and EMAQ are taken from the original papers. We implement OAMPI and TD3+BC according to the official code. Please refer to Appendix for the detailed implementation of baselines.

## Baselines

To answer the first question posed at the start of this section, we compare IQL+LPD against state-of-the-art model-free and model-based offline methods, including Implicit Q-Learning (IQL) (Kostrikov, Nair, and Levine 2021), Conservative Q-Learning (CQL) (Kumar et al. 2020), Onestep RL(OAMPI) (Brandfonbrener et al. 2021), TD3+BC (Fujimoto and Gu 2021) and EMAQ (Ghasemipour, Schuurmans, and Gu 2021). These prior works have achieved superior results on D4RL.

As for the second question, we compare LPD with OPAL (Ajay et al. 2020), which is the first study to distill primitives from the offline dataset before applying offline methods to learn a primitive-directing high-level policy.

## Main Results

**Experiments with limited representation issue** We conduct experiments to show the limited representation issue. A reasonable approach is to evaluate the performance of IQL+OPAL and IQL+LPD with skill length  $c = 1$ . Experimental results in Figure 2(c) show a large margin between IQL+OPAL and IQL while the performance of IQL+LPD is consistent with IQL, which indicates the OPAL limits the ability of IQL. Furthermore, we visualize the similarity  $\epsilon$  between the agent’s decision  $a$  and the dataset, which is calculated by  $\epsilon = \min_{\hat{a} \in \mathcal{C}} \|a - \hat{a}\|_1$ , where  $\mathcal{C} = \{(\hat{s}, \hat{a}) \mid \|s - \hat{s}\|_1 \leq 20, (\hat{s}, \hat{a}) \in \mathcal{D}\}$  is the similar state set to the agent. The results in Figure 2 show the decoded action of OPAL is limited to the vicinity of the dataset, while LPD can represent more diverse actions. This result does not contradict the reconstruction accuracy because we want to decode the action corresponding to the generalized  $z$ , which does not necessarily exist in the dataset.

**Performance on D4RL** The experimental results in Table 1 show that IQL+LPD achieves state-of-the-art perfor-

antmaze (play)	umaze	medium	large
IQL+LPD	<b>93.0±1.3</b>	<b>74.7±2.2</b>	<b>56.2±3.6</b>
IQL+OPAL	83.5±1.9	48.6±1.0	<b>56.9±3.3</b>
antmaze (diverse)	umaze	medium	large
IQL+LPD	<b>81.6±2.0</b>	<b>83.7±1.6</b>	<b>52.8±1.1</b>
IQL+OPAL	70.2±1.8	42.8±3.9	<b>52.4±2.7</b>
kitchen	partial	mixed	complete
IQL+LPD	<b>72.5±1.1</b>	<b>69.2±1.9</b>	<b>75.0±0.7</b>
IQL+OPAL	<b>74.9±0.3</b>	<b>65.7±3.6</b>	11.5±2.0
adroit (human)	door	hammer	pen
IQL+LPD	<b>15.1±2.5</b>	3.3±0.7	<b>63.1±1.6</b>
IQL+OPAL	12.1±2.2	1.9±0.3	52.0±4.6
adroit (cloned)	door	hammer	pen
IQL+LPD	<b>8.1±1.0</b>	2.1±0.2	<b>65.8±2.7</b>
IQL+OPAL	6.0±1.0	1.1±0.4	46.9±3.6

Table 2: Comparison between LPD and OPAL on D4RL tasks with the normalized score metric averaged over five random seeds.

mance in many tasks and significantly outperforms IQL on nearly all tasks. Most model-free offline methods such as CQL, OAMPI, and TD3+BC cannot perform well on kitchen and antmaze tasks, which strictly require approximate dynamical programming compared with locomotion tasks. Although EMAQ is the exiting offline algorithm that achieves good performance in kitchen-partial and kitchen-mixed tasks, IQL+LPD achieves similar or better performance in most kitchen and antmaze tasks. Moreover, we find TD3+BC and COMBO have poor performance on these challenging tasks, although we have tuned hyperparameters. We suspect this is due to the following reason: while TD3+BC and COMBO are strong in mujoco tasks, they are less effective in more challenging tasks. This is be-

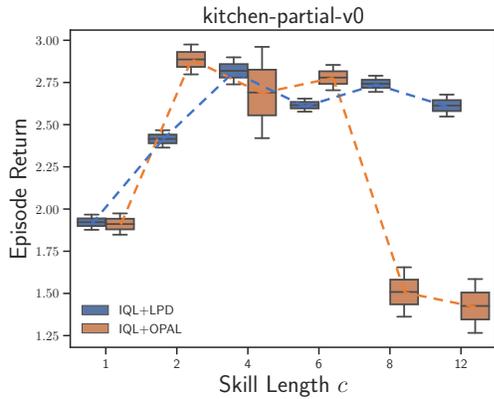


Figure 5: Impact of skill length on performance.

cause under complex environments, BC loss is insufficient to preserve conservatism and the learned model is likely to be inaccurate (Kostrikov, Nair, and Levine 2021). Finally, we are interested in comparing IQL+LPD and IQL, which directly demonstrates the superiority of the temporally-extended primitives. The sources for the scores of the baseline algorithms and the fine-tuning results are shown in Appendix.

**Comparison with OPAL** We conduct a comparison between LPD and OPAL on D4RL tasks. Experimental results in Table 2 show LPD outperforms OPAL in most tasks or achieves the same superior performance. To ensure a fair comparison, we test the performance of IQL+OPAL with the most suitable steps  $c \in \{1, 10\}$  and fine-tune the expectile ratio  $\lambda$  and temperature parameter  $\beta$  in IQL. Then, we select the best parameter combination of IQL+OPAL. We reproduce OPAL with authors’ providing code via email. Although the results of OPAL are a little different from the reported results in the original paper, we argue that the results are convincing and comparable for the following reasons: (1) We incorporate OPAL with IQL. However, this is reasonable since we want the learned primitive policy to be general and can be combined with most offline RL methods. (2) OPAL has the limited representation issue as the above analysis, which will be discussed in the following ablation study.

### Ablation Study

**Impact of skill length on performance** We conduct experiments on kitchen-partial-v0 to evaluate how final performance is impacted as a result of the skill length. As one might expect, the results in Figure 5 show hierarchical offline control achieves better performance than one-level control. Compared with IQL+OPAL, IQL+LPD can accommodate longer steps. Furthermore, we find the optimal skill length in this task is around 5. We also find two methods have similar returns when  $c = 1$ , which does not contradict the above analysis since IQL has limited performance in this task. Complete results of the skill length are shown in Appendix.

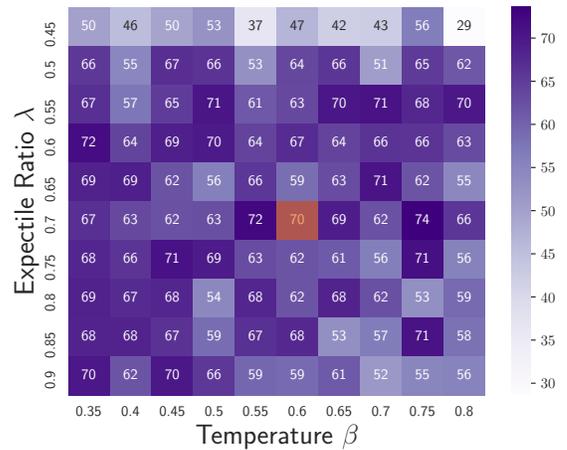


Figure 6: Impact of the hyper-parameters. We evaluate IQL+LPD in kitchen-partial-v0 task with various parameter  $\lambda \in [0.45, 0.9]$  and  $\beta \in [0.35, 0.8]$ . The orange square is the default parameter of IQL and we adopt the normalized score metric.

**Impact of the hyper-parameters** We are concerned about whether the training hyper-parameters of the task policy  $\pi_\phi$  need to be significantly adjusted on hierarchical offline control. For this reason, we ran IQL+LPD on kitchen-partial-v0 with various parameters, such as the expectile ratio  $\lambda \in [0.45, 0.9]$  and the temperature  $\beta \in [0.35, 0.8]$ . The experimental results in Figure 6 show that the performance of IQL+LPD is robust to the changes of the hyper-parameters. Furthermore, we only need to finetune  $\beta$  around the original parameter without tuning  $\lambda$ .

### Conclusion

In this paper, we show that there are provable benefits in learning a hierarchical structure with offline datasets, and it is crucial for low-level primitive skills to be faithful to the original behavior and be expressive to recover the original state space. We empirically show that current skill-based offline RL methods are significantly compromised by the limited representation ability of the learned low-level skills. To solve this issue, we propose the offline Lossless Primitive Discovery (LPD), which learns an invertible function between latent vectors and temporally extended actions. We show that our proposed method has a powerful representation ability to recover the original policy space and it achieves strong results on various D4RL tasks. Learning such expressive skills also enables offline few-shot learning for downstream tasks, and we leave this as an interesting future direction.

### Acknowledgments

This work was funded by the National Natural Science Foundation of China (ID:U1813216, 62192751 and 61425024), National Key Research and Development Project of China under Grant 2017YFC0704100 and Grant 2016YFB0901900, in part by the 111 International Collab-

oration Program of China under Grant BP2018006, BN-Rist Program (BNR2019TD01009) and the National Innovation Center of High Speed Train R&D project (CX/KJ 2020-0006), Science and Technology Innovation 2030 - “New Generation Artificial Intelligence” Major Project (No. 2018AAA0100904) and National Natural Science Foundation of China (62176135).

## References

- Ajay, A.; Kumar, A.; Agrawal, P.; Levine, S.; and Nachum, O. 2020. Opal: Offline primitive discovery for accelerating offline reinforcement learning. *arXiv preprint arXiv:2010.13611*.
- Argenson, A.; and Dulac-Arnold, G. 2020. Model-based offline planning. *arXiv preprint arXiv:2008.05556*.
- Brandfonbrener, D.; Whitney, W.; Ranganath, R.; and Bruna, J. 2021. Offline rl without off-policy evaluation. *Advances in Neural Information Processing Systems*, 34: 4933–4946.
- Dabney, W.; Rowland, M.; Bellemare, M.; and Munos, R. 2018. Distributional reinforcement learning with quantile regression. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.
- Dietterich, T. G.; et al. 1998. The MAXQ Method for Hierarchical Reinforcement Learning. In *ICML*, volume 98, 118–126. Citeseer.
- Dinh, L.; Krueger, D.; and Bengio, Y. 2014. Nice: Non-linear independent components estimation. *arXiv preprint arXiv:1410.8516*.
- Dinh, L.; Sohl-Dickstein, J.; and Bengio, S. 2016. Density estimation using real nvp. *arXiv preprint arXiv:1605.08803*.
- Eysenbach, B.; Gupta, A.; Ibarz, J.; and Levine, S. 2018. Diversity is all you need: Learning skills without a reward function. *arXiv preprint arXiv:1802.06070*.
- Fu, J.; Kumar, A.; Nachum, O.; Tucker, G.; and Levine, S. 2020. D4rl: Datasets for deep data-driven reinforcement learning. *arXiv preprint arXiv:2004.07219*.
- Fujimoto, S.; and Gu, S. S. 2021. A minimalist approach to offline reinforcement learning. *Advances in neural information processing systems*, 34: 20132–20145.
- Ghasemipour, S. K. S.; Schuurmans, D.; and Gu, S. S. 2021. Emaq: Expected-max q-learning operator for simple yet effective offline and online rl. In *International Conference on Machine Learning*, 3682–3691. PMLR.
- Gupta, A.; Kumar, V.; Lynch, C.; Levine, S.; and Hausman, K. 2019. Relay policy learning: Solving long-horizon tasks via imitation and reinforcement learning. *arXiv preprint arXiv:1910.11956*.
- Hu, H.; Yang, Y.; Zhao, Q.; and Zhang, C. 2022. On the Role of Discount Factor in Offline Reinforcement Learning. *arXiv preprint arXiv:2206.03383*.
- Jin, C.; Yang, Z.; Wang, Z.; and Jordan, M. I. 2020. Provably efficient reinforcement learning with linear function approximation. In *Conference on Learning Theory*, 2137–2143. PMLR.
- Kidambi, R.; Rajeswaran, A.; Netrapalli, P.; and Joachims, T. 2020. Morel: Model-based offline reinforcement learning. *Advances in neural information processing systems*, 33: 21810–21823.
- Kostrikov, I.; Nair, A.; and Levine, S. 2021. Offline reinforcement learning with implicit q-learning. *arXiv preprint arXiv:2110.06169*.
- Kulkarni, T. D.; Narasimhan, K.; Saeedi, A.; and Tenenbaum, J. 2016. Hierarchical deep reinforcement learning: Integrating temporal abstraction and intrinsic motivation. *Advances in neural information processing systems*, 29.
- Kumar, A.; Zhou, A.; Tucker, G.; and Levine, S. 2020. Conservative q-learning for offline reinforcement learning. *Advances in Neural Information Processing Systems*, 33: 1179–1191.
- Levine, S.; Kumar, A.; Tucker, G.; and Fu, J. 2020. Offline reinforcement learning: Tutorial, review, and perspectives on open problems. *arXiv preprint arXiv:2005.01643*.
- Lillicrap, T. P.; Hunt, J. J.; Pritzel, A.; Heess, N.; Erez, T.; Tassa, Y.; Silver, D.; and Wierstra, D. 2015. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*.
- Ma, X.; Yang, Y.; Hu, H.; Liu, Q.; Yang, J.; Zhang, C.; Zhao, Q.; and Liang, B. 2021. Offline Reinforcement Learning with Value-based Episodic Memory. *arXiv preprint arXiv:2110.09796*.
- Ma, Y.; Jayaraman, D.; and Bastani, O. 2021. Conservative offline distributional reinforcement learning. *Advances in Neural Information Processing Systems*, 34: 19235–19247.
- Nachum, O.; Gu, S.; Lee, H.; and Levine, S. 2018. Near-optimal representation learning for hierarchical reinforcement learning. *arXiv preprint arXiv:1810.01257*.
- Nair, A.; Dalal, M.; Gupta, A.; and Levine, S. 2020. Accelerating online reinforcement learning with offline datasets. *arXiv preprint arXiv:2006.09359*.
- Nam, T.; Sun, S.-H.; Pertsch, K.; Hwang, S. J.; and Lim, J. J. 2022. Skill-based Meta-Reinforcement Learning. *arXiv preprint arXiv:2204.11828*.
- Peng, X. B.; Kumar, A.; Zhang, G.; and Levine, S. 2019. Advantage-weighted regression: Simple and scalable off-policy reinforcement learning. *arXiv preprint arXiv:1910.00177*.
- Pertsch, K.; Lee, Y.; and Lim, J. J. 2020. Accelerating reinforcement learning with learned skill priors. *arXiv preprint arXiv:2010.11944*.
- Shalev-Shwartz, S.; Shammah, S.; and Shashua, A. 2016. Safe, multi-agent, reinforcement learning for autonomous driving. *arXiv preprint arXiv:1610.03295*.
- Shankar, T.; and Gupta, A. 2020. Learning robot skills with temporal variational inference. In *International Conference on Machine Learning*, 8624–8633. PMLR.
- Sharma, A.; Gu, S.; Levine, S.; Kumar, V.; and Hausman, K. 2019. Dynamics-aware unsupervised discovery of skills. *arXiv preprint arXiv:1907.01657*.

- Singh, A.; Liu, H.; Zhou, G.; Yu, A.; Rhinehart, N.; and Levine, S. 2020a. Parrot: Data-driven behavioral priors for reinforcement learning. *arXiv preprint arXiv:2011.10024*.
- Singh, A.; Yu, A.; Yang, J.; Zhang, J.; Kumar, A.; and Levine, S. 2020b. Cog: Connecting new skills to past experience with offline reinforcement learning. *arXiv preprint arXiv:2010.14500*.
- Sutton, R. S.; Precup, D.; and Singh, S. 1999. Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning. *Artificial intelligence*, 112(1-2): 181–211.
- Swaminathan, A.; and Joachims, T. 2015. Batch learning from logged bandit feedback through counterfactual risk minimization. *The Journal of Machine Learning Research*, 16(1): 1731–1755.
- Yang, L.; and Wang, M. 2019. Sample-optimal parametric q-learning using linearly additive features. In *International Conference on Machine Learning*, 6995–7004. PMLR.
- Yang, Y.; Ma, X.; Chenghao, L.; Zheng, Z.; Zhang, Q.; Huang, G.; Yang, J.; and Zhao, Q. 2021. Believe what you see: Implicit constraint approach for offline multi-agent reinforcement learning. *Advances in Neural Information Processing Systems*, 34: 10299–10312.
- Ye, D.; Liu, Z.; Sun, M.; Shi, B.; Zhao, P.; Wu, H.; Yu, H.; Yang, S.; Wu, X.; Guo, Q.; et al. 2020. Mastering complex control in moba games with deep reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 6672–6679.
- Yu, T.; Kumar, A.; Rafailov, R.; Rajeswaran, A.; Levine, S.; and Finn, C. 2021. Combo: Conservative offline model-based policy optimization. *Advances in neural information processing systems*, 34: 28954–28967.