

---

# Leveraging Non-uniformity in First-order Non-convex Optimization

---

Jincheng Mei<sup>1,2\*</sup> Yue Gao<sup>1\*</sup> Bo Dai<sup>2</sup> Csaba Szepesvári<sup>3,1</sup> Dale Schuurmans<sup>2,1</sup>

## Abstract

Classical global convergence results for first-order methods rely on uniform smoothness and the Łojasiewicz inequality. Motivated by properties of objective functions that arise in machine learning, we propose a non-uniform refinement of these notions, leading to *Non-uniform Smoothness* (NS) and *Non-uniform Łojasiewicz inequality* (NL). The new definitions inspire new geometry-aware first-order methods that are able to converge to global optimality faster than the classical  $\Omega(1/t^2)$  lower bounds. To illustrate the power of these geometry-aware methods and their corresponding non-uniform analysis, we consider two important problems in machine learning: policy gradient optimization in reinforcement learning (PG), and generalized linear model training in supervised learning (GLM). For PG, we find that normalizing the gradient ascent method can accelerate convergence to  $O(e^{-c \cdot t})$  (where  $c > 0$ ) while incurring less overhead than existing algorithms. For GLM, we show that geometry-aware normalized gradient descent can also achieve a linear convergence rate, which significantly improves the best known results. We additionally show that the proposed geometry-aware gradient descent methods escape landscape plateaus faster than standard gradient descent. Experimental results are used to illustrate and complement the theoretical findings.

## 1. Introduction

The optimization of non-convex objective functions is a topic of key interest in modern-day machine learning. Recent, intriguing results show that simple gradient-based optimization can achieve *globally* optimal solutions in certain non-convex problems arising in machine learning, such as

---

\*Equal contribution <sup>1</sup>University of Alberta <sup>2</sup>Google Research, Brain Team <sup>3</sup>DeepMind. Correspondence to: Jincheng Mei <jmei2@ualberta.ca>, Yue Gao <gao12@ualberta.ca>.

in reinforcement learning (RL) (Agarwal et al., 2020), supervised learning (SL) (Hazan et al., 2015), and deep learning (Allen-Zhu et al., 2019). While gradient-based algorithms remain the method of choice in machine learning, the convergence of such algorithms to global minimizers has still only been established in restrictive settings where one can assert two strong assumptions about the objective function: (1) that the objective is smooth, and (2) that the objective satisfies a gradient dominance over sub-optimality such as the Łojasiewicz inequality. We will find it beneficial to recall the definitions of these properties. For the remainder of this paper let  $\Theta = \mathbb{R}^d$ .

**Definition 1** (Smoothness). *The function  $f : \Theta \rightarrow \mathbb{R}$  is  $\beta$ -smooth ( $\beta > 0$ ) if it is differentiable and for all  $\theta, \theta' \in \Theta$ ,*

$$\left| f(\theta') - f(\theta) - \left\langle \frac{df(\theta)}{d\theta}, \theta' - \theta \right\rangle \right| \leq \frac{\beta}{2} \cdot \|\theta' - \theta\|_2^2. \quad (1)$$

**Definition 2.** (Łojasiewicz, 1963; Polyak, 1963; Kurdyka, 1998) *The differentiable function  $f : \Theta \rightarrow \mathbb{R}$  satisfies the  $(C, \xi)$ -Łojasiewicz inequality if for all  $\theta \in \Theta$ ,*

$$\left\| \frac{df(\theta)}{d\theta} \right\|_2 \geq C \cdot \left( f(\theta) - \inf_{\theta \in \Theta} f(\theta) \right)^{1-\xi}, \quad (2)$$

where  $C > 0$  and  $\xi \in [0, 1]$ .

In particular, if an objective function  $f$  satisfies both assumptions, gradient-based optimization can be shown to converge to a global minimizer by noting first that uniform smoothness Eq. (1) ensures the gradient updates achieve monotonic improvement with an appropriate step size (i.e.,  $f(\theta_{t+1}) \leq f(\theta_t) - \frac{1}{2\beta} \cdot \|\nabla f(\theta_t)\|_2^2$ , if  $\theta_{t+1} \leftarrow \theta_t - \frac{1}{\beta} \cdot \nabla f(\theta_t)$ ), while the Łojasiewicz inequality Eq. (2) ensures the gradient does not vanish before a global minimizer is reached. Several global convergence results have recently been achieved in the machine learning literature by exploiting assumptions of this kind. For example, in reinforcement learning it has recently been shown that policy gradient (PG) methods converge to a globally optimal policy (Agarwal et al., 2020; Mei et al., 2020b); in supervised learning it has been shown that gradient descent (GD) methods converge to global minimizers of certain non-convex problems (Hazan et al., 2015); and in deep learning theory it has been shown that (stochastic) GD can converge to a global minimizer with an over-parameterized neural network (Allen-Zhu et al., 2019).

However, previous work that relies on the two *uniform* conditions in Definitions 1 and 2 assumes *universal constants*

$\beta$  and  $C$ , which ignores important problem structure and limits both the applicability of the results and the strength of the results that can be obtained.

In this paper, we expand the class of problems for which gradient-based optimization is globally convergent, develop novel gradient-based methods that better exploit local structure, and improve the convergence rate analysis. We achieve these results by first defining then investigating a new set of *non-uniform* smoothness and Łojasiewicz inequalities, which generalize the classical definitions and allow a refined characterization of the space of objectives. Given these refined notions, we then tailor novel gradient-based algorithms that improve previous methods for these new problem classes, and extend the analysis to exploit these new forms of non-uniformity, achieving significantly stronger convergence rates in many cases. Importantly, these improvements are achieved in non-convex optimization problems that arise in relevant machine learning problems.

The remainder of the paper is organized as follows. First, in Section 2 we illustrate how natural optimization problems, including those in machine learning, exhibit interesting *local* structure that cannot be adequately captured by the uniform smoothness and Łojasiewicz inequalities. Then, Section 3 introduces the the Non-uniform Smoothness (NS) property and the Non-uniform Łojasiewicz (NŁ) inequality, based on which Section 4 provides non-uniform analyses. Sections 5 and 6 then present new results for policy gradient and generalized linear models respectively. Section 7 concludes the paper and discusses some future directions. Due to space limits, we relegate most of the proofs to the appendix.

## 2. Motivation

To illustrate the significance of non-uniformity in machine learning problems, we consider examples motivated by recent theoretical (Zhang et al., 2019; Wilson et al., 2019; Mei et al., 2020b) and empirical studies (Cohen et al., 2021).

Regarding smoothness, it is clear that a uniform smoothness constant  $\beta$  cannot always adequately characterize an objective over its entire domain. For example, the convex function  $f : x \mapsto x^4$  cannot be informatively characterized by a uniform smoothness constant  $\beta$  because its Hessian  $f'' : x \mapsto 12 \cdot x^2$  has the property that  $f''(x) \rightarrow \infty$  as  $|x| \rightarrow \infty$ , and  $f''(x) \rightarrow 0$  as  $|x| \rightarrow 0$ . Varying smoothness of this kind has motivated the study of alternative definitions to explain, for example, the effectiveness of gradient clipping in training neural networks and normalization in optimization (Zhang et al., 2019; Wilson et al., 2019). Meanwhile Cohen et al. (2021) present neural network training results that cannot be well explained using the standard smoothness condition of Definition 1.

Regarding the Łojasiewicz inequality, a recent study of

policy gradient optimization in reinforcement learning has shown that, with the standard softmax parameterization, the expected return objective cannot satisfy *any* Łojasiewicz inequality with a universal constant  $C$  (Mei et al., 2020b), which removes the possibility of using Definition 2 to prove convergence. By introducing a *non-uniform* version of the Łojasiewicz inequality, the same authors were able to show a global convergence rate for the same problem.

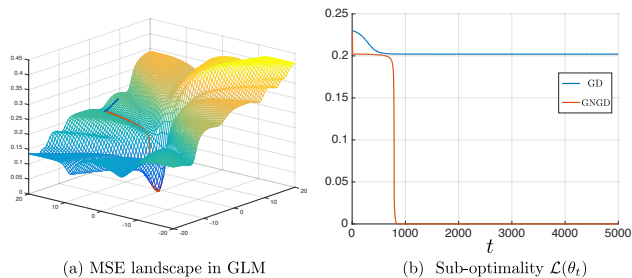


Figure 1. Non-uniform landscape of non-convex function.

Figure 1 illustrates another example of a non-convex objective, which arises in supervised learning. Subfigure 1(a) visualizes the mean squared error (MSE) of a generalized linear model (GLM) (Hazan et al., 2015), which is not only non-convex but also highly non-uniform. As a “teaser”, subfigure 1(b) compares the convergence behavior of two algorithms: standard gradient descent (GD), which suffers from slow convergence on the plateaus due to the non-uniformity of the objective, and an alternative algorithm (GNGD), soon to be introduced. This figure previews how proper handling of non-uniformity in the optimization landscape can enable significant acceleration of optimization progress, including a quick escape from plateaus.

## 3. Non-uniform Properties and Algorithms

The main results in this paper depend on two core concepts, Non-uniform Smoothness (NS) and Non-uniform Łojasiewicz (NŁ) inequality. The NS property is a new, intuitive generalization of smoothness. The NŁ inequality is a recent proposal of Mei et al. (2020b) and generalizes previous Łojasiewicz inequalities. Our key contribution is to show that the *combination* of these two non-uniform concepts is particularly powerful, applicable to important non-convex objectives in machine learning, and allows the development of improved algorithms and analysis.

### 3.1. Non-uniform Smoothness

The first main concept we leverage is a generalized notion of smoothness that depends on the parameters non-uniformly.

**Definition 3** (Non-uniform Smoothness (NS)). *The function  $f : \Theta \rightarrow \mathbb{R}$  satisfies  $\beta(\theta)$  non-uniform smoothness if  $f$  is*

differentiable and for all  $\theta, \theta' \in \Theta$ ,

$$\left| f(\theta') - f(\theta) - \left\langle \frac{df(\theta)}{d\theta}, \theta' - \theta \right\rangle \right| \leq \frac{\beta(\theta)}{2} \cdot \|\theta' - \theta\|_2^2,$$

where  $\beta$  is a positive valued function:  $\beta : \Theta \rightarrow (0, \infty)$ .

We will refer to  $\beta(\theta)$  in Definition 3 as the *NS coefficient*. This alternative definition generalizes and unifies several smoothness concepts from the recent literature. First, NS clearly reduces to Eq. (1) with  $\beta(\theta) = \beta$ . However, NS also generalizes the notion of  $(L_0, L_1)$  smoothness from Zhang et al. (2019) by using  $\beta(\theta) = L_0 + L_1 \cdot \|\nabla f(\theta)\|_2$ .

By using  $\beta(\theta) = c \cdot \|\nabla f(\theta)\|_2^{\frac{p-2}{p-1}}$ , NS also reduces to the notion of strong smoothness of order  $p$  proposed in Wilson et al. (2019). Finally, with  $\beta(\theta) = c/\|\theta\|_p^2$ , NS reduces to a special form of non-uniform smoothness considered in Mei et al. (2020a). We will show later that NS also covers other previously unstudied smoothness variants. Below we will demonstrate the key benefits of Definition 3 in terms of its *generality, better convergence results, and practical implications* in conjunction with the NŁ inequality.

### 3.2. Non-uniform Łojasiewicz Inequality

The second main concept we leverage is a generalized Łojasiewicz inequality introduced by Mei et al. (2020b):

**Definition 4** (Non-uniform Łojasiewicz (NŁ)). *The differentiable function  $f : \Theta \rightarrow \mathbb{R}$  satisfies the  $(C(\theta), \xi)$  non-uniform Łojasiewicz inequality if for all  $\theta \in \Theta$ ,*

$$\left\| \frac{df(\theta)}{d\theta} \right\|_2 \geq C(\theta) \cdot |f(\theta) - f(\theta^*)|^{1-\xi}, \quad (3)$$

where  $\xi \in (-\infty, 1]$ , and  $C(\theta) : \Theta \rightarrow \mathbb{R} > 0$  holds for all  $\theta \in \Theta$ . In this definition, either  $\theta^* = \arg \min_{\theta \in \Theta} f(\theta)$ , or  $f(\theta^*)$  is replaced with  $\inf_{\theta} f(\theta)$  if the global optimum is not achieved within the domain  $\Theta$ .

Definition 4 extends the classical “uniform” Łojasiewicz inequalities in optimization literature, such as the Polyak-Łojasiewicz (PŁ) inequality with  $C(\theta) = C > 0$  and  $\xi = 1/2$  (Łojasiewicz, 1963; Polyak, 1963); and the Kurdyka-Łojasiewicz (KŁ) inequality<sup>1</sup> by setting  $C(\theta) = C > 0$  (Kurdyka, 1998). Following (Mei et al., 2020b;a), we refer to  $\xi$  as the *NŁ degree* and  $C(\theta)$  as the *NŁ coefficient*. Generally speaking, a larger NŁ degree  $\xi$  and NŁ coefficient  $C(\theta)$  indicate faster convergence for gradient based algorithms. Appendix E provides an overview of remarkable non-convex functions that satisfy the NŁ inequality for various  $\xi$  and  $C(\theta)$ . As stated, our main contribution in this paper is to show how, when *combined* with NS, NŁ becomes a powerful tool for both algorithm design and analysis, which is a novel direction of investigation.

<sup>1</sup>The KŁ inequality is violated at bad local optima, since vanishing gradient norm cannot dominate non-zero sub-optimality gap. Therefore Definition 4 actually recovers global KŁ inequality.

### 3.3. Geometry-aware Gradient Descent

A key benefit of the non-uniform definitions is that we can introduce stepsize rules that make gradient descent adapt to the local “geometry” of the optimization objective. First consider the classical gradient decent update.

**Definition 5** (Gradient Descent (GD)).

$$\theta_{t+1} \leftarrow \theta_t - \eta \cdot \nabla f(\theta_t). \quad (4)$$

The key challenge with deploying GD is choosing the step size  $\eta$ ; if  $\eta$  is too large, instability ensues, if too small, progress becomes slow. Recall from the introduction that  $\eta = 1/\beta$  is a canonical choice for assuring convergence in *uniformly*  $\beta$  smooth objectives. This suggests that in the presence of *non-uniform* smoothness  $\beta(\theta)$  given in NS, the stepsize should be adapted to  $1/\beta(\theta)$ . This leads to a new variant of normalized gradient descent.

**Definition 6** (Geometry-aware Normalized GD (GNGD)).

$$\theta_{t+1} \leftarrow \theta_t - \eta \cdot \frac{\nabla f(\theta_t)}{\beta(\theta_t)}. \quad (5)$$

Key to making this approach practical will be efficient ways to measure (or bound)  $\beta(\theta)$ . Below we will show how in the context of NS and NŁ properties, GNGD can be made both practical and extremely efficient at solving various global optimization problems in machine learning. These results also broaden our fundamental knowledge of the set of objectives that admit efficient global optimization.

## 4. Non-uniform Analysis

Our first main contribution is an analysis for GD and GNGD based in the presence of non-uniform properties. For minimization problems, we assume  $\inf_{\theta} f(\theta) > -\infty$  ( $\sup_{\theta} f(\theta) < \infty$  for maximization problems).

**Theorem 1.** *Suppose  $f : \Theta \rightarrow \mathbb{R}$  satisfies NS with  $\beta(\theta)$  and the NŁ inequality with  $(C(\theta), \xi)$ . Suppose  $C := \inf_{t \geq 1} C(\theta_t) > 0$  for GD and GNGD. Let  $\delta(\theta) := f(\theta) - f(\theta^*)$  be the sub-optimality gap. The following hold:*

- (1a) *if  $\beta(\theta) \leq c \cdot \delta(\theta)^{1-2\xi}$  with  $\xi \in (-\infty, 1/2)$ , then the conclusions of (1b) hold;*
- (1b) *if  $\beta(\theta) \leq c \cdot \|\nabla f(\theta)\|_2^{\frac{1-2\xi}{1-\xi}}$  with  $\xi \in (-\infty, 1/2)$ , then GD with  $\eta \in O(1)$  achieves  $\delta(\theta_t) \in O(1/t^{\frac{1}{1-2\xi}})$ , and GNGD achieves  $\delta(\theta_t) \in O(e^{-t})$ .*
- (2a) *if  $\beta(\theta) \leq L_0 + L_1 \cdot \|\nabla f(\theta)\|_2$ , then the conclusions of (2b) hold;*
- (2b) *if  $\beta(\theta) \leq L_0 \cdot \frac{\|\nabla f(\theta)\|_2^2}{\delta(\theta)^{2-2\xi}} + L_1 \cdot \|\nabla f(\theta)\|_2$ , then GD and GNGD both achieve  $\delta(\theta_t) \in O(1/t^{\frac{1}{1-2\xi}})$  when  $\xi \in (-\infty, 1/2)$ , and  $O(e^{-t})$  when  $\xi = 1/2$ . GNGD has strictly better constant than GD ( $1 > C > C^2$ ).*
- (3a) *if  $\beta(\theta) \leq c \cdot \|\nabla f(\theta)\|_2^{\frac{1-2\xi}{1-\xi}}$  with  $\xi \in (1/2, 1)$ , then the conclusions of (3b) hold;*




Figure 2. Different function classes for  $\beta(\theta^*) < \infty$ . We use a label notation where, e.g., C denotes the set of all functions that satisfy property C, and  $ACE := A \cap C \cap E$ . The two largest ellipsoids correspond to  $A \cup B$  and C. We study the following four non-convex function classes in  $(A \cup B) \cap C$ , i.e.,  $\mathcal{W} := AC \setminus (AD \cup ACE)$ ,  $\mathcal{X} := AD \setminus ADE$ ,  $\mathcal{Y} := BC \setminus (BD \cup BCE)$ , and  $\mathcal{Z} := BD \setminus (BDE \cup BF)$ .

(3b) if  $\beta(\theta) \leq c \cdot \delta(\theta)^{1-2\xi}$  with  $\xi \in (1/2, 1)$ , then GD with  $\eta \in \Theta(1)$  does not converge, while GNGD achieves  $\delta(\theta_t) \in O(e^{-t})$ .

**Remark 1.** The cases (1)-(3) cover all three possibilities of  $\beta(\theta^*)$ . Since  $\theta^*$  is the global minimum,  $\nabla^2 f(\theta^*)$  is positive semi-definite (negative if  $\theta^*$  is maximum) if it exists.

- (1) If  $\nabla^2 f(\theta^*) = \mathbf{0}$ , then  $\beta(\theta) \rightarrow 0$  as  $\theta, \theta' \rightarrow \theta^*$ , which means the landscape around  $\theta^*$  is flat;
- (2) If  $\nabla^2 f(\theta^*)$  has at least one strictly positive (negative) eigenvalue, then  $\beta(\theta) \rightarrow \beta > 0$  as  $\theta, \theta' \rightarrow \theta^*$ .

The cases (1)-(2) also cover the situations where the Hessian  $\nabla^2 f(\theta^*)$  does not exist but one can find a finite  $\beta(\theta^*) > 0$  to upper bound the l.h.s. of Definition 3.

- (3) The case (3) is for blow-up type non-existence of  $\nabla^2 f(\theta^*)$ , where  $\beta(\theta^*)$  is unbounded.

**Remark 2.** In Theorem 1,  $C := \inf_{t \geq 1} C(\theta_t)$  is related to the early optimization and plateau escaping behavior studied in Mei et al. (2020a). It remains open to study whether GNGD can be combined with different parameterizations in Mei et al. (2020a) to further improve C.

Note that (1b) recovers the strong smoothness of order  $p$  with  $p = 1/\xi$  in Wilson et al. (2019), and (2a) recovers the  $(L_0, L_1)$  smoothness of Zhang et al. (2019). The results here consider more general NL functions and establish faster rates of convergence. The other cases have not been studied in literature to our knowledge. In Sections 5 and 6 below we study practical machine learning examples that are covered by cases (1) and (2) in Theorem 1. Other cases of different  $\beta(\theta)$  and  $\xi$  are discussed in Appendix A for completeness.

#### 4.1. Function Classes and Existing Lower Bounds

Before applying these results to problems in machine learning, we first provide a refined characterization of function classes organized by their NS and NL properties. This also clarifies the relation between the non-uniform properties and standard notions of convexity and smoothness; see Figure 2.

**Proposition 1.** The following hold for an objective  $f$ :

- (1)  $D \subseteq C$ . If  $f$  satisfies NL with degree  $\xi$ , it satisfies NL

with degree  $\xi' < \xi$ ;

- (2)  $F \subseteq D$ . A strongly convex  $f$  satisfies NL with  $\xi = 1/2$ ;
- (3)  $F \cap A = \emptyset$ . A strongly convex  $f$  cannot satisfy NS with  $\beta(\theta) \rightarrow 0$  as  $\theta, \theta' \rightarrow \theta^*$ ;
- (4)  $E \subseteq C$ . A (not strongly) convex  $f$  satisfies NL with  $\xi = 0$ .

The next proposition provides concrete examples for each convex function class in  $(A \cup B) \cap C$  in Figure 2.

**Proposition 2.** The following results hold: (1)  $ACE \neq \emptyset$ . (2)  $ADE \neq \emptyset$ . (3)  $BCE \neq \emptyset$ . (4)  $BDE \neq \emptyset$ . (5)  $BF \neq \emptyset$ .

A more interesting result considers examples in the classes of non-convex functions  $(A \cup B) \cap C$  in Figure 2. The non-uniform analysis above largely still applies to these problems, even when standard convex analysis cannot apply.

**Proposition 3.** The following results hold: (1)  $\mathcal{W} := AC \setminus (AD \cup ACE) \neq \emptyset$ . (2)  $\mathcal{X} := AD \setminus ADE \neq \emptyset$ . (3)  $\mathcal{Y} := BC \setminus (BD \cup BCE) \neq \emptyset$ . (4)  $\mathcal{Z} := BD \setminus (BDE \cup BF) \neq \emptyset$ .

We next apply the techniques to a class of convex functions, achieving results that cannot be explained by classical convex-smooth analysis.

**Proposition 4.** The convex function  $f : x \mapsto |x|^p$  with  $p > 1$  satisfies the NL inequality with  $\xi = 1/p$  and the NS property with  $\beta(x) \leq c_1 \cdot \delta(x)^{1-2\xi}$ .

Consider any  $p > 2$ , such as  $p = 4$ , where it follows that  $f$  satisfies NL with degree  $\xi = 1/4 < 1/2$ . According to (1a) in Theorem 1, GD will achieve  $\delta(x_t) \in \Theta(1/t^2)$ , while GNGD attains  $\delta(x_t) \in O(e^{-c \cdot t})$  (where  $c > 0$ ). Note that standard convex analysis can only give a  $O(1/t)$  rate on (not strongly) convex smooth functions. The  $\Theta(1/t^2)$  rate for GD here follows from using NL degree  $\xi = 1/4$ , which improves on  $\xi = 0$  from mere convexity ((4) in Proposition 1). The  $O(e^{-c \cdot t})$  rate has also been observed for this example by exploiting strong smoothness ((1b), as noted) (Wilson et al., 2019). Figure 2 provides a more general understanding of when this happens.

**$\Omega(1/t^2)$  lower bound for convexity-smoothness.** Note that GNGD satisfies  $x_{t+1} = x_t - \sum_{i=1}^t \frac{\eta}{\beta(x_i)} \cdot \nabla f(x_i) \in \text{Span}\{x_1, \nabla f(x_1), \dots, \nabla f(x_t)\}$ , which is a first-order ora-

cle (Nesterov, 2003). Thus there exists a worst-case objective in the convex-smooth class that forces  $\delta(x_t) \in \Omega(1/t^2)$  for  $t \in O(n)$ , where  $n$  is the parameter dimension (Nemirovski & Yudin, 1983; Nesterov, 2003; Bubeck, 2014). This is not a contradiction, since the lower bound is established by constructing a convex smooth function with a *constant*  $\beta > 0$  (Bubeck, 2014), and  $\beta(x) \rightarrow \beta > 0$  as  $x, x' \rightarrow x^*$  in Definition 3. Hence, the  $\Omega(1/t^2)$  result covers *some* functions in BCE in Figure 2. Meanwhile  $f : x \mapsto |x|^p$  with  $p > 2$  satisfies  $\beta(x) \rightarrow 0$  as  $x, x' \rightarrow 0$  in Definition 3 (ACE in Figure 2), which implies that the standard convex-smooth class consists of two subclasses. One subclass (BCE) admits first-order sub-linear lower bounds, while the other (ACE) allows linear convergence using first-order methods. This illustrates the *necessity* of non-uniformity in subdividing the NS class as  $A \cup B$  in Figure 2. This partition also inspires geometry-aware GD.

$\Omega(1/\sqrt{t})$  **lower bound for  $(L_0, L_1)$ -smoothness.** For  $\beta(\theta) = L_0 + L_1 \cdot \|\nabla f(\theta)\|_2$  ((2a) in Theorem 1) with  $L_0, L_1 \geq 1$ , standard normalized GD is subject to a  $\Omega(1/\sqrt{t})$  lower bound (Zhang et al., 2019). However, in Section 5, we will show that normalized policy gradient (PG) method achieves a linear rate of  $O(e^{-c \cdot t})$ . Again, this is not a contradiction for similar reasons. With  $L_0 \geq 1$ ,  $\beta(\theta) \rightarrow L_0 > 0$  as  $\theta, \theta' \rightarrow \theta^*$ , the  $\Omega(1/\sqrt{t})$  lower bound will hold for *some* functions in  $BCE \cup \mathcal{Y}$  in Figure 2. While in Section 5 the objective satisfies  $L_0 = 0$  and  $L_1 > 0$ , hence  $\beta(\theta) \rightarrow 0$  as  $\theta, \theta' \rightarrow \theta^*$  (ACE  $\cup \mathcal{W}$  in Figure 2). This shows a similar separation of rates for first-order methods will also occur based on NS conditions. Furthermore, in Section 6, we will show that both GD and GNGD achieve a  $O(e^{-c \cdot t})$  rate for GLM, but here the objective is in  $\mathcal{Z}$  in Figure 2 so the lower bounds do not apply.

**Unbounded Hessian.** Consider any  $p \in (1, 2)$ , such as  $p = 3/2$  where  $f$  satisfies  $\xi = 2/3$ . According to Theorem 1(3a), GD diverges since the Hessian is unbounded near 0. This makes it *necessary* to introduce geometry-aware normalization to ensure convergence, which is verified in Appendix F. This has practical implications for RL, for example ensuring exploration using state distribution entropy, which has unbounded Hessian near probability simplex boundary (Hazan et al., 2019) and alternative probability transforms (Mei et al., 2020a).

## 5. Policy Gradient

Our second main contribution is to show that the expected return objective considered in direct policy optimization in RL falls under the function class  $\mathcal{W}$  in Figure 2, in particular satisfying case (1) of Theorem 1 with NŁ degree  $\xi = 0$ . The key point is that value functions in Markov decision processes (MDPs) satisfy NS properties with coefficient

being the PG norm (Lemmas 2 and 6). This novel finding not only provides a much more precise characterization than existing standard smoothness results (Agarwal et al., 2020; Mei et al., 2020b), but also enables PG with normalization to use the NŁ inequalities (Lemmas 1 and 5) differently than for standard PG ( $\|\nabla f(\theta)\|_2$  vs.  $\|\nabla f(\theta)\|_2^2$ ), which leads to faster convergence as well as plateau escaping.

### 5.1. RL Settings and Notations

For a finite set  $\mathcal{N}$ , let  $\Delta(\mathcal{N})$  denote the set of all probability distributions on  $\mathcal{N}$ . A finite MDP  $\mathcal{M} := (\mathcal{S}, \mathcal{A}, \mathcal{P}, r, \gamma)$  is determined by a finite state space  $\mathcal{S}$ , a finite action space  $\mathcal{A}$ , a transition function  $\mathcal{P} : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$ , a scalar reward function  $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ , and a discount factor  $\gamma \in [0, 1)$ .

In policy-based RL, an agent interacts with the environment, i.e., the MDP  $\mathcal{M}$ , using a policy  $\pi : \mathcal{S} \rightarrow \Delta(\mathcal{A})$ . Given a state  $s_t$ , the agent takes an action  $a_t \sim \pi(\cdot|s_t)$ , receives a one-step scalar reward  $r(s_t, a_t)$  and a next-state  $s_{t+1} \sim \mathcal{P}(\cdot|s_t, a_t)$ . The long-term expected reward, also known as the value function of  $\pi$  under  $s$ , is defined as

$$V^\pi(s) := \mathbb{E}_{\substack{s_0=s, a_t \sim \pi(\cdot|s_t), \\ s_{t+1} \sim \mathcal{P}(\cdot|s_t, a_t)}} \left[ \sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) \right]. \quad (6)$$

The state-action value of  $\pi$  at  $(s, a) \in \mathcal{S} \times \mathcal{A}$  is defined as  $Q^\pi(s, a) := r(s, a) + \gamma \sum_{s'} \mathcal{P}(s'|s, a) V^\pi(s')$ , and  $A^\pi(s, a) := Q^\pi(s, a) - V^\pi(s)$  is the advantage function of  $\pi$ . The state distribution of  $\pi$  is defined as,

$$d_{s_0}^\pi(s) := (1 - \gamma) \sum_{t=0}^{\infty} \gamma^t \Pr(s_t = s | s_0, \pi, \mathcal{P}). \quad (7)$$

Given an initial state distribution  $\rho \in \Delta(\mathcal{S})$ , we denote  $V^\pi(\rho) := \mathbb{E}_{s \sim \rho} [V^\pi(s)]$  and  $d_\rho^\pi(s) := \mathbb{E}_{s_0 \sim \rho} [d_{s_0}^\pi(s)]$ . There exists an optimal policy  $\pi^*$  such that  $V^{\pi^*}(\rho) = \sup_{\pi: \mathcal{S} \rightarrow \Delta(\mathcal{A})} V^\pi(\rho)$ . For convenience, we denote  $V^* := V^{\pi^*}$ . Consider a tabular representation, i.e.,  $\theta(s, a) \in \mathbb{R}$  for all  $(s, a)$ , so that the policy  $\pi_\theta$  can be parameterized by  $\theta$  as  $\pi_\theta(\cdot|s) = \text{softmax}(\theta(s, \cdot))$ ; that is, for all  $(s, a)$ ,

$$\pi_\theta(a|s) = \frac{\exp\{\theta(s, a)\}}{\sum_{a' \in \mathcal{A}} \exp\{\theta(s, a')\}}. \quad (8)$$

When there is only one state the policy  $\pi_\theta = \text{softmax}(\theta)$  is defined as  $\pi_\theta(a) = \exp\{\theta(a)\} / \sum_{a' \in \mathcal{A}} \exp\{\theta(a')\}$ . The problem of policy-based RL is then to find a policy  $\pi_\theta$  that maximizes the value function, i.e.,

$$\sup_{\theta: \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}} V^{\pi_\theta}(\rho). \quad (9)$$

For convenience, and without loss of generality, we assume  $r(s, a) \in [0, 1]$  for all  $(s, a) \in \mathcal{S} \times \mathcal{A}$ .

### 5.2. One-state MDPs

We first illustrate some key insights for one-state MDPs with  $K$  actions and  $\gamma = 0$ . The value function Eq. (6)



reduces to expected reward  $\pi_\theta^\top r$ , where  $r \in [0, 1]^K$ ,  $\theta \in \mathbb{R}^K$ , and  $\pi_\theta = \text{softmax}(\theta)$ . Mei et al. (2020b) have shown that even though  $\max_\theta \pi_\theta^\top r$  is a non-concave maximization, global convergence can be achieved with a  $O(1/t)$  rate using uniform smoothness and the NŁ inequality:

**Lemma 1** (NŁ, Mei et al. (2020b), Lemma 3). *Let  $a^*$  be the optimal action. Denote  $\pi^* = \arg \max_{\pi \in \Delta} \pi^\top r$ . Then,*

$$\left\| \frac{d\pi_\theta^\top r}{d\theta} \right\|_2 \geq \pi_\theta(a^*) \cdot (\pi^* - \pi_\theta)^\top r. \quad (10)$$

Note that Lemma 1 is not improvable in terms of the coefficients  $C(\theta) = \pi_\theta(a^*)$  and  $\xi = 0$  (Mei et al., 2020b, Remark 1 and Lemma 17). However, this result is based on only using a *uniform* smoothness coefficient  $\beta = 5/2$  (Mei et al., 2020b, Lemma 2), which even empirical evidence suggests can be significantly refined. To illustrate, we run standard policy gradient (PG) on a 3-action one-state MDP. As shown in Figure 3(a), PG first goes through a long suboptimal plateau, and then eventually escapes to approach  $\pi^*$ . Figure 3(b) presents the spectral radius of the Hessian and the PG norm  $3 \cdot \left\| \frac{d\pi_{\theta_t}^\top r}{d\theta_t} \right\|_2$  as functions of time  $t$ . It is evident that the smoothness behaves non-uniformly: it is close to zero at the suboptimal plateau and near  $\pi^*$ , highly aligned with the PG norm. Compared to any universal constant  $\beta$ , the PG norm characterizes the non-uniform landscape information far more precisely. We formalize this observation by proving the following key result:

**Lemma 2** (NS). *Denote  $\theta_\zeta := \theta + \zeta \cdot (\theta' - \theta)$  with some  $\zeta \in [0, 1]$ . For any  $r \in [0, 1]^K$ ,  $\theta \mapsto \pi_\theta^\top r$  satisfies  $\beta(\theta_\zeta)$  non-uniform smoothness with  $\beta(\theta_\zeta) = 3 \cdot \left\| \frac{d\pi_{\theta_\zeta}^\top r}{d\theta_\zeta} \right\|_2$ .*

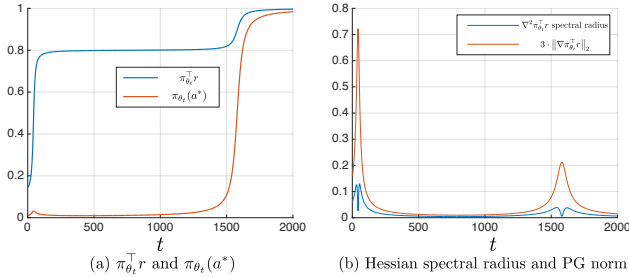


Figure 3. PG results on  $r = (1.0, 0.8, 0.1)^\top$ .

Comparing Lemma 2 with (1b) in Theorem 1, we have  $\xi = 0$ , and GNGD requires normalizing  $\beta(\theta_\zeta)$ , which is the PG norm of  $\theta_\zeta$  rather than  $\theta$ . However,  $\zeta$  is unknown. Fortunately, the next lemma shows that, if we still normalize the PG norm of  $\theta$ , the  $\beta(\theta_\zeta)$  in Lemma 2 can be upper bounded by  $\left\| \frac{d\pi_\theta^\top r}{d\theta} \right\|_2$ , given the learning rate is small enough:

**Lemma 3.** *Let  $\theta' = \theta + \eta \cdot \frac{d\pi_\theta^\top r}{d\theta} / \left\| \frac{d\pi_\theta^\top r}{d\theta} \right\|_2$ . Denote  $\theta_\zeta := \theta + \zeta \cdot (\theta' - \theta)$  with some  $\zeta \in [0, 1]$ . We have, for all*

$\eta \in (0, 1/3)$ ,

$$\left\| \frac{d\pi_{\theta_\zeta}^\top r}{d\theta_\zeta} \right\|_2 \leq \frac{1}{1 - 3\eta} \cdot \left\| \frac{d\pi_\theta^\top r}{d\theta} \right\|_2. \quad (11)$$

Next, the NŁ coefficient  $\pi_\theta(a^*)$  is bounded away from 0, which provides constants in the convergence rate results.

**Lemma 4** (Non-vanishing NŁ coefficient). *Using normalized policy gradient method, we have  $\inf_{t \geq 1} \pi_{\theta_t}(a^*) > 0$ .*

To this point, we demonstrate that the non-concave function  $\pi_\theta^\top r$  satisfies (1b) in Theorem 1 with  $\xi = 0$  in each iteration of normalized PG<sup>2</sup>: Lemmas 2 and 3 show that the NS coefficient  $\beta(\theta_{\zeta_t}) \leq c_1 \cdot \left\| \frac{d\pi_{\theta_t}^\top r}{d\theta_t} \right\|_2$ , while Lemmas 1 and 4 guarantee  $\left\| \frac{d\pi_{\theta_t}^\top r}{d\theta_t} \right\|_2 \geq c_2 \cdot (\pi^* - \pi_{\theta_t})^\top r$ . Therefore, combining Lemmas 1 to 4, we prove the global linear convergence rate  $O(e^{-c \cdot t})$  of normalized PG:

**Theorem 2.** *Using normalized PG  $\theta_{t+1} = \theta_t + \eta \cdot \frac{d\pi_{\theta_t}^\top r}{d\theta_t} / \left\| \frac{d\pi_{\theta_t}^\top r}{d\theta_t} \right\|_2$  with  $\eta = 1/6$ , for all  $t \geq 1$ , we have,*

$$(\pi^* - \pi_{\theta_t})^\top r \leq e^{-\frac{c \cdot (t-1)}{12}} \cdot (\pi^* - \pi_{\theta_1})^\top r, \quad (12)$$

where  $c = \inf_{t \geq 1} \pi_{\theta_t}(a^*) > 0$  is from Lemma 4, and  $c$  is a constant that depends on  $r$  and  $\theta_1$ , but not on the time  $t$ .

**Remark 3.** *If  $\pi_{\theta_1}$  is uniform, i.e.,  $\pi_{\theta_1}(a) = 1/K$ ,  $\forall a \in [K]$ , then we have  $c \geq 1/K$  in Theorem 2. This can be proved by showing that  $\pi_{\theta_{t+1}}(a^*) \geq \pi_{\theta_t}(a^*)$ , similar to Mei et al. (2020b, Proposition 2).*

### 5.3. Geometry-aware Normalized PG (GNPG)

Next, we generalize from one-state to finite MDPs, using the GNPG<sup>3</sup> on value function, as shown in Algorithm 1.

#### Algorithm 1 Geometry-aware Normalized Policy Gradient

**Input:** Learning rate  $\eta > 0$ .

Initialize parameter  $\theta_1(s, a)$  for all  $(s, a)$ .

**while**  $t \geq 1$  **do**

$$\theta_{t+1} \leftarrow \theta_t + \eta \cdot \frac{\partial V^{\pi_{\theta_t}}(\mu)}{\partial \theta_t} / \left\| \frac{\partial V^{\pi_{\theta_t}}(\mu)}{\partial \theta_t} \right\|_2.$$

**end while**

<sup>2</sup>This essentially means we prove that a uniform Łojasiewicz inequality holds for the entire sequence  $\{\theta_t\}_{t \geq 1}$ , but this does not imply that the NŁ condition is unnecessary. As shown in Mei et al. (2020b, Remark 1), Łojasiewicz-type inequalities with constant  $C > 0$  cannot hold. It can only become uniform after specifying an initialization  $\theta_1$  and an algorithm (in this case, PG). Otherwise, uniform Łojasiewicz cannot hold since initialization can make the NŁ coefficient  $\pi_\theta(a^*)$  arbitrarily close to 0.

<sup>3</sup>We use GNPG as the name of Algorithm 1, since NPG is usually used to refer to the natural PG algorithm in RL literature.

#### 5.4. General MDPs

For general finite MDPs, we assume ‘‘sufficient exploration’’ for the initial state distribution  $\mu$ , which is also adapted in literature (Agarwal et al., 2020; Mei et al., 2020b).

**Assumption 1** (Sufficient exploration). *The initial state distribution satisfies  $\min_s \mu(s) > 0$ .*

Given Assumption 1, Agarwal et al. (2020) prove asymptotic global convergence for PG on the non-concave  $\max_\theta V^{\pi_\theta}(\rho)$  problem, while Mei et al. (2020b) strengthens this to a  $O(1/t)$  rate using uniform smoothness and the following NŁ inequality that generalizes Lemma 1.

**Lemma 5** (NŁ, Mei et al. (2020b), Lemma 8). *Denote  $S := |\mathcal{S}|$  as the total number of states. We have,  $\forall \theta \in \mathbb{R}^{S \times A}$ ,*

$$\left\| \frac{\partial V^{\pi_\theta}(\mu)}{\partial \theta} \right\|_2 \geq \frac{\min_s \pi_\theta(a^*(s)|s)}{\sqrt{S} \cdot \left\| d_{\mu^*}^{\pi_\theta} / d_{\mu^*}^{\pi_\theta} \right\|_\infty} \cdot (V^*(\rho) - V^{\pi_\theta}(\rho)),$$

where  $a^*(s)$  is the action that  $\pi^*$  selects in state  $s$ .

Here, the NŁ degree  $\xi = 0$  is not improvable Mei et al. (2020b, Lemma 28). In one-state MDPs with  $S = 1$ , Lemma 5 recovers Lemma 1 with the same NŁ coefficient  $C(\theta) = \pi_\theta(a^*)$ , indicating that  $C(\theta)$  in Lemma 5 might also be unimprovable. On the other hand, the uniform smoothness considered in (Agarwal et al., 2020; Mei et al., 2020b)  $\beta = 8/(1-\gamma)^3$  is too conservative, particularly when  $\gamma$  is close to 1. Our next key result shows that the policy value also satisfies a stronger NS property, with the NS coefficient being the PG norm, generalizing Lemma 2:

**Lemma 6** (NS). *Let Assumption 1 hold and denote  $\theta_\zeta := \theta + \zeta \cdot (\theta' - \theta)$  with some  $\zeta \in [0, 1]$ .  $\theta \mapsto V^{\pi_\theta}(\mu)$  satisfies  $\beta(\theta_\zeta)$  non-uniform smoothness with*

$$\beta(\theta_\zeta) = \left[ 3 + \frac{2 \cdot (C_\infty - (1-\gamma))}{(1-\gamma) \cdot \gamma} \right] \cdot \sqrt{S} \cdot \left\| \frac{\partial V^{\pi_{\theta_\zeta}}(\mu)}{\partial \theta_\zeta} \right\|_2,$$

where  $C_\infty := \max_\pi \left\| \frac{d_\mu^\pi}{\mu} \right\|_\infty \leq \frac{1}{\min_s \mu(s)} < \infty$ .

In one-state MDPs with  $\gamma = 0$  and  $S = 1$ , we have  $C_\infty = 1 - \gamma$ . Thus Lemma 6 reduces to Lemma 2 with the same NS coefficient  $\beta(\theta_\zeta) = 3 \cdot \left\| \frac{d_{\theta_\zeta}^{\pi_\theta} r}{d\theta_\zeta} \right\|_2$ . Similar to Lemma 3, if we use Algorithm 1 with small enough learning rate, then  $\beta(\theta_\zeta)$  in Lemma 6 is upper bounded by the PG norm of  $\theta$ :

**Lemma 7.** *Let  $\eta = \frac{(1-\gamma) \cdot \gamma}{6 \cdot (1-\gamma) \cdot \gamma + 4 \cdot (C_\infty - (1-\gamma))} \cdot \frac{1}{\sqrt{S}}$  and  $\theta' = \theta + \eta \cdot \frac{\partial V^{\pi_\theta}(\mu)}{\partial \theta} / \left\| \frac{\partial V^{\pi_\theta}(\mu)}{\partial \theta} \right\|_2$ . Denote  $\theta_\zeta := \theta + \zeta \cdot (\theta' - \theta)$  with some  $\zeta \in [0, 1]$ . We have,*

$$\left\| \frac{\partial V^{\pi_{\theta_\zeta}}(\mu)}{\partial \theta_\zeta} \right\|_2 \leq 2 \cdot \left\| \frac{\partial V^{\pi_\theta}(\mu)}{\partial \theta} \right\|_2. \quad (13)$$

Next, the NŁ coefficient  $\min_s \pi_\theta(a^*(s)|s)$  in Lemma 5 is lower bounded away from 0:

**Lemma 8** (Non-vanishing NŁ coefficient). *Let Assumption 1 hold. We have,  $c := \inf_{s \in \mathcal{S}, t \geq 1} \pi_{\theta_t}(a^*(s)|s) > 0$ ,*

where  $\{\theta_t\}_{t \geq 1}$  is generated by Algorithm 1.

Now we have the non-concave function  $V^{\pi_\theta}(\rho)$  satisfies (1b) in Theorem 1 with  $\xi = 0$  in each iteration of Algorithm 1. Therefore, combining Lemmas 5 to 8, we prove the global linear convergence rate  $O(e^{-c \cdot t})$  of Algorithm 1:

**Theorem 3.** *Let Assumption 1 hold and let  $\{\theta_t\}_{t \geq 1}$  be generated using Algorithm 1 with  $\eta = \frac{(1-\gamma) \cdot \gamma}{6 \cdot (1-\gamma) \cdot \gamma + 4 \cdot (C_\infty - (1-\gamma))} \cdot \frac{1}{\sqrt{S}}$ , where  $C_\infty := \max_\pi \left\| \frac{d_\mu^\pi}{\mu} \right\|_\infty$ . Denote  $C'_\infty := \max_\pi \left\| \frac{d_\mu^\pi}{\mu} \right\|_\infty$ . Let  $c$  be the positive constant from Lemma 8. We have, for all  $t \geq 1$ ,*

$$V^*(\rho) - V^{\pi_{\theta_t}}(\rho) \leq \frac{(V^*(\mu) - V^{\pi_{\theta_1}}(\mu)) \cdot C'_\infty}{1 - \gamma} \cdot e^{-C \cdot (t-1)},$$

$$\text{where } C = \frac{(1-\gamma)^2 \cdot \gamma \cdot c}{12 \cdot (1-\gamma) \cdot \gamma + 8 \cdot (C_\infty - (1-\gamma))} \cdot \frac{1}{S} \cdot \left\| \frac{d_\mu^{\pi^*}}{\mu} \right\|_\infty^{-1}.$$

Not only the  $O(e^{-c \cdot t})$  rate in Theorem 3 is faster than  $O(1/t)$  for standard PG without normalization, but also the constant is better than Mei et al. (2020b, Theorem 4). The strictly better dependence  $c \gg c^2$  in PG is related to faster escaping plateaus as shown later (Mei et al., 2020a).

**Remark 4.** *The conclusion of GNPG has better constants than PG ( $c \gg c^2$ ) arises from upper bounds (Theorem 3 and Mei et al. (2020b, Theorem 4)), which is also supported by empirical evidence. In fact, there exists a lower bound that shows  $c$  cannot be removed for PG (Mei et al., 2020a, Theorem 1) under one-state MDP settings. For finite MDPs, very recently, Li et al. (2021) show that for softmax PG (without normalization),  $c$  can be very small in terms of the number of states. It remains open to consider whether  $c$  is reasonably large for GNPG.*

**Remark 5.** *To our knowledge, existing PG variants can achieve linear convergence  $O(e^{-c \cdot t})$  only if using at least one of the following techniques: (a) **regularization**; Mei et al. (2020b) prove that entropy regularized PG enjoys  $O(e^{-c \cdot t})$  convergence toward the regularized optimal policy. (b) **natural gradient**; Cen et al. (2020) prove that entropy regularized natural PG achieves linear convergence. (c) **exact line-search**; Bhandari & Russo (2020) prove that without parameterization, PG variants with exact line-search achieve linear rates by approximating policy iteration.*

*Among the above techniques, regularization changes the problem to regularized MDPs, while natural PG and line-search require solving expensive optimization problems to do updates, since each update is an arg max.*

On the contrary, Algorithm 1 enjoys global  $O(e^{-c \cdot t})$  rate (i) without using regularization, since Algorithm 1 directly works on the original MDPs; (ii) without solving optimization problems in each iteration, and the normalized PG update is cheap. The strong results rely on the NS and NŁ

properties, and also the geometry-aware normalization that takes advantage of the non-uniform properties.

**Remark 6.** *Standard softmax PG with bounded learning rate follows  $\Omega(1/t)$  lower bound (Mei et al., 2020b), which is consistent with the case (1) in Theorem 1. Algorithm 1 achieves faster linear convergence rates, indicating that the adaptive update stepsize  $\eta / \|\nabla V^{\pi_{\theta_t}}(\rho)\|_2$  is asymptotically unbounded, since  $\|\nabla V^{\pi_{\theta_t}}(\rho)\|_2 \rightarrow 0$  as  $t \rightarrow \infty$ .*




Figure 4. PG and GNPG on  $r = (1.0, 0.8, 0.1)^\top$ .

We compare PG and GNPG on the one-state MDP problem as shown in Figure 4. Figure 4(a) shows that GNPG escapes from the sub-optimal plateau significantly faster than PG, while Figure 4(b) shows that GNPG follows linear convergence  $O(e^{-c \cdot t})$  of sub-optimality, verifying the theoretical results. Similar experimental results on synthetic tree MDPs with multiple states are presented in Appendix F.

## 6. Generalized Linear Models

Next, we investigate the generalized linear model (GLM) with quasi-maximum likelihood estimate (quasi-MLE), which applied widely in supervised learning. We show that the mean squared error (MSE) of GLM (Hazan et al., 2015) is in the non-convex function class  $\mathcal{Z}$  in Figure 2, and it satisfies the case (2) in Theorem 1 with  $\xi = 1/2$ . As a result, both GD and GNGD achieve global linear convergence rates  $O(e^{-c \cdot t})$ , significantly improving the best existing results of  $O(1/\sqrt{t})$  (Hazan et al., 2015). We also provide new understandings of using normalization in GLM based on our non-uniform analysis.

### 6.1. Basic Settings and Notations

Given a training data set  $\mathcal{D} = \{(x_i, y_i)\}_{i \in [N]}$ , which consists of  $N$  data points, there is a feature map  $x_i \mapsto \phi(x_i) \in \mathbb{R}^d$  for each pair  $(x_i, y_i) \in \mathcal{D}$ . We denote  $\phi_i := \phi(x_i)$  for conciseness. For each data point  $x_i$ , we have  $y_i \in [0, 1]$  as the ground truth likelihood. Following Hazan et al. (2015), our model is parameterized by a weight vector  $\theta \in \mathbb{R}^d$  as ,

$$\pi_i = \sigma(\phi_i^\top \theta) = \frac{1}{1 + \exp\{-\phi_i^\top \theta\}}, \quad (14)$$

where  $\sigma : \mathbb{R} \rightarrow (0, 1)$  is the sigmoid activation. The problem is to minimize the mean squared error (MSE),

$$\min_{\theta} \mathcal{L}(\theta) = \min_{\theta \in \mathbb{R}^d} \frac{1}{N} \cdot \sum_{i=1}^N (\pi_i - y_i)^2. \quad (15)$$

We assume  $y_i = \pi_i^* := \sigma(\phi_i^\top \theta^*)$ , where  $\theta^* \in \mathbb{R}^d$ , and  $\|\theta^*\|_2 < \infty$ , which means the target  $y_i$  is realizable and non-deterministic. According to Hazan et al. (2015), the MSE in Eq. (15) is not quasi-convex (thus not convex). Fortunately, Hazan et al. (2015) manage to show that Eq. (15) satisfies a weaker Strictly-Locally-Quasi-Convex (SLQC) property, based on which they prove the following result:

**Theorem 4** (Hazan et al. (2015)). *With diminishing learning rate  $\eta_t \in \Theta(1/\sqrt{t})$ , the normalized gradient descent (NGD) update  $\theta_{t+1} \leftarrow \theta_t - \eta_t \cdot \frac{\partial \mathcal{L}(\theta_t)}{\partial \theta_t} / \left\| \frac{\partial \mathcal{L}(\theta_t)}{\partial \theta_t} \right\|_2$  satisfies,*

$$\delta(\theta_t) := \mathcal{L}(\theta_t) - \mathcal{L}(\theta^*) \in O(1/\sqrt{t}), \quad (16)$$

where  $\theta^* := \arg \min_{\theta} \mathcal{L}(\theta)$  is the global optimal solution.

### 6.2. Fast Convergence using Non-uniform Analysis

Based on the  $O(1/\sqrt{t})$  rate for NGD in Theorem 4, Hazan et al. (2015) propose to normalize gradient norm in MSE minimization. However, there is no lower bound for other methods including GD on GLM, and thus it is not clear if there exists a faster rate for GLM optimization.

Surprisingly, we prove that both GD and GNGD actually achieve much faster rates of  $O(e^{-c \cdot t})$  using the non-uniform analysis. Our first key finding is to show that the MSE in GLM satisfies a new NŁ inequality with  $\xi = 1/2$ :

**Lemma 9** (NŁ). *Denote  $u(\theta) := \min_{i \in [N]} \{\pi_i \cdot (1 - \pi_i)\}$ , and  $v := \min_{i \in [N]} \{\pi_i^* \cdot (1 - \pi_i^*)\}$ . We have,*

$$\left\| \frac{\partial \mathcal{L}(\theta)}{\partial \theta} \right\|_2 \geq C(\theta, \phi) \cdot \left[ \frac{1}{N} \cdot \sum_{i=1}^N (\pi_i - \pi_i^*)^2 \right]^{\frac{1}{2}}, \quad (17)$$

holds for all  $\theta \in \mathbb{R}^d$ , where

$$C(\theta, \phi) = 8 \cdot u(\theta) \cdot \min\{u(\theta), v\} \cdot \sqrt{\lambda_\phi}, \quad (18)$$

and  $\lambda_\phi$  is the smallest positive eigenvalue of  $\frac{1}{N} \cdot \sum_{i=1}^N \phi_i \phi_i^\top$ .

**Remark 7.** *It is not clear if results similar to Lemma 9 hold without assuming: (i) realizable optimal prediction  $y_i = \pi_i^* := \sigma(\phi_i^\top \theta^*)$ ; (ii) non-deterministic optimal prediction  $\|\theta^*\|_2 < \infty$ . We leave it as an open question to study non-uniformity of GLM without the above assumptions.*

In Lemma 9,  $\lambda_\phi$  is determined by the feature  $\phi$ , and  $u(\theta)$  shows that the gradient is vanishing when  $\pi_i$  is near deterministic, which is consistent with the fact that the sigmoid saturates and provides uninformative gradient as the parameter magnitude becomes large.

We run GD on one example with  $N = 10$  and  $d = 2$ . As shown in Figure 5, the gradient norm  $\|\nabla \mathcal{L}(\theta_t)\|_2$  is close to



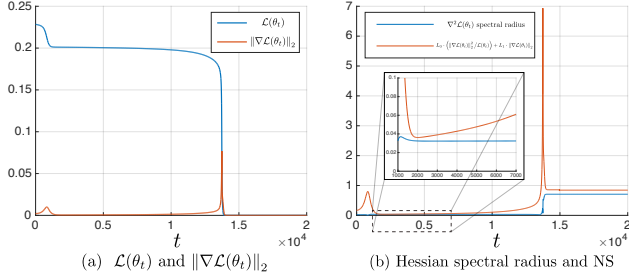


Figure 5. Experiments on GLM using GD.

zero at plateaus and near optimum. However, unlike the PG, the spectral radius of the Hessian  $\nabla^2 \mathcal{L}(\theta_t)$  is only close to zero at plateaus, while it approaches positive constant near optimum. This indicates a different NS condition other than Lemmas 2 and 6 is needed, since only gradient norm  $\rightarrow 0$  cannot upper bound the spectral radius of Hessian  $\rightarrow c > 0$ . With some calculations, we prove the following key results:

**Lemma 10** (Smoothness and NS).  $\mathcal{L}(\theta)$  satisfies  $\beta$  smoothness with

$$\beta = \frac{3}{8} \cdot \max_{i \in [N]} \|\phi_i\|_2^2, \quad (19)$$

and  $\beta(\theta)$  NS with

$$\beta(\theta) = L_1 \cdot \left\| \frac{\partial \mathcal{L}(\theta)}{\partial \theta} \right\|_2 + L_0 \cdot \left( \left\| \frac{\partial \mathcal{L}(\theta)}{\partial \theta} \right\|_2^2 / \mathcal{L}(\theta) \right).$$

At the optimal solution  $\theta^*$ , the spectral radius of the Hessian  $\frac{\partial^2 \mathcal{L}(\theta^*)}{\partial(\theta^*)^2}$  is strictly positive. Therefore, the MSE objective of Eq. (15) is in the non-convex function class  $\mathcal{Z}$  in Figure 2, and it satisfies the case (2) in Theorem 1 with  $\xi = 1/2$ . Combining Lemmas 9 and 10 and applying Theorem 1, we have the following global linear convergence result:

**Theorem 5.** With  $\eta = 1/\beta$ , GD update satisfies for all  $t \geq 1$ ,  $\mathcal{L}(\theta_t) \leq \mathcal{L}(\theta_1) \cdot e^{-C^2 \cdot (t-1)}$ . With  $\eta \in \Theta(1)$ , GNGD update satisfies for all  $t \geq 1$ ,  $\mathcal{L}(\theta_t) \leq \mathcal{L}(\theta_1) \cdot e^{-C \cdot (t-1)}$ , where  $C \in (0, 1)$ , i.e., GNGD is strictly faster than GD.

Theorem 5 significantly improves the  $O(1/\sqrt{t})$  rate in Theorem 4. The key difference is that we discovered a new NŁ inequality of Lemma 9 that is satisfied by GLMs. The linear convergence rates are verified in Appendix F.

In Theorem 5, we have  $C = \inf_{t \geq 1} C(\theta_t, \phi)$ , which is very close to zero if  $\pi_i$  is near deterministic, and GD suffers sub-optimality plateaus as shown in Figure 1. GNGD has strictly (orders of magnitudes) better constant dependence  $C \gg C^2$ , and escapes plateaus significantly faster than GD. Intuitively, for the GLM in Figure 1,  $C$  in Theorem 1 is lower bounded reasonably if  $\theta_1$  is initialized within some finite distance of the central valley containing  $\theta^*$ .

Combining the NŁ and NS properties (Lemmas 9 and 10), we provide new understandings of using normalization in GLM: (i) **First**, using standard NGD (Hazan et al., 2015) for

all  $t \geq 1$  is not a good choice. By examining the asymptotic behaviour as  $\theta \rightarrow \theta^*$ , we have  $\beta(\theta) \rightarrow \beta > 0$ . However, the normalization  $\left\| \frac{\partial \mathcal{L}(\theta)}{\partial \theta} \right\|_2$  in standard NGD gives incremental updates with adaptive stepsize  $\rightarrow \infty$ . To guarantee convergence, it is necessary to use  $\eta_t \rightarrow 0$ , which counteracts normalization and slows down the learning, since it could be not easy to find a learning rate scheme. This is consistent with the  $O(e^{-c \cdot t})$  result for GD with  $\eta > 0$  and without normalization in Theorem 5. (ii) **Second**, using geometry-aware normalization  $\beta(\theta_t)$  is a better choice than normalizing the gradient norm  $\|\nabla \mathcal{L}(\theta_t)\|_2$ . We elaborate this point by investigating *both the asymptotic and the early-stage behaviours* using NS-NŁ. Since  $\beta(\theta_t) \rightarrow \beta > 0$  asymptotically, GNGD is approaching GD as  $\theta_t \rightarrow \theta^*$ , which makes GNGD enjoy the same  $O(e^{-c \cdot t})$  rate. On the other hand, at early-stage optimization (e.g., close to initialization in Figure 1), when  $\theta_t$  is far from  $\theta^*$ , we have thus  $\beta(\theta_t) \leq c \cdot \left\| \frac{\partial \mathcal{L}(\theta_t)}{\partial \theta_t} \right\|_2$ . Then GNGD is close to NGD, which guarantees strictly better progresses than GD. This is because of the progress of GNGD in each iteration at this time is about  $\|\nabla \mathcal{L}(\theta_t)\|_2$ , while the progress of GD is  $\|\nabla \mathcal{L}(\theta_t)\|_2^2$ , and the gradient norm is close to 0 on plateaus. Using NŁ of Lemma 9, GNGD will have strictly better constant dependence  $C$  than  $C^2$  in GD.

## 7. Conclusions and Future Work

The main contributions of this paper concern a general characterization and analysis based on non-uniform properties, which are not only sufficiently general to cover concrete examples, but also significantly improve convergence rates over previous work and even over classical lower bounds. The most exciting part is the techniques apply to important applications in machine learning that involve non-convex optimization problems. One direction is to further push the analysis to other domains with more complex function approximators, including neural networks (Allen-Zhu et al., 2019). Another valuable future work is to incorporate stochastic gradient (Karimi et al., 2016) and other adaptive gradient-based methods (Kingma & Ba, 2014) in the analysis. Finally, applying other cases of non-uniform properties beyond those mentioned in this paper would be interesting.

## Acknowledgements

The authors would like to thank anonymous reviewers for their valuable comments. Jincheng Mei would like to thank Ziyi Chen for pointing that NŁ recovers global KŁ property. Jincheng Mei and Bo Dai would like to thank Nicolas Le Roux for providing feedback on a draft of this manuscript. Csaba Szepesvári and Dale Schuurmans gratefully acknowledge funding from the Canada CIFAR AI Chairs Program, Amii and NSERC.

## References

- Agarwal, A., Kakade, S. M., Lee, J. D., and Mahajan, G. Optimality and approximation with policy gradient methods in markov decision processes. In *Conference on Learning Theory*, pp. 64–66. PMLR, 2020.
- Allen-Zhu, Z., Li, Y., and Song, Z. A convergence theory for deep learning via over-parameterization. In *International Conference on Machine Learning*, pp. 242–252. PMLR, 2019.
- Bhandari, J. and Russo, D. A note on the linear convergence of policy gradient methods. *arXiv preprint arXiv:2007.11120*, 2020.
- Bubeck, S. Convex optimization: Algorithms and complexity. *arXiv preprint arXiv:1405.4980*, 2014.
- Cen, S., Cheng, C., Chen, Y., Wei, Y., and Chi, Y. Fast global convergence of natural policy gradient methods with entropy regularization. *arXiv preprint arXiv:2007.06558*, 2020.
- Cohen, J., Kaur, S., Li, Y., Kolter, J. Z., and Talwalkar, A. Gradient descent on neural networks typically occurs at the edge of stability. In *International Conference on Learning Representations*, 2021.
- Hazan, E., Levy, K., and Shalev-Shwartz, S. Beyond convexity: Stochastic quasi-convex optimization. *Advances in neural information processing systems*, 28:1594–1602, 2015.
- Hazan, E., Kakade, S., Singh, K., and Van Soest, A. Provably efficient maximum entropy exploration. In *International Conference on Machine Learning*, pp. 2681–2691. PMLR, 2019.
- Kakade, S. and Langford, J. Approximately optimal approximate reinforcement learning. In *ICML*, volume 2, pp. 267–274, 2002.
- Karimi, H., Nutini, J., and Schmidt, M. Linear convergence of gradient and proximal-gradient methods under the polyak-łojasiewicz condition. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pp. 795–811. Springer, 2016.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Kurdyka, K. On gradients of functions definable in o-minimal structures. In *Annales de l’institut Fourier*, volume 48, pp. 769–783, 1998.
- Li, G., Wei, Y., Chi, Y., Gu, Y., and Chen, Y. Softmax policy gradient methods can take exponential time to converge. *arXiv preprint arXiv:2102.11270*, 2021.
- Łojasiewicz, S. Une propriété topologique des sous-ensembles analytiques réels. *Les équations aux dérivées partielles*, 117:87–89, 1963.
- Mei, J., Xiao, C., Dai, B., Li, L., Szepesvári, C., and Schuurmans, D. Escaping the gravitational pull of softmax. *Advances in Neural Information Processing Systems*, 33, 2020a.
- Mei, J., Xiao, C., Szepesvari, C., and Schuurmans, D. On the global convergence rates of softmax policy gradient methods. In *International Conference on Machine Learning*, pp. 6820–6829. PMLR, 2020b.
- Nemirovski, A. S. and Yudin, D. B. Problem complexity and method efficiency in optimization. 1983.
- Nesterov, Y. *Introductory lectures on convex optimization: A basic course*, volume 87. Springer Science & Business Media, 2003.
- Polyak, B. T. Gradient methods for minimizing functionals. *Zhurnal Vychislitel’noi Matematiki i Matematicheskoi Fiziki*, 3(4):643–653, 1963.
- Sutton, R. S., McAllester, D. A., Singh, S. P., and Mansour, Y. Policy gradient methods for reinforcement learning with function approximation. In *Advances in neural information processing systems*, pp. 1057–1063, 2000.
- Wilson, A., Mackey, L., and Wibisono, A. Accelerating rescaled gradient descent: Fast optimization of smooth functions. *arXiv preprint arXiv:1902.08825*, 2019.
- Zhang, J., He, T., Sra, S., and Jadbabaie, A. Why gradient clipping accelerates training: A theoretical justification for adaptivity. In *International Conference on Learning Representations*, 2019.