

STO-RL: Offline RL under Sparse Rewards via LLM-Guided Subgoal Temporal Order

Chengyang Gu

Hong Kong University of Science and Technology
(Guangzhou)
Hong Kong, China
cgu893@connect.hkust-gz.edu.cn

Hui Xiong

Hong Kong University of Science and Technology
(Guangzhou)
Hong Kong, China
xionghui@hkust-gz.edu.cn

Yuxin Pan

Hong Kong University of Science and Technology
Hong Kong, China
yuxin.pan@connect.ust.hk

Yize Chen

University of Alberta
Edmonton, Canada
yize.chen@ualberta.ca

ABSTRACT

Offline reinforcement learning (RL) enables policy learning from pre-collected datasets, avoiding costly and risky online interactions, but it often struggles with long-horizon tasks involving sparse rewards. Existing goal-conditioned and hierarchical offline RL methods decompose such tasks and generate intermediate rewards to mitigate limitations of traditional offline RL, but usually overlook temporal dependencies among subgoals and rely on imprecise reward shaping, leading to suboptimal policies. To address these issues, we propose STO-RL (Offline RL using LLM-Guided Subgoal Temporal Order), an offline RL framework that leverages large language models (LLMs) to generate temporally ordered subgoal sequences and corresponding state-to-subgoal-stage mappings. Using this temporal structure, STO-RL applies potential-based reward shaping to transform sparse terminal rewards into dense, temporally consistent signals, promoting subgoal progress while avoiding suboptimal solutions. The resulting augmented dataset with shaped rewards enables efficient offline training of high-performing policies. Evaluations on four discrete and continuous sparse-reward benchmarks demonstrate that STO-RL consistently outperforms state-of-the-art offline goal-conditioned and hierarchical RL baselines, achieving faster convergence, higher success rates, and shorter trajectories. Ablation studies further confirm STO-RL’s robustness to imperfect or noisy LLM-generated subgoal sequences, demonstrating that LLM-guided subgoal temporal structures combined with theoretically grounded reward shaping provide a practical and scalable solution for long-horizon offline RL.

KEYWORDS

Offline RL, Temporal order, Large Language Models

ACM Reference Format:

Chengyang Gu, Yuxin Pan, Hui Xiong, and Yize Chen. 2026. STO-RL: Offline RL under Sparse Rewards via LLM-Guided Subgoal Temporal Order. In



This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). <https://doi.org/10.65109/EHJQ3648>

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 10 pages. <https://doi.org/10.65109/EHJQ3648>

1 INTRODUCTION

In recent years, Deep Reinforcement Learning (DRL) has emerged as a powerful paradigm for solving sequential decision-making problems, achieving remarkable success in domains ranging from game-playing [19] to robotics [52]. Despite these achievements, DRL methods usually require extensive online interaction with the environment, which can be costly, risky, or infeasible in real-world applications such as autonomous driving or medical treatment planning [26, 86]. To overcome this limitation, offline reinforcement learning (offline RL) [34] has been investigated, where agents learn policies entirely from pre-collected datasets without further interactions with targeting environments.

However, to date offline RL faces significant challenges in long-horizon, goal-reaching tasks with sparse rewards [78], which are practical when agents interacting with the real world. In these settings, rewards are often provided only at the final state, and in limited offline datasets, the lack of intermediate feedback also hinders efficient reward propagation. These make RL agents very challenging to learn effective behaviors [24, 88]. To mitigate this issue, goal-conditioned offline RL extends traditional offline RL by conditioning policies or value functions on goal representations, enabling more flexible and generalizable goal-directed behaviors [14, 45, 58]. Complementarily, hierarchical reinforcement learning (HRL) has been proposed to tackle sparse-reward and long-horizon problems by introducing a multi-level decision-making hierarchy [49, 50, 55, 67]. In such a framework, a high-level policy generates subgoals to guide long-term planning, while a low-level policy learns to achieve these subgoals using primitive actions. By decomposing complex tasks into sequences of subgoal-conditioned subtasks, HRL can improve learning efficiency and exploration, but at an additional cost of algorithm design and complexity [89].

Despite these advances, both offline goal-conditioned RL and offline hierarchical RL continue to face two fundamental challenges: (1) **Learning Effective Subgoals** [9, 30, 47]. Generating informative and achievable subgoals is critical for their success, because high-level and goal-conditioned policies rely on them to guide

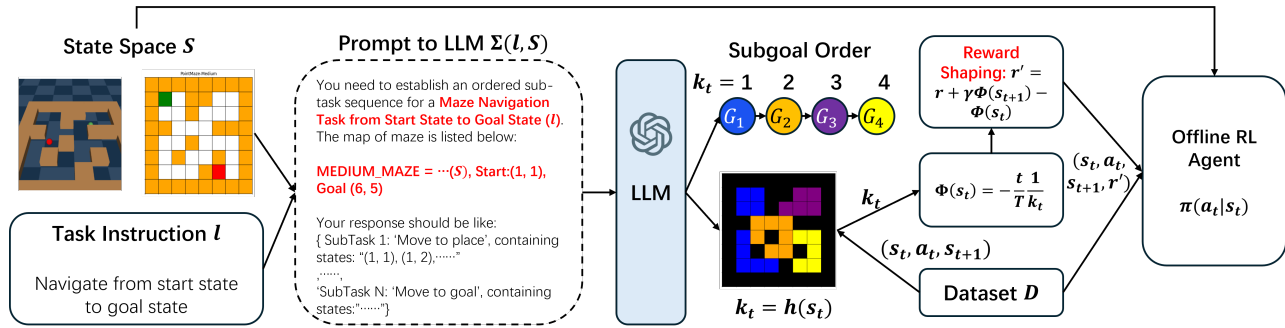


Figure 1: Overview of STO-RL. The agent first prompts an LLM to decompose task instructions into temporally ordered subgoals $\mathcal{G} = \{G_1, G_2, \dots, G_K\}$ with a mapping function h linking each state s_t to its subgoal index k_t ; then applies potential-based reward shaping to assign dense rewards to each transition (s_t, a_t, s_{t+1}) in dataset D based on the mapping function h , followed by training an offline RL agent on the augmented dataset to learn the optimal policy.

low-level actions toward long-horizon goals, directly influencing learning efficiency and task success. However, discovering such subgoals remains difficult due to limited data coverage, sparse rewards, and the presence of suboptimal trajectories and noisy transitions in offline datasets [33, 59, 65, 69, 74]. Furthermore, existing approaches often fail to model the **temporal dependencies among subgoals** [3, 6, 82], making it challenging to learn temporally coherent subgoal sequences. This limitation leads high-level policies to generate subgoals that are inconsistent with task progression and hinder effective hierarchical coordination. (2) **Reward Shaping** [18, 48, 53]. To alleviate the challenge of sparse rewards and enhance learning efficiency, reward shaping is commonly employed to provide denser feedback signals that guide agents toward meaningful, intermediate progress. While shaping schemes such as distance-based potential functions can introduce unintended local optima and misguide policy learning, ultimately leading to suboptimal performance [41].

In this work, we identify large language model (LLM) as the critical aid, and LLM-based planning can redefine how subgoals and rewards are discovered in long-horizon, sparse-reward problems. By leveraging their powerful natural language understanding and reasoning capabilities, LLMs can decompose complex tasks into structured, temporally ordered subgoal sequences, enabling zero-shot or few-shot plan generation [66, 70, 85]. Interestingly, recent works demonstrate that LLMs can also be applied to design effective and dense reward functions for goal-reaching tasks [8, 60, 81]. However, few studies jointly address both perspectives: LLM-based high-level planners mostly rely on pre-trained low-level agents to execute subgoals, while LLM-based reward-shaping mechanisms also often neglect the temporal dependencies among subgoals, leading to inconsistencies and reduced reliability in hierarchical decision-making.

Motivated by these promises and gaps between LLM planners and reward shaping, in this paper we propose *Offline Reinforcement Learning using LLM-Guided Sub-goal Temporal Order (STO-RL)*, a unified framework designed to simultaneously achieve subgoal discovery and reward shaping in the offline setting. Specifically, STO-RL prompts an LLM with task instructions and environment context to generate a sequence of temporally ordered subgoals,

along with a mapping function that assigns each state to its corresponding subgoal stage. LLM excels at interpreting and generalizing these goal-reaching abstractions and relations. Leveraging this temporal structure, STO-RL applies a potential-based reward shaping scheme that integrates subgoal order into the reward function, effectively transforming sparse terminal rewards into dense signals that promotes subgoal progress to the final goal while theoretically avoiding suboptimal shaping. The resulting augmented dataset with shaped rewards enables efficient offline training of an optimal policy from a limited offline dataset.

To validate the effectiveness of STO-RL, we conduct extensive simulations across both discrete and continuous goal-reaching benchmarks, including *CliffWalking*, *FourRoom*, *PointMaze-UMaze* and *PointMaze-Medium*. The results demonstrate that STO-RL significantly improves both learning efficiency and final success rates compared to state-of-the-art offline goal-conditioned and hierarchical RL baselines. In discrete environments, STO-RL achieves faster convergence and more stable policy improvement, while in continuous tasks, it consistently attains higher success rates and generates shorter, more direct trajectories toward the goals. These findings confirm that STO-RL effectively leverages LLM-guided temporally ordered subgoal structures, and, when combined with theoretically grounded reward shaping, provides a practical and scalable solution to the persistent challenges of sparse-reward goal-reaching tasks. Furthermore, ablation studies indicate that STO-RL exhibits a notable robustness to the quality of LLM-generated subgoal sequences: even when the temporal ordering is partially imperfect or noisy, STO-RL maintains competitive performance, demonstrating strong resilience to variability in LLM outputs.

2 RELATED WORKS

Goal-conditioned RL. Goal-conditioned reinforcement learning (RL) aims to enable agents to achieve arbitrary goals by conditioning policies and value functions on goal representations. In offline RL settings, goal-conditioned imitation learning approaches [14, 43, 64] have been investigated to better exploit suboptimal datasets. To facilitate learning in goal-reaching tasks, goal-conditioned RL often

leverages self-supervised techniques, such as hindsight relabeling [4, 36], vector quantization methods [27, 33, 57] and chain-of-thought temporal abstraction via Transformers [28], to derive informative subgoal representations. Despite these advances, existing goal-conditioned RL approaches can suffer from suboptimal or noisy subgoal selection [84], as learning these goals autonomously can be nontrivial.

Hierarchical RL. Hierarchical Reinforcement Learning (HRL) has achieved notable success in addressing long-horizon goal-reaching problems under sparse reward settings by decomposing complex tasks into subgoals. A key factor in its effectiveness is the learning of appropriate subgoal representations, which significantly influences both performance and stability [40]. Recent studies have explored different strategies for discovering semantically meaningful subgoals efficiently. For example, Guider and HLPS [69, 75] leverage pre-trained probabilistic distributions to predict subgoals, while HIDIO [89] and SEADS [1] adopt self-supervised approaches based on intrinsic rewards. Another line of research, including HIGL [30], SFL [23], constructs task graphs from landmarks to guide subgoal selection [25]. Although these approaches demonstrate promising progress in subgoal discovery and skill learning, they largely overlook the temporal structure among subgoals, which constrains their efficiency and overall effectiveness.

LLM Planning. Large Language Models (LLMs) have recently emerged as powerful tools for high-level planning in sequential decision-making tasks. By leveraging their pretrained world knowledge [72, 87] and step-by-step reasoning capabilities [16, 21], LLMs can capture the causal and temporal dependencies inherent in goal-directed behavior. This enables them to generate structured plans composed of subgoals with explicit temporal order [39, 42, 73, 85], providing effective intermediate guidance for hierarchical reinforcement learning agents. For example, ELLM [15] rewards agents for achieving LLM-suggested goals based on their current state, while LDSC [68] and LgTs [70] train specialized RL agents to achieve subgoals generated by language models. Despite these advances, LLMs are limited to high-level planning and cannot directly execute actions, requiring low-level agents to achieve each subgoal along the specified temporal sequence [37, 61, 91]. One intrinsic issues for LLM-based planners are that they lack direct interactions with low-level policies. If the pre-trained fine-grained policies are not well-aligned with the LLM-generated high-level plan, agents may fail to follow the intended subgoal progression. To address this, we introduce a temporal-order-aware reward shaping framework that integrates LLM-generated subgoal sequences into offline RL, allowing agents to autonomously learn subgoal completion in the correct temporal order toward the final goal.

Reward Shaping. Reward shaping augments the ground-truth reward with an additional shaping term, typically in an additive form. Several studies employ manually designed shaping rewards to accelerate policy learning by incorporating domain expertise [11, 63]. However, such methods require considerable problem-specific design, and human knowledge fails to consistently yield gains during policy training. In contrast, numerous methods adopt exploration bonuses or intrinsic rewards to expedite policy convergence

without relying on domain knowledge [7, 12, 44, 46, 56, 90]. However, these approaches often cause the learned policy to deviate from the original optimal policy. To guarantee policy invariance, potential-based reward shaping (PBRS) is formulated as the temporal difference of a potential function [51]. Building on this foundation, subsequent work has focused on the theoretical analysis of PBRS [35, 76] and proposed variants that offer more informative guidance. These include action-dependent PBRS [77], dynamic PBRS [13], dynamic action-dependent PBRS [22], mode-based PBRS [5], dynamics-aware PBRS [31], and self-shaped rewards [2]. Compared to prior PBRS algorithms, our method investigates PBRS within the hierarchical RL setting, where the shaping reward incorporates domain knowledge extracted from a LLM.

3 METHOD

3.1 Problem Formulation

In this work, we focus on the offline reinforcement learning (RL) for goal-reaching tasks with sparse rewards. It can be formulated as a finite-horizon Markov Decision Process (MDP): $M = (\mathcal{S}, \mathcal{A}, P, r, \gamma, T)$. We also set up the problem with a collaborative language-based task description l and a pre-collected dataset \mathcal{D} . Here, \mathcal{S} , \mathcal{A} denotes the state space and the action space, respectively; $P(s_{t+1}|s_t, a_t)$ represents the transition dynamics; $r : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathcal{R}$ is the reward function; $\gamma \in (0, 1)$ is the discount factor; and T is the episode horizon length. The reward function r is given by

$$r(s_t, a_t, s_{t+1}) = \begin{cases} 1 & \text{when } s_t = g, \\ 0 & \text{otherwise.} \end{cases} \quad (1)$$

with $g \in \mathcal{S}$ denoting the target goal state.

The dataset \mathcal{D} consists of a collection of trajectories:

$$\tau = \{(s_0, a_0), (s_1, a_1), \dots, (s_{T-1}, a_{T-1}), s_T\},$$

And our goal is to learn an optimal policy agent $\pi(s_t|a_t) : \mathcal{S} \rightarrow \mathcal{A}$, that maximizes the expected discounted return:

$$V^\pi = \mathbb{E}_\pi \left[\sum_{t=0}^{T-1} \gamma^t r(s_t, a_t, s_{t+1}) \right].$$

3.2 LLMs-Guided Temporal Ordered Subgoals

Standard offline RL methods often struggle to perform effectively on goal-reaching tasks due to sparse rewards and long planning horizons, which hinder the identification of meaningful gradients during policy optimization. A natural approach to mitigate these challenges is to decompose the task into a sequence of subgoals that are easier to achieve, allowing the agent to receive denser feedback and learn structured representations of progress. Such idea of hierarchical decomposition has long been central to reinforcement learning, inspiring the development of subgoal discovery and hierarchical policy learning frameworks.

Due to its strong reasoning and understanding capabilities, LLMs have been increasingly applied to high-level planning in robotics and embodied decision-making [10, 38, 79]. When prompted with natural language task instructions together with necessary state and environment descriptions, LLMs can decompose tasks into structured subgoal sequences. Owing to their strong common sense reasoning and causal interpretations, these LLM-generated plans inherently capture the **temporal dependencies** among subgoals.

However, prior research often overlook such temporal structure, which limits their effectiveness in RL tasks.

In this paper, we assume that for any goal-reaching task, there exists an ordered subgoal sequence of K **temporally ordered subgoals**: $\{G_1, G_2, \dots, G_K\} \subseteq \mathcal{G}$, such that every successful trajectory reaching the final goal g must achieve these subgoals sequentially. The formal definition of such a trajectory is stated in *Definition 1*.

Definition 1. (Properties of Successful Trajectory) Consider a trajectory $\tau = \{(s_0, a_0, k_0), \dots, (s_{H-1}, a_{H-1}, k_{H-1})\}$ successfully reaching the final goal g in H timesteps ($H \leq T$). τ should satisfy: (1) $k_0 = 1, k_{H-1} = K$; (2) There exists $K - 1$ timesteps $0 < t^1 < \dots < t^{K-1} < H - 1$, where $k_{t^i} = i, k_{t^{i+1}} = i + 1$.

We define the *progress index* as $k_t \in \{1, \dots, K\}$, where k_t represents the temporal order of the corresponding subgoal at timestep t . Given the language-based task instruction l and state space \mathcal{S} for environment description, we construct a language prompt $\Sigma(l, \mathcal{S})$ that is fed into a large language model (LLM). Prior studies [39, 61, 71] demonstrate that LLMs can reliably generate subgoals from such prompts. An example of prompt is illustrated in Figure 1. Unlike prior works, we let LLM output a subgoal sequence \mathcal{G} along with an additional mapping function $h : \mathcal{S} \rightarrow \{1, 2, \dots, K\}$ which associates each state s_t with its corresponding progress index $k_t = h(s_t)$. This mapping provides a foundation for incorporating subgoal temporal order into the reward structure and policy learning process, enabling more efficient and interpretable offline RL in complex, sparse-reward environments.

3.3 Subgoal-Temporal-Order-Aware Potential-based Reward Shaping

Built upon the temporally ordered subgoal sequences generated by LLMs, we propose a novel reward shaping mechanism that transforms the original sparse reward signal into a dense one, effectively accelerating learning while avoiding suboptimal solutions in goal-reaching tasks.

Definition 2. (Positive Progress) Let $\{G_1, G_2, \dots, G_K\}$ denote the sequence of subgoals arranged in prescribed temporal order. Define $k_t \in \{1, 2, \dots, K\}$ as the corresponding progress index at time step t . A transition (s_t, a_t, s_{t+1}) is said to constitute **positive progress** if and only if $k_t < k_{t+1}$.

The key insight of our method is to exploit the *temporal structure* inherent in the ordered subgoal sequence, closely aligned with LLM’s strong capabilities. As established in *Definition 1*, any trajectory that successfully reaches the final goal must pass through these subgoals sequentially. Consequently, an optimal policy should reach each subgoal as early as possible within the episode. To formalize this intuition, we first consider the reward design at the transition level. Specifically, following the insights of previous works [6, 62], transitions that **promote progress toward the next subgoal** are considered more desirable. We formally bring out a definition of such transitions as *positive progress* transitions in *Definition 2*.

We introduce a dynamic *potential-based reward shaping* (PBRs) function [13, 51] that encodes both temporal urgency and ordered subgoal progress: $\Phi(s_t, t) = -\frac{t}{T} \cdot \frac{1}{k_t}$, where T denotes the maximum episode length. The shaped reward r' is then defined as:

$$r'(s_t, a_t, s_{t+1}) = r(s_t, a_t, s_{t+1}) + \gamma\Phi(s_{t+1}, t+1) - \Phi(s_t, t). \quad (2)$$

By introducing this potential-based shaping, we not only preserve the optimality of the original policy, but also effectively identify and reward actions that guide the agent more efficiently toward the goal. As is shown in *Theorem 1*, transitions that achieve positive subgoal progress receive strictly higher shaped rewards, and they are therefore preferred during learning, enabling more efficient policy optimization.

Theorem 1. (Shaped Reward Preference for Positive Progress) Consider the shaped reward in Eq. 2: $r'(s_t, a_t, s_{t+1}) = r(s_t, a_t, s_{t+1}) + \gamma\Phi(s_{t+1}, t+1) - \Phi(s_t, t)$ with potential function $\Phi(s_t, t) = -\frac{t}{T} \cdot \frac{1}{k_t}$. Then for any transition (s_t, a_t, s_{t+1}) , the shaped reward assigns strictly higher value to transitions that achieve positive progress compared to those that do not, namely $\Delta r := r'(s_t, a_t^c, s_{t+1}^c) - r'(s_t, a_t^n, s_{t+1}^n) > 0$, where (s_t, a_t^c, s_{t+1}^c) denote a transition that achieves **positive progress**, and (s_t, a_t^n, s_{t+1}^n) a transition that does **not**.

PROOF. Substituting $\Phi(s_t, t) = -\frac{t}{T} \cdot \frac{1}{k_t}$ to the definition of reward shaping, we have:

$$r'(s_t, a_t^c, s_{t+1}^c) = r(s_t, a_t^c, s_{t+1}^c) - \gamma\left(\frac{t+1}{T} \cdot \frac{1}{k_{t+1}^c}\right) + \frac{t}{T} \cdot \frac{1}{k_t}, \quad (3)$$

$$r'(s_t, a_t^n, s_{t+1}^n) = r(s_t, a_t^n, s_{t+1}^n) - \gamma\left(\frac{t+1}{T} \cdot \frac{1}{k_{t+1}^n}\right) + \frac{t}{T} \cdot \frac{1}{k_t}. \quad (4)$$

Taking the difference we have:

$$\begin{aligned} \Delta r &= r'(s_t, a_t^c, s_{t+1}^c) - r'(s_t, a_t^n, s_{t+1}^n) \\ &= r(s_t, a_t^c, s_{t+1}^c) - r(s_t, a_t^n, s_{t+1}^n) + \gamma \frac{t+1}{T} \cdot \left(\frac{1}{k_{t+1}^n} - \frac{1}{k_{t+1}^c}\right). \end{aligned} \quad (5)$$

Since the task has sparse rewards, indicating

$$r(s_t, a_t^c, s_{t+1}^c) = r(s_t, a_t^n, s_{t+1}^n) = 0. \quad (6)$$

By the definition of positive progress in *Definition 1*:

$$k_{t+1}^n \leq k_t < k_{t+1}^c. \quad (7)$$

Therefore:

$$\Delta r = \gamma \frac{t+1}{T} \left(\frac{1}{k_{t+1}^n} - \frac{1}{k_{t+1}^c}\right) > 0. \quad (8)$$

showing that shaped rewards are strictly larger for transitions achieving positive progress. \square

In addition to promoting faster learning, our temporally ordered reward shaping also enhances the agent’s ability to avoid suboptimal solutions. Specifically, it provides a mechanism to differentiate between successful trajectories of varying lengths, favoring shorter ones that reach the goal more efficiently.

Lemma 1. (Equivalence of Returns for Equal-Length Successful Trajectories) Consider a goal-reaching task with K temporally ordered subgoals. Any two successful trajectories with same episode length T_L ($T_L \leq T$) yield equal discounted return $\sum_{t=0}^{T_L-1} \gamma^t r'(s_t, a_t, s_{t+1})$ under $r'(s_t, a_t, s_{t+1}) = r(s_t, a_t, s_{t+1}) + \gamma\Phi(s_{t+1}, t+1) - \Phi(s_t, t)$, where $\Phi(s_t, t) = -\frac{t}{T} \cdot \frac{1}{k_t}$.

PROOF. See Appendix A.1 in our arXiv version of paper [20]. \square

Based on such potential-based reward shaping properties as shown in *Lemma 1*, any two successful trajectories with the same length will yield the same discounted return under the shaped reward, regardless of the specific sequence of intermediate subgoal progress indices. Furthermore, under the condition (discount factor




Figure 2: A toy example illustrating how the subgoal-temporal-order-aware PBRs penalizes longer successful trajectories. A virtual trajectory τ^3 is constructed with the same length T and discounted return as longer trajectory τ^2 ($R(\tau^3) = R(\tau^2)$), consisting of the first $T - L$ steps sharing subgoal indices and return with shorter τ^1 , followed by L non-progress transitions that yield negative returns under *Theorem 2*.

γ satisfies $\gamma \geq (T - 1)/T$, *Theorem 2* demonstrates that transitions which do not contribute to positive subgoal progress receive strictly negative rewards. Consequently, for two successful trajectories of lengths T and $T - L$ ($L > 0$), the discounted return of the longer trajectory can be expressed as the sum of the shorter trajectory’s return and the cumulative rewards of L non-progressive transitions, as illustrated in Figure 2. Since these additional transitions yield negative contributions, the return of the longer trajectory is strictly smaller than that of the shorter one, as proven in *Theorem 3*.

Theorem 2. (Strict Negative Reward for Non-Positive Progress): If $\gamma > \frac{T-1}{T}$, potential-based penalty $\gamma\Phi(s_{t+1}, t+1) - \Phi(s_t, t) < 0$ holds for any transition (s_t, a_t, s_{t+1}) that does not achieve positive progress: $k_{t+1} \leq k_t$.

PROOF. Consider the potential function defined as $\Phi(s_t, t) = -\frac{t}{T} \cdot \frac{1}{k_t}$, where $t \in \{0, \dots, T - 1\}$ is the timestep and k_t is the progress index.

For a transition (s_t, a_t, s_{t+1}) that does **not achieve positive progress**, we have $k_{t+1} \leq k_t$.

The potential-based shaping term is

$$\begin{aligned} \Delta\Phi_t &:= \gamma\Phi(s_{t+1}, t+1) - \Phi(s_t, t) \\ &= \gamma \left(-\frac{t+1}{T} \cdot \frac{1}{k_{t+1}} \right) - \left(-\frac{t}{T} \cdot \frac{1}{k_t} \right) \\ &= -\gamma \frac{t+1}{T} \cdot \frac{1}{k_{t+1}} + \frac{t}{T} \cdot \frac{1}{k_t}. \end{aligned} \quad (9)$$

For any non-progress transition, Since $k_{t+1} \leq k_t$, we have

$$\frac{1}{k_{t+1}} \geq \frac{1}{k_t}. \quad (10)$$

Hence

$$\Delta\Phi_t \leq -\gamma \frac{t+1}{T} \cdot \frac{1}{k_t} + \frac{t}{T} \cdot \frac{1}{k_t} = \frac{1}{k_t} \left(\frac{t}{T} - \gamma \frac{t+1}{T} \right) = \frac{1}{k_t} \cdot \frac{t - \gamma(t+1)}{T}. \quad (11)$$

Since $t \leq T - 1$, we have

$$\gamma > \frac{T-1}{T} > \frac{t}{t+1}. \quad (12)$$

Then the numerator $t - \gamma(t+1) < 0$, which implies

$$\Delta\Phi_t < 0. \quad (13)$$

This proves that the potential-based reward penalizes transitions that fail to achieve positive progress when $\gamma > \frac{T-1}{T}$. \square

This property ensures that our reward shaping framework inherently discourages redundant or inefficient behaviors, guiding the agent to prioritize shorter, optimal trajectories toward the goal.

Theorem 3. (Shorter Successful Trajectories Yield Higher Returns): Let τ^1 and τ^2 be two successful trajectories with length $T_L - L$ and T_L ($L > 0$) respectively. If $\gamma > \frac{T-1}{T}$, then the discounted return of the shorter trajectory is strictly larger, i.e. $R(\tau^1) > R(\tau^2)$, where $R(\tau^1) = \sum_{t=0}^{T_L-L-1} \gamma^t r'(s_t^1, a_t^1, s_{t+1}^1)$, $R(\tau^2) = \sum_{t=0}^{T_L-1} \gamma^t r'(s_t^2, a_t^2, s_{t+1}^2)$.

PROOF. Suppose there exists a successful trajectory τ^3 with length T_L . In τ^3 , $k_{T_L-L} = k_{T_L} = K$, and for $t = T_L - L, \dots, T_L - 1$, $k_{t+1} = k_t = K$. Hence, the discounted return of τ^3 is:

$$\begin{aligned} R(\tau^3) &= \sum_{t=0}^{T_L-1} \gamma^t r'(s_t^3, a_t^3, s_{t+1}^3) \\ &= \sum_{t=0}^{T_L-L-1} \gamma^t r'(s_t^3, a_t^3, s_{t+1}^3) + \sum_{t=T_L-L}^{T_L-1} \gamma^t r'(s_t^3, a_t^3, s_{t+1}^3). \end{aligned} \quad (14)$$

Since $k_{T_L-L} = k_{T_L} = K$, according to *Lemma 1* we have:

$$R(\tau^3) = R(\tau^2), \quad (15)$$

$$R(\tau^1) = \sum_{t=0}^{T_L-L-1} \gamma^t r'(s_t^3, a_t^3, s_{t+1}^3). \quad (16)$$

Hence:

$$R(\tau^2) = R(\tau^1) + \sum_{t=T_L-L}^{T_L-1} \gamma^t r'(s_t^3, a_t^3, s_{t+1}^3); \quad (17)$$

Since $k_{t+1} = k_t$ for $t = T_L - L, \dots, T_L - 1$, according to *Theorem 2*:

$$r'(s_t, a_t, s_{t+1}) < 0, \quad t = T_L - L, \dots, T_L - 1. \quad (18)$$

Therefore, $R(\tau^2) < R(\tau^1)$. \square

3.4 STO-RL

Built upon the temporally ordered subgoal sequences generated by LLMs and the potential-based reward shaping introduced in Section 3.3, we propose *Offline Reinforcement Learning using LLM-Guided Sub-goal Temporal Order (STO-RL)*, a framework for offline goal-conditioned reinforcement learning.

To ensure the RL agent is receiving rich and meaningful rewards, STO-RL first prompts an LLM to extract information from the language-based task instruction and environment description. LLM aids in decomposing the overall task into a temporally ordered subgoal sequence. Using this sequence, the framework applies potential-based reward shaping incorporating subgoal temporal order, to assign a reward signal r'_t for every transition (s_t, a_t, s_{t+1}) in the offline dataset D . With the augmented dataset containing the dense rewards, an offline RL agent is trained to learn the optimal policy $\pi(s_t|a_t)$. This process allows the agent to efficiently leverage subgoal temporal information while performing policy optimization in a purely offline setting.

The complete STO-RL algorithm is summarized in Algorithm 1.

Algorithm 1 STO-RL

Require: Offline dataset \mathcal{D} , task instruction l , state space S , discount factor γ , maximum episode length T

Ensure: Optimal policy $\pi^*(s_t|a_t)$

- 1: **Generate Prompt:** Design prompt $\Sigma = (l, S)$ for LLM.
- 2: **Subgoal Extraction:** LLM outputs temporally ordered subgoal sequence G and mapping function $h: (G, h) = LLM(\Sigma)$.
- 3: **for** each transition $(s_t, a_t, s_{t+1}) \in \mathcal{D}$ **do**
- 4: $\Phi(s_t, t) = -\frac{t}{T} \cdot \frac{1}{k_t}$ where $k_t = h(s_t)$
- 5: $r'_t = r(s_t, a_t, s_{t+1}) + \gamma\Phi(s_{t+1}, t + 1) - \Phi(s_t, t)$
- 6: **end for**
- 7: **Augment dataset:** Replace original rewards in \mathcal{D} with shaped rewards r'_t .
- 8: **Offline RL Training:** Train an offline RL agent to learn optimal policy $\pi^*(s_t|a_t)$.

4 EXPERIMENTS

To thoroughly assess the effectiveness of the proposed STO-RL method, we conduct experiments across *four* distinct goal-reaching tasks under sparse reward settings. Our evaluation is designed to address the following research questions:

- (1) Does STO-RL achieve higher **success rates** on sparse-reward tasks in comparison with existing (hierarchical) offline reinforcement learning baselines?
- (2) Will STO-RL accelerate the RL agent learning process, compared to prior methods?
- (3) Can STO-RL help avoid suboptimal solutions, achieving shorter trajectory solutions?

To generate ordered subgoal sequences in our STO-RL, we majorly employ *ChatGPT 5.0* as the default LLM model (See Sec. 4.3 for other LLM models); For the offline RL component in our STO-RL framework, we adopt common-used *Implicit Q-Learning (IQL)* [32], which learns value functions to infer an implicit policy via advantage-weighted regression. Across all experiments, we employ a two-layer multilayer perceptron (MLP) architecture with 128 hidden units per layer for the learning of value, Q, and policy networks in IQL. We find this architecture effective to learn Q value.

Baselines. We compare our STO-RL framework with 3 baseline approaches:

- **Implicit Q-Learning (IQL)** [32]: a classical offline RL algorithm without hierarchical subgoal decomposition.
- **Goal-Conditioned Behavioral Cloning (GC-BC)** [14]: a behavioral cloning style approach that learns a direct mapping from states, subgoals to actions via supervised imitation. In our experiments, GC-BC utilizes the same LLM-generated subgoals as STO-RL for a fair comparison.
- **Hierarchical Implicit Q-Learning (HIQL)** [59]: a hierarchical extension of offline goal-conditioned RL that derives a representation function, high-level policy, and low-level policy jointly from a single goal-conditioned value function.

4.1 Tasks with Finite, Discrete Action Space

To answer the three questions, we begin by evaluating STO-RL on two simple but classical sparse-reward, goal-reaching tasks with

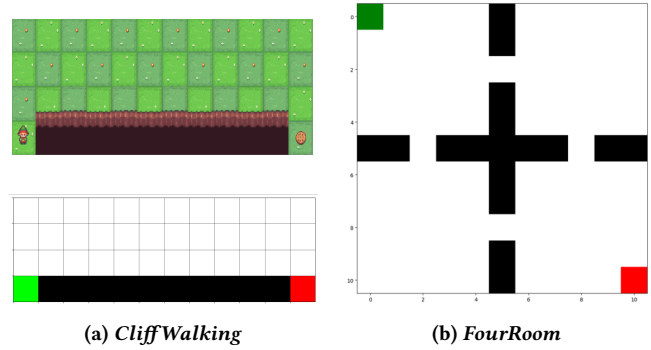


Figure 3: Environments with finite, discrete action spaces.

discrete, finite state and action spaces: (1) *CliffWalking* and (2) *FourRoom*. The configurations of these environments are illustrated in Figure 3.

In the *CliffWalking* task, the agent must navigate a 4×12 grid world from a designated start state to a goal state while avoiding falling off a cliff. The task begins with the agent positioned at $[3, 0]$ (green state in Figure 3a), and the goal is located at $[3, 11]$. A cliff occupies cells $[3, 1 - 10]$ (black region in Figure 3a). If the agent steps onto a cliff cell, it is immediately returned to the start state; Similarly, in the *FourRoom* task, the agent navigates an 11×11 grid world divided into four connected rooms with narrow doorways. The agent must move from a fixed start location in the upper-left room (green state in Figure 3b) to a goal location in the bottom-right room (red state in Figure 3b). Walls (black cells in Figure 3b) separate the rooms, and the agent can transition between rooms only through the designated doorways. For both tasks, the agent operates in a discrete action space consisting of four actions: move up, move down, move right, move left. The agent receives a reward of 1 only upon reaching the goal, and an episode terminates either when the goal is reached or when the maximum number of steps ($T = 100$) is exceeded.

Table 1: Evaluation results on CliffWalking and FourRoom. Success rates and average steps after 1,000 training iterations.

| Method | CliffWalking | | FourRoom | |
|---------------|--------------|-------------|--------------|------------|
| | Success Rate | Steps | Success Rate | Steps |
| Datasets | 0.50 | 69.9 ± 32.8 | 0.14 | 97.2 ± 9.2 |
| IQL | 1.00 | 13.0 ± 0.0 | 1.00 | 20.0 ± 0.0 |
| GC-BC | 1.00 | 13.0 ± 0.0 | 1.00 | 20.0 ± 0.0 |
| HIQL | 1.00 | 13.0 ± 0.0 | 1.00 | 20.0 ± 0.0 |
| STO-RL | 1.00 | 13.0 ± 0.0 | 1.00 | 20.0 ± 0.0 |

Following the similar settings of [29], we generate offline datasets for both tasks, each containing 1,000 trajectories, and evaluate STO-RL alongside the baseline methods on these datasets. Following *Theorem 2*, the discount factor γ is set to $0.991 (> (T - 1)/T)$ for both environments. The results in Table 1, averaged over 100 trials, demonstrate that STO-RL achieves competitive success rates compared with the baselines. STO-RL converges to the same optimal path with shortest length and 100% success rate as all other

methods, indicating no losses in baseline offline RL performance. Furthermore, the learning curves in Figure 4 show that our STO-RL always requires fewer iterations to converge to the optimal solution, indicating a **faster learning process** relative to the baselines.

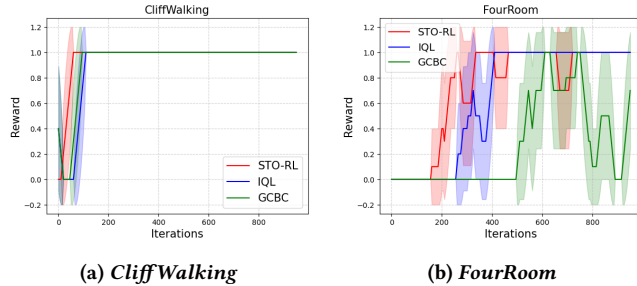


Figure 4: Learning curves for the *CliffWalking* and *FourRoom* tasks, smoothed using a moving average over 50 iterations.

The value maps for the *FourRoom* environment, presented in Figure 5, illustrate why STO-RL achieves a faster learning process. Compared with the IQL value map in Figure 5a, the subgoal temporal order based reward shaping in STO-RL produces a more accurate value landscape (Figure 5b), which better differentiates unrewarded states in the early stages of the task. In Figure 5b, the gradually increasing values along the paths to each room’s doorway guide the agent to efficiently move toward the goal, enabling rapid identification of the optimal trajectory.

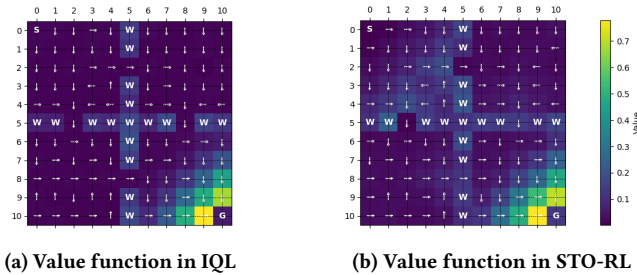


Figure 5: Value functions for the *FourRoom* task. ‘S’ denotes the start state, ‘G’ the goal state, and ‘W’ the walls. Arrows indicate the optimal actions under the learned policy.

4.2 Tasks with Continuous Action Space

In the *CliffWalking* and *FourRoom* tasks, we demonstrate the faster learning capability of STO-RL as addressed in Question (2). To better evaluate answer Question (1) and (3), we assess STO-RL on two additional, more challenging sparse-reward, goal-reaching tasks from D4RL [17]: *PointMaze-UMaze* and *PointMaze-Medium*, which feature continuous state and action spaces. The configurations of these environments are illustrated in Figure 6.

In both *PointMaze-UMaze* and *PointMaze-Medium*, the agent is modeled as a 2-DoF ball actuated by forces in the Cartesian x and y directions. The task requires the agent to navigate a closed maze to reach a designated goal. In *UMaze*, the maze has a U-shaped layout,

Table 2: Evaluation results on *PointMaze-UMaze* and *PointMaze-Medium*. Success rates and average steps after 2000 training iterations.

| Method | UMaze | | Medium | |
|---------------|--------------|--------------------|--------------|----------------------|
| | Success Rate | Steps | Success Rate | Steps |
| Datasets | 0.49 | 191.2 ± 12.7 | 0.51 | 415.1 ± 93.4 |
| IQL | 0.24 | 199.3 ± 0.0 | 0.17 | 467.7 ± 74.7 |
| GC-BC | 0.21 | 198.1 ± 4.4 | 0.58 | 406.8 ± 88.2 |
| HIQL | 0.47 | 197.0 ± 4.0 | 0.38 | 416.6 ± 62.4 |
| STO-RL | 0.68 | 195.1 ± 5.3 | 0.55 | 380.7 ± 109.0 |




Figure 6: Environments with continuous state, action spaces.

while *Medium* presents a more complex structure with multiple corridors and obstacles, increasing the difficulty of planning and exploration. To maintain consistency, the start and goal states are fixed in both environments (See details in Appendix A.3 in our arXiv version [20]). In both tasks, the agent has a 4-dimensional state space (comprising the x and y positions and velocities) and a 2-dimensional action space corresponding to the linear forces applied along x and y directions. A sparse reward of 1 is provided only when the agent reaches the goal (i.e., when the Euclidean distance between the agent and goal is less than 0.5). Episodes terminate either upon reaching the goal or when the maximum number of steps is exceeded ($T = 200$ for *UMaze*, and $T = 500$ for *Medium*) is exceeded, requiring agents to efficiently explore and identify optimal trajectories.

To facilitate the generation of ordered subgoal sequences by the LLMs, the continuous maze maps in the *PointMaze* tasks are discretized into cells with a side length of 1. These discretized maps are then vectorized in matrix form then provided as part of the prompt to the LLMs (See the details Appendix A.2 in [20]).

We generate offline datasets for both tasks, each containing 1,000 trajectories, and evaluate STO-RL alongside the baseline methods on these datasets. Following *Theorem 2*, the discount factor γ is set to 0.996 for *UMaze* and 0.999 for *Medium* ($> (T - 1)/T$). The results in Table 2, averaged over 100 trials, demonstrate that our STO-RL outperforms the baselines in both *PointMaze* tasks, by achieving the most competitive performance with almost highest success rates (0.68 for *UMaze* and 0.55 for *Medium*) and shortest solution trajectories under given offline datasets. In *Medium*, although GC-BC attains competitive success rates compared to our STO-RL, its average trajectory length is approximately 20 steps longer. Such result indicates that our STO-RL framework have better capability

in avoiding suboptimal solutions. Moreover, the learning curves in Figure 7 indicate that our STO-RL maintains its faster learning priorities in the *PointMaze* tasks. We exclude HIQL from the learning curve comparison because it trains a separate high-level policy on additional data, unlike GC-BC, IQL, and STO-RL which only train low-level agents, making direct comparison unfair. The *PointMaze* tasks feature longer horizons and larger, continuous action spaces, which make exploration and learning effective actions challenging. As a result, baselines like GC-BC and IQL achieve poorer performance with limited offline data. In contrast, our STO-RL consistently attains high performance, demonstrating superior adaptability to complex offline RL tasks. Additionally, we show that STO-RL also consistently achieves superior performance across different LLM-generated subgoal sequences, as discussed in detail in Section 4.3.

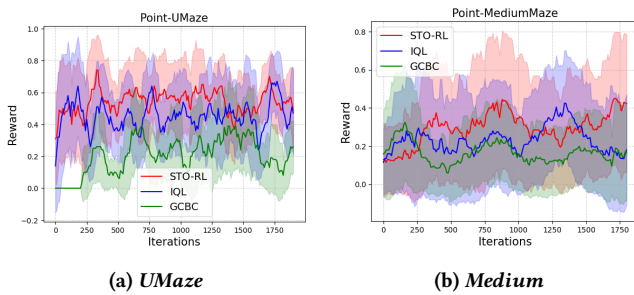


Figure 7: Learning curves for the *PointMaze-UMaze* and *PointMaze-Medium* tasks, smoothed using a moving average over 50 iterations.

4.3 Ablation Study: Subgoal Sequences

An important factor that could significantly affect proposed STO-RL’s performance is the **quality** of the subgoal sequences generated by the LLMs. Different types of language models may produce varying subgoal sequences, potentially influencing task performance. A key question, therefore, is *whether the language model can consistently generate effective subgoal sequences with correct temporal ordering, particularly in complex tasks involving environmental inference.*

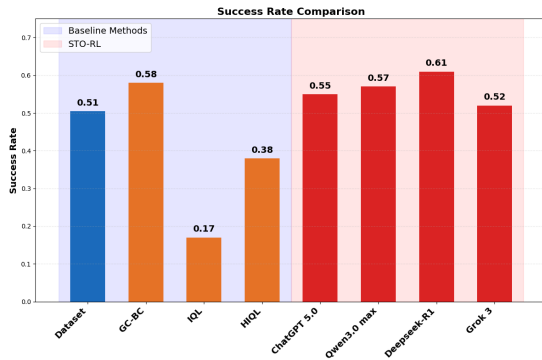


Figure 8: Success rates of STO-RL under subgoal sequences generated by different LLMs, for *PointMaze-Medium*.

To evaluate the general effectiveness of LLM-generated subgoals, we conduct an ablation study on *PointMaze-Medium* task. In this study, four widely used language models—*ChatGPT 5.0* [54], *Grok 3* [80], *DeepSeek-R1* [21], and *Qwen3-Max* [83]—are prompted to generate distinct subgoal sequences using same instructions, which are then individually applied within our STO-RL framework. Figure 9 presents the subgoal states on the maze map corresponding to the sequences generated by the four language models. We observe that, although LLMs can occasionally be misled by environmental distractions such as dead ends (as shown in Figure 9b and Figure 9d), the overall temporal ordering of the subgoals remains largely correct across all LLMs models we tested, even in complex environments like *PointMaze-Medium*. Figure 8 reports the averaged success rates over 100 trials for each of the LLM-generated subgoal sequences. These results demonstrate that, even when the subgoal sequence only partially preserves the correct temporal order, our STO-RL framework still achieves competitive performance. Therefore, our method exhibits robust general effectiveness across different language models. With the continued advancement of LLM capabilities, we anticipate further improvements in the performance of STO-RL.

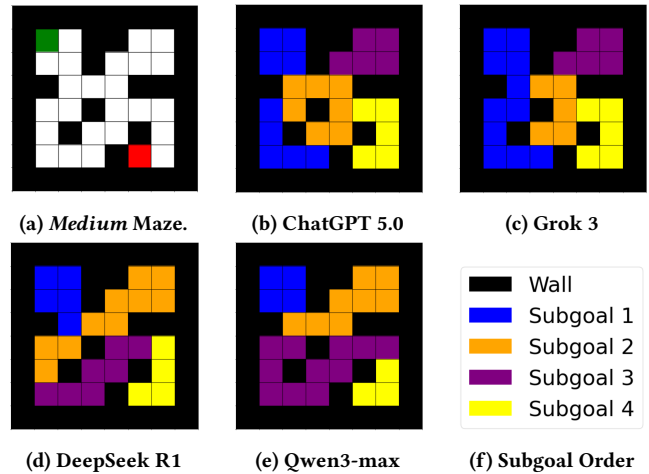


Figure 9: Subgoal Sequences generated by different LLMs. In discretized *PointMaze-Medium* (Figure 9a), green state is start, red state is goal.

Beyond those maze-based tasks, STO-RL can be extended to more complex and realistic scenarios; details are provided in the arXiv version of our paper [20].

5 CONCLUSION

In this work, we propose STO-RL, a novel framework for offline reinforcement learning guided by LLM-generated subgoal temporal order. Leveraging the reasoning capabilities of large language models (LLMs), STO-RL automatically decomposes language-based task instructions into temporally ordered subgoal sequences. Extensive experiments on discrete and continuous benchmarks (*FourRoom*, *CliffWalking*, *PointMaze-UMaze*, and *PointMaze-Medium*) show that STO-RL consistently accelerates learning, improves success rates, and yields shorter, more direct trajectories than state-of-the-art offline goal-conditioned and hierarchical RL baselines.

ACKNOWLEDGMENTS

This work was partly supported by Natural Sciences and Engineering Research Council of Canada (NSERC).

REFERENCES

- [1] Jan Achterhold, Markus Krimmel, and Joerg Stueckler. 2023. Learning temporally extended skills in continuous domains as symbolic actions for planning. In *Conference on Robot Learning*. PMLR, 225–236.
- [2] Jacob Adamczyk, Volodymyr Makarenko, Stas Tiomkin, and Rahul V Kulkarni. 2025. Bootstrapped Reward Shaping. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 39. 15302–15310.
- [3] Hongjoon Ahn, Heewoong Choi, Jisu Han, and Taesup Moon. 2025. Option-aware Temporally Abstracted Value for Offline Goal-Conditioned Reinforcement Learning. *arXiv preprint arXiv:2505.12737* (2025).
- [4] Marcin Andrychowicz, Filip Wolski, Alex Ray, Jonas Schneider, Rachel Fong, Peter Welinder, Bob McGrew, Josh Tobin, OpenAI Pieter Abbeel, and Wojciech Zaremba. 2017. Hindsight experience replay. *Advances in neural information processing systems* 30 (2017).
- [5] John Asmuth, Michael L Littman, and Robert Zinkov. 2008. Potential-based Shaping in Model-based Reinforcement Learning. In *AAAI*. 604–609.
- [6] Seungho Baek, Taegeon Park, Jongchan Park, Seungjun Oh, and Yusung Kim. 2025. Graph-Assisted Stitching for Offline Hierarchical Reinforcement Learning. *arXiv preprint arXiv:2506.07744* (2025).
- [7] Marc Bellemare, Sriram Srinivasan, Georg Ostrovski, Tom Schaul, David Saxton, and Remi Munos. 2016. Unifying count-based exploration and intrinsic motivation. *Advances in neural information processing systems* 29 (2016).
- [8] Siddhant Bhambri, Amrita Bhattacharjee, Durgesh Kalwar, Lin Guan, Huan Liu, and Subbarao Kambhampati. 2024. Extracting Heuristics from Large Language Models for Reward Shaping in Reinforcement Learning. *arXiv preprint arXiv:2405.15194* (2024).
- [9] Elliot Chane-Sane, Cordelia Schmid, and Ivan Laptev. 2021. Goal-conditioned reinforcement learning with imagined subgoals. In *International conference on machine learning*. PMLR, 1430–1440.
- [10] Yaran Chen, Wenbo Cui, Yuanwen Chen, Mining Tan, Xinyao Zhang, Jinrui Liu, Haoran Li, Dongbin Zhao, and He Wang. 2025. Robogpt: an llm-based long-term decision-making embodied agent for instruction following tasks. *IEEE Transactions on Cognitive and Developmental Systems* (2025).
- [11] Marco Colombetti and Marco Dorigo. 1992. *Robot shaping: developing situated agents through learning*. International Computer Science Institute.
- [12] Rati Devidze, Parameswaran Kamalaruban, and Adish Singla. 2022. Exploration-guided reward shaping for reinforcement learning under sparse rewards. *Advances in Neural Information Processing Systems* 35 (2022), 5829–5842.
- [13] Sam Michael Devlin and Daniel Kudenko. 2012. Dynamic potential-based reward shaping. In *11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*. IFAAMAS, 433–440.
- [14] Yiming Ding, Carlos Florensa, Pieter Abbeel, and Mariano Phielipp. 2019. Goal-conditioned imitation learning. *Advances in neural information processing systems* 32 (2019).
- [15] Yuqing Du, Olivia Watkins, Zihan Wang, Cédric Colas, Trevor Darrell, Pieter Abbeel, Abhishek Gupta, and Jacob Andreas. 2023. Guiding pretraining in reinforcement learning with large language models. In *International Conference on Machine Learning*. PMLR, 8657–8677.
- [16] Mohamed Amine Ferrag, Norbert Tihanyi, and Merouane Debbah. 2025. From llm reasoning to autonomous ai agents: A comprehensive review. *arXiv preprint arXiv:2504.19678* (2025).
- [17] Justin Fu, Aviral Kumar, Ofir Nachum, George Tucker, and Sergey Levine. 2020. D4rl: Datasets for deep data-driven reinforcement learning. *arXiv preprint arXiv:2004.07219* (2020).
- [18] Xiaozhu Gao, Jinhui Liu, Bo Wan, and Lingling An. 2024. Hierarchical reinforcement learning from demonstration via reachability-based reward shaping. *Neural Processing Letters* 56, 3 (2024), 184.
- [19] Adrian Goldwaser and Michael Thielscher. 2020. Deep reinforcement learning for general game playing. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 34. 1701–1708.
- [20] Chengyang Gu, Yuxin Pan, Hui Xiong, and Yize Chen. 2026. STO-RL: Offline RL under Sparse Rewards via LLM-Guided Subgoal Temporal Order. *arXiv preprint arXiv:2601.08107* (2026).
- [21] Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirog Ma, Peiyi Wang, Xiao Bi, et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948* (2025).
- [22] Anna Harutyunyan, Sam Devlin, Peter Vranx, and Ann Nowé. 2015. Expressing arbitrary reward functions as potential-based advice. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 29.
- [23] Christopher Hoang, Sungryull Sohn, Jongwook Choi, Wilka Carvalho, and Honglak Lee. 2021. Successor feature landmarks for long-horizon goal-conditioned reinforcement learning. *Advances in neural information processing systems* 34 (2021), 26963–26975.
- [24] Ian Holmes and Min Chi. 2025. Attention-Based Reward Shaping for Sparse and Delayed Rewards. *arXiv preprint arXiv:2505.10802* (2025).
- [25] Zhiao Huang, Fangchen Liu, and Hao Su. 2019. Mapping state space using landmarks for universal goal reaching. *Advances in Neural Information Processing Systems* 32 (2019).
- [26] David Isele, Alireza Nakhaei, and Kikuo Fujimura. 2018. Safe reinforcement learning on autonomous vehicles. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 1–6.
- [27] Riashat Islam, Hongyu Zang, Anirudh Goyal, Alex M Lamb, Kenji Kawaguchi, Xin Li, Romain Laroche, Yoshua Bengio, and Remi Tachet des Combes. 2022. Discrete compositional representations as an abstraction for goal conditioned reinforcement learning. *Advances in Neural Information Processing Systems* 35 (2022), 3885–3899.
- [28] Zhiwei Jia, Vineet Thummuluri, Fangchen Liu, Linghao Chen, Zhiao Huang, and Hao Su. 2024. Chain-of-thought predictive control. In *Proceedings of the 41st International Conference on Machine Learning*. 21768–21790.
- [29] Li Jiang, Sijie Cheng, Jieli Qiu, Haoran Xu, Wai Kin Chan, and Zhao Ding. 2023. Offline reinforcement learning with imbalanced datasets. *arXiv preprint arXiv:2307.02752* (2023).
- [30] Junsu Kim, Younggyo Seo, and Jinwoo Shin. 2021. Landmark-guided subgoal generation in hierarchical reinforcement learning. *Advances in neural information processing systems* 34 (2021), 28336–28349.
- [31] Cevahir Koprulu, Po-Han Li, Tianyu Qiu, Ruihan Zhao, Tyler Westbroek, David Fridovich-Keil, Sandeep Chinchali, and Ufuk Topcu. 2025. Dense Dynamics-Aware Reward Synthesis: Integrating Prior Experience with Demonstrations. In *7th Annual Learning for Dynamics & Control Conference*. PMLR, 894–906.
- [32] Ilya Kostrikov, Ashvin Nair, and Sergey Levine. 2021. Offline reinforcement learning with implicit q-learning. *arXiv preprint arXiv:2110.06169* (2021).
- [33] Kalle Kujanpää, Joni Pajarinen, and Alexander Ilin. 2023. Hierarchical imitation learning with vector quantized models. In *International Conference on Machine Learning*. PMLR, 17896–17919.
- [34] Aviral Kumar, Aurick Zhou, George Tucker, and Sergey Levine. 2020. Conservative q-learning for offline reinforcement learning. *Advances in neural information processing systems* 33 (2020), 1179–1191.
- [35] Adam Laud and Gerald DeJong. 2003. The influence of reward on the speed of reinforcement learning: An analysis of shaping. In *Proceedings of the 20th International Conference on Machine Learning (ICML-03)*. 440–447.
- [36] Alexander Li, Lerrel Pinto, and Pieter Abbeel. 2020. Generalized hindsight for reinforcement learning. *Advances in neural information processing systems* 33 (2020), 7754–7767.
- [37] Lin Li, Runjia Tan, Jianwu Fang, Jianru Xue, and Chen Lv. 2025. LLM-augmented hierarchical reinforcement learning for human-like decision-making of autonomous driving. *Expert Systems with Applications* 294 (2025), 128736.
- [38] Manling Li, Shiyu Zhao, Qineng Wang, Kangrui Wang, Yu Zhou, Sanjana Srivastava, Cem Gokmen, Tony Lee, Erran Li Li, Ruohan Zhang, et al. 2024. Embodied agent interface: Benchmarking llms for embodied decision making. *Advances in Neural Information Processing Systems* 37 (2024), 100428–100534.
- [39] Bo Liu, Yuqian Jiang, Xiaohan Zhang, Qiang Liu, Shiqi Zhang, Joydeep Biswas, and Peter Stone. 2023. Llm+ p: Empowering large language models with optimal planning proficiency. *arXiv preprint arXiv:2304.11477* (2023).
- [40] Chenghao Liu, Fei Zhu, Quan Liu, and Yuchen Fu. 2021. Hierarchical reinforcement learning with automatic sub-goal identification. *IEEE/CAA journal of automatica sinica* 8, 10 (2021), 1686–1696.
- [41] Minghuan Liu, Menghui Zhu, and Weinan Zhang. 2022. Goal-conditioned reinforcement learning: Problems and solutions. *arXiv preprint arXiv:2201.08299* (2022).
- [42] Lajanugen Logeswaran, Yao Fu, Moontae Lee, and Honglak Lee. 2022. Few-shot subgoal planning with language models. *arXiv preprint arXiv:2205.14288* (2022).
- [43] Fei Ma, Guanjun Liu, and Kaiwen Zhang. 2021. Goal-conditioned Behavioral Cloning with Prioritized Sampling. In *2021 IEEE International Conference on Networking, Sensing and Control (ICNSC)*, Vol. 1. IEEE, 1–6.
- [44] Haozhe Ma, Fangling Li, Jing Yu Lim, Zhengding Luo, Thanh Vinh Vo, and Tze-Yun Leong. 2025. Catching Two Birds with One Stone: Reward Shaping with Dual Random Networks for Balancing Exploration and Exploitation. In *Forty-second International Conference on Machine Learning*.
- [45] Jason Yecheng Ma, Jason Yan, Dinesh Jayaraman, and Osbert Bastani. 2022. Offline goal-conditioned reinforcement learning via f -advantage regression. *Advances in neural information processing systems* 35 (2022), 310–323.
- [46] Farzan Memarian, Wonjoon Goo, Rudolf Lioutikov, Scott Niekum, and Ufuk Topcu. 2021. Self-supervised online reward shaping in sparse-reward environments. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2369–2375.
- [47] Amirhossein Mesbah, Reshad Hosseini, Seyed Pooya Shariatpanahi, and Majid Nili Ahmadabadi. 2024. Subgoal Discovery Using a Free Energy Paradigm and State Aggregations. *arXiv preprint arXiv:2412.16687* (2024).

- [48] Lina Mezghani, Sainbayar Sukhbaatar, Piotr Bojanowski, Alessandro Lazaric, and Karteek Alahari. 2023. Learning goal-conditioned policies offline with self-supervised reward shaping. In *Conference on robot learning*. PMLR, 1401–1410.
- [49] Ofir Nachum, Shixiang Shane Gu, Honglak Lee, and Sergey Levine. 2018. Data-efficient hierarchical reinforcement learning. *Advances in neural information processing systems* 31 (2018).
- [50] Ofir Nachum, Haoran Tang, Xingyu Lu, Shixiang Gu, Honglak Lee, and Sergey Levine. 2019. Why does hierarchy (sometimes) work so well in reinforcement learning? *arXiv preprint arXiv:1909.10618* (2019).
- [51] Andrew Y Ng, Daishi Harada, and Stuart Russell. 1999. Policy invariance under reward transformations: Theory and application to reward shaping. In *ICML*, Vol. 99. Citeseer, 278–287.
- [52] Hai Nguyen and Hung La. 2019. Review of deep reinforcement learning for robot manipulation. In *2019 Third IEEE international conference on robotic computing (IRC)*. IEEE, 590–595.
- [53] Takato Okudo and Seiji Yamada. 2021. Subgoal-based reward shaping to improve efficiency in reinforcement learning. *IEEE Access* 9 (2021), 97557–97568.
- [54] OpenAI. 2025. GPT-5 System Card. <https://openai.com/research/gpt-5-system-card>. Accessed: 2025-10-08.
- [55] Takayuki Osa, Voot Tangkaratt, and Masashi Sugiyama. 2019. Hierarchical reinforcement learning via advantage-weighted information maximization. *arXiv preprint arXiv:1901.01365* (2019).
- [56] Georg Ostrovski, Marc G Bellemare, Aaron Oord, and Rémi Munos. 2017. Count-based exploration with neural density models. In *International conference on machine learning*. PMLR, 2721–2730.
- [57] Sherjil Ozair, Yazhe Li, Ali Razavi, Ioannis Antonoglou, Aaron Van Den Oord, and Oriol Vinyals. 2021. Vector quantized models for planning. In *international conference on machine learning*. PMLR, 8302–8313.
- [58] Seohong Park, Kevin Frans, Benjamin Eysenbach, and Sergey Levine. 2024. Ogbench: Benchmarking offline goal-conditioned rl. *arXiv preprint arXiv:2410.20092* (2024).
- [59] Seohong Park, Dibya Ghosh, Benjamin Eysenbach, and Sergey Levine. 2023. Higl: Offline goal-conditioned rl with latent states as actions. *Advances in Neural Information Processing Systems* 36 (2023), 34866–34891.
- [60] Thomas Pouplin, Katarzyna Kobalczuk, Hao Sun, and Mihaela van der Schaar. 2024. The Synergy of LLMs & RL Unlocks Offline Learning of Generalizable Language-Conditioned Policies with Low-fidelity Data. *arXiv preprint arXiv:2412.06877* (2024).
- [61] Bharat Prakash, Tim Oates, and Tinoosh Mohsenin. 2023. LLM augmented hierarchical agents. *arXiv preprint arXiv:2311.05596* (2023).
- [62] Wenjie Qiu, Wensen Mao, and He Zhu. 2023. Instructing goal-conditioned reinforcement learning agents with temporal logic objectives. *Advances in Neural Information Processing Systems* 36 (2023), 39147–39175.
- [63] Jette Randløv and Preben Alstrøm. 1998. Learning to Drive a Bicycle Using Reinforcement Learning and Shaping. In *ICML*, Vol. 98. 463–471.
- [64] Moritz Reuss, Maximilian Li, Xiaogang Jia, and Rudolf Lioutikov. 2023. Goal-conditioned imitation learning using score-based diffusion policies. *arXiv preprint arXiv:2304.02532* (2023).
- [65] Erick Rosete-Beas, Oier Mees, Gabriel Kalweit, Joschka Boedecker, and Wolfram Burgard. 2023. Latent plans for task-agnostic offline reinforcement learning. In *Conference on Robot Learning*. PMLR, 1838–1849.
- [66] Swarnadeep Saha, Archiki Prasad, Justin Chih-Yao Chen, Peter Hase, Elias Stengel-Eskin, and Mohit Bansal. 2024. System-1. x: Learning to balance fast and slow planning with language models. *arXiv preprint arXiv:2407.14414* (2024).
- [67] Devin Schwab and Soumya Ray. 2017. Offline reinforcement learning with task hierarchies. *Machine Learning* 106, 9 (2017), 1569–1598.
- [68] Chak Lam Shek and Pratap Tokekar. 2025. Option Discovery Using LLM-guided Semantic Hierarchical Reinforcement Learning. *arXiv preprint arXiv:2503.19007* (2025).
- [69] Wonchul Shin and Yusung Kim. 2023. Guide to Control: Offline Hierarchical Reinforcement Learning Using Subgoal Generation for Long-Horizon and Sparse-Reward Tasks. In *IJCAI*. 4217–4225.
- [70] Yash Shukla, Wenchang Gao, Vasanth Sarathy, Alvaro Velasquez, Robert Wright, and Jivko Sinapov. 2023. Lgts: Dynamic task sampling using llm-generated sub-goals for reinforcement learning agents. *arXiv preprint arXiv:2310.09454* (2023).
- [71] Tom Silver, Varun Hariprasad, Reece S Shuttleworth, Nishanth Kumar, Tomás Lozano-Pérez, and Leslie Pack Kaelbling. 2022. PDDL planning with pretrained large language models. In *NeurIPS 2022 foundation models for decision making workshop*.
- [72] Shivam Singh, Karthik Swaminathan, Nabanita Dash, Ramandeep Singh, Snehasis Banerjee, Mohan Sridharan, and Madhava Krishna. 2025. AdaptBot: Combining LLM with knowledge graphs and human input for generic-to-specific task decomposition and knowledge refinement. *arXiv preprint arXiv:2502.02067* (2025).
- [73] Chan Hee Song, Jiaman Wu, Clayton Washington, Brian M Sadler, Wei-Lun Chao, and Yu Su. 2023. Llm-planner: Few-shot grounded planning for embodied agents with large language models. In *Proceedings of the IEEE/CVF international conference on computer vision*. 2998–3009.
- [74] Vivienne Huiling Wang, Joni Pajarinen, Tinghuai Wang, and Joni-Kristian Kämäräinen. 2023. State-conditioned adversarial subgoal generation. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 37. 10184–10191.
- [75] Vivienne Huiling Wang, Tinghuai Wang, Wenyang Yang, Joni-Kristian Kämäräinen, and Joni Pajarinen. 2024. Probabilistic subgoal representations for hierarchical reinforcement learning. *arXiv preprint arXiv:2406.16707* (2024).
- [76] Eric Wiewiora. 2003. Potential-based shaping and Q-value initialization are equivalent. *Journal of Artificial Intelligence Research* 19 (2003), 205–208.
- [77] Eric Wiewiora, Garrison W Cottrell, and Charles Elkan. 2003. Principled methods for advising reinforcement learning agents. In *Proceedings of the 20th international conference on machine learning (ICML-03)*. 792–799.
- [78] Mingkang Wu, Umer Siddique, Abhinav Sinha, and Yongcan Cao. 2024. Offline reinforcement learning with failure under sparse reward environments. In *2024 IEEE 3rd International Conference on Computing and Machine Intelligence (ICMI)*.
- [79] Zhenyu Wu, Ziwei Wang, Xiuwei Xu, Jiwen Lu, and Haibin Yan. 2023. Embodied task planning with large language models. *arXiv preprint arXiv:2307.01848* (2023).
- [80] xAI. 2025. Grok 3 Beta: The Age of Reasoning Agents. <https://x.ai/news/grok-3>. Accessed: 2025-10-08.
- [81] Tianbao Xie, Siheng Zhao, Chen Henry Wu, Yitao Liu, Qian Luo, Victor Zhong, Yanhao Yang, and Tao Yu. 2023. Text2reward: Reward shaping with language models for reinforcement learning. *arXiv preprint arXiv:2309.11489* (2023).
- [82] Duo Xu and Faramarz Fekri. 2024. Learning Hidden Subgoals under Temporal Ordering Constraints in Reinforcement Learning. *arXiv preprint arXiv:2411.01425* (2024).
- [83] An Yang, Anpeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. 2025. Qwen3 technical report. *arXiv preprint arXiv:2505.09388* (2025).
- [84] Rui Yang, Yiming Lu, Wenzhe Li, Hao Sun, Meng Fang, Yali Du, Xiu Li, Lei Han, and Chongjie Zhang. 2022. Rethinking goal-conditioned supervised learning and its connection to offline rl. *arXiv preprint arXiv:2202.04478* (2022).
- [85] Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. 2023. React: Synergizing reasoning and acting in language models. In *International Conference on Learning Representations (ICLR)*.
- [86] Gregory Yauney and Pratik Shah. 2018. Reinforcement learning with action-derived rewards for chemotherapy and clinical trial dosing regimen selection. In *Machine Learning for Healthcare Conference*. PMLR, 161–226.
- [87] Ilker Yildirim and LA Paul. 2024. From task structures to world models: what do LLMs know? *Trends in Cognitive Sciences* 28, 5 (2024), 404–415.
- [88] David Yunis, Justin Jung, Falcon Dai, and Matthew Walter. 2024. Subwords as Skills: Tokenization for Sparse-Reward Reinforcement Learning. *Advances in Neural Information Processing Systems* 37 (2024), 67663–67688.
- [89] Jesse Zhang, Haonan Yu, and Wei Xu. 2021. Hierarchical reinforcement learning by discovering intrinsic options. *arXiv preprint arXiv:2101.06521* (2021).
- [90] Zeyu Zheng, Junhyuk Oh, and Satinder Singh. 2018. On learning intrinsic rewards for policy gradient methods. *Advances in neural information processing systems* 31 (2018).
- [91] Tianxing Zhou, Zhirui Wang, Haojia Ao, Guangyan Chen, Boyang Xing, Jingwen Cheng, Yi Yang, and Yufeng Yue. 2025. STEP Planner: Constructing cross-hierarchical subgoal tree as an embodied long-horizon task planner. *arXiv preprint arXiv:2506.21030* (2025).