Quasipseudometric Value Functions with Dense Rewards

Khadichabonu Valieva

 $Khadichabonu.\ Valieva@usm.\ edu$

School of Computing Sciences & Computer Engineering The University of Southern Mississippi

Bikramjit Banerjee

Bikramjit.Banerjee@usm.edu

 $School\ of\ Computing\ Sciences\ \mathscr{C}\ Computer\ Engineering$ $The\ University\ of\ Southern\ Mississippi$

Reviewed on OpenReview: https://openreview.net/forum?id=4LqOl6pDUe

Abstract

As a generalization of reinforcement learning (RL) to parametrizable goals, goal conditioned RL (GCRL) has a broad range of applications, particularly in challenging tasks in robotics. Recent work has established that the optimal value function of GCRL $Q^*(s, a, q)$ has a quasipseudometric structure, leading to targetted neural architectures that respect such structure. However, the relevant analyses assume a sparse reward setting—a known aggravating factor to sample complexity. We show that the key property underpinning a quasipseudometric, viz., the triangle inequality, is preserved under a dense reward setting as well, specifically identifying the key condition necessary for triangle inequality. Contrary to earlier findings where certain types of dense rewards were shown to be detrimental to GCRL, we conjecture that dense reward functions that satisfy this condition can only improve, never worsen, sample complexity. We evaluate this proposal in 12 standard benchmark environments in GCRL featuring challenging continuous control tasks. Our empirical results confirm that training a quasipseudometric value function in our dense reward setting indeed either improves upon, or preserves, the sample complexity of training with sparse rewards. This opens up opportunities to train efficient neural architectures with dense rewards, compounding their benefits to sample complexity.

Keywords: Goal conditioned reinforcement learning, reward shaping, quasipseudometric value functions

1 Introduction

Reinforcement learning (RL) is a popular class of techniques for training autonomous agents to behave (near) optimally, often without requiring a model of the task or environment. In goal-achieving tasks, traditional RL learns policies that reach a single goal at the minimum (maximum) expected cost (value) from any state. Contrastingly in multi-task settings, a goal conditioned value function models the cost-to-go to a set of goal states, not just one. This generalization from a single-goal case to goal-conditioned RL (GCRL) yields effective representations—powered by deep neural networks—for value functions capable of capturing abstract concepts underlying goal achievement in many complex tasks (Wang et al., 2023; Liu et al., 2022; Plappert et al., 2018).

Recent work has established that the true optimal value function in GCRL is always a quasipseudometric, i.e., a metric without the constraint of being symmetric, but crucially respecting the triangle inequality (Pitis et al., 2020; Wang & Isola, 2022; Liu et al., 2023). This allows the search for value functions to be naturally restricted to the space of quasipseudometrics. Additionally, such functions are designed to be universal value function approximators (UVFA), i.e., capable of approximating arbitrarily complex value functions. Accordingly, Liu et al. (2023) propose the metric residual network (MRN) architecture for GCRL value functions that explicitly accommodate an asymmetric component while maintaining the UVFA property and the triangle inequality. This and other similar approaches search a smaller subset of the space of value

functions, yet the true optimal value function is guaranteed to reside in it. This has led to significant gains in terms of sample efficiency in recent GCRL advancements (Liu et al., 2023; Wang & Isola, 2022; Wang et al., 2023).

In this paper, we review some of the theoretical analyses underlying much of the work cited above. In particular, the proof of the key property of triangle inequality in Liu et al. (2023) is established for a sparse reward setting that is easy to design but hard to learn from, i.e., sparse rewards lead to poor sample efficiency. By contrast, dense reward settings using various mechanisms, e.g., reward shaping, intrinsic motivation, human feedback etc., are generally known to improve sample efficiency. If dense reward-based value functions were to satisfy the triangle inequality, then their reward bias could be combined with the representational bias of quasipseudometrics to deliver a double punch to sample complexity. Furthermore in GCRL, since a goal is given, it may be easy to design a dense reward scheme parametrized by the goal that neither requires onerous, intricate reward engineering, nor is computationally demanding. This could bring the benefit of dense rewards-viz., reduced sample complexity-to GCRL with less human effort than in other RL tasks. However, existing negative results (Plappert et al., 2018) specifically in GCRL show that certain classes of dense rewards significantly deteriorate the performance of state-of-the-art RL methods, and might appear to foreclose a discussion on their efficacy in GCRL. Contradictorily, we show that a class of dense rewards, viz., potential shaped rewards, can indeed bring their benefit to bear in GCRL as long as they satisfy a condition under which the triangle inequality is preserved for the *optimal* value function. Furthermore, we establish a condition under which the triangle inequality is preserved for intermediate value functions that may be encountered during RL iterations. This result adds nuance to recent contradictory finding (Wang et al., 2023) that such value functions do not satisfy the triangle inequality. We show experiments in 12 benchmark GCRL tasks to establish that our prescribed dense rewards indeed improve sample complexity in some tasks, but does not deteriorate sample efficiency in any task, while certain types of dense rewards that violate the key condition deteriorate sample complexity consistent with the findings of Plappert et al. (2018). Our main contributions can be summarized as:

- We show that using rewards shaped with potential functions that serve as admissible heuristics, the optimal value function does satisfy the triangle inequality. This extends similar results from Liu et al. (2023) in a sparse reward setting to a commonly used dense reward setting that is more general and beneficial to sample efficiency;
- We define and delineate a progressive criterion for GCRL policies and show that under such policies the intermediate value functions also satisfy the triangle inequality. This adds nuance to a prior negative result in Wang et al. (2023) which showed that there exist intermediate value functions that violate the triangle inequality;
- Via experiments in 12 standard benchmark GCRL tasks, we show that our prescribed dense rewards
 improve sample complexity as well as the learned policy in several tasks, while not deteriorating
 performance in any task. This result informs us how to overcome a prior negative result in Plappert et al. (2018) which showed that GCRL with shaped rewards had significantly worse sample
 complexity than with sparse rewards.

2 Background

This section covers the preliminaries on goal conditioned RL, the prevalent solution approaches for GCRL, and the recent architecture of metric residual networks that we use in this paper.

2.1 Goal-conditioned RL

Goal conditioned RL is modeled by goal-conditioned Markov decision process, $M = (S, A, G, T, R, \gamma, \rho_0, \rho_G)$. While S, A, T, ρ_0 define the state action spaces, the transition function and the initial state distribution just like a standard MDP, G gives the space of goal states, and G0 is the distribution from which a goal is sampled at the beginning of an episode. Further, the reward function G1 is additionally parametrized by the goal,

 $R: \mathcal{S} \times \mathcal{A} \times \mathcal{G} \mapsto \Re$. In the sparse reward setting, R is often defined as

$$R(s, a, g) = \begin{cases} 0 & \text{if } M(s, a) = g \\ -1 & \text{otherwise} \end{cases}$$
 (1)

where $M: \mathcal{S} \times \mathcal{A} \mapsto \mathcal{G}$ maps the product space of \mathcal{S} and \mathcal{A} to \mathcal{G} , specifying the goal state reached by taking action a in state s. As opposed to the common assumption $\mathcal{G} \subset \mathcal{S}$, M allows action a to decide whether the goal is reached (Liu et al., 2023) when testing whether $M(s,a) \stackrel{?}{=} g$. The agent's decision function is its policy, π , which can be either deterministic ($\pi: \mathcal{S} \times \mathcal{G} \mapsto \mathcal{A}$) or stochastic ($\pi: \mathcal{S} \times \mathcal{A} \times \mathcal{G} \mapsto [0,1]$) yielding the probability of selecting actions. We refer to trajectories or episodes as time series of states and actions that a learner encounters, $s_0, a_0, s_1, a_1, \ldots, s_t, a_t, \ldots$

2.2 Solution Approach: DDPG+HER

A popular approach to solving GCRL is a combination of off-policy actor-critic, e.g., DDPG (Lillicrap et al., 2016) with hindsight experience replay (HER) (Andrychowicz et al., 2017). DDPG in GCRL estimates a goal conditioned critic for policy π

$$Q^{\pi}(s, a, g) = \mathbb{E}_{\pi, T, R} \left[\sum_{t=0}^{\infty} \gamma^{t} r_{t,g} | s_{0} = s, a_{0} = a, g \right]$$

where $r_{t,g} = R(s_t, a_t, g)$ and the expectation is taken over future steps of rewards generated by the policy (π) , and the T, R functions. The critic is updated by minimizing the mean squared TD error over samples (s_t, a_t, s_{t+1}, g) drawn from a replay buffer D,

$$L(Q) = \mathbb{E}_D\left[(r_{t,g} + \gamma Q(s_{t+1}, \pi(s_{t+1}), g) - Q(s_t, a_t, g))^2 \right]. \tag{2}$$

By ensuring that Q is differentiable w.r.t. actions a, the actor policy π is updated in the direction of the gradient $\mathbb{E}[\nabla_{a_t}Q(s_t,a_t,g)]$, where the expectation is again evaluated using samples drawn from D. As these samples are drawn from state distributions generated by policies different from π , DDPG is an off-policy method.

Hindsight experience replay (HER) (Andrychowicz et al., 2017) mitigates the sparse reward problem by relabeling failed trajectories. Instead of treating all experience traces where the agent failed to achieve a goal as is, HER changes the goal in some of them to match a step of the trace in hindsight—essentially pretending as if the agent's goal all along was to reach the state that it actually did. This transforms some of the failed episodes into successful experiences that are informative about goal achievement, and allows the agent to generalize, eventually, to the true goal distribution $\rho_{\mathcal{G}}$. A more recent method, goal-conditioned supervised learning (Ghosh et al., 2021), uses the idea of HER but in the context of offline RL. Both DDPG and HER remain popular and are commonly used in online GCRL (Liu et al., 2023).

2.3 Metric Residual Network

Liu et al. (2023) propose a novel neural architecture for GCRL critic based on the insight that the optimal negated action-value function, $-Q^*(s, a, g)$, satisfies the triangle inequality in the sparse reward setting of Eq. 1. Consequently, they introduce the metric residual network (MRN) that decomposes -Q into the sum of a metric and an asymmetric residual component that provably approximates any quasipseudometric. Specifically,

$$Q(s, a, g) = -(d_{sym}(h_{sa}, h_{sg}) + d_{asym}(h_{sa}, h_{sg}))$$
(3)

where h_{sa} and h_{sg} are latent encodings of concatenated (s, a) and (s, g), d_{sym} and d_{asym} are symmetric and asymmetric distance components given by

$$d_{sym}(x,y) = \|\mu_1(x) - \mu_1(y)\|_2,\tag{4}$$

$$d_{asym}(x,y) = \max_{i} (\mu_{2i}(x) - \mu_{2i}(y))_{+}, \tag{5}$$

 $||.||_2$ is the Euclidean distance, $(x)_+ = \max(x, 0)$ is also known as the rectified linear unit (ReLU) function, μ_1 and μ_2 are neural networks that map their inputs to vectors, and μ_{2i} is the *i*-th component of μ_2 's output

vector. Note from Eqs. 4, 5 that indeed $d_{sym}(x,y) = d_{sym}(y,x)$ but possibly $d_{asym}(x,y) \neq d_{asym}(y,x)$. The main idea is that Eq. 3 is a valid decomposition since Q(s,a,g) is essentially a distance and any distance can be (in general) split into symmetric and asymmetric components, where either components can be 0 in limit cases. On the one hand, the Max-ReLU formulation of d_{asym} in Eq. 5 is proven to ensure the UVFA property of Eq. 3 (Liu et al., 2023). On the other hand, d_{sym} improves sample efficiency due to its symmetry. We use DDPG+HER with MRN critic architecture as the base GCRL method for this paper.

2.4 The Quasipseudometric Property

A space \mathcal{X} with a distance measure d defines a quasipseudometric, (d, \mathcal{X}) , if $(1) \ \forall x \in \mathcal{X}$, d(x,x) = 0 and $(2) \ \forall x,y,z \in \mathcal{X}$, $d(x,y) + d(y,z) \geq d(x,z)$, also known as the triangle inequality. A metric is a special case of a quasipeudometric that is also symmetric, but symmetry is not necessary in robotics problems, e.g., a robotic arm could rotate in one direction but not in the other leading to $d(x,y) \neq d(y,x)$. Additionally, a quasimetric is a special case of a quasipseudometric where it is necessary that $x \neq y \implies d(x,y) > 0$. This is also not necessarily true in robotics problems, for instance, when rotating a robotic arm, different quaternions $x \neq y$ could represent the same rotation, i.e., d(x,y) = 0. Liu et al. (2023) argue that in general GCRL problems, -Q(s,a,g) may not be purely symmetric (i.e., it may have an asymmetric component), or > 0 for different states. As a result, the most general form of -Q(s,a,g) for GCRL is neither a metric nor a quasimetric, but a quasipseudometric as constructed via Eqs. 3, 4, 5.

Liu et al. (2023) establish the critical property (2) of a quasipseudometric, i.e., the triangle inequality for optimal $-Q^*(s, a, g)$. This property is the key focus of this paper.

3 Triangle Inequality of Value Functions

In this section, we establish that both the optimal value function as well as intermediate value functions satisfy the triangle inequality under novel conditions, with the following road-map. In Sec. 3.1, we analyze the quasipseudometric property of optimal value functions that use potential-shaped dense rewards. After establishing the main condition for such value functions to satisfy the triangle inequality property, we prove this property formally for two subcases in Secs. 3.1.1 and 3.1.2. In Sec. 3.1.3, we prove that this property is a not a trivial consequence of a well-known property (viz., policy invariance) of potential-shaped rewards. We conclude Sec. 3.1 with discussions on associated topics: value function clipping via an informed constraint (Sec. 3.1.4) and selection of a key parameter of the main condition (Sec. 3.1.5). In Sec. 3.2, we show that another commonly used dense reward that does not use potential functions may violate the triangle inequality. Finally in Sec. 3.3, we establish a novel condition under which even intermediate value functions (with dense or sparse rewards) satisfy the triangle inequality, and formally prove this property.

3.1 Optimal Value Function

Our primary claim is that $-Q^*$ satisfies the triangle inequality not only in the sparse reward setting, but also in the presence of a common type of dense rewards, viz., potential shaped rewards (Ng et al., 1999). This observation lends GCRL to improved sample efficiency when approximating $-Q^*$ using a combination of MRN and potential shaped rewards.

We use the standard potential based shaping rewards (Ng et al., 1999) expressed as the difference between discounted potential (ϕ) of the next state-action pair and the potential of the current state-action pair, for the same goal:

$$F(s, a, s', a', g) = \gamma \phi(s', a', g) - \phi(s, a, g)$$
(6)

and a simple potential function

$$\phi(s, a, g) = -\left(\frac{1 - \gamma^{d(s, a, g)/\eta}}{1 - \gamma}\right)$$

where d is a distance measure between the state and the goal, and η is a measure of the atomicity of actions—distance covered per time step. The specific form of ϕ shown above is explained later in the context of Obs. 1. The function d must be such that

$$M(s,a) = g \implies d(s,a,g) = 0 \implies \phi(s,a,g) = 0$$
 (7)

i.e., the potential is maximized when the goal is reached. Therefore, Eq. 6 can be viewed as a heuristic measure of the agent's step-improvement in the desirability of its state with respect to its goal, where "desirability" is roughly the proximity of the goal. Thus F is a suitable candidate to be an immediate step-reward, and is typically added to the default reward. An important property of such potential shaped reward function is that it leaves the optimal policy unchanged (Ng et al., 1999). Now, in the reward regime of Eq. 1,

$$Q^*(s, a, g) = -\left(\frac{1 - \gamma^{L^*(s, a, g)}}{1 - \gamma}\right)$$

where $L^*(s, a, g)$ is the *optimal* expected number of steps required to reach the goal g from state s. Note the similarity of this expression of Q^* to the potential function ϕ shown above. Our comparable choice of ϕ leads to a simple form of the condition that we show in the next section (Eq. 8) to be the key to satisfying the triangle inequality. Below, we formally specify its premise as an observation.

Observation 1. If
$$d(s, a, g) \leq \eta L^*(s, a, g)$$
, then $\phi(s, a, g) \geq Q^*(s, a, g), \forall s, a, g$

In other words, if $d(s,a,g)/\eta$ is an underestimate of L^* then the above condition will be satisfied. Thus, d acts as an admissible heuristic. In this paper, we study two simple forms of d: the arc-cosine distance $d_{ac}(s,a,g) = \cos^{-1}\left(\frac{M(s,a)\cdot g}{\|M(s,a)\|\|g\|}\right)/\pi$, and the Euclidean distance $d_E(s,a,g) = \|M(s,a)-g\|_2$, both of which are metrics and satisfy Eq. 7 although d_E is not necessarily bounded. These choices are simple and avoid intricate, environment-specific reward engineering, while producing underestimates of L^* (for appropriate values of η). However, these choices are not necessary for our theoretical results to hold and any definition of d that ensures $\phi \geq Q^*$ could be used. Many of the experimental tasks contain state representations that include angles mixed with non-angular features. To keep the design of F simple and domain-independent, we do not segregate these feature categories. Consequently, d_E or d_{ac} applied to the full state/goal vectors may not be rigorous distance measures, but may suffice as $admissible\ heuristics$. When angular features are present, d_{ac} may be relatively more accurate than d_E , yielding a dominant admissible heuristic.

We distinguish $Q^*(s, a, g)$ —the optimal action values with unshaped sparse rewards—from $Q_F^*(s, a, g)$ which corresponds to action values with rewards shaped by F in Eq. 6. Specifically,

$$Q^*(s, a, g) = \max_{\pi} \mathbb{E}_{\pi, T, R} \left[\sum_{t=0}^{\infty} \gamma^t r_{t,g} | s_0 = s, a_0 = a, g \right]$$
$$Q_F^*(s, a, g) = \max_{\pi} \mathbb{E}_{\pi, T, R} \left[\sum_{t=0}^{\infty} \gamma^t (r_{t,g} + F_{t,g}) | s_0 = s, a_0 = a, g \right],$$

where $F_{t,g} = F(s_t, a_t, s_{t+1}, a_{t+1}, g)$. Next we establish the validity of triangle inequality with Q_F^* in two cases: (i) $\mathcal{G} \equiv \mathcal{S} \times \mathcal{A}$ and (ii) $\mathcal{G} \not\equiv \mathcal{S} \times \mathcal{A}$.

3.1.1 Case I: $\mathcal{G} \equiv \mathcal{S} \times \mathcal{A}$

In this setting, M is the identity mapping. We use the notation $x_t = (s_t, a_t)$, which leads to equivalent notations for Q^* : $Q^*(x_1, x_2) \equiv Q^*(s, a, g)$ whenever $x_1 = (s, a)$ and $M(x_2) = g$. The main result is:

Proposition 1. Consider the shaped, goal-conditioned MDP $M_{GCF} = (S, A, G, T, R + F, \gamma, \rho_0, \rho_g)$, with $G \equiv S \times A$. The optimal universal value function Q_F^* satisfies the triangle inequality: $\forall x^1, x^2, x^3 \in \mathcal{X}$,

$$Q_F^*(x^1, x^2) + Q_F^*(x^2, x^3) \le Q_F^*(x^1, x^3),$$

The only condition ϕ must satisfy is

$$\phi(s, a, g) \ge Q^*(s, a, g), \forall s, a, g \tag{8}$$

w.r.t. the unshaped value function, for which a sufficient condition is established in Obs. 1.

Proof: As in (Liu et al., 2023), consider the Markov policies π_1, π_2, π_3 that are optimal w.r.t. $Q_F^*(x^1, x^2)$, $Q_F^*(x^2, x^3)$, $Q_F^*(x^1, x^3)$ and the (non-Markov) policy $\pi_{1\to 2}$ defined for t>0 as:

$$\pi_{1\to 2}(a|s_t) = \begin{cases} \pi_1(a|s_t), & x^2 \notin x_{< t} \text{ (set of states encountered prior to step t)} \\ \pi_2(a|s_t), & \text{otherwise.} \end{cases}$$

Let τ be the random variable that indicates the first time $\pi_{1\to 2}$ reaches x^2 . In the steps below, we notate $F(s_t, a_t, s_{t+1}, a_{t+1}, g)$ as $F_{t,g}$, and $\mathbb{E}_{(x_t, r_t) \sim \pi, T, R, \tau}$ as $\mathbb{E}_{\pi, \cdot}$ for brevity. Then define

$$q_{1\to 2}^{1} = \mathbb{E}_{\pi_{1\to 2},.} \left[\sum_{t=0}^{\tau} \gamma^{t} (r_{t,g} + F_{t,g}) | x_{0} = x^{1}, g = x^{2} \right],$$

$$q_{2\to 3}^{2} = \mathbb{E}_{\pi_{1\to 2},.} \left[\sum_{t=\tau}^{\infty} \gamma^{t} (r_{t,g} + F_{t,g}) | x_{\tau} = x^{2}, g = x^{3} \right]$$

$$= \mathbb{E}_{\pi_{1\to 2},.} \left[\sum_{t=\tau}^{\infty} (\gamma^{t} r_{t,g}) + 0 - \gamma^{\tau} \phi_{\tau} \right].$$

$$\text{Now, } Q_{F}^{*}(x^{1}, x^{2}) = \mathbb{E}_{\pi_{1},.} \left[\sum_{t=0}^{\tau} \gamma^{t} (r_{t,g} + F_{t,g}) | x_{0} = x^{1}, g = x^{2} \right] +$$

$$\mathbb{E}_{\pi_{1},.} \left[\sum_{t=\tau+1}^{\infty} \gamma^{t} (r_{t,g} + F_{t,g}) | x_{\tau+1} = x^{2}, g = x^{2} \right], \tag{9}$$

and $\pi_1 \equiv \pi_{1 \to 2}$ for the first τ steps. Therefore, $Q_F^*(x^1, x^2) - q_{1 \to 2}^1$

$$= \mathbb{E}_{\pi_{1},.} \left[\sum_{t=\tau+1}^{\infty} \gamma^{t}(r_{t,g} + F_{t,g}) | x_{\tau} = x^{2}, g = x^{2} \right],$$

$$= \mathbb{E}_{\pi_{1},.} \left[\sum_{t=\tau+1}^{\infty} (\gamma^{t} r_{t,g}) + \gamma^{\infty} \phi(.) - \gamma^{\tau+1} \phi_{\tau+1} | x^{2}, x^{2} \right]$$

$$= \mathbb{E}_{\pi_{1},.} \left[\sum_{t=\tau+1}^{\infty} (\gamma^{t} r_{t,g}) - \gamma^{\tau+1} \phi_{\tau+1} | x^{2}, x^{2} \right]$$

$$= \mathbb{E}_{\tau} [\gamma^{\tau+1}] [Q^{*}(x^{2}, x^{2}) - \phi(x^{2}, x^{2})]$$

$$\leq 0 \text{ by assumption (Eq. 8)}.$$

Similarly,

$$Q_F^*(x^2, x^3) = \mathbb{E}_{\pi_2, .} \left[\sum_{t=0}^{\infty} \gamma^t (r_{t,g} + F_{t,g}) | x_0 = x^2, g = x^3 \right]$$

$$\leq \gamma^{\tau} \mathbb{E}_{\pi_2, .} \left[\sum_{t=0}^{\infty} \gamma^t (r_{t,g} + F_{t,g}) | x^2, x^3 \right]$$

$$= \mathbb{E}_{\pi_2, .} \left[\sum_{t=0}^{\infty} \gamma^{t+\tau} (r_{t,g} + F_{t,g}) | x^2, x^3 \right]$$

$$= \mathbb{E}_{\pi_2, .} \left[\sum_{k=\tau}^{\infty} \gamma^k r_{k,g} + 0 - \gamma^{\tau} \phi_{\tau} | x^2, x^3 \right]$$

$$= \mathbb{E}_{\pi_{1\to 2, .}} \left[\sum_{k=\tau}^{\infty} \gamma^k r_{k,g} - \gamma^{\tau} \phi_{\tau} | x^2, x^3 \right]$$

$$= q_{2\to 3}^2, \tag{10}$$

since $\pi_2 \equiv \pi_{1\to 2}$ after τ . Therefore, $Q_F^*(x^2, x^3) - q_{2\to 3}^2 \leq 0$. Consequently,

$$Q_F^{\pi_1 \to 2}(x^1, x^3) = q_{1 \to 2}^1 + q_{2 \to 3}^2 \ge Q_F^*(x^1, x^2) + Q_F^*(x^2, x^3).$$

But since the optimal $Q_F^*(x^1, x^3) \geq Q_F^{\pi_1 \to 2}(x^1, x^3)$, we arrive at the triangle inequality.

3.1.2 Case II: $\mathcal{G} \not\equiv \mathcal{S} \times \mathcal{A}$

The proof of this case closely resembles (Liu et al., 2023); we highlight the main difference in blue color but also provide the rest of the proof for completeness. In this case, M is an onto mapping. Given a goal g, the GCRL problem effectively reduces to a standard MDP and there exists a deterministic optimal policy π^* for reaching the goal g from an initial state $x = (s_0, a_0)$. Then, under deterministic dynamics,

$$Q_F^*(x,g) = \sup_{x':M(x')=g} Q_F^*(x,x').$$

Assuming the supremum is attainable, let

$$x_g = \arg\max_{x':M(x')=g} Q_F^*(x, x'),$$
 (11)

then $Q_F^*(x,g) = Q_F^*(x,x_g)$. Assume for contradiction that this is not the case, i.e., $Q_F^*(x,g) \neq Q_F^*(x,x_g)$. There are two possibilities:

- If $Q_F^*(x, x_g) > Q_F^*(x, g)$: This would imply that by using a policy that selects x_g rather than g, one could achieve a higher return. This contradicts the definition of π^* as the optimal policy, thus $Q_F^*(x, x_g) > Q_F^*(x, g)$ cannot be true.
- If $Q_F^*(x, x_g) < Q_F^*(x, g)$: Let

$$\tau = \min_{t} (M(x_t) = g),$$

such that x_{τ} is the first (s, a) pair along the optimal trajectory that achieves the goal. There are two further cases:

1. After reaching x_{τ} , π^* will repeatedly return to x_{τ} . In this case, we have $Q_F^*(x, x_g) \geq Q_F^*(x, x_{\tau})$ by the definition of x_g (Eq. 11) and

$$Q_F^*(x, x_\tau) = Q_F^*(x, g) - \gamma^{\tau+1} Q_F^*(x_\tau, g)$$

$$= Q_F^*(x, g) - \gamma^{\tau+1} [Q^*(x_\tau, g) - \phi(x_\tau, g)]$$

$$\geq Q_F^*(x, g), \text{ by Eq. 8.}$$
(12)

Combining the two, we get $Q_F^*(x, x_g) \geq Q_F^*(x, g)$ which contradicts our assumption that $Q_F^*(x, g) > Q_F^*(x, x_g)$.

- 2. π^* never returns to x_{τ} after reaching it for the first time. In this case, one can find the next $\tau' = \min_{t>\tau}(M(x_t) = g)$, such that $x_{\tau'}$ is another (s,a) along the optimal trajectory. Again, there are two sub-cases:
 - (a) If π^* repeatedly visits $x_{\tau'}$, then the argument in the first case applies.
 - (b) Otherwise, recursively find the next τ'' , and so on. Eventually, we may have a last state x_{ζ} such that no $t > \zeta$ satisfies $M(x_t) = g$. Then, $Q_F^*(x, x_g) \ge Q_F^*(x, x_{\zeta}) \ge Q_F^*(x, g)$. The last inequality is derived in the same way as Eq. 12. Alternatively, there may exist an infinite sequence of such $\{x_{\tau}\}$. Following this sequence, the claim remains true but the supremum is not attainable. However, in this case an x_{τ} can be found in the sequence such that $Q_F^*(x, x_{\tau})$ is arbitrarily close to $Q_F^*(x, g)$.

3.1.3 Insufficiency of Policy Invariance

A well-known property of potential-based reward shaping is that it does not change the optimal policy (Ng et al., 1999). Could the preservation of the quasipseudometric property of the optimal value function under this strategy for reward densification be a trivial consequence of this property? We show in this section that that is not the case. Specifically, we show that violation of Eq. 8 can break triangle inequality without affecting policy invariance.

For a counterexample, consider the simple MDP shown in Fig. 1(a). It has 3 states, $S_1 - S_3$ navigable by action a and an absorbing state S_A that can only be reached by action b from any other state.

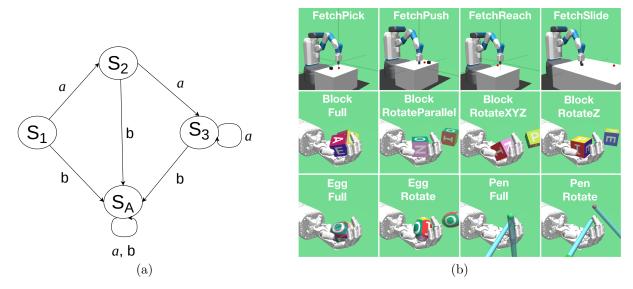


Figure 1: (a) Counterexample used in Secs. 3.1.3 and 3.2. (b) GCRL benchmark environments (Plappert et al., 2018). Figure from (Liu et al., 2023).

Proposition 2. In the counterexample of Fig. 1(a), the potential function for various states with goal S_2 or S_3 can be set such that policy invariance is preserved, yet the assumption of Eq. 8 as well as the quasipseudometric property are violated.

Proof: First note that the optimal policy for goal S_2 or S_3 is to take action a, as these goals are unreachable via action b. For some $\alpha > 0$, we set:

$$\begin{split} \phi(S_1, a, S_3) &= Q^*(S_1, a, S_3) + \alpha & \text{(satisfies Eq. 8)} \\ \phi(S_2, a, S_3) &= Q^*(S_2, a, S_3) - \alpha & \text{(violates Eq. 8)} \\ \phi(S_1, a, S_2) &= Q^*(S_1, a, S_2) + \alpha & \text{(satisfies Eq. 8)} \end{split}$$

These yield: $Q_F^*(S_1, a, S_3) = Q_F^*(S_1, a, S_2) = -\alpha$, $Q_F^*(S_2, a, S_3) = \alpha$. Since the goals are unreachable via action b, it is straightforward to set $\phi(*, b, *)$ such that the optimal policy stays unchanged for Q_F^* . However, now the triangle inequality is violated:

$$Q_F^*(S_1, a, S_2) + Q_F^*(S_2, a, S_3) = -\alpha + \alpha \not\leq Q_F^*(S_1, a, S_3)$$

This proves the criticality of Eq. 8, and that policy invariance of potential based shaping is insufficient to ensure the quasipseudometric property.

3.1.4 Projection

 Q_F^* has the same upper bound as Q^* , since $Q_F^*(s,a,g) = Q^*(s,a,g) - \phi(s,a,g) \le 0$ by Eq. 8. Consequently, the MRN architecture needs no modification, specifically to Eq. 3, as the critic output is guaranteed to be non-positive despite potentially positive shaping rewards. However, Q_F^* has a more informed lower bound:

$$Q_F^*(s, a, g) = Q^*(s, a, g) - \phi(s, a, g)$$

$$\geq -\frac{1}{1 - \gamma} - \phi(s, a, g)$$

$$= -\frac{\gamma^{d(s, a, g)/\eta}}{1 - \gamma}$$
(13)

which we impose on the critic with d_{ac} since it is bounded. Recent analyses (Gupta et al., 2022) have shown that projection informed by shaping effectively reduces the size of the state space for exploration, leading to improved regret bounds. We show the full algorithm for clarity in Alg. 1.

3.1.5 Estimation of η

As both the design of F and the above projection operation needs an estimate of η , it is useful to consider how one might estimate η . From Obs. 1, we want to ensure that d/η is an underestimate of L^* by controlling η . For Euclidean tasks, η can be set to the average distance that the agent covers between successive clock ticks. If this information is not available, it can be approximated via simple experiments. For non-Euclidean tasks (e.g., those with quaternions in states), this is not straightforward. Moreover, we have an added constraint that we do not separate the quaternions from the Euclidean parts of a state to keep the reward design simple as mentioned in Sec. 3.1. In such cases, one could perform data-driven calibration to estimate the ratio d/L^* by learning to optimally complete simpler tasks, measuring their L^* , and finally taking the maximum ratio over those tasks. Sun et al. (2019) propose an approach to estimate L^* using the notion of k-step solvability by learning k-step subpolicies recursively; one of their solvability tests uses an approach similar to the calibration approach suggested above, although it requires the environment to have a reset operator. Either of these approaches could defeat our original purpose of keeping the reward design simple in non-Euclidean tasks, and is a limitation. For our experiments, we select η empirically via a simple parameter search.

Algorithm 1 GCRL with Potential-Shaped Dense Rewards

```
1: Randomly initialize MRN (or another quasimetric) critic Q_F and actor \pi
 2: for episode \leftarrow 1, 2, \dots do
 3:
        Select goal g \sim \rho_{\mathcal{G}}, initial state s \sim \rho_0
        Select action a \sim \pi(\cdot|s, g)
 4:
         for t \leftarrow 1, 2, \dots do
 5:
             Execute action a in state s and receive reward r, next state s'
 6:
             Select next action a' \sim \pi(\cdot|s',g)
 7:
             Store tuple (s, a, r, s', a', g) in replay memory \mathcal{D}
 8:
             a, s \leftarrow a', s'
9:
         end for
10:
        Sample minibatch (s, a, r, s', a', g) from \mathcal{D}
11:
         g', r' \leftarrow \text{Relabel goals } g \text{ to } g' \text{ in the minibatch by HER with original reward } r, \text{ then update } r \text{ to } r'
12:
         Calculate target using shaped reward (r' + F(s, a, s', a', g'))
13:
         Clip target according to Eq. 13 if using arc-cosine distance d_{ac}
14:
         Use (clipped) target in Eq. 2 to update MRN/quasimetric critic Q_F
15:
         Update actor \pi
16:
17: end for
```

3.2 Failure of a Potential-Free, Dense Reward Function

A common dense reward function is $F(s,a,s',a',g) = -d_E(s,a,g) = -\|M(s,a) - g\|_2$ without involving a potential function. Although Euclidean distance generally satisfies the triangle inequality, we argue that $-\|M(s,a) - g\|_2$ may not, using the same counterexample of Fig. 1(a). Since (S_1,a) transitions to (i.e., achieves) S_2 and (S_2,a) to (achieves) S_3 , $\|M(S_1,a) - S_2\|_2 = \|M(S_2,a) - S_3\|_2 = 0$, but $\|M(S_1,a) - S_3\|_2 > 0$ (= δ , say). Hence, $-\|M(S_1,a) - S_2\|_2 - \|M(S_2,a) - S_3\|_2 \not\leq -\|M(S_1,a) - S_3\|_2$. Also, $R(S_1,a,S_2) = R(S_2,a,S_3) = 0$, but $R(S_1,a,S_3) = -1$ by Eq. 1. Noting that the agent can never reach back to state S_2 once it is in S_2 , $Q_F^*(S_1,a,S_2) = -\frac{\gamma}{1-\gamma}$, $Q_F^*(S_2,a,S_3) = 0$ and $Q_F^*(S_1,a,S_3) = -(1+\delta)$. For an appropriate choice of vector representation of states (particularly S_2 and S_3) such that $\delta > \frac{2\gamma-1}{1-\gamma}$, we have $Q_F^*(S_1,a,S_2) + Q_F^*(S_2,a,S_3) \not\leq Q_F^*(S_1,a,S_3)$. Thus, the optimal value function shaped in this way may violate the triangle inequality.

3.3 Intermediate Value Functions

In their critique of on-policy Q-function estimation methods for GCRL such as DDPG in continuous control tasks, Wang et al. (2023) show that *on-policy* Q-function may not be a quasipseudometric, even though the *optimal* Q-function is. However, their counterexample is an extreme policy that is unlikely to be encountered during on-policy iterations. First, we adjust the terminology to refer to these value functions as *intermediate*

value functions, i.e., the value functions learned en-route to the optimal. This adjustment is warranted as the Q-functions may not be on-policy when HER relabelling of goals is utilized. Note however, we use HER in Alg. 1 and for experimentation, but it is not necessary for any theoretical result in this paper.

In this section, we establish that intermediate Q-functions do indeed satisfy the triangle inequality (and hence meet the quasipseudometric criterion) if the policy makes a minimal progress toward the goal. We call such policies progressive policies and believe they are more relevant to intermediate Q-function estimation in GCRL than the counterexample in Wang et al. (2023) (Theorem 1). We first formalize the notion of progressive policies, specify our assumption, and finally show that the corresponding value functions satisfy the triangle inequality.

For notational convenience, we write $\mathbb{E}_{s' \sim T(.|s,a),a' \sim \pi(s')}$ simply as $\mathbb{E}_{s',a'}$. Note that the value function for a policy π satisfies

$$Q^{\pi}(s, a, g) = R(s, a, g) + \gamma \mathbb{E}_{s', a'} \left[Q^{\pi}(s', a', g) \right].$$
 (14)

Definition 1. For any $(s, a, g) \in \mathcal{S} \times \mathcal{A} \times \mathcal{G}$, the progress of a GCRL policy π is given by

$$\Delta^{\pi}(s, a, g) = \mathbb{E}_{s', a'} \left[Q^{\pi}(s', a', g) \right] - Q^{\pi}(s, a, g)$$

We refer to Δ^{π} for the optimal policy as Δ^* . We assume that the progress of π is not unboundedly different from that of the optimal policy, i.e., the following holds for some $0 < \epsilon < \infty$

$$\epsilon \le \Delta^*(s, a, g) - \Delta^{\pi}(s, a, g) \le 2\epsilon. \tag{15}$$

Note that (i) ϵ does not need to be small, just finite; (ii) the counterexample in Wang et al. (2023) does not satisfy this assumption. Our main result of this section is:

Proposition 3. Consider the goal-conditioned MDP $M_{GC} = (S, A, G, T, R, \gamma, \rho_0, \rho_g)$. The value function Q^{π} defined in Eq. 14 for any policy π that satisfies Eq. 15 also satisfies the triangle inequality: $\forall x^1, x^2, x^3 \in \mathcal{X}$,

$$Q^{\pi}(x^1, x^2) + Q^{\pi}(x^2, x^3) \le Q^{\pi}(x^1, x^3).$$

Proof: From Eq. 14 we have,

$$\mathbb{E}_{s',a'} \left[Q^{\pi}(s',a',g) \right] = (Q^{\pi}(s,a,g) - R(s,a,g)) / \gamma.$$

Then, using Eq. 15 and Def. 1, for either $z \equiv (x^1, x^2)$ or $z \equiv (x^2, x^3)$, the following holds:

$$\Delta^*(z) - \Delta^{\pi}(z) = \frac{Q^*(z) - R(z)}{\gamma} - Q^*(z) - \frac{Q^{\pi}(z) - R(z)}{\gamma} + Q^{\pi}(z)$$
$$= (\frac{1}{\gamma} - 1)[Q^*(z) - Q^{\pi}(z)]$$
$$\geq \epsilon \text{ (by Eq. 15)}.$$

Adding for $z \equiv (x^1, x^2)$ and $z \equiv (x^2, x^3)$, we get

$$Q^{\pi}(x^1, x^2) + Q^{\pi}(x^2, x^3) \le Q^*(x^1, x^2) + Q^*(x^2, x^3) - \frac{2\epsilon\gamma}{1 - \gamma}$$
(16)

But similarly for $z \equiv (x^1, x^3)$,

$$\Delta^*(z) - \Delta^{\pi}(z) = (\frac{1}{\gamma} - 1)[Q^*(z) - Q^{\pi}(z)] \le 2\epsilon$$

by Eq. 15. This gives $Q^*(x^1, x^3) \leq Q^{\pi}(x^1, x^3) + \frac{2\epsilon\gamma}{1-\gamma}$. Finally, the result is obtained by combining this with Eq. 16 and noting that the triangle inequality holds for the optimal Q-value function, i.e., $Q^*(x^1, x^2) + Q^*(x^2, x^3) \leq Q^*(x^1, x^3)$.

This result relies on the triangle inequality of the optimal value function as established before in (Liu et al., 2023) for sparse rewards and in Sec. 3.1 for potential-shaped dense rewards. But it does not have any dependence on whether M is one-to-one or onto, hence the two cases $\mathcal{G} \equiv \mathcal{S} \times \mathcal{A}$ and $\mathcal{G} \not\equiv \mathcal{S} \times \mathcal{A}$ do not need to be distinguished. The result also does not assume any specific form of, or bounds on, the reward function. Hence it extends readily to shaped rewards as well, as long as the shaped value function respects the same upper bound (Sec. 3.1.4), $Q_F^{\pi}(.) \leq 0$.

4 Related Work

Several value function representations have been proposed for GCRL over the last decade. Schaul et al. (2015) introduced the bilinear decomposition, later generalized to bilinear value networks (Yang et al., 2022) with better learning efficiency. Pitis et al. (2020) proposed the deep norm (DN) and wide norm (WN) families of neural representations that respect the triangle inequality. However, they are restricted to norm-induced functions, and are generally unable to represent all functions that respect the triangle inequality. By contrast, Possion Quasi-metric Embedding (PQE) (Wang & Isola, 2022) can universally approximate any quasipseudometric, thus improving upon DN/WN. However, as Liu et al. (2023) argue, PQE captures the restrictive form of first hitting-time when applied to GCRL, whereas MRNs capture the more general setting of repeated return to goal $(Q^*(g,g) \neq 0)$, while preserving the UVFA property of PQEs. Durugkar et al. (2021) introduced a quasipseudometric that estimates the Wasserstein-1 distance between state visitation distributions, minimizing which is equivalent to policy optimization in GCRL tasks with deterministic transition dynamics. While they use the Wasserstein discriminator as a potential for reward shaping (as intrinsic motivation), our goal is different. We prove that dense rewards via shaping preserves the triangle inequality for the general class of potential based shaping, not just for the Wasserstein based quasipseudometric. Other recent architectures for GCRL use contrastive representation (Eysenbach et al., 2022) but without regard to quasipseudometric architecture, and quasimetric RL (QRL) (Wang et al., 2023) where temporal distances are learned, although it is unclear if it respects the triangle inequality in stochastic settings.

While the above literature on representation learning has been centered on expressive and flexible representations for GCRL, their analyses are generally restricted to sparse reward settings. In fact, past experimentation with certain types of dense rewards in GCRL have yielded negative results (Plappert et al., 2018). Plappert et al. (2018) argue that dense reward signals are hard to learn from because (i) arbitrary distance measures (e.g., Euclidean distance and quaternions for rotations) are highly non-linear; (ii) dense rewards bias the policy toward specific strategies that may be sub-optimal. Similar arguments also appear in (Liu et al., 2022). However, our setting overcomes these objections. First, we establish the sufficient condition (Eq. 8) for the triangle inequality that may not be satisfied by arbitrary distance measures, ϕ , providing guidance on the contrary. And second, we use potential based reward shaping (Ng et al., 1999) which is policy invariant, hence strategically unbiased. However, we acknowledge the large body of work on reward shaping (Tang et al., 2017; Brys et al., 2014; Devlin & Kudenko, 2012; Knox & Stone, 2009; Van Seijen et al., 2017) of various types (e.g., count-based, intrinsic motivation, human advice, etc.) where careful, heuristic reward design is often employed to explicitly bias the policies.

5 Experimental Results

We use GCRL benchmark manipulation tasks with the Fetch robot and Shadow-hand domains (Plappert et al., 2018); see Fig. 1(b). We use two critic architectures to demonstrate generality: MRN (Liu et al., 2023) and PQE (Wang & Isola, 2022). We experimentally evaluate the following hypotheses:

Hypothesis 1: Potential-based dense rewards satisfying Eq. 8 may improve GCRL's sample complexity when using quasimetric critic architectures. Specifically, the property of Q^* function that they capture—that it satisfies the triangle inequality—is preserved in the presence of shaped rewards with the new value function Q_F^* . Such dense rewards may enable the less restrictive Q_F^* to be learned more efficiently than Q^* .

Hypothesis 2: Plappert et al. (2018) found that certain types of dense rewards hurt RL performance in GCRL robot manipulation tasks. This negative result contradicts our Hypothesis 1. We conjecture that their application of dense rewards did not satisfy the required structure—specifically Eq. 8—which is why it failed. To confirm this contradiction, we verify that our dense reward setting does not deteriorate RL performance

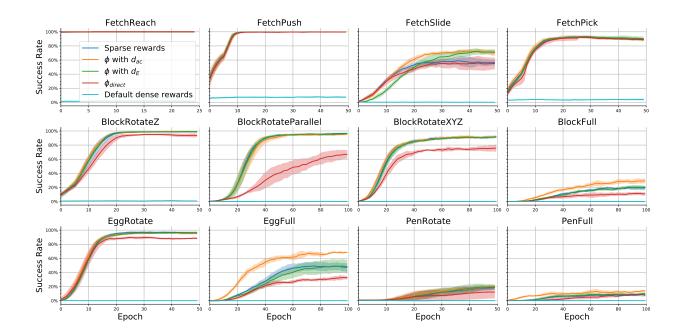


Figure 2: Comparison of MRN with sparse rewards vs. potential-shaped dense rewards with three definitions of potential ϕ and potential-free default dense rewards. Learning curves are averaged over five independent trials, and one standard deviation bands are included. We see statistically significant improvement of performance due to potential-based dense rewards with d_{ac} in 4 of the 12 environments, viz., FetchSlide, BlockFull, Eggfull and PenFull, and in 1 environment (FetchSlide) for d_E . There is no statistically significant deterioration in any environment for either d_{ac} or d_E .

in any task using potential function based on either d_{ac} or d_E . Additionally, to investigate the consequence of violating Eq. 8, we test a direct potential function $\phi_{\text{direct}}(s, a, g) = -\|M(s, a) - g\|_2$ that bypasses the distance function d, and the potential-free default dense rewards of these environments (Gymnasium-Robotics, 2018) that are based on Euclidean distance and may violate the triangle inequality as argued in Sec. 3.2. Note that as ϕ_{direct} is no longer comparable to Q^* , it may easily violate Obs. 1 and Eq. 8. Moreover, as both ϕ_{direct} and the default dense rewards lead to potentially unbounded value functions, we do not clip the corresponding critics (Eq. 13).

Table 1: How often the three potential functions satisfy Eq. 8. We use the learned Q-values with MRN critics from the final training epoch as estimates of optimal Q^* .

Environment	ϕ with d_{ac} (%)	ϕ with d_E (%)	ϕ_{direct} (%)
FetchReach	59.098 ± 1.444	86.345 ± 4.098	8.523 ± 0.203
FetchPush	98.984 ± 0.480	46.860 ± 1.071	28.664 ± 0.599
FetchSlide	99.726 ± 0.089	95.851 ± 0.248	63.053 ± 2.149
FetchPick	97.907 ± 0.819	71.256 ± 0.932	26.057 ± 0.889
BlockRotateZ	96.453 ± 1.052	82.453 ± 0.692	4.693 ± 0.182
BlockRotateParallel	98.927 ± 0.387	81.996 ± 0.681	5.751 ± 0.306
BlockRotateXYZ	98.786 ± 0.289	91.542 ± 0.371	10.269 ± 0.262
BlockFull	99.705 ± 0.115	92.574 ± 0.259	40.727 ± 2.390
EggRotate	97.842 ± 0.908	85.228 ± 1.104	6.674 ± 0.272
EggFull	99.647 ± 0.201	95.662 ± 0.560	56.173 ± 3.040
PenRotate	97.955 ± 1.066	79.251 ± 1.168	20.911 ± 0.964
PenFull	97.994 ± 0.760	76.868 ± 1.044	21.576 ± 0.586

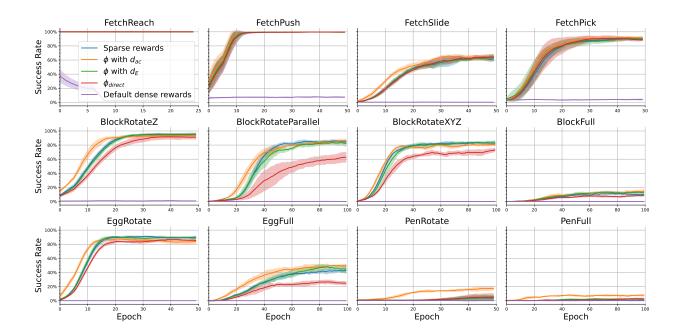


Figure 3: Comparison of PQE critic architecture with sparse rewards vs. potential-shaped dense rewards with three definitions of potential ϕ and potential-free default dense rewards. Learning curves are averaged over five independent trials, and one standard deviation bands are included. We see statistically significant improvement of performance due to potential-based dense rewards with d_{ac} in 3 environments, viz., BlockFull, Eggfull and PenRotate and PenFull, but none with d_E . There is statistically significant deterioration in EggRotate with d_{ac} which we attribute to not tuning η for PQE critics.

We use the MRN code repository publicly available at: https://github.com/Cranial-XIX/metric-residual-network which implements DDPG+HER with MRN and PQE (among others) critics for the GCRL manipulation tasks. We make two simple modifications: (1) add Eq. 6 to the default (sparse) reward function (Eq. 1) as shown on line 13 in Alg. 1; (2) use Eq. 13 to clip the critic's output as shown on line 14 in Alg. 1, but only for d_{ac} . No other changes were made to any algorithm or neural architecture. The newly added parameter η was set to 0.01 for d_{ac} and 0.2 for d_{E} . These values were selected from the sets $\{0.01, 0.02, 0.03, 0.04, 0.05\}^{-1}$ for d_{ac} and $\{0.02, 0.2, 2.0\}$ for d_{E} , using performance improvement with MRN architecture as the criterion. Our modified versions of the code and scripts are also available at: https://github.com/khadimon/GCRL-Dense-Rewards. All experiments were run on NVIDIA Quadro RTX 6000 GPUs with 24 GiB of memory each and running on Ubuntu 24.04.

We show the success rates of the learned policies after every training epoch for MRN and PQE critic architectures in Figs. 2 and 3 respectively. We see from Fig. 2 that indeed potential-based dense rewards with d_{ac} and d_E improve the sample complexity in some environments, to an extent that is statistically significant as shown with standard deviation bands. In particular, there is statistically significant improvement in 4 of the 12 environments, viz., FetchSlide, BlockFull, Eggfull and PenFull with d_{ac} , but only one environment (FetchSlide) with d_E . Not only is the sample complexity improved, but also higher quality policies are learned. This confirms Hypothesis 1. Furthermore, no statistically significant deterioration is observed in any environment with d_{ac} or d_E , confirming Hypothesis 2. In BlockRotateParallel, although the mean of ϕ with d_{ac} appears worse, the difference is not statistically significant as the other methods are subsumed within its error-band. The sample complexity of ϕ with d_E could conceivably improve with a more informed choice of η . We did not conduct a separate parameter search for PQE critics; instead, we used the same η values as above. Fig. 3 with PQE critics shows similar statistically significant improvements in BlockFull,

¹We show sensitivity results for several values of η with MRN in Fig. 4, which indicates that further reducing η is detrimental and possibly fails admissibility of the heuristic.

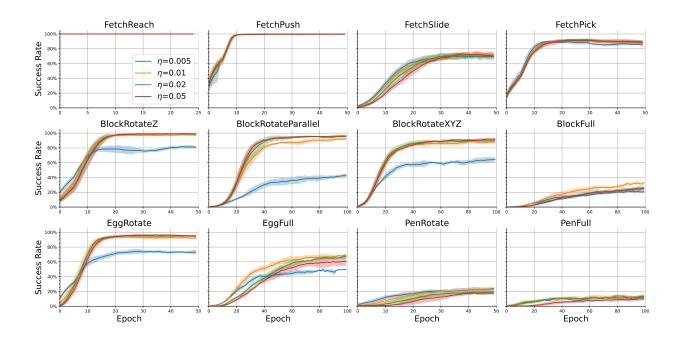


Figure 4: Comparison of MRN with potential-shaped dense rewards using d_{ac} with various values of η . $\eta = 0.01$ is used in experiments reported in Fig. 2.

EggFull, Penrotate and PenFull with d_{ac} , but none with d_E . Interestingly, d_{ac} also shows an initial speed-up in FetchSlide, BlockRotateZ, BlockRotateParallel, and EggRotate compared to the sparse reward setting, but this advantage does not persist when the policies asymptote. Contrarily to Hypothesis 2, however, we see statistically significant deterioration with d_{ac} in EggRotate (and to a lesser extent in BlockRotateXYZ), although none with d_E . We attribute this to the fact that η was not tuned for PQE critics.

Fig. 2 shows that sample complexity is significantly worsened in 6 of 12 environments with ϕ_{direct} , viz., BlockRotateZ, BlockRotateParallel, BlockRotateXYZ, BlockFull, EggRotate and EggFull. Similarly, Fig. 3 with PQE critics shows worse sample complexity for ϕ_{direct} in BlockRotateZ, BlockRotateParallel, BlockRotateXYZ, EggRotate and EggFull environments. These are consistent with Tab. 1 which shows that ϕ_{direct} has a high propensity of violating Eq. 8 in most of these cases, which is further evidence of the criticality of Eq. 8. Figs. 2 and 3 also show significantly worse performance from the default dense rewards in all environments, which is consistent with our finding in Sec. 3.2. In case Plappert et al. (2018) used this default dense reward, these results would also be consistent with their findings.

6 Conclusion

We have presented generalizations of previous results on triangle inequality in the context of value functions in GCRL. Specifically, we have shown that the optimal value function satisfies the triangle inequality even when the reward function is densified with a particular class of shaping functions. Additionally, we have shown that the intermediate value functions also satisfy the triangle inequality if the underlying policy satisfies a certain progressive criterion. Both of these findings contradict previously published results in some ways, which emphasizes the importance of the nuanced conditions behind our results. Experiments in 12 benchmark GCRL tasks confirm that our prescribed dense rewards either improve upon or match the sample efficiency of the sparse reward setting, but could worsen it if the key condition is violated. Future investigations could focus on more general classes of reward functions that preserve the quasipseudometric property of value functions and/or lend themselves to other, potentially more effective, architectures.

7 Acknowledgments

We would like to acknowledge constructive feedback from the anonymous reviewers which helped to substantially improve this paper. An earlier version of this paper was published as the first author's Honors thesis at: https://aquila.usm.edu/honors_theses/1002/. The experiments reported in this paper were conducted on a GPU computing cluster purchased under an Air Force Research Lab grant FA8750-20-1-0105.

References

- Marcin Andrychowicz, Filip Wolski, Alex Ray, Jonas Schneider, Rachel Fong, Peter Welinder, Bob McGrew, Josh Tobin, OpenAI Pieter Abbeel, and Wojciech Zaremba. Hindsight experience replay. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), Advances in Neural Information Processing Systems, volume 30. Curran Associates, Inc., 2017. URL https://proceedings.neurips.cc/paper_files/paper/2017/file/453fadbd8a1a3af50a9df4df899537b5-Paper.pdf.
- Tim Brys, Ann Nowé, Daniel Kudenko, and Matthew Taylor. Combining multiple correlated reward and shaping signals by measuring confidence, volume 28, pp. 1687–1693. AAAI Press, Quec City, Qubec, Canada, Jul 2014.
- Sam Devlin and Daniel Kudenko. *Dynamic potential-based reward shaping*, volume 1, pp. 433–440. International Foundation for Autonomous Agents and Multiagent Systems, Valencia, Spain, 2012.
- Ishan Durugkar, Mauricio Tec, Scott Niekum, and Peter Stone. Adversarial intrinsic motivation for reinforcement learning. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan (eds.), Advances in Neural Information Processing Systems, 2021. https://openreview.net/forum?id=GYr3qnFKgU.
- Benjamin Eysenbach, Tianjun Zhang, Sergey Levine, and Ruslan Salakhutdinov. Contrastive learning as goal-conditioned reinforcement learning. *Advances in Neural Information Processing Systems*, pp. 35603–35620, Dec 2022.
- Dibya Ghosh, Abhishek Gupta, Ashwin Reddy, Justin Fu, Coline Manon Devin, Benjamin Eysenbach, and Sergey Levine. Learning to Reach Goals via Iterated Supervised Learning. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=rALAOXo6yNJ.
- Abhishek Gupta, Aldo Pacchiano, Yuexiang Zhai, Sham Kakade, and Sergey Levine. Unpacking reward shaping: Understanding the benefits of reward engineering on sample complexity. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (eds.), Advances in Neural Information Processing Systems, volume 35, pp. 15281–15295. Curran Associates, Inc., 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/6255f22349da5f2126dfc0b007075450-Paper-Conference.pdf.
- Gymnasium-Robotics. Gymnasium-Robotics Documentation. In Farama Foundation, 2018. https://robotics.farama.org/.
- W Bradley Knox and Peter Stone. Interactively shaping agents via human reinforcement: The tamer framework. In *Proceedings of the Fifth International Conference on Knowledge Capture*, pp. 9–16. ACM, 2009.
- Timothy P. Lillicrap, Jonathan J. Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. In Yoshua Bengio and Yann LeCun (eds.), 4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings, 2016. URL http://arxiv.org/abs/1509.02971.
- Bo Liu, Yihao Feng, Qiang Liu, and Peter Stone. Metric residual networks for sample efficient goal-conditioned reinforcement learning. In *Thirty-Seventh AAAI Conference on Artificial Intelligence (AAAI)*, Februray 2023.

- Minghuan Liu, Menghui Zhu, and Weinan Zhang. Goal-conditioned reinforcement learning: Problems and solutions. In Lud De Raedt (ed.), *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence*, *IJCAI-22*, pp. 5502–5511. International Joint Conferences on Artificial Intelligence Organization, 7 2022. doi: 10.24963/ijcai.2022/770. URL https://doi.org/10.24963/ijcai.2022/770. Survey Track.
- Andrew Y. Ng, Daishi Harada, and Stuart Russell. Policy invariance under reward transformations: Theory and application to reward shaping. In *Proc. 16th International Conf. on Machine Learning*, pp. 278–287. Morgan Kaufmann, 1999.
- Silviu Pitis, Harris Chan, Kiarash Jamali, and Jimmy Ba. An inductive bias for distances: Neural nets that respect the triangle inequality. In *International Conference on Learning Representations*, 2020. https://openreview.net/forum?id=HJeiDpVFPr.
- Matthias Plappert, Marcin Andrychowicz, Alex Ray, Bob McGrew, Bowen Baker, Glenn Powell, Jonas Schneider, Josh Tobin, Maciek Chociej, Peter Welinder, Vikash Kumar, and Wojciech Zaremba. Multi-Goal Reinforcement Learning: Challenging Robotics Environments and Request for Research, 2018. https://arxiv.org/abs/1802.09464.
- Tom Schaul, Daniel Horgan, Karol Gregor, and David Silver. Universal value function approximators. In Francis R. Bach and David M. Blei (eds.), *Proc. International Conference on Machine Learning (ICML)*, volume 37 of *JMLR Workshop and Conference Proceedings*, pp. 1312–1320. JMLR.org, 2015. URL http://dblp.uni-trier.de/db/conf/icml/icml2015.html#SchaulHGS15.
- Hao Sun, Zhizhong Li, Xiaotong Liu, Dahua Lin, and Bolei Zhou. Policy continuation with hindsight inverse dynamics. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, Red Hook, NY, USA, 2019. Curran Associates Inc.
- Haoran Tang, Rein Houthooft, Davis Foote, Adam Stooke, Xi Chen, Yan Duan, John Schulman, Filip DeTurck, and Pieter Abbeel. #exploration: A study of count-based exploration for deep reinforcement learning. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pp. 2753–2762. Curran Associates, Inc., 2017.
- Harm Van Seijen, Mehdi Fatemi, Joshua Romoff, Romain Laroche, Tom Barnes, and Jonathan Tsang. *Hybrid reward architecture for reinforcement learning*, volume 30, pp. 1–10. Curran Associates, Inc., 2017.
- Tongzhou Wang and Phillip Isola. On the learning and learnability of quasimetrics. In *International Conference on Learning Representations*, 2022. https://openreview.net/forum?id=y0VvIg25yk.
- Tongzhou Wang, Antonio Torralba, Phillip Isola, and Amy Zhang. Optimal goal-reaching reinforcement learning via quasimetric learning. In *Proceedings of the 40th International Conference on Machine Learning*, pp. 36411–36430. PMLR, 2023.
- Ge Yang, Zhang-Wei Hong, and Pulkit Agrawal. Bi-linear value networks for multi-goal reinforcement learning. In *International Conference on Learning Representations*, 2022. https://openreview.net/forum?id=LedObtLmCjS.