

Beyond Convexity: Stochastic Quasi-Convex Optimization

Elad Hazan* Kfir Y. Levy† Shai Shalev-Shwartz‡

May 2014

Abstract

Stochastic convex optimization is a basic and well studied primitive in machine learning. It is well known that convex and Lipschitz functions can be minimized efficiently using Stochastic Gradient Descent (SGD).

The Normalized Gradient Descent (NGD) algorithm, is an adaptation of Gradient Descent, which updates according to the direction of the gradients, rather than the gradients themselves. In this paper we analyze a stochastic version of NGD and prove its convergence to a *global* minimum for a wider class of functions: we require the functions to be *quasi-convex* and *locally-Lipschitz*. Quasi-convexity broadens the concept of unimodality to multidimensions and allows for certain types of saddle points, which are a known hurdle for first-order optimization methods such as gradient descent. Locally-Lipschitz functions are only required to be Lipschitz in a small region around the optimum. This assumption circumvents gradient explosion, which is another known hurdle for gradient descent variants.

Interestingly, unlike the vanilla SGD algorithm, the stochastic normalized gradient descent algorithm provably requires a minimal minibatch size.

1 Introduction

The benefits of using the Stochastic Gradient Descent (SGD) scheme for learning could not be stressed enough. For convex and Lipschitz objectives, SGD is guaranteed to find an ϵ -optimal solution within $O(1/\epsilon^2)$ iterations and requires only an unbiased estimator for the gradient, which is obtained with only one (or a few) data samples. However, when applied to non-convex problems several drawbacks are revealed. In particular, SGD is widely used for deep learning [Bengio \(2009\)](#), one of the most interesting fields where stochastic non-convex optimization problems arise. Often, the objective in these kind of problems demonstrates

*Princeton University; ehazan@cs.princeton.edu.

†Technion; kfiryl@tx.technion.ac.il.

‡The Hebrew University; shais@cs.huji.ac.il.

two extreme phenomena: on the one hand plateaus, i.e., regions with very small gradients; and on the other hand sharp cliffs, i.e., exceedingly high gradients. As is expected, applying SGD to these problems is often reported to yield unsatisfactory results.

In this paper we analyze a stochastic version of the Normalized Gradient Descent (NGD) algorithm, which we denote by SNGD. Each iteration of SNGD is as simple and efficient as SGD, but is much more appropriate for non-convex optimization problems, overcoming some of the pitfalls that SGD may encounter. Particularly, we define a family of *locally-quasi-convex* and *locally-Lipschitz* functions, and prove that SNGD is suitable for optimizing such objectives.

Local-Quasi-convexity is a generalization of unimodal functions to multidimensions, which includes quasi-convex, and convex functions as a subclass. Locally-Quasi-convex functions allow for certain types of plateaus and saddle points which are difficult for SGD and other gradient descent variants. Local-Lipschitzness is a generalization of Lipschitz functions that only assumes Lipschitzness in a small region around the minima, whereas farther away the gradients may be unbounded. Gradient explosion is, thus, another difficulty that is successfully tackled by SNGD and poses difficulties for other stochastic gradient descent variants.

Our contributions:

- We introduce *local-quasi-convexity*, a property that extends quasi-convexity and captures unimodal functions which are not quasi-convex. We prove that NGD finds an ϵ -optimal minimum for such functions within $O(1/\epsilon^2)$ iterations. As a special case, we show that the above rate can be attained for quasi-convex functions that are Lipschitz in an $\Omega(\epsilon)$ -region around the optimum (gradients may be *unbounded* outside this region). For objectives that are also smooth in an $\Omega(\sqrt{\epsilon})$ -region around the optimum, we prove a faster rate of $O(1/\epsilon)$.
- We introduce a new setup: stochastic optimization of locally-quasi-convex functions; and show that this setup captures Generalized Linear Models (GLM) regression, [McCullagh and Nelder \(1989\)](#). For this setup, we devise a stochastic version of NGD (SNGD), and show that it converges within $O(1/\epsilon^2)$ iterations to an ϵ -optimal minimum.
- The above positive result requires that at each iteration of SNGD, the gradient should be estimated using a minibatch of a minimal size. We provide a negative result showing that if the minibatch size is too small then the algorithm might indeed diverge.
- We report experimental results supporting our theoretical guarantees and demonstrate an accelerated convergence attained by SNGD.

1.1 Related Work

Quasi-convex optimization problems arise in numerous fields, spanning economics [Varian \(1985\)](#); [Laffont and Martimort \(2009\)](#), industrial organization [Wolfstetter \(1999\)](#), and computer vision [Ke and Kanade \(2007\)](#). It is well known that quasi-convex optimization tasks

can be solved by a series of convex feasibility problems [Boyd and Vandenberghe \(2004\)](#); However, generally solving such feasibility problems may be very costly [Goffin et al. \(1996\)](#).

There exists a rich literature concerning quasi-convex optimization in the offline case, [Polyak \(1967\)](#); [Zabotin et al. \(1972\)](#); [Khabibullin \(1977\)](#); [Sikorski \(1986\)](#). A pioneering paper by [Nesterov \(1984\)](#), was the first to suggest an efficient algorithm, namely Normalized Gradient Descent, and prove that this algorithm attains ϵ -optimal solution within $O(1/\epsilon^2)$ iterations given a differentiable quasi-convex objective. This work was later extended by [Kiwiel \(2001\)](#), showing that the same result may be achieved assuming upper semi-continuous quasi-convex objectives. In [Konnov \(2003\)](#) it was shown how to attain faster rates for quasi-convex optimization, but they assume to know the optimal value of the objective, an assumption that generally does not hold in practice.

Among the deep learning community there have been several attempts to tackle gradient-explosion/plateaus. Ideas spanning gradient-clipping [Pascanu et al. \(2013\)](#), smart initialization [Doya \(1993\)](#), and more, [Martens and Sutskever \(2011\)](#), have shown to improve training in practice. Yet, non of these works provides a theoretical analysis showing better convergence guarantees.

To the best of our knowledge, there are no previous results on stochastic versions of NGD, neither results regarding locally-quasi-convex/locally-Lipschitz functions.

Gradient descent with fixed step sizes, including its stochastic variants, is known to perform poorly when the gradients are too small in a plateau area of the function, or alternatively when the other extreme happens: gradient explosions. These two phenomena have been reported in certain types of non-convex optimization, such as training of deep networks.

Figure 1 depicts a one-dimensional family of functions for which GD behaves provably poorly. With a large step-size, GD will hit the cliffs and then oscillate between the two boundaries. Alternatively, with a small step size, the low gradients will cause GD to miss the middle valley which has constant size of $1/2$. On the other hand, this exact function is quasi-convex and locally-Lipschitz, and hence the NGD algorithm provably converges to the optimum quickly.

2 Definitions and Notations

We use $\|\cdot\|$ to denote the Euclidean norm. $\mathbb{B}_d(\mathbf{x}, r)$ denotes the d dimensional Euclidean ball of radius r , centered around \mathbf{x} , and $\mathbb{B}_d := \mathbb{B}_d(0, 1)$. $[N]$ denotes the set $\{1, \dots, N\}$.

For simplicity, throughout the paper we always assume that functions are *differentiable* (but if not stated explicitly, we do not assume any bound on the norm of the gradients).

Definition 2.1. (*Local-Lipschitzness and Local-Smoothness*) Let $\mathbf{z} \in \mathbb{R}^d$, $G, \epsilon \geq 0$. A function $f : \mathcal{K} \mapsto \mathbb{R}$ is called $(G, \epsilon, \mathbf{z})$ -Locally-Lipschitz if for every $\mathbf{x}, \mathbf{y} \in \mathbb{B}_d(\mathbf{z}, \epsilon)$, we have

$$|f(\mathbf{x}) - f(\mathbf{y})| \leq G\|\mathbf{x} - \mathbf{y}\| .$$

Similarly, the function is $(\beta, \epsilon, \mathbf{z})$ -locally-smooth if for every $\mathbf{x}, \mathbf{y} \in \mathbb{B}_d(\mathbf{z}, \epsilon)$ we have,

$$|f(\mathbf{y}) - f(\mathbf{x}) - \langle \nabla f(\mathbf{y}), \mathbf{x} - \mathbf{y} \rangle| \leq \frac{\beta}{2}\|\mathbf{x} - \mathbf{y}\|^2 .$$

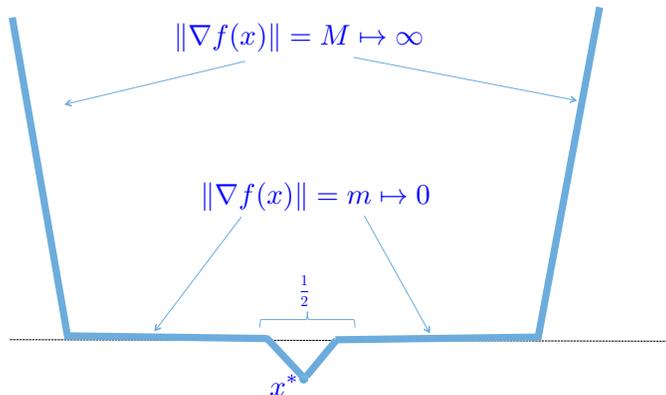


Figure 1: A quasi-convex Locally-Lipschitz function with plateaus and cliffs.

Next we define quasi-convex functions:

Definition 2.2. (Quasi-Convexity) We say that a function $f : \mathbb{R}^d \mapsto \mathbb{R}$ is quasi-convex if $\forall \mathbf{x}, \mathbf{y} \in \mathbb{R}^d$, such that $f(\mathbf{y}) \leq f(\mathbf{x})$, it follows that

$$\langle \nabla f(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle \leq 0 .$$

We further say that f is strictly-quasi-convex, if it is quasi-convex and its gradients vanish only at the global minima, i.e., $\forall \mathbf{y} : f(\mathbf{y}) > \min_{\mathbf{x} \in \mathbb{R}^d} f(\mathbf{x}) \Rightarrow \|\nabla f(\mathbf{y})\| > 0$.

Informally, the above characterization states that the (opposite) gradient of a quasi-convex function directs us in a *global* descent direction. Following is an equivalent (more common) definition:

Definition 2.3. (Quasi-Convexity) We say that a function $f : \mathbb{R}^d \mapsto \mathbb{R}$ is quasi-convex if any α -sublevel-set of f is convex, i.e., $\forall \alpha \in \mathbb{R}$ the set

$$\mathcal{L}_\alpha(f) = \{\mathbf{x} : f(\mathbf{x}) \leq \alpha\} \text{ is convex.}$$

The equivalence between the above definitions can be found in [Boyd and Vandenberghe \(2004\)](#), for completeness we provide a proof in [Appendix A](#). During this paper we denote the sublevel-set of f at \mathbf{x} by

$$S_f(\mathbf{x}) = \{\mathbf{y} : f(\mathbf{y}) \leq f(\mathbf{x})\} . \tag{1}$$

3 Local-Quasi-Convexity

Quasi-convexity does not fully capture the notion of unimodality in several dimension. As an example let $\mathbf{x} = (x_1, x_2) \in [-10, 10]^2$, and consider the function

$$g(\mathbf{x}) = (1 + e^{-x_1})^{-1} + (1 + e^{-x_2})^{-1} . \tag{2}$$

It is natural to consider g as unimodal since it acquires no local minima but for the unique global minima at $\mathbf{x}^* = (-10, -10)$. However, g is not quasi-convex: consider the points $\mathbf{x} = (\log 16, -\log 4)$, $\mathbf{y} = (-\log 4, \log 16)$, which belong to the 1.2-sub-level set, their average does not belong to the same sub-level-set since $g(\mathbf{x}/2 + \mathbf{y}/2) = 4/3$.

Quasi-convex functions always enable us to *explore*, meaning that the gradient always directs us in a global descent direction. Intuitively, from an optimization point of view, we only need such a direction whenever we do not *exploit*, i.e., whenever we are not approximately optimal.

In what follows we define local-quasi-convexity, a property that enables us to either *explore/exploit*. This property ¹ captures a wider class of unimodal function (such as g above) rather than mere quasi-convexity. Later we justify this definition by showing that it captures Generalized Linear Models (GLM) regression, see [McCullagh and Nelder \(1989\)](#); [Kalai and Sastry \(2009\)](#).

Definition 3.1. (*Local-Quasi-Convexity*) Let $\mathbf{x}, \mathbf{z} \in \mathbb{R}^d$, $\kappa, \epsilon > 0$. We say that $f : \mathbb{R}^d \mapsto \mathbb{R}$ is $(\epsilon, \kappa, \mathbf{z})$ -Strictly-Locally-Quasi-Convex (SLQC) in \mathbf{x} , if at least one of the following applies:

1. $f(\mathbf{x}) - f(\mathbf{z}) \leq \epsilon$.
2. $\|\nabla f(\mathbf{x})\| > 0$, and for every $\mathbf{y} \in \mathbb{B}(\mathbf{z}, \epsilon/\kappa)$ it holds that $\langle \nabla f(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle \leq 0$.

Note that if f is G -Lipschitz and strictly-quasi-convex function, then $\forall \mathbf{x}, \mathbf{z} \in \mathbb{R}^d$, $\forall \epsilon > 0$, it holds that f is $(\epsilon, G, \mathbf{z})$ -SLQC in \mathbf{x} . Recalling the function g that appears in Equation (2), then it can be shown that $\forall \epsilon \in (0, 1]$, $\forall \mathbf{x} \in [-10, 10]^2$ then this function is $(\epsilon, 1, \mathbf{x}^*)$ -SLQC in \mathbf{x} , where $\mathbf{x}^* = (-10, -10)$ (see Appendix B).

3.1 Generalized Linear Models (GLM)

3.1.1 The Idealized GLM

In this setup we have a collection of m samples $\{(\mathbf{x}_i, y_i)\}_{i=1}^m \in \mathbb{B}_d \times [0, 1]$, and an activation function $\phi : \mathbb{R} \mapsto \mathbb{R}$. We are guaranteed to have $\mathbf{w}^* \in \mathbb{R}^d$ such that: $y_i = \phi\langle \mathbf{w}^*, \mathbf{x}_i \rangle$, $\forall i \in [m]$ (we denote $\phi\langle \mathbf{w}, \mathbf{x} \rangle := \phi(\langle \mathbf{w}, \mathbf{x} \rangle)$). The performance of a predictor $\mathbf{w} \in \mathbb{R}^d$, is measured by the average square error over all samples.

$$\widehat{\text{err}}_m(\mathbf{w}) = \frac{1}{m} \sum_{i=1}^m (y_i - \phi\langle \mathbf{w}, \mathbf{x}_i \rangle)^2. \quad (3)$$

In [Kalai and Sastry \(2009\)](#) it is shown that the Perceptron problem with γ -margin is a private case of GLM regression.

The sigmoid function $\phi(z) = (1 + e^{-z})^{-1}$ is a popular activation function in the field of deep learning. The next lemma states that in the idealized GLM problem with sigmoid

¹Definition 3.1 can be generalized in a manner that captures a broader range of scenarios (e.g. the Perceptron problem), we defer this definition to Appendix H.

activation, then the error function is SLQC (but not quasi-convex). As we will see in Section 4 this implies that Algorithm 1 finds an ϵ -optimal minima of $\widehat{\text{err}}_m(\mathbf{w})$ within $\text{poly}(1/\epsilon)$ iterations.

Lemma 3.1. *Consider the idealized GLM problem with the sigmoid activation, and assume that $\|\mathbf{w}^*\| \leq W$. Then the error function appearing in Equation (3) is $(\epsilon, e^W, \mathbf{w}^*)$ -SLQC in \mathbf{w} , $\forall \epsilon > 0, \forall \mathbf{w} \in \mathbb{B}_d(0, W)$ (But it is not generally quasi-convex).*

We defer the proof to Appendix C

3.1.2 The Noisy GLM

In the noisy GLM setup (see McCullagh and Nelder (1989); Kalai and Sastry (2009)), we may draw i.i.d. samples $\{(\mathbf{x}_i, y_i)\}_{i=1}^m \in \mathbb{B}_d \times [0, 1]$, from an unknown distribution \mathcal{D} . We assume that there exists a predictor $\mathbf{w}^* \in \mathbb{R}^d$ such that $\mathbf{E}_{(\mathbf{x}, y) \sim \mathcal{D}}[y|\mathbf{x}] = \phi\langle \mathbf{w}^*, \mathbf{x} \rangle$, where ϕ is an activation function. Given $\mathbf{w} \in \mathbb{R}^d$ we define its expected error as follows:

$$\mathcal{E}(\mathbf{w}) = \mathbf{E}_{(\mathbf{x}, y) \sim \mathcal{D}}(y - \phi\langle \mathbf{w}, \mathbf{x} \rangle)^2,$$

and it can be shown that \mathbf{w}^* is a global minima of \mathcal{E} . We are interested in schemes that obtain an ϵ -optimal minima to \mathcal{E} , within $\text{poly}(1/\epsilon)$ samples and optimization steps. Given m samples from \mathcal{D} , their empirical error $\widehat{\text{err}}_m(\mathbf{w})$, is defined as in Equation (3).

The following lemma states that in this setup, letting $m = \Omega(1/\epsilon^2)$, then $\widehat{\text{err}}_m$ is SLQC with high probability. This property will enable us to apply Algorithm 2, to obtain an ϵ -optimal minima to \mathcal{E} , within $\text{poly}(1/\epsilon)$ samples from \mathcal{D} , and $\text{poly}(1/\epsilon)$ optimization steps.

Lemma 3.2. *Let $\delta, \epsilon \in (0, 1)$. Consider the noisy GLM problem with the sigmoid activation, and assume that $\|\mathbf{w}^*\| \leq W$. Given a fixed point $\mathbf{w} \in \mathbb{B}(0, W)$, then w.p. $\geq 1 - \delta$, after $m \geq \frac{8e^{2W}(W+1)^2}{\epsilon^2} \log(1/\delta)$ samples, the empirical error function appearing in Equation (3) is $(\epsilon, e^W, \mathbf{w}^*)$ -SLQC in \mathbf{w} .*

Note that if we had required the SLQC to hold $\forall \mathbf{w} \in \mathbb{B}(0, W)$, then we would need the number of samples to depend on the dimension, d , which we would like to avoid. Instead, we require SLQC to hold for a fixed \mathbf{w} . This satisfies the conditions of Algorithm 2, enabling us to find an ϵ -optimal solution with a sample complexity that is independent of the dimension. We defer the proof of Lemma 3.2 to Appendix D

4 NGD for Locally-Quasi-Convex Optimization

Here we present the NGD algorithm, and prove the convergence rate of this algorithm for SLQC objectives. Our analysis is simple, enabling us to extend the convergence rate presented in Nesterov (1984) beyond quasi-convex functions. We then show that quasi-convex and *locally-Lipschitz* objective are SLQC, implying that NGD converges even if the gradients are unbounded outside a small region around the minima. For quasi-convex and *locally-smooth* objectives, we show that NGD attains a faster convergence rate.

Algorithm 1 Normalized Gradient Descent (NGD)

Input: #Iterations T , $\mathbf{x}_1 \in \mathbb{R}^d$, learning rate η

for $t = 1 \dots T$ **do**

 Update:

$$\mathbf{x}_{t+1} = \mathbf{x}_t - \eta \hat{g}_t \quad \text{where } g_t = \nabla f(\mathbf{x}_t), \hat{g}_t = \frac{g_t}{\|g_t\|}$$

end for

Return: $\bar{\mathbf{x}}_T = \arg \min_{\{\mathbf{x}_1, \dots, \mathbf{x}_T\}} f(\mathbf{x}_t)$

NGD is presented in Algorithm 1. NGD is similar to GD, except we normalize the gradients. It is intuitively clear that to obtain robustness to plateaus (where the gradient can be arbitrarily small) and to exploding gradients (where the gradient can be arbitrarily large), one must ignore the size of the gradient. It is more surprising that the information in the direction of the gradient suffices to guarantee convergence.

Following is the main theorem of this section:

Theorem 4.1. *Fix $\epsilon > 0$, let $f : \mathbb{R}^d \mapsto \mathbb{R}$, and $\mathbf{x}^* \in \arg \min_{\mathbf{x} \in \mathbb{R}^d} f(\mathbf{x})$. Given that f is $(\epsilon, \kappa, \mathbf{x}^*)$ -SLQC in every $\mathbf{x} \in \mathbb{R}^d$. Then running the NGD algorithm with $T \geq \kappa^2 \|\mathbf{x}_1 - \mathbf{x}^*\|^2 / \epsilon^2$, and $\eta = \epsilon / \kappa$, we have that: $f(\bar{\mathbf{x}}_T) - f(\mathbf{x}^*) \leq \epsilon$.*

Theorem 4.1 states that $(\cdot, \cdot, \mathbf{x}^*)$ -SLQC functions admit $\text{poly}(1/\epsilon)$ convergence rate using NGD. The intuition behind this lies in Definition 3.1, which asserts that at a point \mathbf{x} either the (opposite) gradient points out a global optimization direction, or we are already ϵ -optimal. Note that the requirement of (ϵ, \cdot, \cdot) -SLQC in any \mathbf{x} is not restrictive, as we have seen in Section 3, there are interesting examples of functions that admit this property $\forall \epsilon \in [0, 1]$, and for any \mathbf{x} .

For simplicity we have presented NGD for unconstrained problems. Using projections we can easily extend the algorithm and its analysis for constrained optimization over convex sets. This will enable to achieve convergence of $O(1/\epsilon^2)$ for the objective presented in Equation (2), and the idealized GLM problem presented in Section 3.1.1.

We are now ready to prove Theorem 4.1:

Proof of Theorem 4.1. First note that if the gradient of f vanishes at \mathbf{x}_t , then by the SLQC assumption we must have that $f(\mathbf{x}_t) - f(\mathbf{x}^*) \leq \epsilon$. Assume next that we perform T iterations and the gradient of f at \mathbf{x}_t never vanishes in these iterations. Consider the update rule of NGD (Algorithm 1), then by standard algebra we get,

$$\|\mathbf{x}_{t+1} - \mathbf{x}^*\|^2 = \|\mathbf{x}_t - \mathbf{x}^*\|^2 - 2\eta \langle \hat{g}_t, \mathbf{x}_t - \mathbf{x}^* \rangle + \eta^2 .$$

Assume that $\forall t \in [T]$ we have $f(\mathbf{x}_t) - f(\mathbf{x}^*) > \epsilon$. Take $\mathbf{y} = \mathbf{x}^* + (\epsilon/\kappa) \hat{g}_t$, and observe that $\|\mathbf{y} - \mathbf{x}^*\| \leq \epsilon/\kappa$. The $(\epsilon, \kappa, \mathbf{x}^*)$ -SLQC assumption implies that $\langle \hat{g}_t, \mathbf{y} - \mathbf{x}_t \rangle \leq 0$, and therefore

$$\langle \hat{g}_t, \mathbf{x}^* + (\epsilon/\kappa) \hat{g}_t - \mathbf{x}_t \rangle \leq 0 \Rightarrow \langle \hat{g}_t, \mathbf{x}_t - \mathbf{x}^* \rangle \geq \epsilon/\kappa .$$

Setting $\eta = \epsilon/\kappa$, the above implies,

$$\begin{aligned}\|\mathbf{x}_{t+1} - \mathbf{x}^*\|^2 &\leq \|\mathbf{x}_t - \mathbf{x}^*\|^2 - 2\eta\epsilon/\kappa + \eta^2 \\ &= \|\mathbf{x}_t - \mathbf{x}^*\|^2 - \epsilon^2/\kappa^2.\end{aligned}$$

Thus, after T iterations for which $f(\mathbf{x}_t) - f(\mathbf{x}^*) > \epsilon$ we get

$$0 \leq \|\mathbf{x}_{T+1} - \mathbf{x}^*\|^2 \leq \|\mathbf{x}_1 - \mathbf{x}^*\|^2 - T\epsilon^2/\kappa^2,$$

Therefore, we must have $T \leq \kappa^2\|\mathbf{x}_1 - \mathbf{x}^*\|^2/\epsilon^2$. □

4.1 Locally-Lipschitz/Smooth Quasi-Convex Optimization

It can be shown that strict-quasi-convexity and $(G, \epsilon/G, \mathbf{x}^*)$ -local-Lipschitzness of f implies that f is $(\epsilon, G, \mathbf{x}^*)$ -SLQC $\forall \mathbf{x} \in \mathbb{R}^d, \forall \epsilon \geq 0$, and $\mathbf{x}^* \in \arg \min_{\mathbf{x} \in \mathbb{R}^d} f(\mathbf{x})$ (see Appendix E). Therefore the following is a direct corollary of Theorem 4.1:

Corollary 4.1. *Fix $\epsilon > 0$, let $f : \mathbb{R}^d \mapsto \mathbb{R}$, and $\mathbf{x}^* \in \arg \min_{\mathbf{x} \in \mathbb{R}^d} f(\mathbf{x})$. Given that f is strictly quasi-convex and $(G, \epsilon/G, \mathbf{x}^*)$ -locally-Lipschitz. Then running the NGD algorithm with $T \geq G^2\|\mathbf{x}_1 - \mathbf{x}^*\|^2/\epsilon^2$, and $\eta = \epsilon/G$, we have that: $f(\bar{\mathbf{x}}_T) - f(\mathbf{x}^*) \leq \epsilon$.*

In case f is also locally-smooth, we state an even faster rate:

Theorem 4.2. *Fix $\epsilon > 0$, let $f : \mathbb{R}^d \mapsto \mathbb{R}$, and $\mathbf{x}^* \in \arg \min_{\mathbf{x} \in \mathbb{R}^d} f(\mathbf{x})$. Given that f is strictly quasi-convex and $(\beta, \sqrt{2\epsilon/\beta}, \mathbf{x}^*)$ -locally-smooth. Then running the NGD algorithm with $T \geq \beta\|\mathbf{x}_1 - \mathbf{x}^*\|^2/2\epsilon$, and $\eta = \sqrt{2\epsilon/\beta}$, we have that: $f(\bar{\mathbf{x}}_T) - f(\mathbf{x}^*) \leq \epsilon$.*

We prove Theorem 4.2 in Appendix F.

Remark 1. *The above corollary (resp. theorem) implies that f could have arbitrarily large gradients and second derivatives outside $\mathbb{B}(\mathbf{x}^*, \epsilon/G)$ (resp. $\mathbb{B}(\mathbf{x}^*, \sqrt{2\epsilon/\beta})$), yet NGD is still ensured to output an ϵ -optimal point within $G^2\|\mathbf{x}_1 - \mathbf{x}^*\|^2/\epsilon^2$ (resp. $\beta\|\mathbf{x}_1 - \mathbf{x}^*\|^2/2\epsilon$) iterations. We are not familiar with a similar guarantee for GD even in the convex case.*

5 SNGD for Stochastic SLQC Optimization

Here we describe the setting of stochastic SLQC optimization. Then we describe our SNGD algorithm which is ensured to yield an ϵ -optimal solution within $\text{poly}(1/\epsilon)$ queries. We also show that the (noisy) GLM problem, described in Section 3.1.2 is an instance of stochastic SLQC optimization, allowing us to provably solve this problem within $\text{poly}(1/\epsilon)$ samples and optimization steps using SNGD.

Algorithm 2 Stochastic Normalized Gradient Descent (SNGD)

Input: #Iterations T , $\mathbf{x}_1 \in \mathbb{R}^d$, learning rate η , minibatch size b

for $t = 1 \dots T$ **do**

Sample: $\{\psi_i\}_{i=1}^b \sim \mathcal{D}^b$, and define,

$$f_t(\mathbf{x}) = \frac{1}{b} \sum_{i=1}^b \psi_i(\mathbf{x})$$

Update:

$$\mathbf{x}_{t+1} = \mathbf{x}_t - \eta \hat{g}_t \quad \text{where } g_t = \nabla f_t(\mathbf{x}_t), \hat{g}_t = \frac{g_t}{\|g_t\|}$$

end for

Return: $\bar{\mathbf{x}}_T = \arg \min_{\{\mathbf{x}_1, \dots, \mathbf{x}_T\}} f_t(\mathbf{x}_t)$

The stochastic SLQC optimization Setup: Consider the problem of minimizing a function $f : \mathbb{R}^d \mapsto \mathbb{R}$, and assume there exists a distribution over functions \mathcal{D} , such that:

$$f(\mathbf{x}) := \mathbf{E}_{\psi \sim \mathcal{D}}[\psi(\mathbf{x})].$$

We assume that we may access f by randomly sampling minibatches of size b , and querying the gradients of these minibatches. Thus, upon querying a point $\mathbf{x}_t \in \mathbb{R}^d$, a random minibatch $\{\psi_i\}_{i=1}^b \sim \mathcal{D}^b$ is sampled, and we receive $\nabla f_t(\mathbf{x}_t)$, where $f_t(\mathbf{x}) = \frac{1}{b} \sum_{i=1}^b \psi_i(\mathbf{x})$. We make the following assumption regarding the minibatch averages:

Assumption 5.1. *Let $T, \epsilon, \delta > 0$, $\mathbf{x}^* \in \arg \min_{\mathbf{x} \in \mathbb{R}^d} f(\mathbf{x})$. There exists $\kappa > 0$, and a function $b_0 : \mathbb{R}^3 \mapsto \mathbb{R}$, that for $b \geq b_0(\epsilon, \delta, T)$ then w.p. $\geq 1 - \delta$ and $\forall t \in [T]$, the minibatch average $f_t(\mathbf{x}) = \frac{1}{b} \sum_{i=1}^b \psi_i(\mathbf{x})$ is $(\epsilon, \kappa, \mathbf{x}^*)$ -SLQC in x_t . Moreover, we assume $|f_t(\mathbf{x})| \leq M$, $\forall t \in [T], \mathbf{x} \in \mathbb{R}^d$.*

Note that we assume that $b_0 = \text{poly}(1/\epsilon, \log(T/\delta))$.

Justification of Assumption 5.1 Noisy GLM regression (see Section 3.1.2), is an interesting instance of stochastic optimization problem where Assumption 5.1 holds. Indeed according to Lemma 3.2, given $\epsilon, \delta, T > 0$, then for $b \geq \Omega(\log(T/\delta)/\epsilon^2)$ samples², the average minibatch function is $(\epsilon, \kappa, \mathbf{x}^*)$ -SLQC in x_t , $\forall t \in [T]$, w.p. $\geq 1 - \delta$.

Local-quasi-convexity of minibatch averages is a plausible assumption when we optimize an expected sum of quasi-convex functions that share common global minima (or when the different global minima are close by). As seen from the Examples presented in Equation (2),

²In fact, Lemma 3.2 states that for $b = \Omega(\log(1/\delta)/\epsilon^2)$, then the error function is SLQC in a *single decision point*. Using the union bound we can show that for $b = \Omega(\log(T/\delta)/\epsilon^2)$ it holds for T decision points

and in Sections 3.1.1, 3.1.2, this sum is generally not quasi-convex, but is more often locally-quasi-convex.

Note that in the general case when the objective is a sum of quasi-convex functions, the number of local minima of such objective may grow *exponentially* with the dimension d , see Auer et al. (1996). This might imply that a general setup where each $\psi \sim \mathcal{D}$ is quasi-convex may be generally hard.

5.1 Main Results

SNGD is presented in Algorithm 2. SNGD is similar to SGD, except we normalize the gradients. The normalization is crucial in order to take advantage of the SLQC assumption, and in order to overcome the hurdles of plateaus and cliffs. Following is our main theorem:

Theorem 5.1. *Fix $\delta, \epsilon, G, M, \kappa > 0$. Suppose we run SNGD with $T \geq \kappa^2 \|\mathbf{x}_1 - \mathbf{x}^*\|^2 / \epsilon^2$ iterations, $\eta = \epsilon / \kappa$, and $b \geq \max\{\frac{M^2 \log(4T/\delta)}{2\epsilon^2}, b_0(\epsilon, \delta, T)\}$. Assume that for $b \geq b_0(\epsilon, \delta, T)$ then w.p. $\geq 1 - \delta$ and $\forall t \in [T]$, the function f_t defined in the algorithm is M -bounded, and is also $(\epsilon, \kappa, \mathbf{x}^*)$ -SLQC in \mathbf{x}_t . Then, with probability of at least $1 - 2\delta$, we have that $f(\bar{\mathbf{x}}_T) - f(\mathbf{x}^*) \leq 3\epsilon$.*

We prove of Theorem 5.1 at the end of this section.

Remark 2. *Since strict-quasi-convexity and $(G, \epsilon/G, \mathbf{x}^*)$ -local-Lipschitzness are equivalent to SLQC (App. E), the theorem implies that f could have arbitrarily large gradients outside $\mathbb{B}(\mathbf{x}^*, \epsilon/G)$, yet SNGD is still ensured to output an ϵ -optimal point within $G^2 \|\mathbf{x}_1 - \mathbf{x}^*\|^2 / \epsilon^2$ iterations. We are not familiar with a similar guarantee for SGD even in the convex case.*

Remark 3. *Theorem 5.1 requires the minibatch size to be $\Omega(1/\epsilon^2)$. In the context of learning, the number of functions, n , corresponds to the number of training examples. By standard sample complexity bounds, n should also be order of $1/\epsilon^2$. Therefore, one may wonder, if the size of the minibatch should be order of n . This is not true, since the required training set size is $1/\epsilon^2$ times the VC dimension of the hypothesis class. In many practical cases, the VC dimension is more significant than $1/\epsilon^2$, and therefore n will be much larger than the required minibatch size. The reason our analysis requires a minibatch of size $1/\epsilon^2$, without the VC dimension factor, is because we are just “validating” and not “learning”.*

In SGD and for the case of convex functions, even a minibatch of size 1 suffices for guaranteed convergence. In contrast, for SNGD we require a minibatch of size $1/\epsilon^2$. The theorem below shows that the requirement for a large minibatch is not an artifact of our analysis but is truly required.

Theorem 5.2. *Let $\epsilon \in (0, 0.1]$; There exists a distribution over convex loss functions, such that running SNGD with minibatch size of $b = \frac{0.2}{\epsilon}$, with a high probability, it never reaches an ϵ -optimal solution*

We prove Theorem 5.2 in Section 5.2.3. The gap between the upper bound of $1/\epsilon^2$ and the lower bound of $1/\epsilon$ remains as an open question.

We now provide a sketch for the proof of Theorem 5.1:

Proof of Theorem 5.1. Theorem 5.1 is a consequence of the following two lemmas. In the first we show that whenever all f_t 's are SLQC, there exists some t such that $f_t(\mathbf{x}_t) - f_t(\mathbf{x}^*) \leq \epsilon$. In the second lemma, we show that for a large enough minibatch size b , then for any $t \in [T]$ we have $f(\mathbf{x}_t) \leq f_t(\mathbf{x}_t) + \epsilon$, and $f(\mathbf{x}^*) \geq f_t(\mathbf{x}^*) - \epsilon$. Combining these two lemmas we conclude that $f(\bar{\mathbf{x}}_T) - f(\mathbf{x}^*) \leq 3\epsilon$.

Lemma 5.1. *Let $\epsilon, \delta > 0$. Suppose we run SNGD for $T \geq \kappa^2 \|\mathbf{x}_1 - \mathbf{x}^*\|^2 / \epsilon^2$ iterations, $b \geq b_0(\epsilon, \delta, T)$, and $\eta = \epsilon / \kappa$. Assume that w.p. $\geq 1 - \delta$ all f_t 's are $(\epsilon, \kappa, \mathbf{x}^*)$ -SLQC in x_t , whenever $b \geq b_0(\epsilon, \delta, T)$. Then w.p. $\geq 1 - \delta$ we must have some $t \in [T]$ for which $f_t(\mathbf{x}_t) - f_t(\mathbf{x}^*) \leq \epsilon$.*

Lemma 5.1 is proved similarly to Theorem 4.1, we defer the proof to Section 5.2.1.

The second Lemma relates $f_t(\mathbf{x}_t) - f_t(\mathbf{x}^*) \leq \epsilon$ to a bound on $f(\mathbf{x}_t) - f(\mathbf{x}^*)$.

Lemma 5.2. *Suppose $b \geq \frac{M^2 \log(4T/\delta)}{2} \epsilon^{-2}$ then w.p. $\geq 1 - \delta$ and for every $t \in [T]$:*

$$f(\mathbf{x}_t) \leq f_t(\mathbf{x}_t) + \epsilon, \quad \text{and also,} \quad f(\mathbf{x}^*) \geq f_t(\mathbf{x}^*) - \epsilon.$$

Lemma 5.2 is a direct consequence of Hoeffding's bound (see Section 5.2.2). Using the definition of $\bar{\mathbf{x}}_T$ (Alg. 2), together with Lemma 5.2 gives:

$$f(\bar{\mathbf{x}}_T) - f(\mathbf{x}^*) \leq f_t(\mathbf{x}_t) - f_t(\mathbf{x}^*) + 2\epsilon, \quad \forall t \in [T]$$

Combining the latter with Lemma 5.1, establishes Theorem 5.1. □

5.2 Remaining Proofs

5.2.1 Proof of Lemma 5.1

Proof. First note that if the gradient of f_t vanishes at \mathbf{x}_t , then by the SLQC assumption we must have that $f_t(\mathbf{x}_t) - f_t(\mathbf{x}^*) \leq \epsilon$. Assume next that we perform T iterations and the gradient of f_t at \mathbf{x}_t never vanishes in these iterations. Consider the update rule of SNGD (Algorithm 2), then by standard algebra we get:

$$\|\mathbf{x}_{t+1} - \mathbf{x}^*\|^2 = \|\mathbf{x}_t - \mathbf{x}^*\|^2 - 2\eta \langle \hat{g}_t, \mathbf{x}_t - \mathbf{x}^* \rangle + \eta^2.$$

Assume that $\forall t \in [T]$ we have $f_t(\mathbf{x}_t) - f_t(\mathbf{x}^*) > \epsilon$. Take $\mathbf{y} = \mathbf{x}^* + (\epsilon / \kappa) \hat{g}_t$, and observe that $\|\mathbf{y} - \mathbf{x}^*\| \leq \epsilon / \kappa$. Hence the $(\epsilon, \kappa, \mathbf{x}^*)$ -SLQC assumption implies that $\langle \hat{g}_t, \mathbf{y} - \mathbf{x}_t \rangle \leq 0$, thus,

$$\langle \hat{g}_t, \mathbf{x}^* + (\epsilon / \kappa) \hat{g}_t - \mathbf{x}_t \rangle \leq 0 \Rightarrow \langle \hat{g}_t, \mathbf{x}_t - \mathbf{x}^* \rangle \geq \epsilon / \kappa.$$

This implies that, if we set $\eta = \epsilon / \kappa$ then

$$\begin{aligned} \|\mathbf{x}_{t+1} - \mathbf{x}^*\|^2 &\leq \|\mathbf{x}_t - \mathbf{x}^*\|^2 - 2\eta\epsilon / \kappa + \eta^2 \\ &= \|\mathbf{x}_t - \mathbf{x}^*\|^2 - \epsilon^2 / \kappa^2. \end{aligned}$$

So, after T iterations for which $f_t(\mathbf{x}_t) - f_t(\mathbf{x}^*) > \epsilon$ we get

$$0 \leq \|\mathbf{x}_{T+1} - \mathbf{x}^*\|^2 \leq \|\mathbf{x}_1 - \mathbf{x}^*\|^2 - T\epsilon^2/\kappa^2,$$

Therefore, we must have

$$T \leq \frac{\kappa^2 \|\mathbf{x}_1 - \mathbf{x}^*\|^2}{\epsilon^2}.$$

□

5.2.2 Proof of Lemma 5.2

Proof. At each step t , the minibatch is being sampled after \mathbf{x}_t and \mathbf{x}^* are fixed. The random variables $f_t(\mathbf{x}_t)$ (resp. $f_t(\mathbf{x}^*)$) are an average of b i.i.d. random variables whose expectation is $f(\mathbf{x}_t)$ (resp. $f(\mathbf{x}^*)$). These random variables are bounded, since we assume $\forall t, \mathbf{x}, |f_t(\mathbf{x})| \leq M$ (see Thm. 5.1). Applying Hoeffding's bound to the b random samples mentioned above, together with the union bound over $t \in [T]$, and over both sequences of random variables, the lemma follows. □

5.2.3 Proof of Theorem 5.2

We will require the following lemma, whose proof is given in App. G.

Lemma 5.3 (Absorb probabilities). *Let $\{X_t\}_{t=1}^\infty$ be a Markov chain over states $\{i\}_{i=0}^\infty$, such that 0 is an absorbing state, and the transition distribution elsewhere is as follows:*

$$X_{t+1}|\{X_t = i\} = \begin{cases} i - 1 & \text{w.p. } p \\ i + 1 & \text{w.p. } 1 - p \end{cases}$$

Define the absorb probabilities $\alpha_i := P(\exists t > 0 : X_t = 0 | X_0 = i)$, then:

$$\alpha_i = \left(\frac{p}{1-p}\right)^i, \quad \forall i \geq 1$$

Proof. To prove Theorem 5.2, we construct a counter example in one dimension. Consider the following distribution \mathcal{D} over loss functions:

$$f(\mathbf{x}) = \begin{cases} -0.5\epsilon\mathbf{x} & \text{w.p. } 1 - \epsilon \\ (1 - 0.5\epsilon) \max\{\mathbf{x} + 3, 0\} & \text{w.p. } \epsilon \end{cases} \quad (4)$$

It can be verified that the optimum of $\mathbf{E}_{\mathcal{D}}[f(\mathbf{x})]$ is in $\mathbf{x}^* = -3$, and that the slope of the expected loss in $(-3, \infty)$ is 0.5ϵ . Also notice that all points in the segment $[-5, -1]$ are ϵ -optimal.

Suppose we use SNGD with a batchsize of $b = \frac{0.2}{\epsilon}$, i.e., we sample the gradient b times at any query point, and then average the samples and take the sign. Assume that at time t the queried point is greater than $\mathbf{x}^* = -3$. Let Y_t be the averaged gradient over the

batch received at time t , and define $p = P(Y_t \geq 0)$, i.e., the probability that this sign is non-negative. Then the following is a lower bound on $1 - p$:

$$1 - p := P(Y_t < 0) \geq (1 - \epsilon)^b = (1 - \epsilon)^{0.2/\epsilon},$$

where $(1 - \epsilon)^b$ is the probability that all b samples are negative. Now, consider the function $G(\epsilon) = (1 - \epsilon)^{0.2/\epsilon}$, It can be shown that G is monotonically decreasing in $[0, 1]$, and that $G(0.1) \geq 0.8$. Therefore, for any $\epsilon \in [0, 0.1]$ we have, $p \leq 0.2$.

Now, let $\{X_t\}_{t \in [T]}$, be the random variables describing the queries of SNGD under the distribution over loss functions given in Equation (4). Also assume that we start SNGD with $X_1 = 0$, i.e., at a distance of $D = 3$ from the optimum. Then the points that SNGD queries are on the discrete lattice $\{i\eta\}_{i \in \mathcal{Z}}$, and the following holds:

$$X_{t+1} | \{X_t = i\eta\} = \begin{cases} (i - 1)\eta & \text{w.p. } p \\ (i + 1)\eta & \text{w.p. } 1 - p \end{cases}$$

Let $i_0 = \lceil -1/\eta \rceil$, note that i_0 is the minimal number of steps required by SNGD to arrive from $X_1 = 0$, to an ϵ -optimal solution. Now in order to analyze the probability that SNGD ever arrives at an ϵ -optimal point, it is sufficient to consider the Markov chain over the lattice $\{i\eta\}_{i \in \mathcal{Z}}$ with the state $S_0 = i_0\eta$, as an absorbing state. Using Lemma 5.3 we conclude that if we start at $X_1 = 0$ then the probability that we ever absorb is:

$$\begin{aligned} P(\exists t > 0 : X_t \text{ is } \epsilon\text{-optimal} | X_0 = 0) &\leq \left(\frac{p}{1-p}\right)^{\frac{1}{\eta}-1} \\ &\leq \left(\frac{1}{4}\right)^{\frac{1}{\epsilon}-1} \\ &\leq \left(\frac{1}{4}\right)^9, \end{aligned}$$

where we used $p \leq 0.2$, a bound of $G = 1$ on the gradients of losses; NGD's learning rate $\eta = \epsilon/G$, and $\epsilon \leq 0.1$. \square

6 Experiments

A better understanding of how to train deep neural networks is one of the greatest challenges in current machine learning and optimization. Since learning NN (Neural Network) architectures essentially requires to solve a hard non-convex program, we have decided to focus our empirical study on this type of tasks. As a test case, we train a Neural Network with a single hidden layer of 100 units over the MNIST data set. We use a ReLU activation function, and minimize the square loss. We employ a regularization over weights with a parameter of $\lambda = 5 \cdot 10^{-4}$.

At first we were interested in comparing the performance of SNGD to MSGD (Minibatch Stochastic Gradient Descent), and to a stochastic variant of Nesterov's accelerated gradient

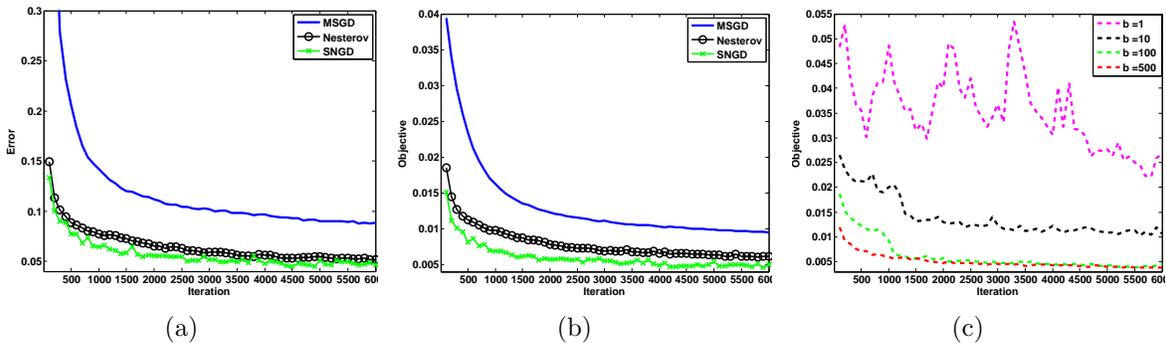


Figure 2: Comparison between optimizations schemes. Left: test error. Middle: objective value (on training set). On the Right we compare the objective of SNGD for different minibatch sizes.

method [Sutskever et al. \(2013\)](#), which is considered to be state-of-the-art. For MSGD and Nesterov’s method we used a step size rule of the form $\eta_t = \eta_0(1 + \gamma t)^{-3/4}$, with $\eta_0 = 0.01$ and $\gamma = 10^{-4}$. For SNGD we used the constant step size of 0.1. In Nesterov’s method we used a momentum of 0.95. The comparison appears in Figures 2(a),2(b). As expected, MSGD converges relatively slowly. Conversely, the performance of SNGD is comparable with Nesterov’s method. All methods employed a minibatch size of 100.

Later, we were interested in examining the effect of minibatch size on the performance of SNGD. We employed SNGD with different minibatch sizes. As seen in Figure 2(c), the performance improves significantly with the increase of minibatch size.

7 Discussion

We have presented the first provable gradient-based algorithm for stochastic quasi-convex optimization. This is a first attempt at generalizing the well-developed machinery of stochastic convex optimization to the challenging non-convex problems facing machine learning, and better characterizing the border between NP-hard non-convex optimization and tractable cases such as the ones studied herein.

Amongst the numerous challenging questions that remain, we note that there is a gap between the upper and lower bound of the minibatch size sufficient for SNGD to provably converge.

References

P. Auer, M. Herbster, and M. K. Warmuth. Exponentially many local minima for single neurons. *Advances in neural information processing systems*, pages 316–322, 1996.

- Y. Bengio. Learning deep architectures for AI. *Foundations and trends in Machine Learning*, 2(1):1–127, 2009.
- S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge university press, 2004.
- K. Doya. Bifurcations of recurrent neural networks in gradient descent learning. *IEEE Transactions on neural networks*, 1:75–80, 1993.
- J.-L. Goffin, Z.-Q. Luo, and Y. Ye. Complexity analysis of an interior cutting plane method for convex feasibility problems. *SIAM Journal on Optimization*, 6(3):638–652, 1996.
- A. T. Kalai and R. Sastry. The isotron algorithm: High-dimensional isotonic regression. In *COLT*, 2009.
- Q. Ke and T. Kanade. Quasiconvex optimization for robust geometric reconstruction. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 29(10):1834–1847, 2007.
- R. F. Khabibullin. A method to find a point of a convex set. *Issled. Prikl. Mat.*, 4:15–22, 1977.
- K. C. Kiwiel. Convergence and efficiency of subgradient methods for quasiconvex minimization. *Mathematical programming*, 90(1):1–25, 2001.
- I. V. Konnov. On convergence properties of a subgradient method. *Optimization Methods and Software*, 18(1):53–62, 2003.
- J.-J. Laffont and D. Martimort. *The theory of incentives: the principal-agent model*. Princeton university press, 2009.
- J. Martens and I. Sutskever. Learning recurrent neural networks with hessian-free optimization. In *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, pages 1033–1040, 2011.
- P. McCullagh and J. Nelder. *Generalised linear models*. London: Chapman and Hall/CRC, 1989.
- Y. E. Nesterov. Minimization methods for nonsmooth convex and quasiconvex functions. *Matekon*, 29:519–531, 1984.
- R. Pascanu, T. Mikolov, and Y. Bengio. On the difficulty of training recurrent neural networks. In *Proceedings of The 30th International Conference on Machine Learning*, pages 1310–1318, 2013.
- B. T. Polyak. A general method of solving extremum problems. *Dokl. Akademii Nauk SSSR*, 174(1):33, 1967.
- J. Sikorski. Quasi subgradient algorithms for calculating surrogate constraints. In *Analysis and algorithms of optimization problems*, pages 203–236. Springer, 1986.

- I. Sutskever, J. Martens, G. Dahl, and G. Hinton. On the importance of initialization and momentum in deep learning. In *Proceedings of the 30th International Conference on Machine Learning (ICML-13)*, pages 1139–1147, 2013.
- H. R. Varian. Price discrimination and social welfare. *The American Economic Review*, pages 870–875, 1985.
- E. Wolfstetter. *Topics in microeconomics: Industrial organization, auctions, and incentives*. Cambridge University Press, 1999.
- Y. I. Zaboltn, A. Korablev, and R. F. Khabibullin. The minimization of quasicomplex functionals. *Izv. Vyssh. Uch. Zaved. Mat.*, (10):27–33, 1972.

A Equivalence Between Definitions 2.2 and 2.3

First let us show that 2.2 \Rightarrow 2.3

Proof of 2.2 \Rightarrow 2.3. Let $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$ such that $f(\mathbf{x}), f(\mathbf{y}) \leq \lambda$. Let $[\mathbf{x}, \mathbf{y}]$ be the line segment connecting these points; we need to show that $\forall \mathbf{z} \in [\mathbf{x}, \mathbf{y}]; f(\mathbf{z}) \leq \lambda$. Assume by contradiction that there exists $\mathbf{z} \in [\mathbf{x}, \mathbf{y}]$ such that $f(\mathbf{z}) > \lambda$. Assume w.l.o.g. that $\langle \nabla f(\mathbf{z}), \mathbf{x} - \mathbf{y} \rangle \neq 0$ (otherwise we could always find $\mathbf{z}' \in [\mathbf{x}, \mathbf{y}]$ such that $f(\mathbf{z}') = f(\mathbf{z})$ and $\langle \nabla f(\mathbf{z}'), \mathbf{x} - \mathbf{y} \rangle \neq 0$), and let $\alpha \in (0, 1)$ such that $\mathbf{z} = \alpha \mathbf{x} + (1 - \alpha) \mathbf{y}$. By Definition 2.2 the following applies:

$$\begin{aligned} 0 &\leq \langle \nabla f(\mathbf{z}), \mathbf{z} - \mathbf{x} \rangle = \langle \nabla f(\mathbf{z}), \alpha \mathbf{x} + (1 - \alpha) \mathbf{y} - \mathbf{x} \rangle = (1 - \alpha) \langle \nabla f(\mathbf{z}), \mathbf{y} - \mathbf{x} \rangle, \\ 0 &\leq \langle \nabla f(\mathbf{z}), \mathbf{z} - \mathbf{y} \rangle = \langle \nabla f(\mathbf{z}), \alpha \mathbf{x} + (1 - \alpha) \mathbf{y} - \mathbf{y} \rangle = -\alpha \langle \nabla f(\mathbf{z}), \mathbf{y} - \mathbf{x} \rangle. \end{aligned}$$

Since $\alpha \in (0, 1)$, we conclude that $\langle \nabla f(\mathbf{z}), \mathbf{y} - \mathbf{x} \rangle \geq 0$ and also $\langle \nabla f(\mathbf{z}), \mathbf{y} - \mathbf{x} \rangle \leq 0$. This is a contradiction since we assumed $\langle \nabla f(\mathbf{z}), \mathbf{y} - \mathbf{x} \rangle \neq 0$. \square

Let us now show that 2.3 \Rightarrow 2.2

Proof of 2.3 \Rightarrow 2.2. Consider the 1 dimensional function $h(\alpha) = f(\mathbf{x} + \alpha(\mathbf{y} - \mathbf{x}))$. The derivative of h at 0 is $h'(0) = \langle \nabla f(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle$. Therefore, we need to show that if $\mathbf{y} \in S_f(\mathbf{x})$ then $h'(0) \leq 0$. By the quasi-convex assumption we have that all the line segment connecting \mathbf{x} to \mathbf{y} is in $S_f(\mathbf{x})$. Therefore, for every $\alpha \in [0, 1]$ we have $h(\alpha) \leq h(0)$. This means that

$$h'(0) = \lim_{\alpha \rightarrow +0} \frac{h(\alpha) - h(0)}{\alpha} \leq 0.$$

\square

B Local Quasi-convexity of g

Here we show that the function g that appears in Equation 2 is SLQC. Denote $\mathbf{x}^* = (-10, -10)$, let $\epsilon \in [0, 1]$ and let \mathbf{x}, \mathbf{y} such that $g(\mathbf{x}) - g(\mathbf{x}^*) \geq \epsilon$ and $\|\mathbf{y} - \mathbf{x}^*\| \leq \epsilon$. In order to prove SLQC it sufficient to show that $\|g_{\mathbf{x}}\| > 0$, and that $\langle g_{\mathbf{x}}, \mathbf{x} - \mathbf{y} \rangle \geq 0$ (we denote $g_{\mathbf{x}} := \nabla g(\mathbf{x})$). Deriving g at x we have:

$$g_{\mathbf{x}} = \nabla g(\mathbf{x}) = (e^{-x_1}/(1 + e^{-x_1})^2, e^{-x_2}/(1 + e^{-x_2})^2).$$

and it is clear that $\|g_{\mathbf{x}}\| > 0, \forall \mathbf{x} \in [-10, 10]^2$, thus strictness always holds. We divide the proof of $\langle g_{\mathbf{x}}, \mathbf{x} - \mathbf{y} \rangle \geq 0$, into cases:

Case 1: Suppose that $x_1 \leq 0, x_2 \leq 0$. In this case it is possible to show that the Hessian of g is positive-semi-definite, thus g is *convex* in $[-10, 0]^2$. Since g is also 1-Lipschitz, then it implies that it is $(\epsilon, 1, \mathbf{x}^*)$ -SLQC in every $\mathbf{x} \in [-10, 0]^2$.

Case 2: Suppose that at least one of x_1, x_2 is positive, w.l.o.g. assume that $x_1 > 0$. In this case:

$$\begin{aligned}
\langle g_{\mathbf{x}}, \mathbf{x} - \mathbf{y} \rangle &\geq \sum_{i=1}^2 \frac{e^{-x_i}(x_i + 10 + (x_i^* - y_i))}{(1 + e^{-x_i})^2} \\
&\geq \sum_{i=1}^2 \frac{e^{-x_i}(x_i + 10 - \epsilon)}{(1 + e^{-x_i})^2} \\
&\geq \frac{19e^{-10}}{(1 + e^{-10})^2} - \frac{\epsilon e^{-10+\epsilon}}{(1 + e^{-10+\epsilon})^2} \\
&\geq \frac{e^{-10}}{(1 + e^{-10})^2} (19 - \epsilon e^\epsilon) \\
&> 0.
\end{aligned}$$

where in the second line we used $\|\mathbf{y} - \mathbf{x}^*\| \leq \epsilon$. In the third line we used $\epsilon \in [0, 1]$, also $10 = \arg \max_{z \in [0, 10]} \frac{e^{-z}(z+9)}{(1+e^{-z})^2}$, and $\min_{z \in [-10, 0]} \frac{e^{-z}(z+10-\epsilon)}{(1+e^{-z})^2} \geq -\frac{\epsilon e^{-10+\epsilon}}{(1+e^{-10+\epsilon})^2}$. The fourth line uses $e^{-10+\epsilon} \geq e^{-10}$, and the last line uses $\epsilon \leq 1$.

The above two cases establish that g is $(\epsilon, 1, \mathbf{x}^*)$ -SLQC in every $\mathbf{x} \in [-10, 10]^2$, $\epsilon \in (0, 1]$.

C Proof of Lemma 3.1

Proof. Given $\epsilon \geq 0$, we will show that $\widehat{\text{err}}_m$ is $(\epsilon, e^W, \mathbf{w}^*)$ -SLQC at every $\mathbf{w} \in \mathbb{B}(0, W)$. Recall $\phi(z) = (1 + e^{-z})^{-1}$, and consider $\|\mathbf{w}\| \leq W$ such that $\widehat{\text{err}}_m(\mathbf{w}) = \frac{1}{m} \sum_{i=1}^m (y_i - \phi(\langle \mathbf{w}, \mathbf{x}_i \rangle))^2 \geq \epsilon$. Also let \mathbf{v} be a point ϵ/e^W close to the minima \mathbf{w}^* , we therefore have:

$$\begin{aligned}
&\langle \nabla \widehat{\text{err}}_m(\mathbf{w}), \mathbf{w} - \mathbf{v} \rangle \\
&= \frac{2}{m} \sum_{i=1}^m \frac{e^{\langle \mathbf{w}, \mathbf{x}_i \rangle}}{(1 + e^{\langle \mathbf{w}, \mathbf{x}_i \rangle})^2} (\phi(\langle \mathbf{w}, \mathbf{x}_i \rangle) - y_i) (\langle \mathbf{w}, \mathbf{x}_i \rangle - \langle \mathbf{v}, \mathbf{x}_i \rangle) \\
&= \frac{2}{m} \sum_{i=1}^m \frac{e^{\langle \mathbf{w}, \mathbf{x}_i \rangle}}{(1 + e^{\langle \mathbf{w}, \mathbf{x}_i \rangle})^2} (\phi(\langle \mathbf{w}, \mathbf{x}_i \rangle) - \phi(\langle \mathbf{w}^*, \mathbf{x}_i \rangle)) (\langle \mathbf{w}, \mathbf{x}_i \rangle - \langle \mathbf{w}^*, \mathbf{x}_i \rangle + \langle \mathbf{w}^* - \mathbf{v}, \mathbf{x}_i \rangle) \\
&\geq \frac{2}{m} \sum_{i=1}^m \frac{4e^{\langle \mathbf{w}, \mathbf{x}_i \rangle}}{(1 + e^{\langle \mathbf{w}, \mathbf{x}_i \rangle})^2} (\phi(\langle \mathbf{w}, \mathbf{x}_i \rangle) - \phi(\langle \mathbf{w}^*, \mathbf{x}_i \rangle))^2 - \frac{\epsilon e^{-W}}{2} \\
&\geq 2e^{-W} \widehat{\text{err}}_m(\mathbf{w}) - \frac{\epsilon e^{-W}}{2} \\
&> 0.
\end{aligned} \tag{5}$$

In the second line we used $y_i = \phi(\langle \mathbf{w}^*, \mathbf{x}_i \rangle)$, which holds for the idealized setup. In the third line we used the fact that $\phi(z)$ is monotonically increasing and 1/4-Lipschitz, and therefore $(\phi(z) - \phi(z'))(z - z') \geq 4(\phi(z) - \phi(z'))^2$. We also used $|\langle \mathbf{w}^* - \mathbf{v}, \mathbf{x}_i \rangle| \leq \|\mathbf{w}^* - \mathbf{v}\| \cdot \|\mathbf{x}_i\| \leq \epsilon e^{-W}$, and $|\phi(\langle \mathbf{w}, \mathbf{x}_i \rangle) - \phi(\langle \mathbf{w}^*, \mathbf{x}_i \rangle)| \leq 1$; Finally we used $\max_z \frac{e^z}{(1+e^z)^2} \leq 1/4$. The fourth line

uses $\min_{\|z\| \leq W} \frac{4e^z}{(1+e^z)^2} \geq e^{-W}$. The last line follows since we assume $\widehat{\text{err}}_m(\mathbf{w}) \geq \epsilon$. The strictness is immediate since $\langle \nabla \widehat{\text{err}}_m(\mathbf{w}), \mathbf{w} - \mathbf{v} \rangle > 0$, therefore, the above establishes SLQC.

We will now show that $\widehat{\text{err}}_m$ is generally not quasi-convex: Consider the idealized setup with two samples $(\mathbf{x}_1, y_1), (\mathbf{x}_2, y_2)$ where $\mathbf{x}_1 = (0, -\log 4), \mathbf{x}_2 = (-\log 4, 0)$ and $y_1 = y_2 = 1/5$, The error function is therefore:

$$\widehat{\text{err}}_m(\mathbf{w}) = \frac{1}{2} \left(\frac{1}{5} - \frac{1}{1 + e^{-\langle \mathbf{w}, \mathbf{x}_1 \rangle}} \right)^2 + \frac{1}{2} \left(\frac{1}{5} - \frac{1}{1 + e^{-\langle \mathbf{w}, \mathbf{x}_2 \rangle}} \right)^2.$$

and it can be verified that the optimal predictor is $\mathbf{w}^* = (1, 1)$, yielding $\widehat{\text{err}}_m(\mathbf{w}^*) = 0$. Now let $\mathbf{w}_1 = (3, 1), \mathbf{w}_2 = (1, 3)$, it can be shown that $\widehat{\text{err}}_m(\mathbf{w}_1) = \widehat{\text{err}}_m(\mathbf{w}_2) \leq 0.018$, yet $\widehat{\text{err}}_m(\mathbf{w}_1/2 + \mathbf{w}_2/2) \geq 0.019$. Thus $\widehat{\text{err}}_m$ is not quasi-convex. \square

D Proof of Lemma 3.2

Proof. Since we are in the noisy idealized setup, and $\forall i, y_i \in [0, 1]$ the following holds

$$y_i = \phi\langle \mathbf{w}^*, \mathbf{x}_i \rangle + \xi_i.$$

where $\{\xi_i\}_{i=1}^m$ are zero mean, independent and bound random variables, $\forall i \in [m], |\xi_i| \leq 1$. Therefore $\widehat{\text{err}}_m$ can be written as follows:

$$\begin{aligned} \widehat{\text{err}}_m(\mathbf{w}) &= \frac{1}{m} \sum_{i=1}^m (y_i - \phi\langle \mathbf{w}, \mathbf{x}_i \rangle)^2 \\ &= \frac{1}{m} \sum_{i=1}^m (\phi\langle \mathbf{w}^*, \mathbf{x}_i \rangle - \phi\langle \mathbf{w}, \mathbf{x}_i \rangle)^2 + \frac{1}{m} \sum_{i=1}^m \xi_i \theta_i(\mathbf{w}) + c. \end{aligned}$$

where $c = \frac{1}{m} \sum_{i=1}^m \xi_i^2$, and $\theta_i(\mathbf{w}) = 2(\phi\langle \mathbf{w}, \mathbf{x}_i \rangle - \phi\langle \mathbf{w}^*, \mathbf{x}_i \rangle)$. We therefore have:

$$\widehat{\text{err}}_m(\mathbf{w}) - \widehat{\text{err}}_m(\mathbf{w}^*) = \frac{1}{m} \sum_{i=1}^m (\phi\langle \mathbf{w}^*, \mathbf{x}_i \rangle - \phi\langle \mathbf{w}, \mathbf{x}_i \rangle)^2 + \frac{1}{m} \sum_{i=1}^m \xi_i \theta_i(\mathbf{w}).$$

Now fix $\epsilon > 0$, and let \mathbf{w} be a fixed point in $\mathbb{B}(0, W)$ such that $\widehat{\text{err}}_m(\mathbf{w}) - \widehat{\text{err}}_m(\mathbf{w}^*) \geq \epsilon$. Also let \mathbf{v} be a point ϵ/e^W close to the minima \mathbf{w}^* , we therefore have:

$$\begin{aligned}
& \langle \nabla \widehat{\text{err}}_m(\mathbf{w}), \mathbf{w} - \mathbf{v} \rangle \\
&= \frac{2}{m} \sum_{i=1}^m \frac{e^{\langle \mathbf{w}, \mathbf{x}_i \rangle}}{(1 + e^{\langle \mathbf{w}, \mathbf{x}_i \rangle})^2} (\phi \langle \mathbf{w}, \mathbf{x}_i \rangle - y_i) (\langle \mathbf{w}, \mathbf{x}_i \rangle - \langle \mathbf{v}, \mathbf{x}_i \rangle) \\
&= \frac{2}{m} \sum_{i=1}^m \frac{e^{\langle \mathbf{w}, \mathbf{x}_i \rangle}}{(1 + e^{\langle \mathbf{w}, \mathbf{x}_i \rangle})^2} (\phi \langle \mathbf{w}, \mathbf{x}_i \rangle - \phi \langle \mathbf{w}^*, \mathbf{x}_i \rangle - \xi_i) (\langle \mathbf{w}, \mathbf{x}_i \rangle - \langle \mathbf{w}^*, \mathbf{x}_i \rangle + \langle \mathbf{w}^* - \mathbf{v}, \mathbf{x}_i \rangle) \\
&\geq \frac{2}{m} \sum_{i=1}^m \frac{4e^{\langle \mathbf{w}, \mathbf{x}_i \rangle}}{(1 + e^{\langle \mathbf{w}, \mathbf{x}_i \rangle})^2} (\phi \langle \mathbf{w}, \mathbf{x}_i \rangle - \phi \langle \mathbf{w}^*, \mathbf{x}_i \rangle)^2 - \frac{\epsilon e^{-W}}{2} + \frac{1}{m} \sum_{i=1}^m \xi_i \lambda_i(\mathbf{w}) \\
&\geq \frac{2}{m} \sum_{i=1}^m \frac{4e^{\langle \mathbf{w}, \mathbf{x}_i \rangle}}{(1 + e^{\langle \mathbf{w}, \mathbf{x}_i \rangle})^2} [(\phi \langle \mathbf{w}, \mathbf{x}_i \rangle - \phi \langle \mathbf{w}^*, \mathbf{x}_i \rangle)^2 + \xi_i \theta_i(\mathbf{w})] - \frac{\epsilon e^{-W}}{2} + \frac{1}{m} \sum_{i=1}^m \xi_i \tilde{\lambda}_i(\mathbf{w}) \\
&\geq 2e^{-W} (\widehat{\text{err}}_m(\mathbf{w}) - \widehat{\text{err}}_m(\mathbf{w}^*)) - \frac{\epsilon e^{-W}}{2} + \frac{1}{m} \sum_{i=1}^m \xi_i \tilde{\lambda}_i(\mathbf{w}) \\
&\geq \frac{3}{2} \epsilon e^{-W} + \frac{1}{m} \sum_{i=1}^m \xi_i \tilde{\lambda}_i(\mathbf{w}). \tag{6}
\end{aligned}$$

where we denote $\lambda_i(\mathbf{w}) = \frac{8e^{\langle \mathbf{w}, \mathbf{x}_i \rangle}}{(1 + e^{\langle \mathbf{w}, \mathbf{x}_i \rangle})^2} (\langle \mathbf{w}^*, \mathbf{x}_i \rangle - \langle \mathbf{w}, \mathbf{x}_i \rangle)$, and $\tilde{\lambda}_i(\mathbf{w}) = \lambda_i(\mathbf{w}) - \frac{8e^{\langle \mathbf{w}, \mathbf{x}_i \rangle}}{(1 + e^{\langle \mathbf{w}, \mathbf{x}_i \rangle})^2} \theta_i(\mathbf{w})$. The argumentation justifying the above inequalities is the same as is done for Equation (5) (see Appendix C). According to Equation (6), the lemma is established if we can show that

$$\frac{1}{m} \sum_{i=1}^m \xi_i \tilde{\lambda}_i(\mathbf{w}) \geq -\epsilon e^{-W}$$

The $\{\xi_i\}_{i=1}^m$ are zero mean and independent, and $|\xi_i \tilde{\lambda}_i(\mathbf{w})| \leq 4(W + 1)$, thus applying Hoeffding's bound we get that the above does hold for $m \geq \frac{8e^{2W}(W+1)^2}{\epsilon^2} \log(1/\delta)$. Note that in bounding $|\xi_i \tilde{\lambda}_i(\mathbf{w})|$, we used $|\xi_i| \leq 1$, also $\mathbf{w}, \mathbf{w}^* \in \mathbb{B}(0, W)$, and $\max_z \frac{e^z}{(1+e^z)^2} \leq \frac{1}{4}$. \square

E Locally-Lipschitz and Strictly Quasi-Convex are SLQC

In order to show that strictly quasi-convex function which is also $(G, \epsilon/G, \mathbf{x}^*)$ -Lipschitz, is SLQC, we require the following lemma:

Lemma E.1. *Let $\mathbf{z} \in \mathbb{R}^d$, and assume that f is $(G, \epsilon/G, \mathbf{z})$ -Locally-Lipschitz. Then, for every \mathbf{x} with $f(\mathbf{x}) - f(\mathbf{z}) > \epsilon$ we have $B(\mathbf{z}, \epsilon/G) \subseteq S_f(\mathbf{x})$*

Proof. Recall the notation $S_f(\mathbf{x}) = \{\mathbf{y} : f(\mathbf{y}) \leq f(\mathbf{x})\}$. By Lipschitzness, for every $\mathbf{y} \in B(\mathbf{z}, \epsilon/G)$ we have $f(\mathbf{y}) \leq f(\mathbf{z}) + \epsilon$. Combining with the assumption that $f(\mathbf{z}) + \epsilon < f(\mathbf{x})$ we obtain that $\mathbf{y} \in S_f(\mathbf{x})$. \square

Therefore, if $f(\mathbf{x}) - f(\mathbf{x}^*) \geq \epsilon$, then $\forall \mathbf{y} \in \mathbb{B}(\mathbf{x}^*, \epsilon/G)$ it holds that $f(\mathbf{x}) - f(\mathbf{y}) \geq 0$, and since f is strictly quasi-convex, the latter means that $\langle \nabla f(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle \leq 0$, and $\|\nabla f(\mathbf{x})\| > 0$. Thus $(\epsilon, G, \mathbf{x}^*)$ -SLQC is established.

F Proof of Theorem 4.2

The key lemma, that enables us to attain faster rates for smooth functions is the following:

Lemma F.1. *Let \mathbf{x}^* be a global minima of f . Also assume that f is $(\beta, \sqrt{2\epsilon/\beta}, \mathbf{x}^*)$ -locally-smooth. Then, for every \mathbf{x} with $f(\mathbf{x}) - f(\mathbf{x}^*) > \epsilon$ we have $B(\mathbf{x}^*, \sqrt{2\epsilon/\beta}) \subseteq S_f(\mathbf{x})$.*

Proof. Combining the definition of local-smoothness (Def. 2.1) together with $\nabla f(\mathbf{x}^*) = 0$ we get

$$|f(\mathbf{y}) - f(\mathbf{x}^*)| \leq \frac{\beta}{2} \|\mathbf{y} - \mathbf{x}^*\|^2, \quad \forall \mathbf{y} \in \mathbb{B}(\mathbf{x}^*, \sqrt{2\epsilon/\beta})$$

Therefore, for every $\mathbf{y} \in B(\mathbf{x}^*, \sqrt{2\epsilon/\beta})$ we have $f(\mathbf{y}) \leq f(\mathbf{x}^*) + \epsilon$. Combining with the assumption that $f(\mathbf{x}^*) + \epsilon < f(\mathbf{x})$ we obtain that $\mathbf{y} \in S_f(\mathbf{x})$. \square

The proof of Theorem 4.2 follows the same lines as the proof of Theorem 4.1. The main difference is that whenever $f(\mathbf{x}_t) - f(\mathbf{x}^*) \geq \epsilon$, we use Lemma F.1 and quasi-convexity to show that for $\mathbf{y} = \mathbf{x}^* + \sqrt{2\epsilon/\beta}\hat{g}_t$ it follows that

$$\langle \nabla f(\mathbf{x}_t), \mathbf{y} - \mathbf{x}_t \rangle \leq 0 .$$

We therefore omit the details of the proof.

G Proof of Lemma 5.3

Proof. Using the stationarity and Markov property of the chain, we can write the following recursive equations for the absorb probabilities:

$$\alpha_i = (1 - p)\alpha_{i+1} + p\alpha_{i-1}, \quad \forall i > 1 \tag{7}$$

$$\alpha_1 = (1 - p)\alpha_2 + p \tag{8}$$

Lets guess a solution of the form, $\alpha_i = c_0\rho^i$, where ρ is the decay parameter of the absorb probabilities. By inserting this solution into equation (7) we get an equation for ρ :

$$(1 - p)\rho^2 - \rho + p = 0 .$$

And it can be validated that the only nontrivial solution is $\rho = \frac{p}{1-p}$, using the latter ρ in equation (8) we get $c_0 = 1$, and therefore we conclude that:

$$\alpha_i = \left(\frac{p}{1-p}\right)^i, \quad \forall i \geq 1$$

\square

H A Broader Notion of Local-Quasi-Convexity

Definition 3.1 describes a rich family of function, as depicted in Section 3.1.1, and 3.1.2. However, it is clear that it does not capture piecewise constant and quasi-convex functions, such as the zero-one loss, or the Perceptron problem.

In some cases, e.g. the Perceptron problem, we may have an access to a *direction oracle*, $\mathcal{G} : \mathbb{R}^d \mapsto \mathbb{R}^d$. This oracle is a proxy to the gradient, aiming us in a global ascent (descent) direction. Following is a broader definition of locally-quasi-convex functions:

Definition H.1. (*Local-Quasi-Convexity*) Let $\mathbf{x}, \mathbf{z} \in \mathbb{R}^d$, $\kappa, \epsilon > 0$. Also let $\mathcal{G} : \mathbb{R}^d \mapsto \mathbb{R}^d$. We say that $f : \mathbb{R}^d \mapsto \mathbb{R}$ is $(\epsilon, \kappa, \mathbf{z})$ -Strictly-Locally-Quasi-Convex (SLQC) in \mathbf{x} , with respect to the direction oracle \mathcal{G} , if at least one of the following applies:

1. $f(\mathbf{x}) - f(\mathbf{z}) \leq \epsilon$.
2. $\|\mathcal{G}(\mathbf{x})\| > 0$, and for every $\mathbf{y} \in \mathbb{B}(\mathbf{z}, \epsilon/\kappa)$ it holds that $\langle \mathcal{G}(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle \leq 0$.

Thus, Definition 3.1, is a private case of the above, which takes the gradient of f to be the direction oracle. Note that we can show that NGD/SNGD and their guarantees still hold for SLQC functions with a direction oracle. The algorithms and proofs are very similar to the ones that appear in the paper, and we therefore omit the details.

In the following section we illustrate a scenario that fits the above definition.

H.1 The γ -margin Perceptron

In this setup we have a collection of m samples $\{(\mathbf{x}_i, y_i)\}_{i=1}^m \in \mathbb{B}_d \times \{0, 1\}$, and we are guaranteed to have $\mathbf{w}^* \in \mathbb{R}^d$ such that: $y_i \langle \mathbf{w}^*, \mathbf{x}_i \rangle \geq \gamma$, $\forall i \in [m]$. Thus, using the sign of $\langle \mathbf{w}^*, \mathbf{x}_i \rangle$ as a predictor, it classifies all the points correctly (with a margin of γ).

Letting ϕ be the zero-one loss $\phi(z) = \mathbb{1}_{z \geq 0}$, we measure the performance of a predictor $\mathbf{w} \in \mathbb{R}^d$, by the average (square) error over all samples,

$$\widehat{\text{err}}_m(\mathbf{w}) = \frac{1}{m} \sum_{i=1}^m (y_i - \phi(\langle \mathbf{w}, \mathbf{x}_i \rangle))^2. \quad (9)$$

Clearly, the gradients of $\widehat{\text{err}}_m(\mathbf{w})$ vanish almost everywhere. Luckily, from the convergence analysis of the Perceptron (see e.g. Kalai and Sastry (2009)), we know the following to be a direction oracle for $\widehat{\text{err}}_m(\mathbf{w})$:

$$\mathcal{G}(\mathbf{w}) = \frac{1}{m} \sum_{i=1}^m (\phi(\langle \mathbf{w}, \mathbf{x}_i \rangle) - y_i) \mathbf{x}_i. \quad (10)$$

The next lemma states that in the above setup, the error function is SLQC with respect to \mathcal{G} . This implies that Algorithm 1 finds an ϵ -optimal minima of $\widehat{\text{err}}_m(\mathbf{w})$ within $\text{poly}(1/\epsilon)$ iterations.

Lemma H.1. *Consider the γ -margin Perceptron problem. Then the error function appearing in Equation (9) is $(\epsilon, 2/\gamma, \mathbf{w}^*)$ -SLQC in \mathbf{w} , $\forall \epsilon \in (0, 1)$, $\forall \mathbf{w} \in \mathbb{R}^d$ with respect to the direction oracle appearing in Equation (10).*

Proof. Given $\epsilon \in (0, 1)$, we will show that $\widehat{\text{err}}_m$ is $(\epsilon, 2/\gamma, \mathbf{w}^*)$ -SLQC at every $\mathbf{w} \in \mathbb{R}^d$. Consider $\mathbf{w} \in \mathbb{R}^d$ such that $\widehat{\text{err}}_m(\mathbf{w}) = \frac{1}{m} \sum_{i=1}^m (y_i - \phi\langle \mathbf{w}, \mathbf{x}_i \rangle)^2 \geq \epsilon$. Also let \mathbf{v} be a point $\gamma\epsilon/2$ close to the minima \mathbf{w}^* , we therefore have:

$$\begin{aligned}
& \langle \mathcal{G}(\mathbf{w}), \mathbf{w} - \mathbf{v} \rangle \\
&= \frac{1}{m} \sum_{i=1}^m (\phi\langle \mathbf{w}, \mathbf{x}_i \rangle - y_i)(\langle \mathbf{w}, \mathbf{x}_i \rangle - \langle \mathbf{v}, \mathbf{x}_i \rangle) \\
&= \frac{1}{m} \sum_{i=1}^m (\phi\langle \mathbf{w}, \mathbf{x}_i \rangle - \phi\langle \mathbf{w}^*, \mathbf{x}_i \rangle)(\langle \mathbf{w}, \mathbf{x}_i \rangle - \langle \mathbf{w}^*, \mathbf{x}_i \rangle + \langle \mathbf{w}^* - \mathbf{v}, \mathbf{x}_i \rangle) \\
&\geq \frac{\gamma}{m} \sum_{i=1}^m (\phi\langle \mathbf{w}, \mathbf{x}_i \rangle - \phi\langle \mathbf{w}^*, \mathbf{x}_i \rangle)^2 - \gamma\epsilon/2 \\
&\geq \gamma\epsilon - \gamma\epsilon/2 \\
&> 0.
\end{aligned} \tag{11}$$

In the second line we used $y_i = \phi\langle \mathbf{w}^*, \mathbf{x}_i \rangle$, which holds by our assumption on \mathbf{w}^* . In the fourth line we used the fact that $(\phi\langle \mathbf{w}, \mathbf{x}_i \rangle - \phi\langle \mathbf{w}^*, \mathbf{x}_i \rangle)(\langle \mathbf{w}, \mathbf{x}_i \rangle - \langle \mathbf{w}^*, \mathbf{x}_i \rangle) \geq \gamma(\phi\langle \mathbf{w}, \mathbf{x}_i \rangle - \phi\langle \mathbf{w}^*, \mathbf{x}_i \rangle)^2$, which holds since \mathbf{w}^* is a minimizer with a γ -margin. We also used $\widehat{\text{err}}_m(\mathbf{w}) \leq 1$, and $|\langle \mathbf{w}^* - \mathbf{v}, \mathbf{x}_i \rangle| \leq \|\mathbf{w}^* - \mathbf{v}\| \cdot \|\mathbf{x}_i\| \leq \gamma\epsilon/2$. Lastly, we use $\widehat{\text{err}}_m(\mathbf{w}) \geq \epsilon$.

The strictness is immediate since $\langle \mathcal{G}(\mathbf{w}), \mathbf{w} - \mathbf{v} \rangle > 0$, therefore, the above establishes SLQC. \square