
Coarse-Grained Smoothness for Reinforcement Learning in Metric Spaces

Omer Gottesman
Brown University

Kavosh Asadi
Amazon Web Services

Cameron Allen
Brown University

Sam Lobel
Brown University

George Konidaris
Brown University

Michael Littman
Brown University

Abstract

Principled decision-making in continuous state-action spaces is impossible without some assumptions. A common approach is to assume Lipschitz continuity of the Q-function. We show that, unfortunately, this property fails to hold in many typical domains. We propose a new coarse-grained smoothness definition that generalizes the notion of Lipschitz continuity, is more widely applicable, and allows us to compute significantly tighter bounds on Q-functions, leading to improved learning. We provide a theoretical analysis of our new smoothness definition, and discuss its implications and impact on control and exploration in continuous domains.

1 INTRODUCTION

Reinforcement learning (RL) aims to develop algorithms that learn optimal policies for sequential decision-making problems [Sutton and Barto, 2018]. In RL, the dynamics of the environment are unknown, and therefore an agent must learn using environmental interactions. While RL has been studied extensively when the set of environmental states is discrete, continuous domains pose a significant challenge: impressive empirical results have been achieved thanks to advances in deep RL, but the theory underlying RL in continuous environments is still poorly developed. A recent body of work addresses this problem by assuming that the Q-function of the optimal policy is Lipschitz continuous [Ni et al., 2019, Tang et al., 2020, Touati et al., 2020].

The Lipschitz assumption is intuitive and allows strong theoretical results to be derived for a variety of algorithms, but

we will show that it is, in fact, overly strong and limiting—and rarely satisfied in domains of interest. In particular, we argue that the key assumption about the smoothness of the Q-function is less about the Q-function itself and more about the ability of the metric to reflect the distance between states under the transition dynamics of the environment. Furthermore, rather than make assumptions on the structure of the Q-function, which are hard to confirm, our assumptions are over the dynamics of the domain, making them much easier to verify empirically or through domain knowledge. We then show that, by choosing a generalized smoothness definition that is qualitatively similar to Lipschitz continuity but designed to properly account for the topology of practical domains, theoretical guarantees similar to those provided by the Lipschitz continuity assumption can be obtained for a much wider range of domains.

The generalization we introduce to Lipschitz smoothness is termed *coarse-grained smoothness*. It is similar to Lipschitz smoothness in that it is (informally) a measure of the largest gradients of the function, but unlike Lipschitz smoothness, gradients that persist only over small scales are ignored. This feature makes coarse-grained smoothness robust to discontinuities and large gradients on small scales. Furthermore, for any function, an infinite number of coarse-grained smoothness measures can be applied, each characterized by the length scale over which gradients are smoothed.

Our main contributions are the following: (1) Identifying that the commonly used Lipschitz assumption on the Q-function is often violated and analyzing the properties of Q-function geometries (Sec. 4 and 5). (2) Proposing a generalization of the Lipschitz assumption (Sec. 6) and analyzing its properties from a purely mathematical perspective (Sec. 7). (3) Providing theoretical analysis of the implications of coarse-grained smoothness to RL (Sec. 8).

2 BACKGROUND AND NOTATION

We use the standard RL formulation of an MDP denoted by $\langle \mathcal{S}, \mathcal{A}, T, R, P_0, \gamma \rangle$, where \mathcal{S} , \mathcal{A} and γ are the state space,

action space, and the discount factor, respectively. The next-state transition probabilities and reward functions are given by $T(s'|s, a)$ and $R(s, a)$, respectively. The initial state distribution is $P_0(s)$. We denote the state-action space as $\mathcal{X} = \mathcal{S} \times \mathcal{A}$ and a point in that space as $x = (s, a)$. The dimensionality of \mathcal{X} is D .

A policy $\pi(a|s)$ gives the probability of taking an action at a given state. The value function is the expected return collected by following the policy from state s : $V^\pi(s) := \mathbb{E}[\sum_{t=0}^T \gamma^t r_t | a_t \sim \pi, s_0 = s]$. The function $Q^\pi(s, a)$ is the expected return for taking action a at state s , and afterwards following π in selecting future actions. An optimal policy, π^* , is defined as a policy that maximizes $V^{\pi^*}(s)$ for all s and its value and Q-functions are denoted $V^*(s)$ and $Q^*(s, a)$, respectively.

We assume \mathcal{X} and \mathcal{S} are metric spaces and denote their distance functions as $d_{\mathcal{X}}(x, x')$ and $d_{\mathcal{S}}(s, s')$, respectively, but will drop the subscript when there is no ambiguity. Note that, in the more intuitive case where metrics are defined over \mathcal{S} and \mathcal{A} , a metric over \mathcal{X} can be easily defined as $d_{\mathcal{X}}(x, x') = d_{\mathcal{S}}(s, s') + C d_{\mathcal{A}}(a, a')$, for any $C \geq 0$. We leave further discussion regarding assumptions on the metric functions to later sections, as one of the main objectives of this paper is to establish what types of metrics are useful in an RL context.

2.1 Lipschitz Continuity and Bounds

The Lipschitz constant of a function f is defined as

$$L \equiv \sup \left\{ \frac{|f(x) - f(x')|}{d(x, x')} \mid x, x' \in \mathcal{X} \right\}. \quad (1)$$

If L is finite, the function is said to be L -Lipschitz continuous. Intuitively, the Lipschitz constant can be thought of as the magnitude of the largest gradient of a continuous function, $f(x)$, over \mathcal{X} , and is infinite for discontinuous functions. If L and the value of $f(x')$ at a set of points $x' \in \mathcal{X}' \subset \mathcal{X}$ are known, an upper and a lower bound for $f(x)$ at any $x \in \mathcal{X}$ can be obtained by

$$\begin{aligned} \hat{f}_{UB}(x) &= \min_{x' \in \mathcal{X}'} f(x') + d(x, x')L \\ \hat{f}_{LB}(x) &= \max_{x' \in \mathcal{X}'} f(x') - d(x, x')L \end{aligned} \quad (2)$$

as shown schematically in Fig. 1(a).

In the RL literature, these bounds are used to make generalizations about the Q-function that support reasoning about which areas of the state-action space should still be explored.

3 RELATED WORK

Numerous papers have employed Lipschitz assumptions to obtain sample-complexity results for exploration. [Pazis and](#)

[Parr \[2013\]](#) introduced a PAC-learning algorithm in metric spaces that maintains a set of known state-action pairs and uses an approximate nearest neighbor approach. Similarly, [Shah and Xie \[2018\]](#) introduced the nearest-neighbor Q-learning algorithm that can output a near-optimal approximation of the Q-function in finite time. [Ortner and Ryabko \[2013\]](#) tackle the undiscounted episodic RL problem and introduced an algorithm whose regret scales sub-linearly with the number of episodes. This algorithm was later improved [[Lakshmanan et al., 2015](#)] and simplified [[Song and Sun, 2019](#)]. [Osband and Van Roy \[2014\]](#) leveraged Lipschitz continuity of the value function to obtain regret bounds that depend on the dimensionality, rather than cardinality, of the underlying MDP.

A more recent line of work focuses on learning adaptive discretization techniques for exploration in metric spaces [[Sinclair et al., 2020](#), [Touati et al., 2020](#)]. These algorithms leverage Lipschitz continuity assumptions to provide upper bounds on the value function of a candidate set of state-action pairs chosen more frequently from promising areas, while maintaining optimism in the face of uncertainty, and are inspired by zooming approaches from the bandit literature [[Kleinberg et al., 2008](#), [Krishnamurthy et al., 2020](#)].

Lipschitz assumptions have been used to tackle numerous other challenges in RL. For example, [Pirotta et al. \[2015\]](#) use Lipschitz assumptions to provide a new policy-gradient algorithm that can adaptively choose a learning rate that yields monotonic improvements in the policy-gradient objective. [Asadi et al. \[2018\]](#) show that having access to Lipschitz models is a key ingredient for combating the compounding-error problem in model-based reinforcement learning. [Berkenkamp et al. \[2017\]](#) give a Lyapunov approach, and [Chandak et al. \[2020\]](#) propose a Seldonian algorithm to maintaining safety in RL that leverages Lipschitz assumptions on the model. [Lecarpentier and Rachelson \[2019\]](#) propose a worst-case approach to addressing non-stationary problems where the rate of non-stationarity with respect to time satisfies a Lipschitz assumption. [Tang et al. \[2020\]](#) introduce Lipschitz Value Iteration, an algorithm that gradually tightens the range of Q-value estimates to perform off-policy evaluation. [Gelada et al. \[2019\]](#) lean on Lipschitz assumptions to build a representation and state abstraction that is conducive to value-function optimization, and [Gottesman et al. \[2019\]](#) use similar Lipschitz assumptions for combining different estimators for off-policy evaluation and applying the resultant estimator to medical applications. Lipschitz continuity has also been leveraged in the context of lifelong RL [[Lecarpentier et al., 2020](#)].

More generally, Lipschitz continuity has garnered attention in other areas of machine learning and in particular in the deep learning literature. Most notably, Lipschitz properties have been used to provide norm-based capacity control [[Neyshabur et al., 2015](#)], learn generative models for matching distributions in high-dimensional spaces [[Ar-](#)

iovsky et al., 2017], and handle adversarial examples to improve the robustness of neural networks against malevolent actors [Finlay et al., 2019].

In the measure theory literature, Lancien and Dalet [2017] have used the same definition we use in generalising the Lipschitz assumption. However, our theoretical analysis of coarse grained smoothness has a different focus than that found in the math literature: while Lancien and Dalet [2017] use coarse-Lipschitzness as a stepping stone to generalize the notion of norm-attaining Lipschitz equivalences, we use it to derive modified bounds, and analyze how they compare with traditional Lipschitz bounds. Furthermore, our analysis of the application of coarse-grained smoothness to the field of RL is novel as well.

Closely related to our work, there exist alternative assumptions in the literature that, akin to Lipschitz continuity, capture the intuitive notion of smoothness in different ways. For example, Kakade et al. [2003] make a local-modeling assumption—even more stringent than the Lipschitz assumption—that they use to extend the E^3 algorithm [Kearns and Singh, 2002] to metric spaces. Additionally, Touati et al. [2020] allow for a generalized variant of the Lipschitz assumption where an additive error term may exist, and show that their exploration algorithm is robust to such mis-specifications. A conceptual advantage of our generalized smoothness assumption compared to Touati et al. [2020], is that our approach is explicit in defining a length scale over which changes of the Q-function are ignored, allowing for a principled way of specifying exactly how the Lipschitz smoothness is violated, as well as quantifying the smoothness of the function over multiple length scales.

4 MOTIVATING EXAMPLE: Q-FUNCTION DISCONTINUITIES IN CONTINUOUS RIVERSWIM

To provide intuition for why alternative smoothness measures to Lipschitz continuity should be explored, we examine the smoothness properties of the Q-function in a simple environment with characteristics common to many continuous control tasks of interest; we formalize these characteristics in Sec. 5. In many domains of interest, such as navigation and robotics, the reward function is derived from attempting to reach a set of goal states, and is thus discontinuous. We will assume the metric function, $d_S(s, s')$, quantifies how quickly an agent can transition between states—we will later argue that this assumption, although strong, can replace the smoothness assumption on the Q-function, and is much easier to verify or enforce at problem-formulation time.

The environment we use for illustration is a continuous version of the riverswim environment [Strehl and Littman, 2008]. The 1D state is a real number, the initial state is

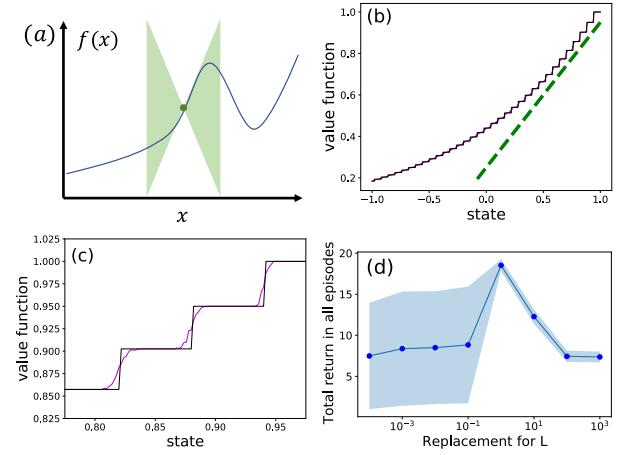


Figure 1: (a) Schematic of Lipschitz bounds. (b) The value function for the optimal policy in the continuous riverswim environment for both deterministic (black) and stochastic (magenta) dynamics. The dashed green line represents a coarse grained slope of the function. (c) Magnification of (b). (d) Performance of the algorithm presented in Ni et al. [2019] for different replacement values for L .

$s_0 = 0$, and the continuous actions are $a \in [-a_{\max}, a_{\max}]$. The transition function is $s_{t+1} = s_t + a_t - c$, where $0 \leq c \leq a_{\max}$. The dynamics can be seen as modeling swimming against a current. An episode ends with a small positive reward, r_{left} , upon reaching $s \leq -1$, and a larger reward, $r_{\text{right}} > r_{\text{left}}$, upon reaching $s \geq 1$. It is a good example of an environment in which an agent must explore effectively, or it will never discover the large reward at $s \geq 1$.

Here, the transition function is continuous with respect to both states and actions. However, the goal-based reward function results in a discontinuous Q-function. The optimal policy is to always take action $a = a_{\max}$.¹ The value function and Q-function for the optimal policy are:

$$V^*(s) = r_{\text{right}} \gamma^{\lceil \frac{1-s}{a_{\max}-c} \rceil}, \quad (3)$$

$$Q^*(s, a) = \begin{cases} r_{\text{left}} & s + a \leq c - 1 \\ \gamma V^*(s + a - c) & c - 1 < s + a < c + 1, \\ r_{\text{right}} & c + 1 \leq s + a \end{cases} \quad (4)$$

where $\lceil \cdot \rceil$ denotes rounding up to the closest integer. Note that there is a discontinuous jump in Q^* at every state where the number of steps required to reach the goal increases by one, as well as when an action will lead to such a state. Thus, Q^* has discontinuities with respect to both states and

¹If $r_{\text{left}} > r_{\text{right}} \gamma^{\lfloor \frac{2}{a_{\max}-c} \rfloor}$ the optimal policy will have states from which it will move left, but for simplicity we ignore that region of the parameter space.

actions and clearly violates the Lipschitz property. In Sec. 5, we argue that this phenomenon is common to a wide variety of domains.

In Fig. 1(b), we plot V^* in black and see that, despite the discontinuous structure of the value function, in any region in the state space, a relatively clear slope can be seen if one ignores the discontinuous steps—qualitatively approximated by the green line. Furthermore, we also plot (in magenta) the empirically estimated state values when Gaussian noise with standard deviation of $0.03a_{\max}$ is added to the transition (Fig. 1(c) zooms in on the upper right part of Fig. 1(b) for clarity). Although adding stochasticity eliminates the discontinuity of the Q-function, it is still the case that its maximal gradient is much larger than its coarse-grained rate of change.

The Lipschitz constant is usually used to bound the Q-function in one region of the state-action space based on observations in other regions. Therefore, it seems intuitive that it would be better to use the effective slope rather than the true “infinite” Lipschitz constant. In Sec. 6, we formalize this idea.

To illustrate why the coarse-grained slope might be a better measure of the smoothness of the Q-function, we apply the control algorithm proposed by Ni et al. [2019] to the river-swim environment with different inputs as the Lipschitz-constant value, and plot the total reward obtained over 50 episodes as a function of this value. Despite the true Lipschitz constant being “infinite”, we see that using large replacement for L leads to poor performance due to over-exploration. The algorithm performs best for an intermediate replacement for $L \approx 1$, which tellingly matches the largest coarse-grained slope shown by the green line in Fig. 1(d).² Note also that using too small of a replacement for L leads to poor exploration as seen by the low mean performance and large variance for these values. These results demonstrate that the Lipschitz constant is not necessarily the smoothness measure we should be using in RL algorithms, and therefore alternative measures should be explored.

5 CONSTRAINTS ON THE GEOMETRY OF Q-FUNCTIONS

Motivated by the observations in the previous section, we now show that discontinuous Q-functions will be the norm rather than the exception, as they arise under conditions very common in RL. Thus, care must be taken when designing algorithms that assume Lipschitz continuity. On the other hand, we show that these discontinuities can be bounded by making some assumptions regarding the metric over \mathcal{S} , which motivates the introduction, in Sec. 6, of an alternative

²Although the Lipschitz constant of the Q-function is defined over the state-action space, here we approximate it by the slope in the $a = a_{\max}$ subspace.

smoothness measure that can handle these discontinuities.

The key assumption we make formalizes the intuitive notion that the metric over the state space quantifies how quickly an agent can transition between states.

Assumption 1. *For any two states $s, s' \in \mathcal{S}$ such that $d(s, s') \leq d_{\min}$, there exists a policy for which within at most k time steps (i) with probability $1 - \delta$, s' can be reached from s , and (ii) with probability 1, the agent remains within a distance of less than d_{\min} from s' .*

This assumption can be applied to deterministic continuous domains (in which case $\delta = 0$) or stochastic discrete domains when a suitable metric is provided (or stochastic continuous domains under discretization). In Appendix B, we analyze how Assumption 1 can be modified for stochastic continuous domains, and how to relax part (ii) of the assumption. Note that such a metric is not always trivial to obtain—consider for example a continuous 2D maze. Two points on opposing sides of a wall may be very close in Euclidean distance, but it may take the agent many steps to move from one to the other. However, such assumptions, or ones similar in nature, are arguably required for any RL approach in continuous domains where learning requires generalization to unobserved states. Furthermore, such assumptions are easier to validate or design into the problem formulation than assumptions on the structure of the Q-function. While it is often difficult to know what the k -step reachable neighborhood for an MDP is, it is often possible to do so for the special case of $k = 1$, allowing for estimation of a d_{\min} value for which Assumption 1 holds.

For deterministic continuous domains, Assumption 1 leads immediately to discontinuities, as we saw in Sec. 4. The simplest example demonstrating this property is to consider the boundary of a region from which a terminal goal state can be reached in one step (such a boundary arises by taking $k = 1$ in Assumption 1). The value of all states within the boundary will be the reward, r , for that goal state, while states at an infinitely small distance outside that boundary will have value at most γr . While adding stochasticity to the dynamics might remove the strict discontinuities, as long as the transition noise is much smaller than the average step sizes, local gradients, while finite, may still be unhelpfully large, as can be seen in Fig. 1(c).

While the Q-function may have discontinuous “jumps” (infinite gradients), Assumption 1 also implies that the *magnitude* of such “jumps” can be bounded.

Theorem 1. *Under Assumption 1, for any two points $x_1, x_2 \in \mathcal{X}$ such that $d(x_1, x_2) \leq d_{\min}$, the difference $\Delta Q^* \equiv |Q^*(x_1) - Q^*(x_2)|$ is bounded by*

$$\Delta Q^* \leq \frac{1 - \gamma^k}{1 - \delta \gamma^k} \left(Q_{\max, (1,2)}^* - \frac{r_{\min}}{1 - \gamma} \right), \quad (5)$$

where $Q_{\max, (1,2)}^* \equiv \max(Q^*(x_1), Q^*(x_2))$.

The proof can be found in Appendix A, but the main idea is to note that, for any two states less than d_{\min} from each other, the agent can transition from the lower value state to the higher value state, and then follow the policy from the higher value state. Thus, all that is needed to complete the proof is to compute the rewards and discount factor accumulated along the way between the two states.

The bound in Theorem 1 is a local one in the sense that it depends on the Q-values of the state-action points compared, and reflects the fact that, if the values of the points are small, the difference between their values will be comparatively small as well. (For example, note that the “steps” in Fig. 1(b) are smaller where the value function is lower.) However, by substituting Q_{\max}^* , the maximum value for states in the domain,³ for $Q_{\max, (1,2)}^*$, a corollary of Theorem 1 provides a looser bound for the difference between any two points in \mathcal{X} that we can obtain without knowledge of their Q-values.

Corollary 1. *For any two points $x_1, x_2 \in \mathcal{X}$ such that $d(s_1, s_2) \leq d_{\min}$, the difference $\Delta Q^* \equiv |Q^*(x_1) - Q^*(x_2)|$ is bounded by*

$$\Delta Q^* \leq \Delta Q_{\max}^* \equiv \frac{1 - \gamma^k}{1 - \delta\gamma^k} \left(Q_{\max}^* - \frac{r_{\min}}{1 - \gamma} \right), \quad (6)$$

where Q_{\max}^* is the maximal value of any state in the environment.

We have demonstrated that, in contrast to the key assumption commonly made in the continuous RL literature, Q-functions are discontinuous for many domains of interest. However, the additional constraints we have proved regarding the geometry of the Q-function allow us to propose a new, more meaningful smoothness measure for use in RL, which we do next.

6 COARSE-GRAINED SMOOTHNESS MEASURES

The results presented in Sec. 5 suggest that, while the Q-function may be discontinuous locally, there are constraints on how much it can change over a given distance in \mathcal{X} . This insight leads us to propose an alternative smoothness measure that ignores discontinuities on the local scale, defined by the length parameter α , with a corresponding coarse-grained smoothness parameter, L_α :

$$L_\alpha \equiv \sup \left\{ \frac{|Q^*(x) - Q^*(x')|}{d(x, x')} \mid d(x, x') \geq \alpha \right\}. \quad (7)$$

Note that, unlike L , for any bounded function over a bounded metric space, L_α exists and is finite for all $\alpha > 0$. Furthermore, L_α is a monotonically decreasing function of α , and $\lim_{\alpha \rightarrow 0} L_\alpha = L$ is the standard Lipschitz constant.

³ Q_{\max}^* is usually known and depends on the dynamics of the MDP. For example, it is r_{\max} in domains where the episode ends when a goal is reached or $r_{\max}/(1-\gamma)$ in infinite horizon domains.

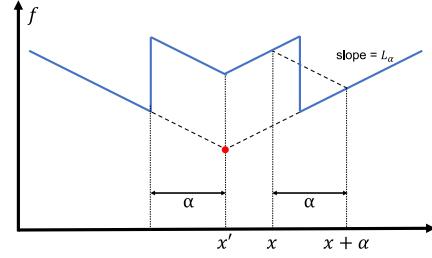


Figure 2: Schematic demonstration how L_α is used to derive the upper bound of a function based on knowledge of its value at a point x .

We can use L_α for upper and lower bounds of a function in a manner similar to the use of the traditional Lipschitz constant (Eq. 2). The upper bound that the function at point x' imposes on the function’s value at point x is simply $f(x') + L_\alpha d(x, x')$ for $d(x, x') > \alpha$.

We will now derive an upper bound on $f(x)$ for $d' \equiv d(x, x') \leq \alpha$. The derivation of these bounds is presented in scalar (1D) notation but is easily applied to any normed metric over \mathbb{R}^n , by applying the 1D derivation along the unit vector pointing from x to x' . For $d' \equiv d(x, x') \leq \alpha$, we can bound $f(x)$ in the following way (Fig. 2 schematically demonstrates how these bounds are derived): By noting that $f(x' + d' + \alpha) \leq f(x') + L_\alpha(d' + \alpha)$, then because $d(x' + d' + \alpha, x) = \alpha$, we get $f(x) \leq f(x' + d' + \alpha) + L_\alpha \alpha \leq f(x') + L_\alpha(d' + \alpha) + L_\alpha \alpha = f(x') + L_\alpha(d' + 2\alpha)$. Applying the same logic to the lower bound we get:

$$\begin{aligned} \hat{f}_{UB}^\alpha(x) &= \min_{x'} f(x') + g_x^\alpha(x') \\ \hat{f}_{LB}^\alpha(x) &= \max_{x'} f(x') - g_x^\alpha(x'), \end{aligned} \quad (8)$$

$$g_x^\alpha(x') \equiv \begin{cases} L_\alpha(d(x, x') + 2\alpha) & d(x, x') \leq \alpha \\ L_\alpha d(x, x') & d(x, x') > \alpha. \end{cases} \quad (9)$$

An important point is that, while this derivation provides bounds on points at any distance from each other, the L_α smoothness property makes an assumption *only* on points that are at least α away from each other. These bounds could be replaced by simpler more intuitive bounds, such as $g_x^\alpha(x') = \alpha L_\alpha$ for $d(x, x') \leq \alpha$, but such bounds would require additional assumptions on the geometry of the Q-function at small length scales, which is exactly what our formulation aims to avoid.

Also of note is that the L_α bounds are discontinuous, which would make them more difficult to work with using gradient-based algorithms. However, most Lipschitz-based methods are non-parametric, and therefore do not use a parameterized representation of the bound, and for such algorithms the discontinuity of the bounds will not pose a problem.

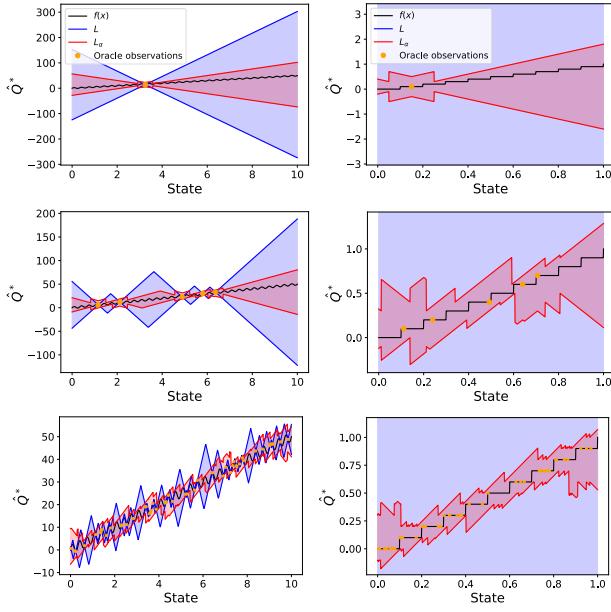


Figure 3: Comparison of L (blue) vs. L_α (red) bounds for a continuous (left) and discontinuous (right) function, with different numbers of sampled points.

Taking the limit $d(x, x') \rightarrow 0$ in Eq. 8 and 9 leads to the following proposition bounding the maximum discontinuity gap in a L_α -smooth function:

Proposition 1. *For an L_α -smooth function f , for any two points x and x' ,*

$$\lim_{d(x, x') \rightarrow 0} |f(x) - f(x')| < 2L_\alpha\alpha. \quad (10)$$

To compare the L_α bounds with the traditional Lipschitz bounds, in Fig. 3, we plot both bounds based on 1, 5, and 50 oracle observations for two functions: a sine function riding on top of a linear slope and a discontinuous “stairs” step function.⁴ The L_α bounds are shown in red (for α equal to the wavelength of the sine function and the width of a step for the stairs function), while the standard Lipschitz bounds are shown in blue for comparison. (Note that for the stairs function L is infinite, resulting in vacuous bounds, indicated as a blue “background”.)

Fig. 3 demonstrates the strength of using a coarse-grained Lipschitz constant: Because L_α could be orders of magnitude smaller than L (even if L is finite), L_α provides significantly tighter bounds for points far away from observations. A trade-off, however, is that L_α is unable to provide tight bounds close to observed points. (For the sine function, the red zone tends to lie inside of the blue zone, except close to the data points.)

⁴In Appendix C, we present the exact functions as well as discuss how their L and L_α values are computed.

A crucial advantage of the L_α formalism is that, unlike the Lipschitz constant, L_α can be upper bounded using empirical data. While empirically upper bounding the Lipschitz constant of a function is impossible, as infinitely large gradients may be “hiding” at infinitely small length scales, L_α smoothness is insensitive to gradients on scales smaller than α , and thus can be upper bounded by sampling the function at a density larger than $\mathcal{O}(\alpha^{-1})$, as we show in Theorem 2 (proof in Appendix E). Doing so requires making an assumption bounding the difference between the values of a function for neighboring points, but because in Corollary 1 we derived exactly such a bound, this assumption is appropriate in RL contexts.

Theorem 2. *Let $f(x)$ be a function s.t. for all $x_i, x_j \in \mathcal{X}$, $d_{ij} \equiv d(x_i, x_j) \leq d_0$ implies $\Delta_{ij} = |f(x_i) - f(x_j)| \leq \Delta_{d_0}$ and define $\Delta_\epsilon \equiv \lceil \frac{\epsilon}{d_0} \rceil \Delta_{d_0}$. Assume also that \mathcal{X} is convex and that $f(x')$ is known at a subset of points $\mathcal{X}' \subset \mathcal{X}$, s.t. \mathcal{X}' is an ϵ -cover of \mathcal{X} , and $d_0 \leq \epsilon \leq \alpha$. Then*

$$L_\alpha \leq \max_{\substack{x_i, x_j \in \mathcal{X}' \\ d_{ij} \geq \alpha - 2\epsilon}} \frac{\Delta_{ij} + 2\Delta_\epsilon}{\max(\alpha, d_{ij} - 2\epsilon)}. \quad (11)$$

7 THEORETICAL PROPERTIES OF L_α BOUNDS

In this section we analyze L_α bounds from a pure function approximation perspective, independently of RL. We will discuss implications to RL in Sec. 8.

7.1 Comparing L and L_α Bounds

We first quantify when using L_α to bound a function will provide tighter bounds than the standard Lipschitz bounds. Obviously, if a function $f(x)$ is discontinuous, as is common for Q-functions, L_α -bounds are trivially better because Lipschitz bounds do not exist. When the Lipschitz bound does exist, Theorem 3 quantifies over what fraction of \mathcal{X} the L_α -bound is tighter than the Lipschitz bound.

Before we can state Theorem 3 (proof in Appendix D), we first define several quantities and constants to make the presentation precise. First, the volume in \mathcal{X} of a ball of radius r is $C_{D,r}^D$. For example, for a Euclidean metric in 2 and 3 dimensions, we have $C_D = \pi$ and $C_D = 4\pi/3$, respectively. To generalize the notion of a radius to non-spherical spaces, we also define the linear size of an arbitrarily shaped metric space \mathcal{X} , $l_{\mathcal{X}}$, as a function of the volume of \mathcal{X} , $V_{\mathcal{X}}$, through the relation $V_{\mathcal{X}} = C_{D_{\mathcal{X}}}(l_{\mathcal{X}})^D$. Thus, if \mathcal{X} is spherical, we have $C_{D_{\mathcal{X}}} = C_D$, but the more general value of $C_{D_{\mathcal{X}}}$ allows for $l_{\mathcal{X}}$ to be defined for spaces of different shapes.

Theorem 3. *Let $f(x)$ be defined over the metric space \mathcal{X} , and assume the value of $f(x)$ is known at N points. Assume $f(x)$ is Lipschitz continuous with Lipschitz constant L . For any value of α and its corresponding L_α over $f(x)$, define*

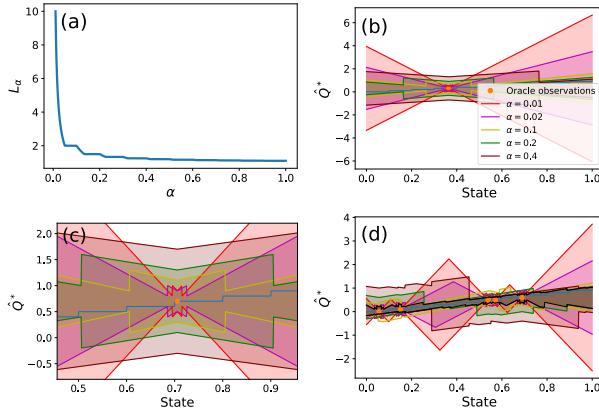


Figure 4: (a) L_α vs. α for the steps function. By combining the L_α bounds for different values of α ((b) and zoomed in (c)), we can obtain a much tighter overall bound (d).

l_α to be

$$l_\alpha = \begin{cases} \alpha & L \geq L_\alpha > L/3 \\ \frac{2L_\alpha\alpha}{L-L_\alpha} & L/3 \geq L_\alpha. \end{cases} \quad (12)$$

Then, if $\frac{C_D}{C_{D,\mathcal{X}}} \left(\frac{l_\alpha}{l_{\mathcal{X}}} \right)^D N < 1$, the fraction of volume of \mathcal{X} over which the L_α bound is tighter than the Lipschitz bound is at least $1 - \frac{C_D}{C_{D,\mathcal{X}}} \left(\frac{l_\alpha}{l_{\mathcal{X}}} \right)^D N$.

Intuitively, Theorem 3 states that, as long as observations are not so numerous as to cover \mathcal{X} such that all observed points are less than a distance of l_α away from each other (note that $l_\alpha < \alpha$), L_α -bounds are tighter than Lipschitz across a large fraction of the space. Importantly, Theorem 3 shows that L_α -bounds are tighter over a larger fraction of the space in the low data regime (small N) and exponentially better with increasing dimensionality. Furthermore, the margin by which L_α bounds are tighter also grows significantly in the low data regime (small N) and in high-dimensions. In that sense, the curse of dimensionality increases the relative advantage of L_α bounds over Lipschitz bounds.

7.2 Overcoming the Shortcomings of L_α Bounds

As mentioned earlier, one drawback of using L_α bounds is that close to observed points, they are looser than Lipschitz bounds, if those exist. In fact, even in the limit of infinite data, the uncertainty gap, $\hat{f}_{UB}^\alpha(x) - \hat{f}_{LB}^\alpha(x)$, will approach $2L_\alpha\alpha$ rather than 0, for all x . In this section we address this problem by providing a method of using multiple values of α and their corresponding L_α .

The trade-off between tight bounds close to observed points for L_α with small values of α , and tight bounds away from observed points for large α values suggests that we can achieve the best of both worlds by combining bounds over

different scales to obtain a *composite-bound*. If we know L_α for multiple values of α , we can compute the tightest possible upper-bound by taking the minimum of the upper-bounds computed using different (α, L_α) pairs, and follow a similar procedure for computing a lower bound, giving:

$$\begin{aligned} \hat{f}_{UB}^{\alpha^*}(x) &= \min_{\alpha} \min_{x'} f(x') + g_x^\alpha(x') \quad (13) \\ \hat{f}_{LB}^{\alpha^*}(x) &= \max_{\alpha} \max_{x'} f(x') - g_x^\alpha(x'). \end{aligned}$$

In Fig. 4, we illustrate how such a procedure would work for bounding a step function $f(x)$ for which we have oracular knowledge of $f(x)$ for sampled points. In Fig. 4(a), we plot L_α vs. α ,⁵ and compare the bounds generated using the value of a single point for five (α, L_α) pairs in Fig. 4(b-c). In Fig. 4(d), we show in black the resulting bounds when all five (α, L_α) pairs are used. Note that, if $f(x)$ is Lipschitz continuous and known, L can also be used for this method.

The main hurdle to using this method is that it requires knowing L_α for multiple values of α (in an ideal world for *all* values of α), which assumes a lot of domain knowledge. However, even a small number of additional (α, L_α) pairs can provide significant improvement to the tightness of Lipschitz bounds, and in an RL setting, estimates for such pairs could be obtained using results such as Prop. 2.

8 IMPLICATIONS FOR RL

We now discuss the implications of using L_α bounds instead of Lipschitz bounds in RL. We first show in Sec. 8.1 that the structure of MDPs allows us to upper bound L_α for one meaningful value of α , and thus the ability to use it for RL is guaranteed for any domain over which a suitable distance metric is defined. This fact stands in contrast to the Lipschitz constant, which is not always known and is hard to estimate. In Sec. 8.2, we discuss how L_α bounds can be incorporated into existing Lipschitz-based algorithms, and what type of theoretical results can be expected to carry over when switching between the bounds.

8.1 Upper Bounding L_α

A crucial concern when applying Lipschitz (or L_α) bounds is knowing the specific value of L (or L_α). For Lipschitz bounds, assuming L is known, as many recent papers do, is a very strong assumption. Estimating the Lipschitz constant is very hard, and, because very steep gradients could potentially exist on very small length scales in \mathcal{X} , it is not possible to derive an upper bound from samples. Such a bound is required for many optimality results to hold. Fortunately, upper bounds *can* be estimated for the L_α constant: See Prop. 2 for an upper bound on L_α for $\alpha = d_{\min}$ (proof in Appendix F). This property not only provides an off-the-shelf

⁵In Appendix C, we show how L_α can be analytically computed for the step function.

value for use in algorithms, it does so for a relevant length scale of the MDP. Consider for simplicity the case of $k = 1$, so for all points $x, x' \in \mathcal{X}$ such that $d(x, x') \leq d_{\min}$ the agent can transition between the points in one step. The motivation for defining L_α was to have a smoothness measure that is insensitive to fluctuations at small length scales of \mathcal{X} , and therefore choosing $\alpha = d_{\min}$ smooths over changes in Q on the scale of single transitions, which, as shown in Sec. 4 and 5, is an important source of discontinuous jumps of the Q-function.

Proposition 2. *For an MDP satisfying Assumption 1, and convex \mathcal{X} , for $\alpha = d_{\min}$, and ΔQ_{\max}^* defined in Corollary 1, we can bound L_α by $L_\alpha \leq \frac{2\Delta Q_{\max}^*}{d_{\min}}$.*

This bound is a marked improvement over the potentially infinite Lipschitz constant, which cannot be bounded empirically. Furthermore, in Appendix F.1, we present an example for which this bound is arbitrarily tight, and give intuition as to the general MDPs in which we can expect such tightness.

8.2 Extending L_α Bounds to Existing Lipschitz Algorithms and Results

Most Lipschitz-based algorithms use the Lipschitz assumption to generate bounds on the Q-function; by substituting the L_α bounds for the Lipschitz bounds, the same algorithms can be used as-is. How would the theoretical guarantees of these algorithms compare with the original algorithms? This discussion can be divided into several cases. First, we have demonstrated that, for a wide range of common RL tasks, the Lipschitz smoothness assumption on the Q-function is violated, making the argument moot. By contrast, for any bounded domain and a bounded Q-function, L_α will always exist for all $\alpha > 0$. Second, even if the Q-function is Lipschitz continuous, L is usually unknown and hard to estimate or even upper bound. Meanwhile, Prop. 2 provides an upper bound for L_α for at least one value of α that can readily be used in algorithms. Furthermore, if L is known, any algorithm that relies on using it to upper bound the Q-function can be modified to also incorporate L_α for $\alpha = d_{\min}$ using Eq. 13 at an additional computational factor of at most 2.

Finally, although we can always use both L_α and L simultaneously (if L exists and is known), we can compare their performance individually. Early in training, when the agent has observed relatively few points, by Theorem 3, L_α provides tighter bounds than L , and can therefore explore more efficiently by ignoring large regions of the state-action space whose values can be bounded below optimality. However, later in training, when the domain has been thoroughly explored, L_α bounds are looser in the vicinity of observed points and are less suitable to fine tuning policies already close to optimality. As a consequence, many of the theoretical results in the literature can be extended to using L_α bounds, but care must be taken to keep track of when the agent has observed enough of the space that the L bounds become

tighter and should be used instead.

As an example, Prop. 3 demonstrates how L_α bounds can be applied to the theoretical results of the recently proposed ZoomingRL algorithm [Touati et al., 2020] (proof in Appendix G).

Proposition 3. *In the ZoomRL algorithm introduced in Touati et al. [2020], for any value of α such that $L_\alpha \leq L/2$, substituting L_α -bounds for L results in improved regret for at least the first $\frac{2}{3\alpha}$ episodes.*

This result underscores a trade-off when using L_α for exploration. When α is small, we obtain better regret using L_α over more episodes, but the regret gap is small. Conversely, with a large value of α , we can obtain larger regret improvement, but that improvement exists for fewer episodes. Note that quite often L_α is much smaller than L even for relatively small values of α (see, for example, Fig. 4(a)), and therefore in practice it would often be easy to find an α such that $L_\alpha \leq L/2$.

The example we provide for a modified bound of existing theory is given for the results of Touati et al. [2020] rather than the algorithm of Ni et al. [2019] which we used in Sec. 4 to provide intuition and motivation, and also comes with theoretical performance bounds. The proof used in Ni et al. [2019] does not lend itself to be modified for coarse grained smoothness, as it relies on the asymptotic case where \mathcal{X} is sampled arbitrarily densely. Since coarse grained smoothness is designed to provide advantage when data is sparse, and L_α -bounds are loose in the vicinity of observed points, proofs which rely on the distance between sampled points going to zero are not amenable to modification for coarse grained smoothness.

9 CONCLUSIONS

Using the smoothness properties of the Q-function is a promising direction for performing RL in continuous state-action spaces. We have shown, however, that care must be taken when making assumptions about smoothness properties, as traditional notions of smoothness do not always hold for such domains. Our coarse-grained smoothness approach takes a first step towards bridging the gap between existing work on Lipschitz-based RL and the realistic Q-function geometry of many MDPs. We have demonstrated that many of the algorithms already presented in the literature can easily be modified to incorporate this new smoothness measure, and that useful theoretical properties of these algorithms can be carried over.

The majority of works exploiting the Lipschitz smoothness of Q-functions such as Touati et al. [2020], Ni et al. [2019] are non-parametric in nature, whereas the overwhelming majority of state-of-the-art algorithms used for high dimensional MDPs rely on parametric neural networks. Therefore an important potential long term contribution of this work

lies not in replacing and outperforming existing parametric methods, but rather in being combined with them to improve performance. For example, the results in our paper, as well previous related works, allow us to make theoretical statements about how much exploration is required in continuous MDPs, and how the infinite search space for exploration can be pruned. These insights could potentially be combined with parametric models to improve our understanding of how exploration should be conducted and to develop more efficient exploration algorithms, a problem widely recognized as one of the biggest challenges in RL.

Acknowledgments

This work was supported in part by an NSF Graduate Research Fellowship under grant #2040433, NSF grants #1717569, #1955361, and CAREER award #1844960, DARPA grant W911NF1820268, and ONR contracts N00014-17-1-2699 and N00014-21-1-2200, and was conducted using computational resources and services at the Center for Computation and Visualization, Brown University. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The content is solely the responsibility of the authors and does not necessarily represent the official views of DARPA, the NSF, the ONR, or the AFOSR.

References

- M. Arjovsky, S. Chintala, and L. Bottou. Wasserstein generative adversarial networks. In *International Conference on Machine Learning*, pages 214–223. PMLR, 2017.
- K. Asadi, D. Misra, and M. Littman. Lipschitz continuity in model-based reinforcement learning. In *International Conference on Machine Learning*, pages 264–273, 2018.
- F. Berkenkamp, M. Turchetta, A. P. Schoellig, and A. Krause. Safe model-based reinforcement learning with stability guarantees. In I. Guyon, U. von Luxburg, S. Bengio, H. M. Wallach, R. Fergus, S. V. N. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 908–918, 2017.
- Y. Chandak, S. M. Jordan, G. Theocharous, M. White, and P. S. Thomas. Towards safe policy improvement for non-stationary MDPs. *arXiv preprint arXiv:2010.12645*, 2020.
- C. Finlay, A. M. Oberman, and B. Abbasi. Improved robustness to adversarial examples using Lipschitz regularization of the loss, 2019. URL <https://openreview.net/forum?id=HkxAisC9FQ>.
- C. Gelada, S. Kumar, J. Buckman, O. Nachum, and M. G. Bellemare. DeepMDP: Learning continuous latent space models for representation learning. *CoRR*, abs/1906.02736, 2019.
- O. Gottesman, Y. Liu, S. Sussex, E. Brunskill, and F. Doshi-Velez. Combining parametric and nonparametric models for off-policy evaluation. In K. Chaudhuri and R. Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pages 2366–2375. PMLR, 2019.
- S. Kakade, M. J. Kearns, and J. Langford. Exploration in metric state spaces. In *Proceedings of the 20th International Conference on Machine Learning (ICML-03)*, pages 306–312, 2003.
- M. Kearns and S. Singh. Near-optimal reinforcement learning in polynomial time. *Machine Learning*, 49(2-3):209–232, 2002.
- R. Kleinberg, A. Slivkins, and E. Upfal. Multi-armed bandits in metric spaces. *CoRR*, abs/0809.4882, 2008.
- A. Krishnamurthy, J. Langford, A. Slivkins, and C. Zhang. Contextual bandits with continuous actions: Smoothing, zooming, and adapting. *J. Mach. Learn. Res.*, 21:137:1–137:45, 2020. URL <http://jmlr.org/papers/v21/19-650.html>.
- K. Lakshmanan, R. Ortner, and D. Ryabko. Improved regret bounds for undiscounted continuous reinforcement learning. In F. R. Bach and D. M. Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, volume 37 of *JMLR Workshop and Conference Proceedings*, pages 524–532. JMLR.org, 2015.
- G. Lancien and A. Dalet. Some properties of coarse lipschitz maps between banach spaces. *HAL*, 2017, 2017.
- E. Lecarpentier and E. Rachelson. Non-stationary Markov decision processes, a worst-case approach using model-based reinforcement learning. In H. M. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. B. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 7214–7223, 2019.
- E. Lecarpentier, D. Abel, K. Asadi, Y. Jinnai, E. Rachelson, and M. L. Littman. Lipschitz lifelong reinforcement learning. *arXiv preprint arXiv:2001.05411*, 2020.
- B. Neyshabur, R. Tomioka, and N. Srebro. Norm-based capacity control in neural networks. In *Conference on Learning Theory*, pages 1376–1401. PMLR, 2015.

C. Ni, L. F. Yang, and M. Wang. Learning to control in metric space with optimal regret. In *2019 57th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 726–733. IEEE, 2019.

R. Ortner and D. Ryabko. Online regret bounds for undiscounted continuous reinforcement learning. *CoRR*, abs/1302.2550, 2013.

I. Osband and B. Van Roy. Model-based reinforcement learning and the eluder dimension. *arXiv preprint arXiv:1406.1853*, 2014.

J. Pazis and R. Parr. PAC optimal exploration in continuous space Markov decision processes. In M. desJardins and M. L. Littman, editors, *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence, July 14-18, 2013, Bellevue, Washington, USA*. AAAI Press, 2013.

M. Pirotta, M. Restelli, and L. Bascetta. Policy gradient in Lipschitz Markov decision processes. *Mach. Learn.*, 100(2-3):255–283, 2015.

D. Shah and Q. Xie. Q-learning with nearest neighbors. In S. Bengio, H. M. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pages 3115–3125, 2018.

S. R. Sinclair, T. Wang, G. Jain, S. Banerjee, and C. L. Yu. Adaptive discretization for model-based reinforcement learning. In H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.

Z. Song and W. Sun. Efficient model-free reinforcement learning in metric spaces. *CoRR*, abs/1905.00475, 2019.

A. L. Strehl and M. L. Littman. An analysis of model-based interval estimation for Markov decision processes. *Journal of Computer and System Sciences*, 74(8):1309–1331, 2008.

R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, 2018.

Z. Tang, Y. Feng, N. Zhang, J. Peng, and Q. Liu. Off-policy interval estimation with Lipschitz value iteration. *Advances in Neural Information Processing Systems*, 33, 2020.

A. Touati, A. A. Taiga, and M. G. Bellemare. Zooming for efficient model-free reinforcement learning in metric spaces. *arXiv preprint arXiv:2003.04069*, 2020.

A Proof of Theorem 1

Theorem 1 Under Assumption 1, for any two points $x_1, x_2 \in \mathcal{X}$ such that $d(s_1, s_2) \leq d_{\min}$, the difference $\Delta Q^* \equiv |Q^*(x_1) - Q^*(x_2)|$ is bounded by

$$\Delta Q^* \leq \frac{1 - \gamma^k}{1 - \delta\gamma^k} \left(Q_{\max,(1,2)}^* - \frac{r_{\min}}{1 - \gamma} \right), \quad (14)$$

where $Q_{\max,(1,2)}^* \equiv \max(Q^*(x_1), Q^*(x_2))$.

Proof. Without loss of generality, assume $Q^*(x_1) > Q^*(x_2)$ and therefore $Q_{\max,(1,2)}^* = Q^*(x_1)$. To lower bound ΔQ^* , we need to find the lowest possible value of $Q^*(x_2)$. Because $d(s_1, s_2) \leq d_{\min}$, the agent can, with probability $1 - \delta$, reach s_1 from s_2 in at most k steps (Assumption 1), and from s_1 can follow the optimal policy.

Let n_k be the number of k -step attempts needed to reach s_1 . In the worst case, the agent will collect a (potentially negative) discounted reward of $\frac{1 - \gamma^{kn_k}}{1 - \gamma} r_{\min}$ along the way. Therefore, for a particular n_k , the discounted expected reward from x_2 is at least $Q_{n_k}^*(x_2) = \frac{1 - \gamma^{kn_k}}{1 - \gamma} r_{\min} + \gamma^{kn_k} Q^*(x_1)$. Part (ii) of Assumption 1 is used to make sure that even if the agent does not reach s_1 in a particular k -step attempt it remains within d_{\min} of it, and can therefore try again until it does (in Appendix B we relax this assumption). The probability of the agent requiring exactly n_k attempts is $\delta^{n_k-1}(1 - \delta)$, and therefore the value of $Q^*(x_2)$ is at least the expectation over n_k of $Q_{n_k}^*(x_2)$:

$$\begin{aligned} Q^*(x_2) &\geq \sum_{n_k=1}^{\infty} \delta^{n_k-1}(1 - \delta) Q_{n_k}^*(x_2) \\ &= \sum_{n_k=1}^{\infty} \delta^{n_k-1}(1 - \delta) \left(\frac{1 - \gamma^{kn_k}}{1 - \gamma} r_{\min} + \gamma^{kn_k} Q^*(x_1) \right) \\ &= \frac{(1 - \delta)r_{\min}}{1 - \gamma} \sum_{n_k=1}^{\infty} \delta^{n_k-1} \\ &\quad - \frac{(1 - \delta)r_{\min}}{1 - \gamma} \sum_{n_k=1}^{\infty} \delta^{n_k-1} \gamma^{kn_k} \\ &\quad + (1 - \delta)Q^*(x_1) \sum_{n_k=1}^{\infty} \delta^{n_k-1} \gamma^{kn_k} \\ &= \frac{(1 - \delta)r_{\min}}{(1 - \gamma)(1 - \delta)} - \frac{\gamma^k(1 - \delta)r_{\min}}{(1 - \gamma)(1 - \delta\gamma^k)} \\ &\quad + \frac{\gamma^k(1 - \delta)Q^*(x_1)}{(1 - \delta\gamma^k)} \\ &= \frac{(1 - \delta)r_{\min}}{(1 - \gamma)} \left(\frac{1}{1 - \delta} - \frac{\gamma^k}{1 - \delta\gamma^k} \right) \\ &\quad + \frac{\gamma^k(1 - \delta)Q^*(x_1)}{(1 - \delta\gamma^k)} \\ &= \frac{(1 - \delta)r_{\min}}{(1 - \gamma)} \frac{(1 - \delta\gamma^k - \gamma^k + \delta\gamma^k)}{(1 - \delta)(1 - \delta\gamma^k)} \\ &\quad + \frac{\gamma^k(1 - \delta)Q^*(x_1)}{(1 - \delta\gamma^k)} \\ &= \frac{r_{\min}(1 - \gamma^k)}{(1 - \gamma)(1 - \delta\gamma^k)} + \frac{\gamma^k(1 - \delta)Q^*(x_1)}{(1 - \delta\gamma^k)} \end{aligned} \quad (15)$$

Now to obtain the largest difference between the values of the states we simply subtract the result of Eq. 15 from $Q^*(x_1)$ to obtain:

$$\begin{aligned}
 Q^*(x_1) - Q^*(x_2) &\leq Q^*(x_1) - \frac{r_{\min}(1 - \gamma^k)}{(1 - \gamma)(1 - \delta\gamma^k)} \\
 &\quad - \frac{\gamma^k(1 - \delta)Q^*(x_1)}{(1 - \delta\gamma^k)} \\
 &= Q^*(x_1) \frac{(1 - \delta\gamma^k - \gamma^k + \delta\gamma^k)}{1 - \delta\gamma^k} \\
 &\quad - \frac{r_{\min}(1 - \gamma^k)}{(1 - \gamma)(1 - \delta\gamma^k)} \\
 &= \frac{1 - \gamma^k}{1 - \delta\gamma^k} \left(Q^*(x_1) - \frac{r_{\min}}{1 - \gamma} \right) \\
 &= \frac{1 - \gamma^k}{1 - \delta\gamma^k} \left(Q_{\max,(1,2)}^* - \frac{r_{\min}}{1 - \gamma} \right),
 \end{aligned} \tag{16}$$

where the last substitution is due to $Q^*(x_1) > Q^*(x_2)$. \square

B Modifications of Assumption 1

Below we present two variations to Assumption 1 which may relax it further and allows application to a wider range of problems. We first discuss a variation of the Assumption which allows for applications to continuous stochastic domains (Appendix B.1). We then present a variation which may alleviate the restrictiveness of part (ii) of Assumption 1, that with probability 1 the agent stays within d_{\min} of the state it is trying to reach (Appendix B.2).

B.1 Extension to stochastic continuous domains

Assumption 1 is sufficient to derive a bound on the difference of the Q^* function between two states within a given distance from each other for deterministic continuous domains or stochastic discrete domains when a suitable metric is provided (or stochastic continuous domains under discretization). However, Assumption 1 is not suitable for fully stochastic continuous domains. To see that consider an environment in which the action of the agent determines the next state of the agent up to some Gaussian noise, i.e. $s_{t+1} = s_t + a_t + \varepsilon_t$, with $\varepsilon_t \sim \mathcal{N}(0, \sigma^2)$. In this case the probability of reaching *any* specific state is zero, as the transition function is now a probability density rather than a probability mass function.

To alleviate this issue we provide the following alternative to Assumption 1:

Assumption 2. *For any two states $s, s' \in \mathcal{S}$ such that $d(s, s') \leq d_{\min}$, we have that (i) with probability $1 - \delta$, a ball with center s' and radius d_{ε} can be reached from s in at most k time-steps, and (ii) with probability 1, the agent can remain within a distance of less than d_{\min} from s' .*

While the assumption above can be applied to stochastic continuous domains, it is not enough to prove an analogous result to Theorem 1, because there is no way to relate the value at a state s' , and the values at states in a ball of radius d_{ε} around it. For that we must make an additional assumption:

Assumption 3. *There exists a real number $\Delta Q_{d_{\varepsilon}}^* > 0$, such that for any two points $x_1, x_2 \in \mathcal{X}$ satisfying $d(x_1, x_2) \leq d_{\varepsilon}$, we have that $\Delta Q_{d_{\varepsilon}}^* \leq |Q^*(x_1) - Q^*(x_2)|$.*

We will shortly discuss the limitations of Assumption 3, but we first note that Assumptions 3 and 2 lead to the following alternative to Theorem 1:

Theorem 4. *Under Assumptions 2 and 3, for any two points $x_1, x_2 \in \mathcal{X}$ such that $d(x_1, x_2) \leq d_{\min}$, the difference $\Delta Q^* \equiv |Q^*(x_1) - Q^*(x_2)|$ is bounded by*

$$\Delta Q^* \leq \frac{1 - \gamma^k}{1 - \delta\gamma^k} \left(Q_{\max,(1,2)}^* - \frac{r_{\min}}{1 - \gamma} \right) + \frac{\gamma^k(1 - \delta)}{1 - \delta\gamma^k} \Delta Q_{d_{\varepsilon}}^*, \tag{17}$$

where $Q_{\max,(1,2)}^* \equiv \max(Q^*(x_1), Q^*(x_2))$.

Proof. The proof is similar to the proof for Theorem 1, with the modification that once the agent reaches a ball with center s_1 and radius d_ε , its value is at least $Q^*(x_1) - \Delta Q_{d_\varepsilon}^*$ (rather than equaling $Q^*(x_1)$ when exactly reaching s_1). \square

Assumption 3 may seem very limiting as it appears to make a local smoothness assumption on the Q -function which is similar to the types of assumptions we criticize in this paper. While this is partly true and unfortunately necessary to obtain results for stochastic continuous domains, note that Assumption 3 is still significantly weaker than a Lipschitz-smoothness assumption, as it still allows for discontinuous jumps of the Q -function, and makes no assumptions on points farther than d_ε away. Generally d_ε is assumed to be small (on the order of the transition noise).

A different way to phrase such critique of Assumption 3 is to note that it makes an assumption to bound $|Q^*(x_1) - Q^*(x_2)|$ for points within a given distance from each other ($d(s_1, s_2) < d_\varepsilon$), to obtain a bound for $|Q^*(x_1) - Q^*(x_2)|$ for points at a different distance from each other ($d(s_1, s_2) < d_{\min}$). This, in fact, leads to demonstrating the value of using Assumption 3 to obtain Theorem 4. Note that the same method of proof used for Lemma 2 in Appendix F can be used to get a more straight forward bound on $|Q^*(x_1) - Q^*(x_2)|$ when $d(s_1, s_2) < d_{\min}$, namely:

$$|Q^*(x_1) - Q^*(x_2)| \leq \Delta Q_{d_\varepsilon}^* \left\lceil \frac{d_{\min}}{d_\varepsilon} \right\rceil. \quad (18)$$

In many practical settings, the length-scale over which actions affect the transition dynamics is much larger than the length-scale of the noise, i.e. $d_{\min} \gg d_\varepsilon$. Under such conditions, $\left\lceil \frac{d_{\min}}{d_\varepsilon} \right\rceil$ will be very large and the bound in Theorem 4 will be much tighter than the bound given by Eq. 18. Thus, Theorem 4 makes use of a local smoothness assumption on the length-scale of the noise to obtain a bound on the length-scale of the dynamics due to actions, which is significantly tighter than the bound which can be obtained without making use of Assumption 2.

B.2 Relaxing part (ii) of Assumption 1

We now provide an alternative to Assumption 1 which replaces the requirement that with probability 1 the agent stays within d_{\min} of the state it is trying to reach. The motivation for the original assumption is that when the agent is trying to reach a state s' , which it can reach within k steps with high probability, if it fails, it should at the very least not get further away from the state. Here we discuss an alternative assumption, which allows for the possibility of the agent unfortunately ending up further away from the state it is trying to reach. We will assume that if the agent is unable to reach state s' from state s within k steps, the state it will end up in, s'' , satisfies $d(s, s'') \leq d_{\max}$. We constrain the value of d_{\max} in Assumption 4, and in the proof of Theorem 5 justify this constraint by showing that if it is violated, the expected time for an agent to transition between two states that are less than d_{\min} away from each other could be infinite. Thus, the constraint on d_{\max} ensures our requirement of the metric to quantify how quickly an agent can transition between states.

Assumption 4. For any two states $s, s' \in \mathcal{S}$ such that $d(s, s') \leq d_{\min}$, we have that (i) with probability $1 - \delta$, s' can be reached from s in at most k time-steps, and (ii) if the agent does not reach s' within k time-steps, its state after k steps, s'' , satisfies $d(s, s'') \leq d_{\max} \equiv \frac{1-\delta}{\delta} d_{\min}$.

Note that δ is assumed to be small, and therefore $d_{\max} > d_{\min}$. Under Assumption 4, we can prove the following equivalent result to Theorem 1:

Theorem 5. Under Assumption 4, if $r_{\min} \leq 0$, for any two points $x_1, x_2 \in \mathcal{X}$ such that $d(s_1, s_2) \leq d_{\min}$, the difference $\Delta Q^* \equiv |Q^*(x_1) - Q^*(x_2)|$ is bounded by

$$\Delta Q^* \leq (1 - \gamma^{\tilde{k}}) \left(Q_{\max, (1,2)}^* - \frac{r_{\min}}{1 - \gamma} \right) \quad (19)$$

where $Q_{\max, (1,2)}^* \equiv \max(Q^*(x_1), Q^*(x_2))$ and $\tilde{k} \equiv \frac{k}{(1-\delta) - \left\lceil \frac{d_{\max}}{d_{\min}} \right\rceil \delta}$.

Proof. Proving Theorem 5 requires a different approach than Theorem 1. If we try to use the same logic, with probability $(1 - \delta)$ the agent reaches state s' within k steps. However, if that is not the case, the agent may end up as far away as $d_{\max} + d(s, s')$ away from s'' . To use the same method of proof, we can break $d(s, s'')$ into $\left\lceil \frac{d_{\max}}{d_{\min}} \right\rceil$ intervals of length d_{\min} , and assume the agent can transition between the edges of these intervals every k steps with probability $(1 - \delta)$. Thus,

the problem can be thought of as a 1D random walker which starts at a distance of one interval away from the goal, and with probability $(1 - \delta)$ moves one step toward the goal (for this random walker, all “steps” are in units of d_{\min} and each time-step is equivalent to k time-steps of the agent), and with probability δ moves (in the worst case) $D_{d_{\max}} \equiv \left\lceil \frac{d_{\max}}{d_{\min}} \right\rceil$ steps away from the goal. Denote the location of the random walker D , with the goal at $D = 0$ and the starting point at $D = 1$ (equivalent to at most d_{\min} away from the goal). The expected number of time steps to reach the goal from position D , denoted $\tau(D)$, obeys the following equations:

$$\tau(D) = (1 - \delta)(1 + \tau(D - 1)) + \delta(1 + \tau(D + D_{d_{\max}})) \quad (20)$$

$$\tau(0) = 0 \quad (21)$$

These equations can be solved by considering the mean field approximation of the random walker as moving in continuous space with velocity $u = (1 - \delta)(-1) + \delta D_{d_{\max}}$. For the agent to have a net velocity in the direction of the goal, we must require $u < 0$, which gives us the constraint on d_{\max} stated in Assumption 4:

$$\begin{aligned} (1 - \delta)(-1) + \delta D_{d_{\max}} &< 0 \Rightarrow D_{d_{\max}} < \frac{1 - \delta}{\delta} \\ &\Rightarrow \left\lceil \frac{d_{\max}}{d_{\min}} \right\rceil < \frac{1 - \delta}{\delta} \end{aligned} \quad (22)$$

which is satisfied when $d_{\max} < \frac{1 - \delta}{\delta} d_{\min}$. Intuitively this constraint ensures that even if there is a possibility of stochasticity transitioning the agent away rather than towards a state it is trying to reach, the probability of this happening and the amount by which the agent is distanced from its desired state are balanced in such a way that the agent can eventually reach its desired state. If that is not the case, the metric does not satisfy the notion we desire—quantifying how quickly states can be reached from other nearby states.

For the continuous mean-field dynamics, $\tau(D)$ is equivalent to the time required to traverse a distance of $(-D)$ and is given by

$$\tau(D) = \frac{D}{1 - \delta - \delta D_{d_{\max}}}. \quad (23)$$

Note that by substitution we can verify that Eq. 23 is indeed a solution to Eq. 20, and therefore the mean field solution is valid for the discrete dynamics as well. The expected number of time-steps in which an agent will reach s' from a state s such that $d(s, s') \leq d_{\min}$ can now be expressed by

$$\tilde{k} \equiv k\tau(1) = \frac{k}{1 - \delta - \delta D_{d_{\max}}} = \frac{k}{(1 - \delta) - \left\lceil \frac{d_{\max}}{d_{\min}} \right\rceil \delta} \quad (24)$$

The logic of the proof for Theorem 1 can be applied here—without loss of generality, assume $Q^*(x_1) > Q^*(x_2)$, and therefore the agent can transition from s_2 to s_1 in \tilde{k} steps, collecting a minimum reward of $r_{\min} \frac{1 - \gamma^{\tilde{k}}}{1 - \gamma}$, and from state s_1 follow the policy for the discounted Q-value of state action pairs in state s_1 of $\gamma^{\tilde{k}} Q^*(x_1)$. We thus obtain that for the expected number of steps to reach the goal, $Q_k^*(x_2) \geq r_{\min} \frac{1 - \gamma^{\tilde{k}}}{1 - \gamma} + \gamma^{\tilde{k}} Q^*(x_1)$. Unfortunately, this is not a true lower bound for $Q^*(x_2)$ because we used an expression lower bounding $Q^*(x_2)$ for a given number of steps to the goal and substituted the expectation of the number of steps, rather than take an expectation over the number of steps for the lower bound of $Q^*(x_2)$. However, for $r_{\min} \leq 0$, the expression $r_{\min} \frac{1 - \gamma^k}{1 - \gamma} + \gamma^k Q^*(x_1)$ is a convex function of k , and therefore by the Jensen inequality

$$Q^*(x_2) \geq Q_k^*(x_2) \geq r_{\min} \frac{1 - \gamma^{\tilde{k}}}{1 - \gamma} + \gamma^{\tilde{k}} Q^*(x_1). \quad (25)$$

Subtracting from $Q^*(x_1)$ we get

$$\begin{aligned}
 |Q^*(x_1) - Q^*(x_2)| &\geq Q^*(x_1) - \left(r_{\min} \frac{1 - \gamma^{\tilde{k}}}{1 - \gamma} + \gamma^{\tilde{k}} Q^*(x_1) \right) \\
 &= (1 - \gamma^{\tilde{k}}) \left(Q_{\max, (1,2)}^* - \frac{r_{\min}}{1 - \gamma} \right)
 \end{aligned} \tag{26}$$

□

C Functional forms and smoothness measures of illustrative functions

Sine-like function. The analytic form of the function illustrated in Figure 3 in the main text (left column) is

$$f(x) = A \sin(2\pi\omega x) + mx. \tag{27}$$

Specifically, in Figure 3, the parameter values used are $A = 3$, $\omega = 2$ and $m = 5$. The Lipschitz constant of the function is the largest slope, or the maximum of the first derivative of $f(x)$,

$$f'(x) = 2\pi\omega A \cos(2\pi\omega x) + m, \tag{28}$$

which is attained for all integer values of x :

$$L = 2\pi\omega A + m. \tag{29}$$

To compute $L_\alpha(\alpha)$ we write the slope between any two points x_1 and x_2 and look for its maximum when the distance between them is larger than α . In other words, we wish to maximize $M(x_1, x_2) = |f(x_1) - f(x_2)|/|x_1 - x_2|$ subject to $|x_1 - x_2| \geq \alpha$:

$$\begin{aligned}
 M(x_1, x_2) &= \\
 &= \left| \frac{A[\sin(2\pi\omega x_1) - \sin(2\pi\omega x_2)] + m(x_1 - x_2)}{x_1 - x_2} \right| \\
 &= \left| \frac{A[\sin(2\pi\omega x_1) - \sin(2\pi\omega x_2)]}{x_1 - x_2} + m \right| \\
 &= \left| \frac{2A[\sin(\pi\omega\{x_1 - x_2\}) \cos(\pi\omega\{x_1 + x_2\})]}{x_1 - x_2} + m \right| \\
 &= \left| \frac{2A\pi\omega[\sin(y_1) \cos(y_2)]}{y_1} + m \right|,
 \end{aligned} \tag{30}$$

where in the last equality we substituted $y_1 = \pi\omega\{x_1 - x_2\}$ and $y_2 = \pi\omega\{x_1 + x_2\}$. For simplicity we will assume $m > 0$, which implies the largest M obtained will be positive. This assumption allows us to remove the absolute value for easier differentiation.

To find the maximum of M we take it's partial derivatives:

$$\frac{\partial M}{\partial y_1} = 2\pi\omega \cos(y_2) \left[\frac{\cos(y_1)}{y_1} - \frac{\sin(y_1)}{y_1^2} \right] \quad (31)$$

$$\frac{\partial M}{\partial y_2} = -2\pi\omega \frac{\sin(y_1) \sin(y_1)}{y_1} \quad (32)$$

Equating both derivatives to zero provides two groups of solutions:

Group I:

$$\sin(y_1) = 0; \cos(y_2) = 0$$

Substituting the solutions for group I in Equation 30 yields $M = m$. Once we examine the solutions for Group II, we will see that for every cycle of the sine, we have a solution from Group II with $M > m$.

Group II:

$$y_1 = \tan y_1; \sin(y_2) = 0$$

Substituting the solutions for group II in Equation 30 yields

$$M = 2A\pi\omega \cos(y_2) \cos(y_1) + m. \quad (33)$$

Because $\sin(y_2) = 0$ implies $\cos(y_2) = \pm 1$, and therefore because we are interested in the maxima of M we choose the sign such that the first term is positive (this can always be done because y_1 and y_2 can be chosen independently). The maxima of M can then be found by substituting the solutions (which can be found numerically) for $y_1 = \tan y_1$ in

$$M = 2A\pi\omega |\cos(y_1)| + m. \quad (34)$$

For larger y_1 , the closest solution of $y_1 = \tan y_1$ is closer to $\pi/2 + n\pi$ for an integer n , and therefore the corresponding $\cos(y_1)$ is closer to zero. This means that

$$L_\alpha = 2A\pi\omega |\cos(y_*)| + m, \quad (35)$$

with y_* being the first solution of $y_1 = \tan y_1$ such that $y_1 \geq \pi\omega\alpha$ (which is equivalent to the constraint $x_1 - x_2 \geq \alpha$). In the example shown in Figure 3 in the main text, where we take $\alpha = \omega^{-1}$, we take the second solution of $y_1 = \tan y_1$ which is $y_* \approx 4.49$ and yields:

$$L_{\alpha=\omega} \approx 2A\pi\omega |\cos(4.49)| + m \approx 13.29. \quad (36)$$

Stairs function The stairs function used for demonstration is a stairs function with a step of magnitude 0.1 every interval of 0.1 in the x direction:

$$f(x) = A \left\lfloor \frac{x}{w} \right\rfloor \quad (37)$$

with $A = 0.1$ and $w = 0.1$. The function is obviously not continuous and therefore L does not exist (or is considered “infinite”).

To find L_α , we must find two points that maximize $M(x_1, x_2) = |f(x_1) - f(x_2)|/|x_1 - x_2|$ subject to $|x_1 - x_2| \geq \alpha$. To understand the possible candidates for such pairs of points, let us first consider that $[x_2, x_1]$ forms an interval that contains n

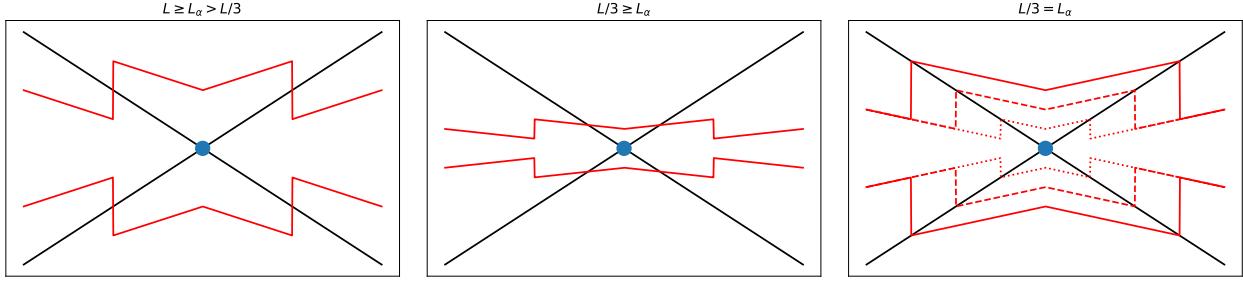


Figure 5: Schematic of how l_α , the maximum distance from a known point over which Lipschitz bounds are tighter than L_α bounds, as a function of the relation between L and L_α (see also Theorem 3).

steps (without loss of generality, assume $x_1 > x_2$, and denote $\Delta \equiv x_1 - x_2$). In such a case we have $f(x_1) - f(x_2) = nA$, and $M = nA/\Delta$. Given an interval which contains n steps, we would like to make Δ as small as possible in order to maximize M . This is achieved for $\Delta = (n-1)w + \epsilon$, with $\epsilon > 0$ as small as possible (to see how this is true, consider that in a width of slightly more than the width of an entire stair we can include two jumps). Then we have

$$M = \lim_{\epsilon \rightarrow 0} \frac{nA}{(n-1)w + \epsilon} = \frac{nA}{(n-1)w}. \quad (38)$$

It is clear from Equation 38 that M decreases with n , and so for a given α , one potential choice for x_1 and x_2 would be choosing them in such a way that they form an interval which contains $\lceil \alpha/w \rceil + 1$ steps, for which we have $M = \frac{A(\lceil \alpha/w \rceil + 1)}{\lceil \alpha/w \rceil w} = \frac{A}{w} \left(1 + \frac{1}{\lceil \alpha/w \rceil} \right)$.

The only other potential choice for the interval $[x_2, x_1]$ is to choose its width to be exactly α . In that case we can fit $\lceil \alpha/w \rceil$ steps in it, and obtain $M = \frac{A\lceil \alpha/w \rceil}{\alpha}$.

Thus, for the stairs function we have:

$$L_\alpha(\alpha) = \max \left[\frac{A}{w} \left(1 + \frac{1}{\lceil \alpha/w \rceil} \right), \frac{A}{\alpha} \lceil \alpha/w \rceil \right]. \quad (39)$$

D Proof of Theorem 3

Theorem 3 Let $f(x)$ be defined over the metric space \mathcal{X} , and assume the value of $f(x)$ is known at N points. Assume $f(x)$ is Lipschitz continuous with Lipschitz constant L . For any value of α and its corresponding L_α over $f(x)$, define l_α to be

$$l_\alpha = \begin{cases} \alpha & L \geq L_\alpha > L/3 \\ \frac{2L_\alpha\alpha}{L-L_\alpha} & L/3 \geq L_\alpha. \end{cases} \quad (40)$$

Then, if $\frac{C_D}{C_{D_\mathcal{X}}} \left(\frac{l_\alpha}{l_\mathcal{X}} \right)^D N < 1$, the fraction of volume of \mathcal{X} over which the L_α bound is tighter than the Lipschitz bound is at least

$$1 - \frac{C_D}{C_{D_\mathcal{X}}} \left(\frac{l_\alpha}{l_\mathcal{X}} \right)^D N. \quad (41)$$

Proof. We will first prove that, if a point $x \in \mathcal{X}$ satisfies $d(x, x') > l_\alpha$ for every x' at which $f(x')$ is known, then the L_α -bound is tighter than the Lipschitz bound. We prove the result for the upper bound, but the proof is similar for the lower bound.

Consider the case in which the value of $f(x)$ is known at only one point, x_0 . When $d(x_0, x) > \alpha$, the upper bounds $\hat{f}_{UB}^\alpha(x_0)$ and $\hat{f}_{UB}(x_0)$ are given by $d(x_0, x)L_\alpha$ and $d(x_0, x)L$, respectively, and because $L > L_\alpha$, we have that $\hat{f}_{UB}^\alpha(x_0) > \hat{f}_{UB}(x_0)$ when $d(x_0, x) > \alpha$, for any value of L_α .

When $L/3 \geq L_\alpha$, we can obtain an even smaller value for l_α by equating Equations 2 and 8 and solving for $d(x_0, x)$ to obtain $l_\alpha = \frac{2L_\alpha\alpha}{L-L_\alpha}$. Figure 5 in the appendix shows a schematic demonstrating the different cases for $L \geq L_\alpha > L/3$ and $L/3 \geq L_\alpha$. Note that the different cases are determined only by comparing L and L_α , but are not affected by α itself, as demonstrated in Figure 5 (right).

Now assume the value of $f(x)$ is known at N points x_n , such that for all points $d(x_n, x) > l_\alpha$. Denote $\hat{f}_{UB}^{\alpha, x_i}(x)$ and $\hat{f}_{UB}^{x_i}(x)$ as the upper bound we would have computed for $f(x)$ if only $f(x_i)$ were known. Because $\hat{f}_{UB}^{\alpha, x_i}(x) > \hat{f}_{UB}^{x_i}(x)$ for all i , we also have that $\min_i \hat{f}_{UB}^{\alpha, x_i}(x) > \min_i \hat{f}_{UB}^{x_i}(x)$ and therefore $\hat{f}_{UB}^\alpha(x) > \hat{f}_{UB}(x)$.

A similar proof can be provided for the lower bound, and the two bounds together imply that Lipschitz estimation may only provide tighter bounds in balls of radius l_α around known points.

The largest volume of the space over which Lipschitz estimation is tighter than L_α estimation is achieved if all balls of radius l_α centered at the N points where $f(x)$ is known, and that volume will be $NC_D l_\alpha^D$, and the fraction that volume is of \mathcal{X} is $\frac{V_{\mathcal{X}} - NC_D l_\alpha^D}{V_{\mathcal{X}}}$, and algebraic manipulation yields Equation 41. \square

E Proof of Theorem 2

As part of the proof, we will start with the following Lemma:

Lemma 1. *Let $f(x)$ be a function such that for all $x_i, x_j \in \mathcal{X}$, $d_{ij} \equiv d(x_i, x_j) \leq d_0$ implies $\Delta_{ij} = |f(x_i) - f(x_j)| \leq \Delta_{d_0}$. Further assume that \mathcal{X} is convex. Then for any two points $x, x' \in \mathcal{X}$, $d(x, x') \leq \epsilon$ implies*

$$|f(x_1) - f(x_2)| \leq \left\lceil \frac{\epsilon}{d_0} \right\rceil \Delta_{d_0}. \quad (42)$$

Proof of Lemma 2. Define $n_* \equiv \left\lceil \frac{\epsilon}{d_0} \right\rceil$. Let $x_n = x + n(x' - x)$ for $n = 0, 1, \dots, n_* - 1$. Because \mathcal{X} is convex, $x_n \in \mathcal{X}$ for all n . Because $d(x_n, x_{n+1}) = d_0$, by assumption, for all n we have that $f(x_{n+1}) \leq f(x_n) + \Delta_{d_0}$. Thus, since $x_{n=0} = x$, we can write $f(x_n) \leq f(x_n) + n\Delta_{d_0}$. Furthermore, because $d(x_{n*-1}, x') \leq d_{\min}$, we have one additional inequality: $f(x') \leq f(x_{n*-1}) + \Delta_{d_0} \leq f(x) + \Delta_{d_0} n_*$. A similar inequality can be derived for the lower bound of $f(x')$. \square

Proof of Theorem 2. Let $x_1, x_2 \in \mathcal{X}$ such that $d_{12} \equiv d(x_1, x_2) \geq \alpha$. We need to upper bound

$$L_\alpha \equiv \max_{\substack{x_1, x_2 \in \mathcal{X} \\ d_{12} \geq \alpha}} \frac{\Delta_{12}}{d_{12}}. \quad (43)$$

Because \mathcal{X}' is an ϵ -cover of \mathcal{X} , we can find $x_i, x_j \in \mathcal{X}'$ such that $d_{i1} \leq \epsilon$ and $d_{j2} \leq \epsilon$. Without loss of generality, assume $f(x_i) \geq f(x_j)$. By Lemma 1 we have $f(x_1) \leq f(x_i) + \Delta_\epsilon$ and $f(x_2) \geq f(x_j) - \Delta_\epsilon$, and therefore $\Delta_{12} \leq \Delta_{ij} + 2\Delta_\epsilon$. Furthermore, $d_{12} \geq d_{ij} - 2\epsilon$ and we are interested in x_1, x_2 such that $d_{12} \geq \alpha$, and we can therefore write

$$L_\alpha \equiv \max_{\substack{x_1, x_2 \in \mathcal{X} \\ d_{12} \geq \alpha}} \frac{\Delta_{12}}{d_{12}} \leq \max_{\substack{x_i, x_j \in \mathcal{X}' \\ d_{ij} \geq \alpha - 2\epsilon}} \frac{\Delta_{ij} + 2\Delta_\epsilon}{\max(\alpha, d_{ij} - 2\epsilon)}. \quad (44)$$

Taking the outer max over all x_i and x_j such that $d_{ij} \geq \alpha - 2\epsilon$ ensures that the procedure of picking an appropriate x_i, x_j pair is possible for all x_1, x_2 pairs satisfying $d_{12} \geq \alpha$. \square

Note that taking the max over α and $d_{ij} - 2\epsilon$ in the denominator is not necessary for the bound to be correct it significantly tightens the bound. If we always take α , we'd get very large ratios for $d_{ij} \gg \alpha$ where the numerator can be very large but the denominator is fixed. Conversely, if we always use $d_{ij} - 2\epsilon$ in the denominator, this could lead to very large ratios for d_{ij} which are only slightly larger than α . Also note that without taking the max in the denominator, we would have needed to deal with the special case of $d_{ij} - 2\epsilon \leq 0$.

F Proof of Proposition 2

We start with the following Lemma:

Lemma 2. For an MDP satisfying the assumptions of Prop. 2, for any $x, x' \in \mathcal{X}$,

$$|Q^*(x) - Q^*(x')| \leq \Delta Q_{\max}^* \left\lceil \frac{d(x, x')}{d_{\min}} \right\rceil. \quad (45)$$

Proof of Lemma 2. Substitute ΔQ_{\max}^* and d_{\min} from Corollary 1 into Lemma 1, and substitute ϵ from Lemma 1 with $d(x, x')$ (the relation $d(x, x') \leq \epsilon$ holds trivially for $d(x, x') = \epsilon$). \square

Proof of Proposition 2. From Equation 7, all we need to show is that for all $x, x' \in \mathcal{X}$ such that $d(x, x') \geq d_{\min}$, the inequality $\frac{|Q^*(x) - Q^*(x')|}{d(x, x')} \leq \frac{2\Delta Q_{\max}^*}{d_{\min}}$ holds. By Lemma 2:

$$\frac{|Q^*(x) - Q^*(x')|}{d(x, x')} \leq \Delta Q_{\max}^* \left\lceil \frac{d(x, x')}{d_{\min}} \right\rceil \frac{1}{d(x, x')}. \quad (46)$$

For $d(x, x') \geq d_{\min}$, the function $\left\lceil \frac{d(x, x')}{d_{\min}} \right\rceil \frac{1}{d(x, x')}$ obtains its supremum at $d(x, x') = \lim_{\epsilon \rightarrow 0} d_{\min} + \epsilon$, and the value of the supremum is $\frac{2}{d_{\min}}$. Substituting this value in Equation 46 completes the proof. \square

F.1 Illustrative example for the potential tightness of the bound in Proposition 2

In this section we give an example for which the bound in Proposition 2 becomes arbitrarily tight in the limit of $\gamma \rightarrow 1$, and provide intuition for other MDPs in which we can expect this bound to be tight. In the riverswim example from Section 4, if we set $c = 0$, then for $k = 1$ we have $d_{\min} = a_{\max}$. Set $r_{\text{right}} = 1$, $r_{\text{left}} = 0$ and $\delta = 0$. Then $Q_{\max} = 1$, and from Corollary 1 we have $\Delta Q_{\max} = (1 - \gamma)$. Consider the following two points (for any $\epsilon < a_{\max}$):

$$\begin{aligned} x_1 &= (1 - a_{\max} + \epsilon, a_{\max}) \\ x_2 &= (1 - 2a_{\max} - \epsilon, a_{\max}). \end{aligned}$$

x_1 is one step away from the right edge of the domain, just ϵ to the right of the first ‘‘step’’ of the Q-function, and therefore $Q(x_1) = 1$. Meanwhile, x_2 is three steps away from the right edge of the domain, ϵ to the left of the second ‘‘step’’ of the Q-function, and therefore $Q(x_2) = \gamma^2$. Note that these two results can also be obtained directly from Eq. 4.

The slope of the Q-function between x_1 and x_2 in the limit of $\epsilon \rightarrow 0$ is

$$\lim_{\epsilon \rightarrow 0} \frac{|Q(x_1) - Q(x_2)|}{|x_1 - x_2|} = \lim_{\epsilon \rightarrow 0} \frac{1 - \gamma^2}{a_{\max} + 2\epsilon} = \frac{1 - \gamma^2}{a_{\max}} = \frac{(1 - \gamma)(1 + \gamma)}{a_{\max}} = \frac{(1 + \gamma)\Delta Q_{\max}}{d_{\min}}$$

Because $|x_1 - x_2| > d_{\min}$, we have shown that $L_{d_{\min}} \geq \frac{(1 + \gamma)\Delta Q_{\max}}{d_{\min}}$ and as the discount factor approaches 1, the bound in Proposition 2 becomes arbitrarily tight. Note that even for small values of γ , the bound will be off by at most a factor of 2, and therefore captures the correct order of magnitude.

Note also that while in the example as presented we chose the values for c , r_{right} and r_{left} such that the math is simplified, the example would hold for any assignment of these values.

This example gives intuition as to the MDPs under which we can expect the bound to be tight. In goal based MDPs, for $k = 1$, when the stochasticity is low, the boundaries between the states from which the goal can be reached in one or two steps, and two or three steps, these boundaries will result in the sharp steps we observed to the Q-function. Therefore we can expect the bound to be tight in MDPs where the majority of the potential reward is gained upon reaching the goal, and will be tighter for low stochasticity and large discount factor.

G Proof of Proposition 3

Proposition 3. In the ZoomRL algorithm introduced in Touati et al. [2020], for any value of α such that $L_\alpha \leq L/2$, substituting L_α -bounds for L results in improved regret for at least the first $\frac{2}{3\alpha}$ episodes.

Proof. The ZoomRL algorithm works by partitioning the state-action space into balls of decreasing radii, and constructing value estimates for each ball. Actions are selected optimistically by using an Lipschitz-based upper bound on the Q values of state-actions in each ball, called the *index* of the ball. Therefore, to prove Prop. 3, we just need to prove that the indices obtained using L_α are all tighter than the ones obtained using L .

The proof will be divided into two parts. We will first show that for at least the first $\frac{2}{3\alpha}$ episodes, the radius of all balls is greater or equal to 3α (part I). In the second part we will show that this lower bound on the radii of all balls, together with the assumption that $L_\alpha \leq L/2$, implies that the L_α bounds in the expression for the indices of balls used by the ZoomRL algorithm results in tighter bounds (part II).

(I) For each step h , and a particular episode number k , let $\text{rad}(B_h^k)$ denote the radius of the corresponding ball. Each ball has a radius of half of its parent, and a new ball is created only when its parent has been visited at least $\frac{1}{\text{rad}(B_h^k)}$ times. When the learner creates new balls as quickly as possible, which happens when each ball is the only child of its parent, to generate a ball of size smaller than 3α we would create balls of radii $(\frac{1}{2}, \frac{1}{2}, \dots, \frac{1}{2^{n-1}}, \frac{1}{2^n})$, where n is chosen such that

$$\frac{1}{2^{n-1}} \geq 3\alpha > \frac{1}{2^n}. \quad (47)$$

The sum of episodes needed to generate such a ball is $K = 1 + 2 + 4 + 2^{n+1} = 2^n - 1 \approx 2^n$. Thus as long as $K < 2^n$, the smallest ball's radius is larger than 3α . By Equation 47, the condition $K < 2^n$ can be written as

$$K < 2^n \leq \frac{2}{3\alpha}.$$

(II) We now show that if $L_\alpha \leq L/2$ and $3\alpha \leq \text{rad}(B_h^k)$ for all balls, then substituting the L_α bounds in the expression for the indices of balls used by the ZoomRL algorithm results in tighter bounds.

The expression for the indices in the ZoomRL algorithm is

$$\begin{aligned} \text{index}(B) \equiv & L \cdot \text{rad}(B) + \\ & \min_{B' \mid \text{rad}(B') \geq \text{rad}(B)} \left[\hat{Q}(B') + L \cdot \text{dist}(B, B') \right] \end{aligned} \quad (48)$$

Where $\hat{Q}(B')$ denotes the value estimate for the center of a ball, and $\text{dist}(B, B')$ the distance between the centers of the balls. For simplicity, we drop the super- and sub- scripts from the B_h^k notation for the balls.

The index of a ball is an upper bound for the value of all state-action pairs contained within it. We wish to write an analogous bound obtained using L_α . We first address the alternative for the $L \cdot \text{rad}(B)$ term. We must make sure that the loosely bounded region at a distance less than α from the center of the ball cannot have violation of the bound. However, because $3\alpha \leq \text{rad}(B)$, we know that $L_\alpha \cdot \text{rad}(B) \geq 3\alpha L_\alpha$, and the $L_\alpha \cdot \text{rad}(B)$ properly upper bounds the difference between the upper bound on the Q-value at the center of ball B and the Q-value at all other points in the ball.

Similarly, to replace L with L_α in the $L \cdot \text{dist}(B, B')$ term, we must make sure that we do not get bound violations when $\text{dist}(B, B') < \alpha$. To do this we will split our analysis to two parts. First, we assume that $\text{dist}(B, B') \geq \alpha$, and then we can safely replace L with L_α . In that case, we can see that

$$\begin{aligned} \text{index}_{L_\alpha}(B) \equiv & L_\alpha \cdot \text{rad}(B) + \\ & \min_{B' \mid \text{rad}(B') \geq \text{rad}(B)} \left[\hat{Q}(B') + L_\alpha \cdot \text{dist}(B, B') \right] \end{aligned} \quad (49)$$

is a valid index (upper bound) for ball B . Then, because $L_\alpha \leq L$, the inequality

$$\begin{aligned} & L_\alpha \cdot \text{rad}(B) + [\hat{Q}(B') + L_\alpha \cdot \text{dist}(B, B')] \\ & \leq L \cdot \text{rad}(B) + [\hat{Q}(B') + L \cdot \text{dist}(B, B')] \end{aligned} \tag{50}$$

holds for all B' , proving that $\text{index}_{L_\alpha}(B) \leq \text{index}_L(B)$, or, in other words, the L_α index is tighter than the index obtained using L .

In the case that $\text{dist}(B, B') < \alpha$, we can upper bound the contribution to the bound of the $L_\alpha \cdot \text{dist}(B, B')$ by $3\alpha L_\alpha$. Using the upper bound to the term in the index, makes it a valid upper bound. Now all that remains is to demand that it is always tighter than the Lipschitz based bound. In other words, we wish to find parameters for which

$$\begin{aligned} & L_\alpha \cdot \text{rad}(B) + [\hat{Q}(B') + L_\alpha \cdot \text{dist}(B, B')] \\ & \leq L_\alpha \cdot \text{rad}(B) + [\hat{Q}(B') + 3\alpha L_\alpha] \\ & \leq L \cdot \text{rad}(B) + [\hat{Q}(B') + L \cdot \text{dist}(B, B')] \end{aligned} \tag{51}$$

for all B' . To do this we will ensure the inequality between the second and third lines hold even when we set $\text{dist}(B, B') = 0$. Eliminating $\hat{Q}(B')$ from both sides leaves us with showing that

$$\begin{aligned} & L_\alpha \cdot \text{rad}(B) + 3\alpha L_\alpha \leq L \cdot \text{rad}(B) \\ & \Rightarrow \frac{L_\alpha}{L} \leq \frac{\text{rad}(B)}{\text{rad}(B) + 3\alpha}. \end{aligned} \tag{52}$$

Because $\text{rad}(B) \geq 3\alpha$, we have that

$$\frac{1}{2} = \frac{3\alpha}{3\alpha + 3\alpha} \leq \frac{\text{rad}(B)}{\text{rad}(B) + 3\alpha} \tag{53}$$

And since we have assumed $L_\alpha \leq L/2$, this completes the proof. □