

Q-WSL:Leveraging Dynamic Programming for Weighted Supervised Learning in Goal-conditioned RL

Xing Lei^a, Xuetao Zhang^{a,*}, Zifeng Zhuang^b and Donglin Wang^b

^a*Institute of Artificial Intelligence and Robotics, Xi'an Jiaotong University, Xi'an, 710049, Shaanxi, China*

^b*Machine Intelligence Laboratory (MiLAB), School of Engineering, Westlake University, Hangzhou, China*

ARTICLE INFO

Keywords:

Goal-conditioned RL
Dynamic Programming
Weighted Supervised Learning

ABSTRACT

A novel class of advanced algorithms, termed Goal-Conditioned Weighted Supervised Learning (GCWSL), has recently emerged to tackle the challenges posed by sparse rewards in goal-conditioned reinforcement learning (RL). GCWSL consistently delivers strong performance across a diverse set of goal-reaching tasks due to its simplicity, effectiveness, and stability. However, GCWSL methods lack a crucial capability known as trajectory stitching, which is essential for learning optimal policies when faced with unseen skills during testing. This limitation becomes particularly pronounced when the replay buffer is predominantly filled with sub-optimal trajectories. In contrast, traditional TD-based RL methods, such as Q-learning, which utilize Dynamic Programming, do not face this issue but often experience instability due to the inherent difficulties in value function approximation. In this paper, we propose Q-learning Weighted Supervised Learning (Q-WSL), a novel framework designed to overcome the limitations of GCWSL by incorporating the strengths of Dynamic Programming found in Q-learning. Q-WSL leverages Dynamic Programming results to output the optimal action of (state, goal) pairs across different trajectories within the replay buffer. This approach synergizes the strengths of both Q-learning and GCWSL, effectively mitigating their respective weaknesses and enhancing overall performance. Empirical evaluations on challenging goal-reaching tasks demonstrate that Q-WSL surpasses other goal-conditioned approaches in terms of both performance and sample efficiency. Additionally, Q-WSL exhibits notable robustness in environments characterized by binary reward structures and environmental stochasticity.

1. Introduction

Deep reinforcement learning (RL) empowers agents to attain sophisticated objectives in complex and uncertain environments, such as computer games (Hou et al., 2022; Liu et al., 2023; Kim et al., 2024), robot control (Xu and Wang, 2022; Bouktif et al., 2023; Xiao et al., 2024), and language processing (Sharifani and Amini, 2023; Zeng et al., 2024). One of the central challenges in deep RL is facilitating efficient learning in environments characterized by sparse rewards. This issue is particularly acute in goal-conditioned RL (GCRL) (Kaelbling, 1993; Schaul et al., 2015; Andrychowicz et al., 2017; Liu et al., 2022), where the agent is tasked with learning generalized policies that can reach a variety of goals.

Recently, we note that several goal-conditioned weighted supervised learning (GCWSL) methods have been proposed (Ghosh et al., 2021; Yang et al., 2022; Ma et al., 2022; Hejna et al., 2023) to tackle the GCRL challenges. In contrast to conventional RL approaches that maximizes discounted cumulative return, GCWSL provides theoretical guarantees that supervised learning from hindsight relabeled experiences optimizes a lower bound on the goal-conditioned RL objective. Specifically, these methods first retrospectively infer which goals would have been fulfilled by a trajectory,

irrespective of whether it successfully reached the desired goal. After this hindsight relabeling, GCWSL aims to imitate the corresponding sub-trajectories of (state, goal) pairs within the same trajectory in the relabeled data. Due to the benefit of the weighted function, GCWSL, unlike behavior cloning, can identify potentially optimal trajectories. Compared with other goal-conditioned RL or self-supervised (SL) methods (Ding et al., 2019; Chen et al., 2020; Lynch et al., 2020; Paster et al., 2020; Eysenbach et al., 2020; Ghosh et al., 2021; Eysenbach et al., 2022), GCWSL has demonstrated outstanding performance across various goal-reaching tasks in a simple, effective, and stable manner.

Despite the successful application of GCWSL in effectively learning sparse rewards for certain goal-reaching tasks, some studies (Brandfonbrener et al., 2022; Yang et al., 2023; Ghugare et al., 2024) indicate that GCWSL may leads to sub-optimal policies when applied to the corresponding sub-trajectories of (state, goal) pairs across different trajectories and identify this issue as lacking of the ability to stitch trajectories. To effectively address this challenge, it is crucial to amalgamate information from multiple trajectories. On the other hand, we demonstrate that Q-learning possesses this ability due to its Dynamic Programming framework, leading to better policy outcomes. However, Q-learning faces certain challenges in environments with sparse rewards. One of the key issues is the difficulty in propagating the value function back to the initial state, primarily due to the use of function approximation (Sutton and Barto, 1998).

*Corresponding author

**Principal corresponding author

✉ leixing@stu.xjtu.edu.cn (X. Lei); xuetaozh@xjtu.edu.cn (X. Zhang); zhuangzifeng@westlake.edu.cn (Z. Zhuang); wangdonglin@westlake.edu.cn (D. Wang)

ORCID(s):

In this work, we introduce Q-WSL, a novel method designed to address the aforementioned challenges by leveraging Dynamic Programming within Q-learning to enhance the performance of GCWSL. Unlike many goal-conditioned RL algorithms that propose new architectures for agents to achieve better performance, Q-WSL focuses on improving the quality of the replay buffer, enabling more effective utilization of existing goal-conditioned RL algorithms. Our approach is motivated by the following insight: while GCWSL often learns suboptimal policies for certain sub-trajectories of (state, goal) pairs across different trajectories, Q-learning can derive better actions. However, Q-learning tends to struggle when dealing with states that require backpropagation over large time steps. In these cases, GCWSL offers a distinct advantage by circumventing the need for backpropagation. By combining the strengths of both methods, Q-WSL mitigates their respective weaknesses, leading to a more robust and efficient performance. To evaluate the effectiveness of Q-WSL, we conduct experiments using challenging GCRL benchmark environments. The results demonstrate that Q-WSL significantly improves sample efficiency compared to previous GCWSL approaches and other competitive GCRL methods, including DDPG+HER (Andrychowicz et al., 2017), ActionModels (Chebotar et al., 2021a), and Model-based HER (Yang et al., 2021). Additionally, Q-WSL exhibits strong robustness to variations in reward functions and environmental stochasticity.

We briefly summarize our contributions:

- While state-of-the-art GCWSL methods tend to yield suboptimal performance for certain unseen skills during testing and lack trajectory stitching capabilities, we demonstrate that Q-learning inherently possesses this ability and can produce optimal actions in such scenarios.
- We introduce a novel and efficient GCRL method, Q-WSL, which harnesses the stitching capability of Q-learning to overcome the limitations of GCWSL. Additionally, we demonstrate that the optimization lower bound of Q-WSL exceeds that of several competitive GCRL algorithms, including WGCSL (Yang et al., 2022) and GCSL (Ghosh et al., 2021).
- Experimental results on several GCRL benchmarks show that Q-WSL obtains better generalization performance and sample efficiency compared to prior GCRL methods. Additionally, Q-WSL is robust to variations in reward forms and environmental stochasticity.

2. Related Works

The Stitching Property The concept of stitching, as discussed by (Ziebart et al., 2008), is a characteristic feature of TD-based RL algorithms, such as those described by (Mnih et al., 2013), (Lillicrap et al., 2015), (Fujimoto et al., 2018), and (Kostrikov et al., 2021), which employ Dynamic Programming techniques. This feature allows these algorithms to integrate data from different trajectories, thereby

enhancing their effectiveness in managing complex tasks by leveraging historical data (Cheikh and Russo, 2023). Recent works have shown that self-supervised learning also exhibits such property, particularly through methods like data augmentation (Shorten and Khoshgoftaar, 2019; Ghugare et al., 2024). However, there remains a significant performance gap between these methods and algorithms that incorporate Dynamic Programming.

This property motivates us to employ simple yet efficient Q-learning in TD-based RL to reach previously unvisited goals across different trajectories by leveraging Dynamic Programming. However, it is not clear that whether Q-learning process this important stitching property in GCRL. Using a straightforward counterexample, we demonstrate that stitching enables a unique property: the ability to infer solutions for composable goal-reaching trajectories during testing time. This includes maneuvering through specific (state, goal) pairs that never appear together during training, although they do appear separately.

Goal-conditioned Weighted Supervised Learning

Goal-conditioned weighted supervised learning methods (GCWSL) (Liu et al., 2021; Yang et al., 2022; Ma et al., 2022; Hejna et al., 2023) are new outstanding goal-conditioned algorithms which enable agent to reach multiple goals in GCRL. These methods repeatedly perform imitation learning on self-collected, relabeled data without any RL updates. Furthermore, these methodologies undergo theoretical validation to confirm that self-supervised learning from relabeled experiences optimizes a lower bound on the GCRL objective. This validation guarantees that the learning process is both efficient and underpinned by a rigorous theoretical foundation.

However, recent researches (Brandfonbrener et al., 2022; Yang et al., 2023; Ghugare et al., 2024) indicate that these methods are sub-optimal for some unseen skills and lack the ability to stitch information from multiple trajectories. In response, our method, Q-WSL, which extends the framework of GCWSL by leveraging Dynamic Programming in Q-learning, can specifically designed to achieve optimal policy outcomes and integrates the capability for trajectory stitching.

3. Preliminaries

3.1. Goal-conditioned Reinforcement Learning

Goal-conditioned Reinforcement Learning (GCRL) can be characterized by the tuple $\langle S, \mathcal{A}, \mathcal{G}, \mathcal{P}, r, \gamma, \rho_0, T \rangle$, where $S, \mathcal{A}, \mathcal{G}, \gamma, \rho_0$ and T respectively represent the state space, action space, goal space, discounted factor, the distribution of initial states and the horizon of the episode. $\mathcal{P} : \mathcal{P}(s'|s, a)$ is the dynamic transition function, and $r : r(s, a, g)$ is typically a simple unshaped binary signal. The sparse reward function employed in GCRL can be expressed as follows:

$$r(s_t, a_t, g) = \begin{cases} 0, & \|\phi(s_t) - g\|_2^2 < \delta \\ -1, & \text{otherwise} \end{cases}, \quad (1)$$

δ is a threshold and $\phi : S \rightarrow \mathcal{G}$ is a known state-to-goal mapping function from states to goals. Agents must develop

a policy $\pi : \mathcal{S} \times \mathcal{G} \rightarrow \mathcal{A}$ that optimizes the expected returns associated with reaching goals sampled from the goal distribution $p(g)$:

$$\mathcal{J}(\pi) = E_{g \sim p(g), a_t \sim \pi, s_{t+1} \sim p(\cdot | s_t, a_t)} \left[\sum_{t=0}^T \gamma^t r(s_t, a_t, g) \right] \quad (2)$$

Q-learning methods optimizes $\mathcal{J}(\pi)$ via the gradient computation:

$$\mathcal{J}(\pi) = E_{g, a_t, s_{t+1}} \left[\sum_{t=0}^T Q^\pi(s_t, a_t, g) \right], \quad (3)$$

where $Q^\pi(s_t, a_t, g)$ is the goal-conditioned Q-function. Note that in $\mathbb{1}[\phi(s_t) = g]$ reward (i.e. If the state s reaches the goal g , then the reward is 1; otherwise, it is 0.), the $Q^\pi(s_t, a_t, g)$ is always zero. This Q-function is not useful for predicting the future goal distribution.

3.2. Goal-conditioned Weighted Supervised Learning

In contrast to goal-conditioned RL methods, which focus on directly optimizing the discounted cumulative return, Goal-Conditioned Weighted Supervised Learning methods (GCWSL) introduce an innovative learning framework by iteratively relabeling and imitating self-generated experiences. This iterative process enhances the model's ability to generalize and efficiently learn from past experiences, thereby improving its overall performance in goal-reaching tasks. Trajectories in the replay buffer are relabeled with hindsight method (Andrychowicz et al., 2017). And then the policy optimization satisfies the following definition during imitating:

Definition 1. Given replay buffer B and hindsight relabeling mechanism, the objective of GCWSL is to mimic those relabeled transitions through supervised learning with weighted function:

$$\mathcal{J}_{GCWSL}(\pi) = E_{(s_t, a_t, g) \sim B_r} [f(A) \cdot \log \pi_\theta(a_t | s_t, g)], \quad (4)$$

where B_r denotes relabeled replay buffer, $g = \phi(s_i)$ denotes the relabeled goals for $i \geq t$ and $f(A)$ is a function about advantage $A := A(s_t, a_t, g)$. The weighted function $f(A)$ exists various forms in GCWSL methods (Ghosh et al., 2021; Yang et al., 2022; Ma et al., 2022; Hejna et al., 2023). (See Appendix A.3 for the specific weight function meaning). Therefore GCWSL includes typical two process, acquiring sub-trajectories corresponding to state-goal (i.e. (s, g)) pairs and imitating them. In the process of imitation, GCWSL first train the related specific weighted function, and then extract the policy with the Eq. 4.

We briefly summarize the GCWSL training procedure in Algorithm 1, where $\beta(a|s, g) := \pi_{old}(a|s, g)$ in practice. And then we have following theorem.

Theorem 1. Suppose there exist a deterministic policy $\pi_{relabel}$ capable of generating relabeled data B_r , GCWSL

Algorithm 1 Goal-conditioned self-supervised learning with weight function

- 1: **Input** behavior policy $\beta(a | s, g)$
- 2: $B_r \leftarrow \text{RELABEL}(\tau)$ where $\tau \sim \beta(\tau)$
- 3: $\mathcal{J}_{GCWSL}(\pi) \leftarrow E_{(s_t, a_t, g) \sim B_r} [f(A) \cdot \log \pi(a_t | s_t, g)]$
- 4: $\pi(a | s, g) \leftarrow \arg \max_\pi \mathcal{J}_{GCWSL}(\pi)$
- 5: **return** $\pi(a | s, g)$

is a variant form of goal-conditioned AWR over B_r :

$$\pi_{k+1} = \arg \max_{\pi \in \Pi} E_{(s, g, s') \sim B_r, a_t \sim \pi(s, g)} [A^{\pi_k}(s, a, g)], \quad (5)$$

s.t. $\mathcal{D}_{KL}(\pi \| \pi_{relabel}) < \epsilon$,

where $A^{\pi_k} = Q^{\pi_k}(s, a, g) - V^{\pi_k}(s, g)$, \mathcal{D}_{KL} is the KL-divergence and g is the goals from relabeled data. See Appendix A.1 for detail proof. This theorem indicate GCWSL guarantees policy improvement when learning from the relabelled data.

4. Q-WSL: Q-learning Weighted Supervised Learning

To generalize GCWSL's skills outside the relabeled data and further improve sample efficiency, we start with a counter example to illustrate why Q-learning in TD-based RL can solve the experience stitching problem in Section 4.1. Then we propose a modified of GCWSL called Q-WSL that leveraging the stitching ability of Q-learning in Section 4.2. In Section 4.2, we also provide theoretical guarantee for Q-WSL. Further, we introduce the overall algorithm implementation process of Q-WSL in Section 4.3. Finally, we summarize and compare Q-WSL with previous related work in Section 4.4.

4.1. Trajectory Stitching Example

From the perspective of GCWSL updating the corresponding (s, g) pairs, GCWSL aims to identify and imitate the optimal sub-trajectory from the combination data generated by the old policy that meets the (s, g) pairs. The optimal trajectory here satisfies the weighted function $f(A)$, such as the the horizon length from s to g in (Yang et al., 2022). However, it has been observed that when data is aggregated from a variety of older policies, certain (s, g) pairs are rarely encountered within the same trajectory, despite being common in different trajectories. In these instances, GCWSL often struggles to produce an optimal policy.

Consider an example illustrated in Fig. 1, where the goals $g \in [g^a, g^b]$ are reachable from the states $s \in [s_0^a, s_0^b, s_t]$ although not explicitly present in the off-policy replay buffer. During training, GCWSL is updated with (state, goal) pairs sampled from the same trajectory, such as (s_0^a, g^a) and (s_0^b, g^b) . However, during testing, when unseen (state, goal) pairs such as (s_0^a, g^b) , (s_0^b, g^a) originating from distinct trajectories are introduced, GCWSL cannot find the correct actions (proof can be found in Figure 1 and Lemma 4.1 of (Ghugare et al., 2024)). In contrast, TD-based RL methods such as Q-learning can output the correct actions




Fig. 1. A counterexample of GCWSL lack stitching ability while Q-learning can output the path across different trajectories. $g \in [g^a, g^b]$ are conveniently expressed as reachable goals. Specifically, the off-policy replay buffer stores history data collected by behavior policy $\beta(a|s, g)$ during training, and then $f(A) \cdot \log \pi(a|s, g)$ is maximized with batch (state, goal) pairs from this buffer. During testing, the $\pi(a|s, g)$ is tested on random batch (state, goal) pairs. GCWSL can consistently learn (state, goal) pairs within the same trajectory (i.e., $(s_0^a, g^a), (s_0^b, g^b)$), which are represented by the same color. However, it lacks the guarantee of generating the correct actions for (state, goal) pairs that originate from different trajectories (i.e., $(s_0^a, g^b), (s_0^b, g^a)$). In contrast, Q-learning can output optimal actions for these (state, goal) pairs across different trajectories thanks to Dynamic Programming.

for $(s_0^a, g^b), (s_0^b, g^a)$ due to the Dynamic Programming component. For instance, consider a value estimation example, $V(s, g) = r + V(s_{t+1}, g)$. For the pair (s_{t+1}, g) , provided that g is accessible and the value $V(s_{t+1}, g)$ is accurately approximated, Q-learning can back-propagate the state-goal (state-action-goal) values to preceding pairs. An example from the off-policy replay buffer illustrates the state s_t as a "hub state" shared by two initial trajectories, from which all goals $g^i \sim [g^a, g^b]$ are reachable. The values of these goal-conditioned states $Q(s_t, a_t^i, g^i)$ can be computed and recursively back-propagated to $Q(s_0, a_0^i, g^i), i \in [a, b]$. Q-learning has the capability to propagate rewards through paths that are connected retrospectively. This allows Q-learning to determine the optimal actions for state-goal pairs that have not been encountered together previously, such as (s_0^a, g^b) and (s_0^b, g^a) . By iteratively back-propagating value estimates through all potential trajectory combinations, Q-learning ensures that values are consistently updated, thereby facilitating the correct action selection. In summary, Q-learning can link (state, goal) pairs across different trajectories, a capability that GCWSL lacks.

4.2. Bounding KL-Constrained Values and Q-WSL

If we replace the advantage function $A^\pi(s, a, g)$ in Eq. 5 by Q-function $Q^\pi(s, a, g)$ we have the policy improvement objective:

$$\pi_{k+1} = \arg \max_{\pi} \mathbb{E} [Q^{\pi_k}(s, \pi_k(s, g), g)], \quad \int_a \pi(a|s, g) da = 1, \text{ s.t. } \text{KL}(\pi \| \pi_{\text{relabel}}) \leq \epsilon. \quad (6)$$

Since minimizing the KL-divergence is equivalent to maximizing the likelihood (LeCun et al., 2015), we have the following Lagrangian:

$$\mathcal{L}(\lambda, \pi^H) = \mathbb{E}_{a \sim \pi(\cdot|s, g)} [Q^\pi(s, a, g)] + \lambda \mathbb{E}_{a, s, g \sim B_r} \log(\pi(a|s, g)). \quad (7)$$

For a deterministic policy π , it can be regarded as a Dirac-Delta function. Consequently, the constraint $\int_a \pi(a|s, g) da = 1$ is always satisfied. Therefore, Equation Eq. 7 simplifies to:

$$\arg \min_{\pi} \mathbb{E}_{a, s, g \sim B_r} [-Q^\pi(s, a, g) + \|\pi(s, g) - a\|_2^2]. \quad (8)$$

It also can be rewritten to maximize the objective form:

$$\mathcal{J}_q(\pi) = \mathbb{E}_{s, a, g \sim B_r} [Q^\pi(s, a, g) + \log(\pi(a|s, g))]. \quad (9)$$

The two objectives optimize the policy through different mechanisms. The objective $\mathcal{J}_{wbc}(\pi)$ represents weighted behavior cloning, which is typically used in prior GCWSL methods. However, $\mathcal{J}_{wbc}(\pi)$ may not yield the optimal goal-conditioned actions (Section 3.2), since the policy constraint is inclined to be more conservative, it tends to align with the actions present in the dataset. On the other hand, the objective $\mathcal{J}(\theta)_q$ optimizes the policy through Q-learning, a method inherent to TD-based RL, as corroborated by (Fujimoto and Gu, 2021).

It is important to note that these two objectives are mutually reinforcing: the weighted behavior cloning guarantees that the policy constraint remains adaptable and selects high-quality samples for the (s, g) pair corresponding to sub-trajectories. Concurrently, the Q-learning steers the policy towards the optimal action for the (s, g) pair.

Based on this discussion and the interoperation of these two objectives, we combine Eq. 6 with Eq. 5 to optimize the policy concurrently. This leads to our Q-WSL policy optimization objective for policy updates:

$$\mathcal{J}_{Q-WSL} = \mathbb{E}_{(s, a, g) \sim B_r} \left[\underbrace{Q^\pi(s, a, g)}_{\text{Q-learning}} + \underbrace{\eta f(A^\pi(s, a, g)) \cdot \log \pi_\theta(s, a, g)}_{\text{Advantage-weighted Regression}} \right], \quad (10)$$

where η is defined to balance the RL term (maximizing Q-function) and imitation term (maximizing the weighted behavior cloning term) like TD3+BC (Fujimoto and Gu, 2021).

4.3. Practice Algorithm

We adopt the DDPG+HER as our baseline framework following (Fang et al., 2019; Yang et al., 2021), therefore for a deterministic policy $\pi(s, g)$, the Eq. 20 and Eq. 8 are used during training procedure. Moreover, we follow WGCSL (Yang et al., 2022) to define $f(A)$ equal to $\gamma^{i-t} \exp_{\text{clip}}(A(s_t, a_t, g)) \cdot \epsilon(A(s_t, a_t, g))$, where $A(s_t, a_t, g) = r(s_t, a_t, g) + \gamma V(s_{t+1}, g) -$

Table 1

Comparison of various goal-conditioned methods policy objectives. F signifies the advantage-weighted function in WGCSL, F' signifies the f -divergence function in GoFAR and $\{B, B_r, B_m\}$ correspond to $\{\text{replay buffer, relabeled data, model-based relabeled data}\}$ respectively.

Goal-conditioned Methods	Policy objective	Optimization		Relabeling
		Q-learning	Advantage-weighted Regression	
DDPG(Paul et al., 2018)	$\min_{\pi_\theta} \mathbb{E}_{(s,g) \sim B} [-Q(s, \pi_\theta(s, g), g)]$	✓	✗	✗
Action Models(Chebatar et al., 2021b)	$\min_{\pi_\theta} \mathbb{E}_{(s,g,Q') \sim B} [-Q'(s, \pi_\theta(s, g), g)]$	✗	✗	✗
DDPG+HER(Andrychowicz et al., 2017)	$\min_{\pi_\theta} \mathbb{E}_{(s,g) \sim B_r} [-Q(s, \pi_\theta(s, g), g)]$	✓	✗	✓
Model-based HER(Yang et al., 2021)	$\min_{\pi_\theta} \mathbb{E}_{(s,a,g) \sim B_m} [-Q(s, \pi(s, g), g)] + \alpha \mathbb{E}_{(s,a,g) \sim B_m} [\ a - \pi_\theta(s, g)\ _2^2]$	✓	✗	✓
GCSL(Ghosh et al., 2021)	$\min_{\pi_\theta} \mathbb{E}_{(s,a,g) \sim B_r} [-\log \pi_\theta(a s, g)]$	✗	✗	✓
WGCSL(Yang et al., 2022)	$\min_{\pi_\theta} \mathbb{E}_{(s_t, a_t, g) \sim B_r} [-F(A(s_t, a_t, g), t) \cdot \log \pi_\theta(a_t s_t, g)]$	✗	✓	✓
GoFar(Ma et al., 2022)	$\min_{\pi_\theta} \mathbb{E}_{(s,a,g) \sim B} [F'_\star(R(s; g) + \gamma \mathcal{T}V^\star(s, a; g) - V(s; g)) \log \pi_\theta(a s, g)]$	✗	✓	✓
DWSL(Hejna et al., 2023)	$\min_{\pi_\theta} \mathbb{E}_{(s_t, a_t, g) \sim B_r} [-e^{A/\omega} \cdot \log \pi_\theta(a_t s_t, g)]$	✗	✓	✓
Q-WSL(ours)	$\min_{\pi_\theta} \mathbb{E}_{(s_t, a_t, g) \sim B_r} [-Q(s_t, \pi_\theta(s_{t+1}, g), g) + \eta F(A(s_t, a_t, g), t) \ \pi_\theta(s_t, g) - a\ _2^2]$	✓	✓	✓

$V(s_t, g)$ and the value function $V(s_t, g)$ is estimated using Q-value function as $V(s_t, g) = Q(s_t, \pi(s_{t+1}, g), g)$. Because in the experiment, we found that the effect of this form of $f(A)$ is better for online goal-conditioned RL, and other forms of $f(A)$ may be more friendly for offline settings. The final overview of Q-WSL is presented in Algorithm 2.

Compared with WGCSL Motivated by (Chane-Sane et al., 2021; Yang et al., 2022), we also compared with WGCSL and have the following boundary:

Theorem 2. Suppose function $\exp_{\text{clip}}(A(s_t, a_t, g) \cdot \epsilon(A(s_t, a_t, g))) \geq 1, g = \phi(s_i), i \geq 1$ over the state-action-goal combination. Consider a finite-horizon discrete Markov Decision Process (MDP), a deterministic policy π that selects actions with non-zero probability, and a sparse reward function $r(s_t, a_t, g) = \mathbb{1}[\phi(s_i) = g]$, where $\mathbb{1}[\phi(s_i) = g]$ is an indicator function. Given trajectories $\tau = (s_1, a_1, \dots, s_T, a_T)$ and a discount factor $\gamma \in (0, 1]$, the following bounds hold:

$$J_{Q\text{-WSL}}(\pi) \geq J_{WGCSL}(\pi) \geq J_{GCSL}(\pi). \quad (11)$$

We defer the proof to Appendix A.1. This theorem indicates that under a reward of $r(s_t, a_t, g) = \mathbb{1}[\phi(s_i) = g]$, the optimization objective of Q-WSL is at least better than that of WGCSL and GCSL.

4.4. Compared with Previous Goal-conditioned Methods

Table 1 compares the policy optimization objectives of the proposed goal-conditioned methods with those of prior works. For this comparison, we selected four actor-critic methods (Silver et al., 2014; Chebotar et al., 2021a; Andrychowicz et al., 2017; Yang et al., 2021), and four state-of-the-art goal-conditioned weighted supervised learning methods (Liu et al., 2021; Yang et al., 2022; Ma et al., 2022; Hejna et al., 2023).

Table 1 demonstrates that most prior weighted self-supervised methods are based on advantage-weighted regression approaches, with the exception of GCSL (Liu et al., 2021), which employs an objective similar to Equation 4. However, methodologies such as GCSL (Liu et al., 2021), WGCSL (Yang et al., 2022), and DWSL (Hejna et al., 2023)

are constrained to utilizing trajectories from successful examples only. GoFAR (Ma et al., 2022) requires the collected dataset to cover all reachable goals, and Action Models (Chebotar et al., 2021a) fails on noisy trajectories. In contrast, our Q-WSL avoids these drawbacks. Our approach shares the principles of prior actor-critic methods such as DDPG (Silver et al., 2014), DDPG+HER (Andrychowicz et al., 2017), and Model-based HER (Yang et al., 2021). To properly recognize the contributions of previous studies, we assert that our proposed method capitalizes on the benefits of both Q-learning optimization and advantage-weighted regression. Consequently, many of the aforementioned techniques, particularly WGCSL, GoFAR, and DWSL, can be seamlessly integrated into our model.

5. Experiments

We begin by introducing the benchmarks and baseline methods, followed by a detailed description of the experimental setup. Subsequently, we present the results and analysis, which confirm our assumptions and theoretical framework.

Benchmarks We use some standard goal-conditioned RL research benchmarks (Plappert et al., 2018) including one task on the *Shadow – hand* and four tasks on the *Fetch* robot (See Fig. 2).

Baselines In this section, we compare our method Q-WSL with prior works on various goal-conditioned policy learning optimization algorithms as show in Table 1. We conducted a series of comparative experiments by implementing the baseline methods within the same framework as our proposed approach. Detailed baselines algorithm implementation is shown in Appendix A.3. We provide a brief overview of our baseline algorithms as below:

- **DDPG** (Paul et al., 2018) is an efficient off-policy actor-critic method for learning continuous actions. In this paper, we adopt it as the basic framework.
- **DDPG+HER** (Andrychowicz et al., 2017) combines DDPG with HER, which learns from failed experiences with sparse rewards.

Algorithm 2 Q-WSL

```

1: Initialize off-policy replay buffer  $\mathcal{B}$ 
2: Initialize policy  $\pi_\theta$ , policy target network weights  $\bar{\theta} \leftarrow \theta$ ,  $Q$ -network  $Q_\psi$ ,  $Q$ -target network weights  $\bar{\psi} \leftarrow \psi$ ,
3: while a fixed number of iteration do
4:   Collect trajectories with the policy  $\pi$  and save to the replay buffer  $\mathcal{B}$ 
5:   for Update goal-conditioned policy step do
6:     Generate relabeled data  $\mathcal{B}_r$  from  $\mathcal{B}$ :  $m \leftarrow \{(s_t, a_t, g, r(s_t, a_t, g)), i \geq t\}$ , where  $g = \phi(s_i), i \geq t$ 
7:     // Sample relabeled data and calculate Q-value function
8:     Calculate  $Q$  targets  $y_q$ :  $y_q = r(s_t, a_t, g) + \gamma V_{\bar{\psi}}(s_{t+1}, \pi_{\bar{\theta}}(s_{t+1}, g), g)$ 
9:     Clip the Q-value function:  $y_q = y_q \cdot \text{clip}(-1/(1-\gamma), 0)$ 
10:    Optimize  $Q$  by minimizing following TD error:  $\mathbb{E}_{(s_t, a_t, \phi(s_i, s_{t+1})) \sim \mathcal{B}_r} [(y_q - Q_\psi(s_t, a_t, g))^2]$ 
11:    Calculate the weight within WGCSL:  $w_A = \gamma^{i-t} \exp_{\text{clip}}(A(s_t, a_t, g)) \cdot \epsilon(A(s_t, a_t, g))$ 
12:    //Policy optimization with Q-WSL
13:    Update  $\pi_\theta$  with maximize goal-conditioned policy optimization with Eq. 20 and Eq. 8
14:  end for
15:  Soft update the target policy,  $Q$ -network and  $V$ -network:  $\bar{\theta} \leftarrow \tau\theta + (1-\tau)\bar{\theta}$ ,  $\bar{\psi} \leftarrow \tau\psi + (1-\tau)\bar{\psi}$ 
16: end while

```




Fig. 2. Examples of goal-conditioned benchmark environments: (a) FetchReach, (b) FetchPush, (c) FetchSlide, (d) FetchPickAndPlace, (e) HandReach.

- **Actionable Models** (Chebotar et al., 2021b) subtly combines goal-chaining with conservative Q -Learning. we re-implement this goal relabeling technique which apply the the Q -value at the final state of trajectories in \mathcal{B} to enable goal chaining, as well as the negative action sampling operation.
- **MHER** (Yang et al., 2021) constructs a dynamics model using historical trajectories and combines current policy to generate future trajectories for goal relabeling.
- **GCSL** (Ghosh et al., 2021) is a typical goal-conditioned supervised learning. GCSL incorporates hindsight relabeling in conjunction with behavior cloning to imitate the suboptimal trajectory.
- **WGCSL** (Yang et al., 2022) introduces an enhanced version of GCSL, incorporating both discounted goal relabeling and advantage-weighted updates into the policy learning process. This integration aims to improve the efficiency and effectiveness of the policy learning updates.
- **GoFar** (Ma et al., 2022) employs advantage-weighted regression with f -divergence regularization based on state-occupancy matching. Additionally, it introduces a hierarchical framework in offline goal-conditioned

reinforcement learning to effectively manage long-horizon tasks.

- **DWSL** (Hejna et al., 2023) initially creates a model to quantify the distance between state and goal. And then DWSL utilizes this learned distance to improve and extract policy.

Architecture Setup We utilize the DDPG (Paul et al., 2018) in conjunction with HER as the foundational algorithm for goal-conditioned RL, against which we benchmark various baselines. Specifically, the critic architecture $Q(s, a, g)$ adheres to the Universal Value Function Approximators (UVFA) model (Schaal et al., 2015). The actor network generates actions parameterized by a Gaussian distribution. The actor and critic networks are structured with three hidden layers of multi-layer perceptrons (MLPs), each comprising 256 neurons and employing ReLU activations. The architectural setup can be denoted as [linear-ReLU] $\times 3$, followed by a linear output layer. We follow the optimal parameter settings for DDPG+HER as specified in (Plappert et al., 2018). All baseline hyperparameters are maintained consistently. Additional details are provided in Appendix A.5.

Evaluation Setup For each baseline and task, we conducted evaluations using random five seeds (e.g, {100, 200, 300, 400, 500}). The policy was trained for 1000 episodes per epoch.

Table 2

Mean success rate (%) for challenging multi-goal robotics tasks after training

Task	Q-WSL(Ours)	Weighted Behavior Cloning				Actor-Critic Methods			
		DWSL	GoFar	WGCSL	GCSL	DDPG+HER	MHER	Action Models	DDPG
FetchReach	100.0 \pm 0.0	96.0 \pm 2.74	100.0 \pm 0.0	99.0 \pm 2.2	93.2 \pm 4.2	99.2 \pm 1.3	99.8 \pm 0.4	97.4 \pm 5.2	100.0 \pm 0.0
FetchPush	99.4 \pm 0.89	7.2 \pm 3.9	88.8 \pm 8.0	94.0 \pm 5.3	5.4 \pm 4.8	91.6 \pm 17.3	96.8 \pm 4.5	6.8 \pm 3.5	7.0 \pm 4.2
FetchSlide	50.6 \pm 8.4	0.2 \pm 0.45	36.6 \pm 10.9	29.0 \pm 14.6	0.4 \pm 0.5	30.2 \pm 14.9	43.6 \pm 8.56	0.2 \pm 0.45	2.2 \pm 3.8
FetchPickAndPlace	96 \pm 2.9	3.8 \pm 2.28	45.2 \pm 4.14	66.4 \pm 12.44	3.8 \pm 2.28	92.6 \pm 1.14	95.4 \pm 2.88	3.4 \pm 2.51	3.8 \pm 2.3
HandReach	59.8 \pm 11.8	0.0 \pm 0.0	17.6 \pm 4.56	19.2 \pm 10.59	0.0 \pm 0.0	64.8 \pm 9.9	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
Average	81.16	21.44	57.64	61.52	20.56	75.68	67.12	21.56	22.6

Upon completing each training epoch, the policy's performance was measured by calculating the average success rate from 100 independent rollouts, each using randomly selected goals. These success rates were averaged across five seeds and plotted over the learning epochs, with the standard deviation illustrated as a shaded region on the graph.

5.1. Performance Evaluation on Multi-goal Benchmark Results

For all baselines, we use 16 CPUs to train the agent for 50 epochs in all tasks. Upon completing the training phase, the most effective policy is evaluated by testing it on the designated tasks. The performance outcomes are then expressed as average success rates. Performance comparisons across training epochs are illustrated in Table 2. For Q-WSL, the policy objective parameter η is set to 0.1.

From Table 2, it is evident that Q-WSL attains significantly higher performance than prior baselines. The results indicate that DDPG and Actionable Models perform poorly across all tasks, while DDPG+HER, MHER, and GCWSL methods, which benefit from HER, demonstrate greater effectiveness. This underscores the importance of HER in multi-goal RL for enhancing learning with sparse rewards and improving sample efficiency. Furthermore, TD-based RL methods built upon HER consistently outperform GCWSL methods under the same conditions, confirming our conclusions in Section 4.1.


Table 2 also shows that compared to the best baseline method, DDPG+HER, Q-WSL attains an improvement of up to 5.48 percentage points. The average improvement across the five tasks is 37.645 percentage points. In summary, Q-WSL constitutes a straightforward yet highly effective approach that attains state-of-the-art performance in goal-conditioned reinforcement learning.

5.2. Sample Efficiency

To evaluate the sample efficiency of Q-WSL compared to baseline methods, we analyze the number of training samples (i.e., $\langle s, a, g', g \rangle$ tuples, g' is the relabeled goal) required to attain a specific mean success rate. For this comparison, we have selected competitive baselines that employ goal relabeling strategies in the HandReach environment.

The results are depicted in Fig. 3. In the HandReach environment, as shown in Fig. 3, Q-WSL requires only

1000 samples to attain the 0.1 mean success rate, demonstrating superior sample efficiency. In this case, Q-WSL is at least twice the sample efficiency of previous works. The reason why a state-of-the-art baseline DWSL is not listed here is that there is basically no success rate in the HandReach environment. Overall, Q-WSL enhances sample efficiency by an average factor of seven (6.8) compared to the baseline methods.

**Fig. 3.** Training samples for HandReach environment under 0.1 success rate.

5.3. Robust to Environmental Stochasticity

To test whether our Q-WSL is robust to random environmental factors, we follow Gofar's settings. Specifically, we examine a modified FetchPush environment characterized by the introduction of Gaussian noise with a zero mean before action execution. This modification generates various environmental conditions with standard deviations of 0.2, 0.5, 1.0, 1.5, allowing us to analyze the robustness and performance of the proposed method under differing levels of stochasticity.

As we see in Fig. 4, Q-WSL is the most robust to stochasticity in the FetchPush environment, also outperforming baseline algorithms in terms of mean success rate under various noise levels. WGCSL exhibits minimal sensitivity to variations in noise levels, whereas DDPG+HER is more sensitive, and GoFar is the most affected by noise fluctuations. Despite DWSL's insensitivity to noise, its overall performance remains sub-optimal. We analyze that the assumption of deterministic dynamics


**Fig. 4.** Mean success rate (%) for FetchPush task under environment stochasticity

Table 3

Mean success rate (%) for challenging multi-goal robotics tasks under indicator reward after training

Task	Q-WSL(Ours)	DWSL	GoFar	WGCSL	DDPG+HER	DDPG
FetchReach	100.0 ± 0.0	98.0 ± 1.58	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0
FetchPush	99.4 ± 0.89	4.4 ± 1.8	9.4 ± 5.1	83.0 ± 8.6	57.4 ± 12.7	4.4 ± 1.8
FetchSlide	44.6 ± 6.3	0.4 ± 0.5	2.0 ± 2.1	31.4 ± 4.3	7.4 ± 1.9	0.2 ± 0.4
FetchPickAndPlace	45.4 ± 14.5	4.4 ± 1.5	16.5 ± 9.2	40.6 ± 6.1	23.6 ± 2.6	4.4 ± 1.7
HandReach	75.6 ± 5.68	0.0 ± 0.0	21.2 ± 4.5	18.8 ± 10.4	40.2 ± 12.7	0.0 ± 0.0
Average	73.0	21.44	29.82	54.76	45.72	21.8

in the self-supervised learning methods such as WGCSL, GoFar, and DWSL may lead to overly optimistic performance estimates in stochastic environments. In contrast, reinforcement learning methods are capable of effectively adapting to these changes.

5.4. Robust to Reward Function

Because the benchmark utilizes a default reward value of Eq. 1, we set it to $\mathbb{1}[\phi(s_t) = g]$ to assess the algorithm's robustness to variations in reward structure. The comparison is shown in Table 3. From Table 3, in the all 5 goal-conditioned RL benchmarks, we can see Q-WSL still attains significantly higher performance than prior baselines and HER principle is also still important to learning sparse reward. Moreover, DDPG+HER is poor performance compared to GCWSL as it is always zero in indicator reward, this means Q-learning is less robust to reward function. This phenomenon and shortcomings of Q-learning have also been demonstrated in (Eysenbach et al., 2020).

5.5. Ablation Studies

To evaluate the significance of incorporating both Q-learning and advantage-weighted regression in Q-WSL, we conducted ablation experiments comparing various Q-WSL variants with HER. For the joint optimization in Equation Eq. 10, the parameter η is set to 0.1 by default. We conduct our experiments under the following conditions:

- **Q-WSL** Q-learning + Goal-conditioned Weighted Supervised Learning (GCWSL).
- **No Q-learning** which is equivalent to WGCSL.
- **No WSL** which is equivalent to Q-learning as described in Eq. 8.

Empirical results are shown in Fig. 5. As depicted in Fig. 5, Q-learning and GCWSL are essential components. Q-WSL demonstrates faster learning than DDPG+HER, whereas the state-of-the-art baseline DWSL fails to learn in FetchPick and FetchPush, suggesting that supervised learning is sub-optimal for relabeled data. The integration of Q-learning and GCWSL significantly enhances performance.

The underlying rationale for these results is that Q-learning and GCWSL mutually reinforce each other within the Q-WSL framework. GCWSL supplies high-quality samples for policy training, while the Q-learning refines the policy further, thereby guiding the curriculum of advantage-weighted regression.

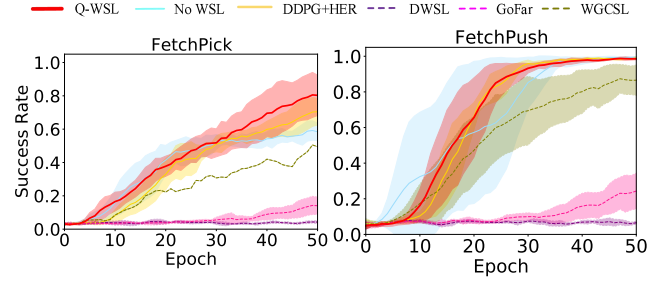


Fig. 5. Ablation studies in FetchPick and FetchPush tasks. The results indicate that Dynamic Programming plays a crucial role in the effectiveness of GCWSL.

6. Conclusion

In this paper, we develop a novel approach termed Q-learning Weighted Supervised Learning (Q-WSL) for Goal-Conditioned Reinforcement Learning (GCRL), which employs the principles of Dynamic Programming from Q-learning to overcome the notable shortcomings of the Goal-Conditioned Weighted Supervised Learning (GCWSL) technique. This innovative framework significantly enhances the capabilities of goal-conditioned reinforcement learning algorithms. Q-WSL specifically targets situations where GCWSL demonstrates suboptimal results due to the necessity of stitching trajectories for task completion. By utilizing Q-learning, our method effectively propagates optimal values backward to the initial states, thereby ensuring optimal decision-making. Experimental evaluations in rigorous goal-conditioned RL scenarios show that Q-WSL not only outperforms existing advanced algorithms but also maintains consistent performance amidst variations in reward dynamics and environmental uncertainties, illustrating its robustness and versatility.

Future research will explore the potential of purely supervised learning and other approaches, such as efficient data augmentation or other Dynamic Programming former, to facilitate faster and more efficient learning of enhanced goal-conditioned policies.

A. Appendix

A.1. PROOF of Theorem 1

Goal-conditioned AWR maximizes the expected improvement over a sampling policy and can implemented as a constrained policy search problem. This optimization problem in goal-conditioned settings can be formulated as (Peng et al., 2020; Nair et al., 2020):

$$\arg \max_{\pi} = \int_g \rho_{B_r}(g) \int_s d_{B_r}(s|g) \int_a [A^{\pi}(s, a, g)] da ds dg, \\ \int_a \pi(a|s, g) da = 1, \text{ s.t. } \text{KL}(\pi_{\theta} \| \pi_{\text{relabel}}) \leq \epsilon, \quad (12)$$

where $\rho_{B_r}(g)$ represents the goal distribution in relabeled data, and $d_{B_r}(s|g) = \sum_{t=0}^{\infty} \gamma^t p(s_t = s | \pi_{\text{relabel}}, g)$ denotes the goal-conditioned un-normalized discounted state distribution

induced by the relabeling policy $\pi_{relabel}$ (Sutton and Barto, 2018). In this context, π_ψ is the learned policy, \mathcal{B}_r refers to the relabeled data, and $\pi_{relabel}$ signifies the behavioral policy responsible for collecting the dataset. KL represents the KL-divergence measurement metric, and ϵ is the threshold. By enforcing the Karush-Kuhn-Tucker (KKT) condition on Eq. 12, we obtain the Lagrangian:

$$\mathcal{L}(\lambda, \pi, c\alpha) \quad (13)$$

$$\begin{aligned} &= \int_g \rho_{\mathcal{B}_r}(g) \int_s d_{\mathcal{B}_r}(s|g) \int_a \pi(a|s, g) [A^\pi(s, a, g)] da ds dg \\ &+ \lambda(\epsilon - \int_g \rho_{\mathcal{B}_r}(g) \int_s d_{\mathcal{B}_r}(s|g) \text{KL}(\pi_\psi \| \pi_{relabel}) ds dg \\ &+ \int_s \alpha(1 - \int_a \pi(a|s, g) da) ds. \end{aligned} \quad (14)$$

Differentiating with respect to π :

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial \pi} &= p_{\mathcal{B}_r}(s, g) A^\pi(s, a, g) - \lambda \cdot p_{\mathcal{B}_r}(s, g) \cdot \log \pi_{relabel}(a|s, g) \\ &+ \lambda p_{\mathcal{B}_r}(s, g) \log \pi(a|s, g) + \lambda p_{\mathcal{B}_r}(s, g) - \alpha, \end{aligned} \quad (15)$$

where $p_{\mathcal{B}_r}(s, g) = \rho_{\mathcal{B}_r}(g) d_{\mathcal{B}_r}(s|g)$. Then set $\frac{\partial \mathcal{L}}{\partial \pi}$ to zero we can have:

$$\pi(a|s, g) = \pi_{relabel}(a|s, g) \exp\left(\frac{1}{\lambda} A^\pi\right) \exp\left(-\frac{1}{p_{\mathcal{B}_r}(s, g)} \frac{\alpha}{\lambda} - 1\right), \quad (16)$$

where $\exp\left(-\frac{1}{p_{\mathcal{B}_r}(s, g)} \frac{\alpha}{\lambda} - 1\right)$ is the partition function $Z(s, g)$ which normalizes the action distribution under state-goal (Peng et al., 2020):

$$\begin{aligned} Z(s, g) &= \exp\left(-\frac{1}{p_{\mathcal{B}_r}(s, g)} \frac{\alpha}{\lambda} - 1\right) \\ &= \int_a' \pi_{\mathcal{B}_r}(a'|s, g) \exp\left(\frac{1}{\lambda} A^\pi(s, a', g)\right) da'. \end{aligned} \quad (17)$$

The closed-form solution can thus be expressed as follows:

$$\pi^*(a|s, g) = \frac{1}{Z(s, g)} \pi_{relabel}(a|s, g) \exp\left(\frac{1}{\lambda} A^\pi(s, a, g)\right) \quad (18)$$

Finally, since π is a parameterized function, translate the solution into these policies, resulting in the following optimization objective:

$$\begin{aligned} &\arg \min_{\pi} \mathbb{E}_{s, g \sim \mathcal{B}_r} [\text{KL}(\pi^*(\cdot|s, g) \| \pi(\cdot|s, g))] \\ &= \arg \min_{\pi} \mathbb{E}_{s, g \sim \mathcal{B}_r} \left[\text{KL}\left(\frac{1}{Z(s, g)} \pi_{relabel}(a|s, g) \right. \right. \\ &\quad \left. \left. \cdot \exp\left(\frac{1}{\lambda} A^\pi(s, a, g)\right) \| \pi(\cdot|s, g)\right) \right] \\ &= \arg \min_{\pi} \mathbb{E}_{s, g \sim \mathcal{B}_r} \mathbb{E}_{a \sim \mathcal{B}_r} \left[-\pi(a|s, g) \exp\left(\frac{1}{\lambda} A^\pi(s, a, g)\right) \right] \end{aligned} \quad (19)$$

The objective can now be interpreted as a weighted maximum likelihood estimation. If the policy is deterministic, the objective can be simplified to:

$$\arg \min_{\pi} \mathbb{E}_{s, a, g \sim \mathcal{B}_r} \left[\exp\left(\frac{1}{\lambda} A^\pi(s, a, g) \cdot \|\pi_\theta(s, g) - a\|_2^2\right) \right] \quad (20)$$

It can be rewritten to maximize the objective form:

$$\mathcal{J}_{wbc}(\pi) = \mathbb{E}_{s, a, g \sim \mathcal{B}_r} \left[\exp\left(\frac{1}{\lambda} A^\pi(s, a, g) \cdot \log(\pi(a|s, g))\right) \right] \quad (21)$$

However, $Z(s)$ is typically ignored in practical implementations (Peng et al., 2020; Nair et al., 2020). Notably, (Nair et al., 2020) suggests that including $Z(s, g)$ often degrades performance. Therefore, for practical purposes, we also omit the term $Z(s)$. If we rewrite $\exp\left(\frac{1}{\lambda} A^\pi(s, a, g)\right)$ to a general function $f(A)$, then GCWSL is the variant form of goal-conditioned AWR which guarantees policy improvement.

A.2. PROOF of Theorem 2

Before proving, we note the following premises: $1 \geq \gamma > 1$, $1 \geq \eta > 0$, $\exp_{clip}(A(s_t, a_t, g) \cdot \epsilon(A(s_t, a_t, g))) \geq 1$, $g = \phi(s_t)$, $i \geq t$.

$$\begin{aligned} \mathcal{J}_{Q-WSL} &= \mathcal{J}_q + \eta \mathcal{J}_{wgcsl} \\ &= \mathbb{E}_{g \sim p(g), \tau \sim \pi_b(\cdot|g), t \sim [1, T], i \sim [t, T]} [\gamma^{i-t} r(s_t, a_t, g) + \eta \gamma^{i-t} \\ &\quad \exp_{clip}(A(s_t, a_t, g)) \cdot \epsilon(A(s_t, a_t, g)) \log \pi(a_t | s_t, g)] \\ &\geq \mathbb{E}_{g \sim p(g), \tau \sim \pi_b(\cdot|g), t \sim [1, T], i \sim [t, T]} [\exp_{clip}(A(s_t, a_t, g)) \\ &\quad \cdot \epsilon(A(s_t, a_t, g)) \log \pi(a_t | s_t, g)] \\ &\geq \mathbb{E}_{g \sim p(g), \tau \sim \pi_b(\cdot|g), t \sim [1, T], i \sim [t, T]} [\exp_{clip}(A(s_t, a_t, g)) \\ &\quad \cdot \epsilon(A(s_t, a_t, g)) \log \pi(a_t | s_t, g)] \triangleq \mathcal{J}_{WGCSL}(\pi) \\ &\geq \mathbb{E}_{g \sim p(g), \tau \sim \pi_b(\cdot|g), t \sim [1, T], i \sim [t, T]} [\log \pi(a_t | s_t, g)] \\ &\triangleq \mathcal{J}_{GCSL}(\pi). \end{aligned} \quad (22)$$

A.3. Baseline Details

In our experiments, all algorithms share the same actor-critic architecture (i.e., DDPG) and hyperparameters. Throughout the data collection phase governed by policy π , Gaussian noise is systematically added with a mean of zero and a constant standard deviation, as detailed in (Fujimoto et al., 2018), to enhance exploration capabilities. Implementations of DDPG, AM, MHER, DDPG+HER, and GCSL were sourced from GoFar's open code. Although WGCSL, GoFar, and DWSL are designed as offline algorithms, they are also applicable in off-policy settings; hence, we re-implemented them into our framework based on the principles outlined in the literature. Here, we specifically focus on methods based on advantage-weighted regression. Before introducing the methods, we will first define \mathcal{B} as the replay buffer and \mathcal{B}_r represent the data relabeled from this buffer. We will evaluate the following GCWSL methods for comparison.

- **GCSL** (Ghosh et al., 2021) is implemented by removing the critic component of the DDPG algorithm and

modifying the policy loss to be based on maximum likelihood estimation:

$$\min_{\pi} -\mathbb{E}_{(s,a,g) \sim \mathcal{B}_r} [\log \pi(a | s, g)]. \quad (23)$$

GCSSL can be regarded as GCWSL when $f(A) = 1$. The following baselines are all on the offline goal-conditioned. We re-implement them on the online.

- **WGCSL** (Yang et al., 2022) is implemented as an extension of GCSSL by incorporating a Q-function. The training of the Q-function employs Temporal Difference (TD) error, similar to the methodology used in the Deep Deterministic Policy Gradient (DDPG) algorithm. Furthermore, this process is enhanced by incorporating advantage weighting into the regression loss function. The advantage term that we compute is denoted as $A(s_t, a_t, g) = r(s_t, a_t, g) + \gamma Q(s_{t+1}, \pi(s_{t+1}, g), g) - Q(s_t, a_t, g)$. By employing this denotes, the policy objective of WGCSL is:

$$\min_{\pi} -\mathbb{E}_{(s_t, a_t, g) \sim \mathcal{B}_r} [\gamma^{i-t} \exp_{clip}(A(s_t, a_t, g)) \cdot \epsilon(A(s_t, a_t, g)) \log \pi(a_t | s_t, g)], \quad (24)$$

where \exp_{clip} is the clip for numerical stability and

$$\epsilon(A(s_t, a_t, g)) = \begin{cases} 1, & A(s_t, a_t, g) > \hat{A} \\ \epsilon_{min}, & \text{otherwise} \end{cases}, \quad (25)$$

(\hat{A} is a threshold, ϵ_{min} is a small positive value).

- **GoFar** (Ma et al., 2022) proposed the adoption of a state-matching objective for multi-goal RL, wherein a reward discriminator and a bootstrapped value function were employed to assign weights to the imitation learning loss. The policy objective of GoFar is:

$$\min_{\pi} -\mathbb{E}_{(s_t, a_t, g) \sim \mathcal{B}} [\max(A(s, a, g) + 1, 0) \log \pi(a_t | s_t, g)]. \quad (26)$$

The advantage function is estimated using discriminator-based rewards. The discriminator, denoted as c , is trained to minimize

$$\mathbb{E}_{g \sim p(g)} [\mathbb{E}_{p(s;g)} [\log c(s, g)] + \mathbb{E}_{(s,g) \sim \mathcal{B}} [\log(1 - c(s, g))]] \quad (27)$$

. The value function V is trained to minimize $(1 - \gamma) \mathbb{E}_{(s,g) \sim \mu_0, p(g)} [V(s, g)] + \frac{1}{2} \mathbb{E}_{(s,a,g,s') \sim \mathcal{B}} [(r(s; g) + \gamma V(s'; g) - V(s; g) + 1)^2]$, where $V \geq 0$. After obtaining the optimal V^* , the calculation method of advantage function $A(s, a, g)$ is the same as WGCSL. We have re-implemented GoFar with an enhanced version, GoFar+HER (i.e., $g = g$), to attain higher sample efficiency.

- **DWSL** (Hejna et al., 2023) presents a cutting-edge supervised learning approach that models the empirical

distribution p_{ζ}^r of discrete state-to-goal distances \hat{d} based on offline data. We adapted this method to an off-policy setting, implementing the following principle for distribution training:

$$\max_{\zeta} \mathbb{E}_{\mathcal{B}_r} [\log p_{\zeta}^r ((j - i - 1) / N | s_i, g)], \quad (28)$$

where $s_i, s_{i+1}, g = \phi(s_j), j > i$ are the relabeled data from the \mathcal{B}_r , and N represents the number of steps. DWSL calculates significant statistics from the distance distribution to determine the shortest path estimates between any two states, thus avoiding the challenges of bootstrapping. Specifically, DWSL utilizes the LogSumExp function for a smooth estimation of the minimum distance. The method calculates distances between any two states as follows:

$$\hat{d}(s_i, g) = -\alpha \log \mathbb{E}_k \sim p_{\zeta}^r(\cdot | (s_i, g)) [e^{-k/(I\alpha)}] \quad (29)$$

and

$$\hat{d}(s_{i+1}, g) = -\alpha \log \mathbb{E}_k \sim p_{\zeta}^r(\cdot | (s_{i+1}, g)) [e^{-k/(I\alpha)}], \quad (30)$$

where I is the bin count in the distribution and α is the temperature. We follow DWSL to set I to default to 100.

Ultimately, DWSL reweights actions according to their effectiveness in reducing the estimated distance to the goal:

$$A = \hat{d}(s_i, \phi(s_j)) - c(s_i, g) - \hat{d}(s_{i+1}, g), \quad (31)$$

where $c(s_i, g) = 1 [\phi(c_{i+1}) \neq g] / I$.

The policy is then extracted as:

$$\max_{\theta} \mathbb{E}_{\mathcal{B}_r} [e^{A/\omega} \log \pi_{\theta}(a_i | s_i, g)], \quad (32)$$

, where ω is the temperature parameter, and π is the policy parameterized by θ .

A.4. Additional Experiments Results

In this section, we examine the robustness of Q-WSL concerning various parameters, including the relabeling ratio. Due to space limitations, detailed ablation studies are provided here instead of the main text.

Relabeling Ratio The multi-goal framework we consider assumes that data is associated with goal labels. We explore the influence of the relabeling ratio on performance, as depicted in Fig. 6. The results in Fig. 6 indicate that Q-WSL remains robust across different relabeling ratios. Furthermore, Q-WSL consistently outperforms competitive methods such as WGCSL and DDPG+HER at various relabeling ratios.

The Impact of Hyperparameter η Since our approach optimizes both supervised learning (SL) and reinforcement learning (RL) concurrently, this section explores the influence of the balancing parameter η . We evaluate η values from

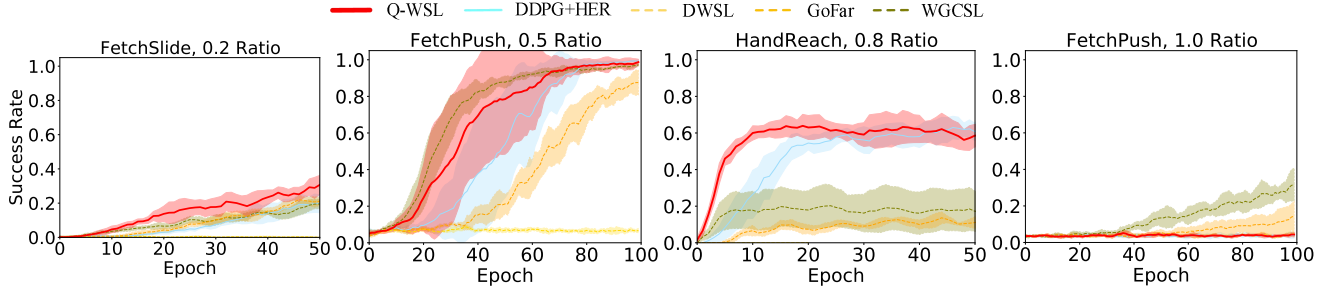


Fig. 6. Relabel ratio ablation studies in such multi-goal tasks. The results demonstrate that Q-WSL outperforms other competitive HER-based methods, except in cases where the labeling rate is 1.0.




Fig. 7. Hyperparameter η ablation studies in such multi-goal tasks. Although the parameter η can take different values, the results demonstrate that Q-WSL remains unaffected by variations in η , consistently achieving state-of-the-art performance.

the set $\{0.1, 0.2, 1.0, 3.0\}$ and compare the results against competitive HER-based algorithms such as WGCSL and DDPG+HER, as shown in Fig. 7. The findings in Fig. 7 reveal that Q-WSL consistently delivers superior performance over the other algorithms, regardless of the η parameter variation. This demonstrates that our method maintains robustness and is not significantly affected by changes in the η parameter.

A.5. Hyperparameters

For all experiments, we utilize the Adam optimizer. To relabel goals, we uniformly sample from all future states within each trajectory. For baseline methods employing discount factors, we set $\gamma = 0.98$ across all Gym Robotics environments. Each algorithm is configured with a consistent set of hyperparameters for all multi-goal tasks. The hyperparameters for DWSL, GoFar, WGCSL, GCSL, MHER, and DDPG have been optimized for our task set, and we use the values reported in previous studies. In our implementation of AM, we use the same network architecture as DDPG and therefore adopt its hyperparameters. All hyperparameters are described in greater detail in (Andrychowicz et al., 2017).

References

- Andrychowicz, M., Wolski, F., Ray, A., Schneider, J., Fong, R., Welinder, P., McGrew, B., Tobin, J., Pieter Abbeel, O., Zaremba, W., 2017. Hindsight experience replay. *Advances in neural information processing systems* 30.
- Bouktif, S., Cheniki, A., Ouni, A., El-Sayed, H., 2023. Deep reinforcement learning for traffic signal control with consistent state and reward design approach. *Knowledge-Based Systems* 267, 110440.
- Brandfonbrener, D., Bietti, A., Buckman, J., Laroché, R., Bruna, J., 2022. When does return-conditioned supervised learning work for offline

Table 4

Hyperparameters for Baselines.

Actor and critic networks	value
Learning rate	1e-3
Initial buffer size	10^6
Polyak-averaging coefficient	0.95
Action L2 norm coefficient	1.0
Observation clipping	$[-200, 200]$
Batch size	256
Rollouts per MPI worker	2
Number of MPI workers	16
Epochs	50
Cycles per epoch	50
Batches per cycle	40
Test rollouts per epoch	10
Probability of random actions	0.3
Gaussian noise	0.2
Probability of hindsight experience replay	0.8
Normalized clipping	$[-5, 5]$
η	0.1

reinforcement learning? *Advances in Neural Information Processing Systems* 35, 1542–1553.

- Chane-Sane, E., Schmid, C., Laptev, I., 2021. Goal-conditioned reinforcement learning with imagined subgoals.
- Chebatar, Y., Hausman, K., Lu, Y., Xiao, T., Kalashnikov, D., Varley, J., Irpan, A., Eysenbach, B., Julian, R., Finn, C., 2021a. Actionable models: Unsupervised offline reinforcement learning of robotic skills.
- Chebatar, Y., Hausman, K., Lu, Y., Xiao, T., Kalashnikov, D., Varley, J., Irpan, A., Eysenbach, B., Julian, R., Finn, C., et al., 2021b. Actionable models: Unsupervised offline reinforcement learning of robotic skills. *arXiv preprint arXiv:2104.07749*.

- Cheikh, D., Russo, D., 2023. On the statistical benefits of temporal difference learning, in: International Conference on Machine Learning, PMLR. pp. 4269–4293.
- Chen, L., Paleja, R., Gombolay, M., 2020. Learning from suboptimal demonstration via self-supervised reward regression. arXiv preprint arXiv:2010.11723 .
- Ding, Y., Florensa, C., Abbeel, P., Phielipp, M., 2019. Goal-conditioned imitation learning. Advances in neural information processing systems 32.
- Eysenbach, B., Salakhutdinov, R., Levine, S., 2020. C-learning: Learning to achieve goals via recursive classification. arXiv preprint arXiv:2011.08909 .
- Eysenbach, B., Zhang, T., Levine, S., Salakhutdinov, R.R., 2022. Contrastive learning as goal-conditioned reinforcement learning. Advances in Neural Information Processing Systems 35, 35603–35620.
- Fang, M., Zhou, T., Du, Y., Han, L., Zhang, Z., 2019. Curriculum-guided hindsight experience replay. Advances in neural information processing systems 32.
- Fujimoto, S., Gu, S.S., 2021. A minimalist approach to offline reinforcement learning. Advances in neural information processing systems 34, 20132–20145.
- Fujimoto, S., Hoof, H., Meger, D., 2018. Addressing function approximation error in actor-critic methods, in: International conference on machine learning, PMLR. pp. 1587–1596.
- Ghosh, D., Gupta, A., Reddy, A., Fu, J., Devin, C.M., Eysenbach, B., Levine, S., 2021. Learning to reach goals via iterated supervised learning, in: International Conference on Learning Representations. URL: <https://openreview.net/forum?id=rALA0Xo6yNJ>.
- Ghugare, R., Geist, M., Berseth, G., Eysenbach, B., 2024. Closing the gap between td learning and supervised learning—a generalisation point of view. arXiv preprint arXiv:2401.11237 .
- Hejna, J., Gao, J., Sadigh, D., 2023. Distance weighted supervised learning for offline interaction data. arXiv preprint arXiv:2304.13774 .
- Hou, X., Guo, Z., Wang, X., Qian, T., Zhang, J., Qi, S., Xiao, J., 2022. Parallel learner: A practical deep reinforcement learning framework for multi-scenario games. Knowledge-Based Systems 236, 107753.
- Kaelbling, L.P., 1993. Learning to achieve goals, in: IJCAI, Citeseer. pp. 1094–8.
- Kim, J., Lee, Y.J., Kwak, M., Park, Y.J., Kim, S.B., 2024. Dynasti: Dynamics modeling with sequential temporal information for reinforcement learning in atari. Knowledge-Based Systems , 112103.
- Kostrikov, I., Nair, A., Levine, S., 2021. Offline reinforcement learning with implicit q-learning. arXiv preprint arXiv:2110.06169 .
- LeCun, Y., Bengio, Y., Hinton, G., 2015. Deep learning. nature 521, 436–444.
- Lillicrap, T.P., Hunt, J.J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., Wierstra, D., 2015. Continuous control with deep reinforcement learning. arXiv preprint arXiv:1509.02971 .
- Liu, M., Zhu, M., Zhang, W., 2022. Goal-conditioned reinforcement learning: Problems and solutions. arXiv preprint arXiv:2201.08299 .
- Liu, P., Zhou, J., Lv, J., 2023. Exploring the first-move balance point of go-moku based on reinforcement learning and monte carlo tree search. Knowledge-Based Systems 261, 110207.
- Liu, S., Cao, J., Wang, Y., Chen, W., Liu, Y., 2021. Self-play reinforcement learning with comprehensive critic in computer games. Neurocomputing 449, 207–213.
- Lynch, C., Khansari, M., Xiao, T., Kumar, V., Tompson, J., Levine, S., Sermanet, P., 2020. Learning latent plans from play, in: Conference on robot learning, PMLR. pp. 1113–1132.
- Ma, J.Y., Yan, J., Jayaraman, D., Bastani, O., 2022. Offline goal-conditioned reinforcement learning via f -advantage regression. Advances in Neural Information Processing Systems 35, 310–323.
- Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., Riedmiller, M., 2013. Playing atari with deep reinforcement learning. arXiv preprint arXiv:1312.5602 .
- Nair, A., Gupta, A., Dalal, M., Levine, S., 2020. Awac: Accelerating online reinforcement learning with offline datasets. arXiv preprint arXiv:2006.09359 .
- Paster, K., McIlraith, S.A., Ba, J., 2020. Planning from pixels using inverse dynamics models. arXiv preprint arXiv:2012.02419 .
- Paul, L.T., James, H.J., David, S., Tom, E., Yuval, T., Otto, H.N.M., Pieter, W.D., Alexander, P., 2018. Continuous control with deep reinforcement learning .
- Peng, X.B., Kumar, A., Zhang, G., Levine, S., 2020. Advantage weighted regression: Simple and scalable off-policy reinforcement learning. URL: <https://openreview.net/forum?id=H1gdF34FvS>.
- Plappert, M., Andrychowicz, M., Ray, A., McGrew, B., Baker, B., Powell, G., Schneider, J., Tobin, J., Chociej, M., Welinder, P., et al., 2018. Multi-goal reinforcement learning: Challenging robotics environments and request for research. arXiv preprint arXiv:1802.09464 .
- Schaul, T., Horgan, D., Gregor, K., Silver, D., 2015. Universal value function approximators, in: International conference on machine learning, PMLR. pp. 1312–1320.
- Sharifani, K., Amini, M., 2023. Machine learning and deep learning: A review of methods and applications. World Information Technology and Engineering Journal 10, 3897–3904.
- Shorten, C., Khoshgoftaar, T.M., 2019. A survey on image data augmentation for deep learning. Journal of big data 6, 1–48.
- Silver, D., Lever, G., Heess, N., Degris, T., Wierstra, D., Riedmiller, M., 2014. Deterministic policy gradient algorithms, in: International conference on machine learning, Pmlr. pp. 387–395.
- Sutton, R.S., Barto, A.G., 1998. Reinforcement Learning. The MIT Press.
- Sutton, R.S., Barto, A.G., 2018. Reinforcement learning: An introduction. MIT press.
- Xiao, H., Chen, C., Zhang, G., Chen, C.P., 2024. Reinforcement learning-driven dynamic obstacle avoidance for mobile robot trajectory tracking. Knowledge-Based Systems 297, 111974.
- Xu, M., Wang, J., 2022. Learning strategy for continuous robot visual control: A multi-objective perspective. Knowledge-Based Systems 252, 109448.
- Yang, R., Fang, M., Han, L., Du, Y., Luo, F., Li, X., 2021. Mher: Model-based hindsight experience replay. arXiv preprint arXiv:2107.00306 .
- Yang, R., Lu, Y., Li, W., Sun, H., Fang, M., Du, Y., Li, X., Han, L., Zhang, C., 2022. Rethinking goal-conditioned supervised learning and its connection to offline rl. arXiv preprint arXiv:2202.04478 .
- Yang, W., Wang, H., Cai, D., Pajarinen, J., Kämäräinen, J.K., 2023. Swapped goal-conditioned offline reinforcement learning. arXiv preprint arXiv:2302.08865 .
- Zeng, H., Wei, B., Liu, J., 2024. Rtrl: Relation-aware transformer with reinforcement learning for deep question generation. Knowledge-Based Systems , 112120.
- Ziebart, B.D., Maas, A.L., Bagnell, J.A., Dey, A.K., et al., 2008. Maximum entropy inverse reinforcement learning., in: Aaai, Chicago, IL, USA. pp. 1433–1438.