

Understanding Notions of Stationarity in Non-Smooth Optimization

Jiajin Li, Anthony Man-Cho So, *Senior Member, IEEE*, and Wing-Kin Ma, *Fellow, IEEE*

Abstract—Many contemporary applications in signal processing and machine learning give rise to structured non-convex non-smooth optimization problems that can often be tackled by simple iterative methods quite effectively. One of the keys to understanding such a phenomenon—and, in fact, one of the very difficult conundrums even for experts—lie in the study of “stationary points” of the problem in question. Unlike smooth optimization, for which the definition of a stationary point is rather standard, there is a myriad of definitions of stationarity in non-smooth optimization. In this article, we give an introduction to different stationarity concepts for several important classes of non-convex non-smooth functions and discuss the geometric interpretations and further clarify the relationship among these different concepts. We then demonstrate the relevance of these constructions in some representative applications and how they could affect the performance of iterative methods for tackling these applications.

Index Terms—non-smooth analysis, subdifferential, stationarity.

I. INTRODUCTION

In recent years, we have witnessed a fast-growing body of literature that utilizes non-convex non-smooth optimization techniques to tackle machine learning and signal processing applications. Although such a development seems to run contrary to the long-held belief that non-convex optimization problems pose serious analytic and algorithmic challenges, it is proven to be practically relevant and opens up an exciting avenue for dealing with contemporary applications. For instance, various low-rank matrix recovery problems admit natural non-convex optimization formulations that can be readily tackled by lightweight first-order methods (e.g., (sub)gradient descent or block coordinate descent) and are more scalable than their convex approximations; see, e.g., [1]–[5]. On the other hand, many modern statistical estimation problems involve non-convex loss functions and/or regularizers. While such problems are non-convex, they possess certain convexity properties (which can be made precise) that can be exploited in computation, and there are algorithms that can compute solutions to these problems with good empirical performance; see, e.g., [6]–[10]. Another example that has drawn immense interest is deep neural networks

with non-smooth activation functions (e.g., the rectified linear unit (ReLU) $x \mapsto \max\{x, 0\}$). To train such networks, one often needs to optimize a loss function that is recursively defined via compositions of linear mappings with nonlinear activation functions. Despite the non-convexity and possible non-smoothness of the loss function, various stochastic algorithms (e.g., stochastic (sub)gradient descent or Adam-type algorithms) for optimizing it can still yield exciting empirical performance on a host of machine learning tasks; see, e.g., [11]–[13]. There are many other applications whose natural optimization formulations are non-convex yet highly structured, such as dictionary learning [14], [15], non-negative matrix factorization [16], [17], and phase retrieval [18], [19]. It is becoming increasingly clear that by carefully exploiting the structure of the non-convex formulation at hand, one can design algorithms that have better empirical performance and runtime than those for solving the corresponding convex approximations.

To better understand such phenomenon, a general approach is to study the “stationary points” of the problem in question and investigate how existing iterative methods behave around these stationary points. For smooth optimization, the definition of a stationary point is rather standard. Indeed, consider the unconstrained minimization problem

$$\inf_{\mathbf{x} \in \mathbb{R}^n} f(\mathbf{x}) \quad (1)$$

with $f : \mathbb{R}^n \rightarrow \mathbb{R}$. Suppose that f is smooth and let $\nabla f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be its gradient. A point $\mathbf{x} \in \mathbb{R}^n$ is said to be stationary if

$$\nabla f(\mathbf{x}) = \mathbf{0},$$

which means that \mathbf{x} is either a local minimum, a local maximum, or a saddle point. However, for non-smooth optimization, one can find a myriad of definitions of a stationary point in the literature; see, e.g., [20], [21] and the references therein. It is far from clear how these different definitions of stationarity are related and, more fundamentally, why they need to be introduced. This not only creates potential confusion among readers but also obscures the nature of the solutions that are being computed by different iterative methods.

In this paper, our main objective is to give an introduction to the theory of subdifferentiation for non-convex non-smooth functions, with a focus on motivating the different constructions of the subdifferential and developing the corresponding stationarity concepts for several important function classes, as well as discussing the geometric interpretations and further clarifying the relationship among the different

This work is supported in part by the Hong Kong Research Grants Council (RGC) General Research Fund (GRF) projects CUHK 14208117 and CUHK 14208819, and in part by the CUHK Research Sustainability of Major RGC Funding Schemes project 3133236.

J. Li and A. M.-C. So is with the Department of Systems Engineering and Engineering Management, The Chinese University of Hong Kong, Shatin, N.T., Hong Kong. E-mail: {jjli, manchoso}@se.cuhk.edu.hk.

W.-K. Ma is with the Department of Electronic Engineering, The Chinese University of Hong Kong, Shatin, N.T., Hong Kong. E-mail: wkma@iee.edu.hk.

constructions. We will also demonstrate the relevance of these constructions in some representative applications and how they could affect the performance of iterative methods for tackling these applications. Readers may just be as intrigued by what classes of iterative algorithms can lead to efficient computation of a stationary point under the aforementioned concepts. Unfortunately, owing to the need for exposition of more sophisticated concepts and also to the page limitation, we decide not to cover algorithms in this introductory article.

II. CONVEX NON-SMOOTH FUNCTIONS

To set the stage for our later developments, let us review the theory of subdifferentiation for convex non-smooth functions. For simplicity, we restrict our discussion to finite-valued convex functions $f : \mathbb{R}^n \rightarrow \mathbb{R}$. Recall that if f is convex and smooth, then its gradient ∇f at $\mathbf{x} \in \mathbb{R}^n$ provides an affine minorant of f at $\mathbf{x} \in \mathbb{R}^n$; i.e.,

$$f(\mathbf{y}) \geq f(\mathbf{x}) + \nabla f(\mathbf{x})^T(\mathbf{y} - \mathbf{x}) \quad \text{for all } \mathbf{y} \in \mathbb{R}^n.$$

In the non-smooth case, a suitable generalization of gradient is the notion of *subgradient*; i.e., a vector $\mathbf{s} \in \mathbb{R}^n$ is a subgradient of f at $\mathbf{x} \in \mathbb{R}^n$ if

$$f(\mathbf{y}) \geq f(\mathbf{x}) + \mathbf{s}^T(\mathbf{y} - \mathbf{x}) \quad \text{for all } \mathbf{y} \in \mathbb{R}^n.$$

Since the subgradient at a point may not be unique, we are led to the notion of *subdifferential*, which is the set

$$\partial f(\mathbf{x}) = \{ \mathbf{s} \in \mathbb{R}^n : f(\mathbf{y}) \geq f(\mathbf{x}) + \mathbf{s}^T(\mathbf{y} - \mathbf{x}) \text{ for all } \mathbf{y} \in \mathbb{R}^n \}. \quad (2)$$

As it turns out, the subdifferential (2) can be constructed by considering the directional derivative of f . Given a point $\mathbf{x} \in \mathbb{R}^n$ and a direction $\mathbf{d} \in \mathbb{R}^n$, the *difference quotient* q of f at \mathbf{x} is defined by

$$t \mapsto q(t) = \frac{f(\mathbf{x} + t\mathbf{d}) - f(\mathbf{x})}{t} \quad \text{for } t > 0. \quad (3)$$

Observe that by the convexity of f , the function q is increasing in t (see, e.g., [22, Chapter 0, Proposition 6.1]) and bounded around 0 (see, e.g., [22, Chapter B, Theorem 3.1.2]). Thus, the *directional derivative* of f at $\mathbf{x} \in \mathbb{R}^n$ in the direction $\mathbf{d} \in \mathbb{R}^n$, which is defined by

$$f'(\mathbf{x}, \mathbf{d}) = \lim_{t \searrow 0} \frac{f(\mathbf{x} + t\mathbf{d}) - f(\mathbf{x})}{t}, \quad (4)$$

exists and is equal to $f'(\mathbf{x}, \mathbf{d}) = \inf_{t > 0} q(t)$. One of the key properties of f' is the following:

Fact 1 ([22, Chapter D, Proposition 1.1.2]) *For any $\mathbf{x} \in \mathbb{R}^n$, the function $\mathbf{d} \mapsto f'(\mathbf{x}, \mathbf{d})$ is finite sublinear (recall that a function $h : \mathbb{R}^n \rightarrow \mathbb{R} \cup \{+\infty\}$ is sublinear if it is convex and satisfies $h(t\mathbf{x}) = t \cdot h(\mathbf{x})$ for all $\mathbf{x} \in \mathbb{R}^n$ and $t > 0$).*

A fundamental result in convex analysis is that there is a correspondence between closed sublinear functions and closed convex sets; see [22, Chapter C]. In particular, upon invoking [22, Chapter C, Theorem 3.1.1], we know that $f'(\mathbf{x}, \cdot)$ is the *support function* of the non-empty closed convex set

$$\partial f(\mathbf{x}) = \{ \mathbf{s} \in \mathbb{R}^n : \mathbf{s}^T \mathbf{d} \leq f'(\mathbf{x}, \mathbf{d}) \text{ for all } \mathbf{d} \in \mathbb{R}^n \}; \quad (5)$$

i.e., $f'(\mathbf{x}, \mathbf{d}) = \sup_{\mathbf{s} \in \partial f(\mathbf{x})} \mathbf{s}^T \mathbf{d}$. By the finiteness of $f'(\mathbf{x}, \cdot)$, the set $\partial f(\mathbf{x})$ is bounded ([22, Chapter C, Proposition 2.1.3]). Hence, $\partial f(\mathbf{x})$ is in fact compact. It can be shown that (5) and (2) describe the same set; see [22, Chapter D, Theorem 1.2.2]. Interestingly, even though we define the set (2) without reference to differentiation, its support function turns out to be the directional derivative $f'(\mathbf{x}, \cdot)$.

In applications we often need to compute an element of the subdifferential of a given function. Let us now give the subdifferentials of some concrete convex functions f .

- **(Smooth function).** Suppose that f is differentiable at \mathbf{x} . Then, $\partial f(\mathbf{x}) = \{ \nabla f(\mathbf{x}) \}$; see [22, Chapter D, Corollary 2.1.4].
- **(Norm).** Let f be a norm on \mathbb{R}^n . Then,

$$\partial f(\mathbf{x}) = \{ \mathbf{s} \in \mathbb{R}^n : \mathbf{s}^T \mathbf{x} = f(\mathbf{x}), f_*(\mathbf{s}) \leq 1 \},$$

where f_* is the dual norm of f defined by $f_*(\mathbf{s}) = \sup_{\mathbf{d} \in \mathbb{R}^n : f(\mathbf{d}) \leq 1} \mathbf{d}^T \mathbf{s}$; see [22, Chapter D, Example 3.1].

In particular, for the ℓ_1 -norm $f(\cdot) = \|\cdot\|_1$, we have $\partial(\|\mathbf{x}\|_1) = \text{Sign}(\mathbf{x})$, where Sign is the element-wise sign function given by

$$[\text{Sign}(\mathbf{x})]_i = \begin{cases} \{x_i/|x_i|\} & \text{if } x_i \neq 0, \\ [-1, 1] & \text{otherwise;} \end{cases}$$

for the ℓ_2 -norm $f(\cdot) = \|\cdot\|_2$, we have

$$\partial(\|\mathbf{x}\|_2) = \begin{cases} \{ \mathbf{x}/\|\mathbf{x}\|_2 \} & \text{if } \mathbf{x} \neq \mathbf{0}, \\ B(\mathbf{0}, 1) & \text{otherwise,} \end{cases}$$

where $B(\mathbf{0}, 1)$ is the unit ball centered at the origin.

- **(Max function).** Suppose that f takes the form $f(\cdot) = \max_{\mathbf{y} \in Y} g(\cdot, \mathbf{y})$, where $Y \subseteq \mathbb{R}^\ell$ is compact and $g : \mathbb{R}^n \times Y \rightarrow \mathbb{R}$ is such that $\mathbb{R}^n \ni \mathbf{x} \mapsto g(\mathbf{x}, \mathbf{y})$ is convex for each $\mathbf{y} \in Y$ and $Y \ni \mathbf{y} \mapsto g(\mathbf{x}, \mathbf{y})$ is continuous for each $\mathbf{x} \in \mathbb{R}^n$. Let $Y(\mathbf{x}) = \{ \mathbf{y} \in Y : f(\mathbf{x}) = g(\mathbf{x}, \mathbf{y}) \}$ be the set of optimal solutions to $\max_{\mathbf{y} \in Y} g(\mathbf{x}, \mathbf{y})$. Then,

$$\partial f(\mathbf{x}) = \text{conv} \left\{ \bigcup_{\mathbf{y} \in Y(\mathbf{x})} \partial g(\mathbf{x}, \mathbf{y}) \right\}; \quad (6)$$

cf. [22, Chapter D, Theorem 4.4.2].

The above result is extremely useful, as many convex functions can be represented as the maximum of a collection of convex functions. For instance, let \mathcal{S}^n denote the set of $n \times n$ real symmetric matrices and consider the largest eigenvalue function $\mathcal{S}^n \ni \mathbf{M} \mapsto \lambda(\mathbf{M})$. By the Courant-Fischer theorem, we have the characterization

$$\lambda(\mathbf{M}) = \max_{\mathbf{u} \in \mathbb{R}^n : \|\mathbf{u}\|_2 = 1} \mathbf{u}^T \mathbf{M} \mathbf{u}.$$

Since the function $\mathbf{M} \mapsto \mathbf{u}^T \mathbf{M} \mathbf{u}$ is linear with gradient $\mathbf{u} \mathbf{u}^T$, it follows from (6) that

$$\partial \lambda(\mathbf{M}) = \text{conv} \{ \mathbf{u} \mathbf{u}^T : \|\mathbf{u}\|_2 = 1, \mathbf{M} \mathbf{u} = \lambda(\mathbf{M}) \mathbf{u} \}.$$

- **(Sum rule).** Suppose that f takes the form $f = \alpha_1 f_1 + \alpha_2 f_2$, where $f_1, f_2 : \mathbb{R}^n \rightarrow \mathbb{R}$ are convex functions and $\alpha_1, \alpha_2 > 0$ are positive scalars. Then, $\partial f = \alpha_1 \partial f_1 + \alpha_2 \partial f_2$; see [22, Chapter D, Theorem 4.1.1].

- **(Composition with affine mapping).** Suppose that f takes the form $f = g \circ A$, where $g : \mathbb{R}^m \rightarrow \mathbb{R}$ is a convex function and $A : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is an affine mapping given by $A(\mathbf{x}) = \mathbf{A}_0 \mathbf{x} + \mathbf{b}$ with $\mathbf{A}_0 \in \mathbb{R}^{m \times n}$ and $\mathbf{b} \in \mathbb{R}^m$. Then,

$$\partial f(\mathbf{x}) = \mathbf{A}_0^T \partial g(A(\mathbf{x})) = \{\mathbf{A}_0^T \mathbf{s} : \mathbf{s} \in \partial g(A(\mathbf{x}))\};$$

see [22, Chapter D, Theorem 4.2.1]. The above result can be viewed as a *chain rule* for subdifferentials. Note that we restrict ourselves to the composition of a convex function with an affine mapping here, as the resulting function is guaranteed to be convex and hence its subdifferential (2) is well defined. To obtain more general chain rules, we need to define a notion of subdifferential for non-convex functions. This will be our objective in subsequent sections.

- **(Indicator).** Although our development so far focuses on finite-valued convex functions, it can be extended to convex functions that take values in $\mathbb{R} \cup \{+\infty\}$. One important example of such functions is the indicator of a closed convex set. Specifically, let $C \subseteq \mathbb{R}^n$ be a closed convex set and define the indicator of C by

$$\mathbb{I}_C(\mathbf{x}) = \begin{cases} 0 & \text{if } \mathbf{x} \in C, \\ +\infty & \text{otherwise.} \end{cases} \quad (7)$$

Using the construction (2) of the subdifferential, it can be verified that

$$\partial \mathbb{I}_C(\mathbf{x}) = \{\mathbf{s} \in \mathbb{R}^n : \mathbf{s}^T(\mathbf{y} - \mathbf{x}) \leq 0 \text{ for all } \mathbf{y} \in C\} \quad (8)$$

if $\mathbf{x} \in C$ and $\partial \mathbb{I}_C(\mathbf{x}) = \emptyset$ otherwise. The set on the right-hand side of (8) is known as the *normal cone* to C at \mathbf{x} and is denoted by $\mathcal{N}_C(\mathbf{x})$. Each element $\mathbf{s} \in \mathcal{N}_C(\mathbf{x})$ is called a *normal direction* to C at \mathbf{x} . The terminology is motivated by the observation that for every $\mathbf{s} \in \mathcal{N}_C(\mathbf{x})$, the set C is completely contained in the halfspace $\{\mathbf{y} \in \mathbb{R}^n : \mathbf{s}^T(\mathbf{y} - \mathbf{x}) \leq 0\}$, whose boundary is a hyperplane that passes through \mathbf{x} and has normal \mathbf{s} ; see the figure below.




Fig. 1. Normal cone of a closed convex set.

Using the notion of subdifferential, we can formulate the optimality condition of the minimization of a convex non-smooth function. Specifically, let $g : \mathbb{R}^n \rightarrow \mathbb{R}$ be a convex function and $C \subseteq \mathbb{R}^n$ be a closed convex set. Consider the problem

$$\inf_{\mathbf{x} \in C} g(\mathbf{x}), \quad (9)$$

which can be put into the form (1) by letting $f = g + \mathbb{I}_C$. We then have the following result:

Fact 2 (cf. [20, Theorem 8.15]) *The following are equivalent:*

- $\bar{\mathbf{x}}$ is an optimal solution to (9).
- $\mathbf{0} \in \partial g(\bar{\mathbf{x}}) + \mathcal{N}_C(\bar{\mathbf{x}})$ (cf. (8)).
- $g'(\bar{\mathbf{x}}, \mathbf{y} - \bar{\mathbf{x}}) \geq 0$ for all $\mathbf{y} \in C$.

Although the main focus of this paper is on notions of stationarity, let us briefly digress and discuss the algorithmic aspects of Problem (9). We say that $\mathbf{d} \in \mathbb{R}^n$ is a *descent direction* of the convex function g at $\bar{\mathbf{x}} \in \mathbb{R}^n$ if there exists a $\bar{t} > 0$ satisfying $g(\bar{\mathbf{x}} + t\mathbf{d}) < g(\bar{\mathbf{x}})$ for all $t \in (0, \bar{t})$. As can be easily verified, this is equivalent to $g'(\bar{\mathbf{x}}, \mathbf{d}) < 0$. In view of Fact 2(c), we are thus motivated to use *feasible descent* methods to solve Problem (9). Roughly speaking, at the current iterate $\mathbf{x}^k \in C$, such methods find a direction \mathbf{d}^k and step size $\alpha_k > 0$ such that the next iterate $\mathbf{x}^{k+1} = \mathbf{x}^k + \alpha_k \mathbf{d}^k$ satisfies $g(\mathbf{x}^{k+1}) < g(\mathbf{x}^k)$ and $\mathbf{x}^{k+1} \in C$. As simple as the above description may seem, there are various subtleties in its implementation. For instance, since $g'(\mathbf{x}^k, \mathbf{d}) < 0$ is equivalent to $\max_{\mathbf{s} \in \partial g(\mathbf{x}^k)} \mathbf{s}^T \mathbf{d} < 0$ (see (5)), one may be tempted to compute the entire subdifferential $\partial g(\mathbf{x}^k)$ in each iteration. However, this could be rather expensive. Moreover, even for the unconstrained minimization of a convex non-smooth function, some natural descent methods (such as a straightforward extension of the steepest descent method for smooth minimization) are not necessarily convergent; see, e.g., [23]. It turns out that the above difficulties can be overcome. We refer the reader to [24] for developments in this direction.

Another idea for solving Problem (9) is to use projected subgradient methods. At the current iterate $\mathbf{x}^k \in C$, such methods proceed by first finding a subgradient $\mathbf{s}^k \in \partial g(\mathbf{x}^k)$ and choosing a step size $\alpha_k > 0$, and then obtaining the next iterate via $\mathbf{x}^{k+1} = \Pi_C(\mathbf{x}^k - \alpha_k \mathbf{s}^k)$, where Π_C is the projector onto C . It should be noted that subgradient methods are generally *not* descent methods. For instance, consider the function $\mathbb{R}^2 \ni (x_1, x_2) \mapsto f(x_1, x_2) = |x_1| + 2|x_2|$, whose contour plot is given in Figure 2. It can be easily seen that $(1, 2) \in \partial f(1, 0)$, but $\mathbf{d} = -(1, 2)$ is not a descent direction.




Fig. 2. Contour plot of $f(x_1, x_2) = |x_1| + 2|x_2|$.

In spite of this, by choosing step sizes that decay at an appropriate rate, it can be shown that subgradient methods will converge to an optimal solution and their convergence

rates can be estimated. We refer the reader to [25], [26] for details.

So far we have only discussed constructions of the subdifferential for convex functions. It should not take long for one to realize that those constructions do not yield much useful information when applied to even some very simple non-convex functions. For instance, if we consider the smooth non-convex function $\mathbb{R} \ni x \mapsto f(x) = -x^2$, then using the definition (2) we have $\partial f(0) = \emptyset$. Another example is the non-smooth non-convex function $\mathbb{R} \ni x \mapsto f(x) = -|x|$, where $\partial f(0) = \{s \in \mathbb{R} : s \geq 1 \text{ and } s \leq -1\} = \emptyset$ according to (2). In view of these examples, it is natural to ask whether one can construct a subdifferential that can better capture the geometry of non-smooth non-convex functions. Before we address this question, let us list some desirable properties that we wish such a *generalized subdifferential* to possess. First, the subdifferential should be a singleton consisting of the gradient (resp. coincide with the usual convex subdifferential) when the function in question is smooth (resp. convex). Second, from a computational point of view, the subdifferential should satisfy some basic calculus rules, particularly the chain rule for composite functions and the sum rule for sum of functions. Without such rules, many concrete non-convex functions that arise in applications cannot be tackled easily. Third, the subdifferential should yield a necessary condition for local optimality; i.e., if f attains a local minimum at \bar{x} , then $\mathbf{0} \in \partial f(\bar{x})$. Fourth, the subdifferential should be *tight*, in the sense that the set $\{\mathbf{x} \in \mathbb{R}^n : \mathbf{0} \in \partial f(\mathbf{x})\}$ of *stationary points* of f should contain as few non-local minima as possible. In summary, we have the following desiderata of a generalized subdifferential:

Desirable Properties of a Generalized Subdifferential

- for smooth f , $\partial f(\mathbf{x}) = \{\nabla f(\mathbf{x})\}$
- coincide with the usual convex subdifferential for convex functions
- basic calculus rules
 - *chain rule*: for $f = g \circ F$, $\partial f(\mathbf{x}) = (JF(\mathbf{x}))^T \partial g(F(\mathbf{x}))$, where JF is the Jacobian of F (see (14) for the definition)
 - *sum rule*: for $f = f_1 + f_2$, $\partial f(\mathbf{x}) = \partial f_1(\mathbf{x}) + \partial f_2(\mathbf{x})$
- necessary condition for local optimality
- tight subdifferential

The above discussion suggests that one can consider an axiomatic approach to constructing subdifferentials with the desired properties for more general functions. Such an approach has been explored, e.g., in [27]. Another approach, which is more geometric in nature and follows our development for convex functions, is to construct a convex set that serves as the generalized subdifferential and take its support function to be the generalized directional derivative. Alternatively, one can define a sublinear function that serves as the generalized directional derivative and take the convex set it supports as the generalized subdifferential. Let us now take this geometric approach as the starting point of our exposition.

III. LOCALLY LIPSCHITZ FUNCTIONS

As we move beyond convex functions, one direction to explore is the class of locally Lipschitz functions. Such a class captures a wide variety of non-convex functions [28] and includes the class of convex functions as a special case [29, Theorem 10.4]. Let us recall the definition:

Definition 1 A function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is *locally Lipschitz* if for any bounded $S \subseteq \mathbb{R}^n$, there exists a constant $L > 0$ such that

$$|f(\mathbf{x}) - f(\mathbf{y})| \leq L \|\mathbf{x} - \mathbf{y}\|_2 \quad \text{for all } \mathbf{x}, \mathbf{y} \in S.$$

By a classic result of Rademacher, a locally Lipschitz function f is differentiable almost everywhere (a.e.) [20, Theorem 9.60]. In particular, every neighborhood of \mathbf{x} contains a point \mathbf{y} for which $\nabla f(\mathbf{y})$ exists, so that there is at least one cluster point due to the Lipschitzian property. This motivates the following construction, which is known as the *Bouligand subdifferential*:

$$\partial_B f(\mathbf{x}) = \left\{ \mathbf{s} \in \mathbb{R}^n : \begin{array}{l} \exists \mathbf{x}^k \rightarrow \mathbf{x}, \nabla f(\mathbf{x}^k) \text{ exists,} \\ \nabla f(\mathbf{x}^k) \rightarrow \mathbf{s} \end{array} \right\}.$$

As a quick illustration, consider the absolute value function $\mathbb{R} \ni x \mapsto f(x) = |x|$. It can be easily verified that $\partial_B f(0) = \{-1, 1\}$. Such an example is instructive, as it highlights two drawbacks of the Bouligand subdifferential. First, the Bouligand subdifferential does not coincide with the usual convex subdifferential when the function in question is convex. Second, the condition $\mathbf{0} \in \partial_B f(\mathbf{x})$ is not even necessary for the local optimality of \mathbf{x} . One possible remedy is to *convexify* the Bouligand subdifferential by considering its convex hull; i.e.,

$$\partial_C f(\mathbf{x}) = \text{conv}(\partial_B f(\mathbf{x})). \quad (10)$$

It can be shown that $\partial_C f(\mathbf{x})$ so defined is a non-empty compact convex set and is called the *Clarke subdifferential* in the literature; see [28, Definition 1.1]. From our discussion in Section II, we know that $\partial_C f(\mathbf{x})$ can also be described by its support function. This leads to the question: What is the support function of $\partial_C f(\mathbf{x})$ when f is locally Lipschitz? The following remarkable result due to Clarke furnishes the answer.

Fact 3 (cf. [28, Proposition 1.4]) Given a point $\mathbf{x} \in \mathbb{R}^n$ and a direction $\mathbf{d} \in \mathbb{R}^n$, the Clarke directional derivative of f at \mathbf{x} in the direction \mathbf{d} is defined by

$$\begin{aligned} f^\circ(\mathbf{x}, \mathbf{d}) &= \limsup_{\mathbf{x}' \rightarrow \mathbf{x}, t \searrow 0} \frac{f(\mathbf{x}' + t\mathbf{d}) - f(\mathbf{x}')}{t} \\ &= \inf_{\substack{\epsilon > 0, \\ \lambda > 0}} \sup_{\substack{\mathbf{x}' \in \mathbf{x} + \epsilon \mathbf{B}(\mathbf{0}, 1), \\ t \in (0, \lambda)}} \frac{f(\mathbf{x}' + t\mathbf{d}) - f(\mathbf{x}')}{t}. \end{aligned} \quad (11)$$

Then, $f^\circ(\mathbf{x}, \cdot)$ is the support function of the set $\partial_C f(\mathbf{x})$ defined in (10); i.e.,

$$f^\circ(\mathbf{x}, \mathbf{d}) = \max_{\mathbf{s} \in \partial_C f(\mathbf{x})} \mathbf{s}^T \mathbf{d}.$$

In particular, we have

$$\partial_C f(\mathbf{x}) = \{ \mathbf{s} \in \mathbb{R}^n : \mathbf{s}^T \mathbf{d} \leq f^\circ(\mathbf{x}, \mathbf{d}) \text{ for all } \mathbf{d} \in \mathbb{R}^n \}, \quad (12)$$

and the function $\mathbf{d} \mapsto f^\circ(\mathbf{x}, \mathbf{d})$ is finite sublinear for all $\mathbf{x} \in \mathbb{R}^n$. Additionally, we have $f'(\mathbf{x}, \mathbf{d}) \leq f^\circ(\mathbf{x}, \mathbf{d})$ if the directional derivative of f exists.

We remark that a locally Lipschitz function may not be directionally differentiable. In other words, the difference quotient in (3) may not have a limit even though it is bounded due to the Lipschitzian property. Here, we give an example to showcase such possibility.

Consider the function

$$\mathbb{R} \ni x \mapsto f(x) = \begin{cases} x \sin(\log(\frac{1}{x})) & \text{if } x > 0, \\ 0 & \text{otherwise.} \end{cases}$$

It is clear that f is smooth on $\mathbb{R} \setminus \{0\}$. Its derivative at any $x > 0$ is given by $f'(x) = \sin(\log(\frac{1}{x})) - \cos(\log(\frac{1}{x}))$, which is bounded by 2. Using this and the structure of f , it can be shown that f is locally Lipschitz. However, the directional derivative of f at $\bar{x} = 0$ does not exist. In fact, the difference quotient $q(t) = \frac{f(t)}{t} = \sin(\log(\frac{1}{t}))$ does not converge, as can be seen by considering the sequence $t_n = e^{-(n+\frac{1}{2})\pi}$ and computing

$$q(t_n) = \sin\left(\left(n + \frac{1}{2}\right)\pi\right) = \begin{cases} 1 & \text{if } n \text{ is even,} \\ -1 & \text{otherwise.} \end{cases}$$

It is instructive to compare the two notions of directional derivatives in (4) and (11) from a geometric point of view. The former considers the variation of f along a ray emanating from \mathbf{x} in the direction \mathbf{d} (i.e., $f(\mathbf{x} + t_k \mathbf{d})$ vs. $f(\mathbf{x})$ with $t_k \searrow 0$), while the latter considers the variation of f in the direction \mathbf{d} for points in the neighborhood of \mathbf{x} (i.e., $f(\mathbf{x}^k + t_k \mathbf{d})$ vs. $f(\mathbf{x}^k)$ with $t_k \searrow 0$ and $\mathbf{x}^k \rightarrow \mathbf{x}$). In particular, the latter is able to explore the behavior of f in a neighborhood of \mathbf{x} rather than just along a ray emanating from \mathbf{x} . Generally, $f^\circ(\mathbf{x}, \mathbf{d})$ is an upper bound on the difference quotient in the neighborhood of \mathbf{x} . As we shall see, such an idea turns out to be very fruitful when studying the local behavior of non-smooth functions.

Our discussion above reveals a fundamental difference in the theory of subdifferentiation for convex functions and non-convex functions. Specifically, in the convex case, subdifferentiation entails linearization of the function at hand; in the non-convex case, subdifferentiation can be seen as a convexification process. This allows the use of concepts from convex analysis to study the subdifferentials of non-convex functions.


Recall that in Section II, we have introduced several properties that the generalized subdifferential should possess. Now, let us check whether the Clarke subdifferential possesses those properties.

- **(Smooth function).** If f is smooth (i.e., continuously differentiable) at \mathbf{x} , then $f^\circ(\mathbf{x}, \mathbf{d}) = \nabla f(\mathbf{x})^T \mathbf{d}$ for all $\mathbf{d} \in \mathbb{R}^n$ and $\partial_C f(\mathbf{x}) = \{\nabla f(\mathbf{x})\}$; see [28, Proposition 1.13].

- **(Convex function).** As mentioned above, convex functions are locally Lipschitz. In this case, the Clarke subdifferential and Clarke directional derivative take on particularly simple forms. Indeed, the Clarke subdifferential coincides with the usual convex subdifferential (2) due to [29, Theorems 17.2 and 25.6]. In addition, the directional derivative of a convex function, which always exists, is equal to the Clarke directional derivative; i.e.,

$$f^\circ(\mathbf{x}, \mathbf{d}) = f'(\mathbf{x}, \mathbf{d}). \quad (13)$$

- **(Sum rule).** The following example demonstrates that the sum rule $\partial_C(f_1 + f_2) = \partial_C f_1 + \partial_C f_2$ does not hold in general. Consider the function $f : \mathbb{R} \rightarrow \mathbb{R}$ given by $f(x) = \max\{x, 0\} + \min\{0, x\}$. Let us compute $\partial_C f_1(0)$, $\partial_C f_2(0)$, and $\partial_C f(0)$:



Observe that

$$\partial_C f(0) = \{1\} \subsetneq \partial_C f_1(0) + \partial_C f_2(0) = [0, 2].$$

The failure of the sum rule is one of the obstacles to computing the Clarke subgradient. Nevertheless, not all is lost, as we still have the following weaker version of the sum rule:


$$\partial_C(f_1 + f_2) \subseteq \partial_C f_1 + \partial_C f_2;$$

see [28, Proposition 1.12].

- **(Tightness).** It is known that if f attains a local minimum at \bar{x} , then $0 \in \partial_C f(\bar{x})$; see [30, Proposition 2.3.2]. By Fact 3, this is equivalent to $f^\circ(\bar{x}, \mathbf{d}) \geq 0$ for all $\mathbf{d} \in \mathbb{R}^n$. However, the Clarke subdifferential may contain stationary points that are not local minima. For instance, consider the function $\mathbb{R} \ni x \mapsto f(x) = -|x|$. It is easy to see that $\partial_C f(0) = [-1, 1]$. It follows that $\bar{x} = 0$ is a stationary point (as $0 \in \partial_C f(0)$). However, the point $\bar{x} = 0$ is clearly not a local minimum (in fact, it is a global maximum). Moreover, observe that the corresponding Clarke directional derivatives are $f^\circ(0, 1) = f^\circ(0, -1) = 1$, which shows that neither $d = 1$ nor $d = -1$ is a descent direction according to Clarke's definition. However, the ordinary directional derivatives exist and are given by $f'(0, 1) = f'(0, -1) = -1$. It follows that both $d = 1$ and $d = -1$ are descent directions. One may argue that the above example is not persuasive enough, as similar phenomena occur in the smooth case (e.g., $\mathbb{R} \ni x \mapsto f(x) = -x^2$). Hence, let us provide another, perhaps more convincing, example:

Consider the function

$$\mathbb{R} \ni x \mapsto f(x) = \begin{cases} x + x^2 \sin(\frac{1}{x}) & \text{if } x > 0, \\ x & \text{otherwise.} \end{cases}$$



For $x > 0$, $f'(x) = 1 + 2x \sin(\frac{1}{x}) - \cos(\frac{1}{x})$ is bounded on compact sets. Using this and the structure of f , it can be shown that f is locally Lipschitz. On one hand, we have $\partial_C f(0) = [0, 2]$, which means that $\bar{x} = 0$ is a stationary point. On the other hand, we have $f'(0, 1) = f'(0, -1) = 1$. Hence, the point $\bar{x} = 0$ is neither a local minimum nor a local maximum.

Observe that in the above examples, the ordinary directional derivative exists but is strictly smaller than the corresponding Clarke directional derivatives (i.e., $f'(x, \mathbf{d}) < f^\circ(x, \mathbf{d})$). This, together with (12), suggests that one may obtain a tighter subdifferential by using other directional derivatives.

In view of the aforementioned drawbacks of the Clarke subdifferential, it is natural to ask whether the notion is useful in applications. As it turns out, the Clarke subdifferential can still be a very powerful tool for studying certain sub-classes of locally Lipschitz functions.

IV. SUBDIFFERENTIALLY REGULAR FUNCTIONS

In this section, we introduce a representative function class called *subdifferentially regular* functions. The Clarke subdifferential for such functions preserves many of the nice properties of the subdifferential for convex functions. This greatly facilitates the manipulation of such functions in computational procedures.

Definition 2 ([30, Definition 2.3.4]) *A locally Lipschitz function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is subdifferentially regular (or simply regular) at $\mathbf{x} \in \mathbb{R}^n$ if for every $\mathbf{d} \in \mathbb{R}^n$, the ordinary directional derivative (4) exists and coincides with the generalized one in (11):*

$$f'(\mathbf{x}, \mathbf{d}) = f^\circ(\mathbf{x}, \mathbf{d}).$$

If f is regular at every $\mathbf{x} \in \mathbb{R}^n$, then we simply say that f is regular.

As a first example, we note that a convex function f is regular. This follows immediately from (13). In this case, we have $\partial_C f = \partial f$. Another important example of a regular function is the max function given by $f = \max_{i \in \{1, \dots, m\}} g_i$, where $g_i : \mathbb{R}^n \rightarrow \mathbb{R}$ ($i = 1, \dots, m$) is smooth; see [20, Example 7.28].

In particular, this implies that a smooth function is regular. Upon letting $I(\mathbf{x}) = \{i \in \{1, \dots, m\} : f(\mathbf{x}) = g_i(\mathbf{x})\}$ be the set of indices whose corresponding functions g_i is active at \mathbf{x} , we have $\partial_C f(\mathbf{x}) = \text{conv}\{\nabla g_i(\mathbf{x}) : i \in I(\mathbf{x})\}$; see [20, Exercise 8.31]. We remark that a similar result holds for max functions involving an infinite collection of smooth functions. The interested reader is referred to [20, Theorem 10.31] for details.

One of the nice properties of regular functions is that they satisfy the following basic calculus rules.

Fact 4 (cf. [20, Theorem 10.6, Corollary 10.9])

- (a) **(Chain Rule).** *Suppose that $f : \mathbb{R}^n \rightarrow \mathbb{R}$ takes the form $f = g \circ F$, where $g : \mathbb{R}^m \rightarrow \mathbb{R}$ is a locally Lipschitz function and $F : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is a smooth mapping. Given a point $\mathbf{x} \in \mathbb{R}^n$, if g is regular at $F(\mathbf{x})$, then f is regular at \mathbf{x} and*

$$\partial_C f(\mathbf{x}) = (JF(\mathbf{x}))^T \partial_C g(F(\mathbf{x})),$$

where JF is the Jacobian of F ; i.e.,

$$JF(\mathbf{x}) = \left[\frac{\partial f_i}{\partial x_j}(\mathbf{x}) \right]_{i,j=1}^{m,n} \in \mathbb{R}^{m \times n}. \quad (14)$$

In particular, for a real-valued function $F : \mathbb{R}^n \rightarrow \mathbb{R}$, the gradient of F at \mathbf{x} is given by $\nabla F(\mathbf{x}) = (JF(\mathbf{x}))^T$.

- (b) **(Sum Rule).** *Suppose that $f = f_1 + f_2 + \dots + f_m$, where $f_i : \mathbb{R}^n \rightarrow \mathbb{R}$ ($i = 1, \dots, m$) are locally Lipschitz functions. Given a point $\mathbf{x} \in \mathbb{R}^n$, if f_1, \dots, f_m are regular at \mathbf{x} , then so is f and*

$$\partial_C f(\mathbf{x}) = \partial_C f_1(\mathbf{x}) + \partial_C f_2(\mathbf{x}) + \dots + \partial_C f_m(\mathbf{x}).$$

We remark that it is possible to develop (possibly weaker) versions of the above calculus rules under weaker assumptions. For instance, a variant of the above chain rule holds in the setting where g is lower semi-continuous¹ and F is a locally Lipschitz mapping (and thus not necessarily smooth), while a variant of the above sum rule holds in the setting where f_1, \dots, f_m are lower semi-continuous. We refer the reader to [20, Theorems 10.6 and 10.49] for details.

To illustrate the usefulness of the above calculus rules, let us turn our attention to another fundamental class of regular functions, namely weakly convex functions. Such functions have recently received much attention, as they arise in many contemporary signal processing and machine learning applications. We begin with the definition.

Definition 3 *A function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is called ρ -weakly convex (with $\rho \geq 0$) if the function $\mathbf{x} \mapsto h(\mathbf{x}) = f(\mathbf{x}) + \frac{\rho}{2} \|\mathbf{x}\|_2^2$ is convex.*

It is immediate from the definition that a convex function is 0-weakly convex. As it turns out, weakly convex functions are locally Lipschitz and regular; see [31, Propositions 4.4 and 4.5]. This implies that the basic calculus rules in Fact 4

¹Recall that a function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is *lower semi-continuous* if $\liminf_{\mathbf{y} \rightarrow \mathbf{x}} f(\mathbf{y}) = f(\mathbf{x})$, or equivalently, the epigraph $\text{epi}(f) = \{(\mathbf{x}, t) \in \mathbb{R}^n \times \mathbb{R} : f(\mathbf{x}) \leq t\}$ of f is closed in $\mathbb{R}^n \times \mathbb{R}$; see [20, Theorem 1.6].

can be applied to weakly convex functions. In particular, we can compute the subdifferential of a weakly convex function f as follows. By definition, the function $\mathbb{R}^n \ni \mathbf{x} \mapsto h(\mathbf{x}) = f(\mathbf{x}) + \frac{\rho}{2} \|\mathbf{x}\|_2^2$ is convex for some $\rho \geq 0$. Using the fact that $\mathbf{x} \mapsto \frac{\rho}{2} \|\mathbf{x}\|_2^2$ is regular and applying the sum rule, we have $\partial_C h(\mathbf{x}) = \partial_C f(\mathbf{x}) + \{\rho \mathbf{x}\}$. Since $\partial_C h$ equals the usual convex subdifferential ∂h of h , we obtain

$$\partial_C f(\mathbf{x}) = \partial h(\mathbf{x}) - \{\rho \mathbf{x}\}.$$

Weakly convex functions are ubiquitous in applications. One prototypical example is the composite function $f = g \circ F$, where $g : \mathbb{R}^m \rightarrow \mathbb{R}$ is convex and Lipschitz continuous on \mathbb{R}^m and $F : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is a smooth map with Lipschitz continuous Jacobian [32]. Note that the chain rule in Fact 4 yields a formula for $\partial_C f$. Below are some concrete examples of such a composite function that arise in applications.

- **(Robust low-rank matrix recovery).** In various signal processing [33] and machine learning [34] applications, a fundamental computational task is to recover a low-rank matrix $\mathbf{X}^* \in \mathbb{R}^{n_1 \times n_2}$ from a small number of noisy linear measurements of the form

$$\mathbf{y} = \mathcal{A}(\mathbf{X}^*) + \mathbf{s}^*,$$

where $\mathcal{A} : \mathbb{R}^{n_1 \times n_2} \rightarrow \mathbb{R}^m$ is a known linear operator, $\mathbf{s}^* \in \mathbb{R}^m$ is a noise vector, and $\mathbf{y} \in \mathbb{R}^m$ is the vector of observed values. For simplicity, let us assume that the ground-truth matrix \mathbf{X}^* is an $n \times n$ symmetric positive semidefinite matrix of rank $r \geq 1$. In the setting where the noise vector represents *outliers* in the measurements, the ℓ_1 -loss function is usually preferred over the ℓ_2 -loss for recovering the ground-truth signal. This gives rise to the following weakly convex formulation for recovering \mathbf{X}^* [5]:

$$\min_{\mathbf{U} \in \mathbb{R}^{n \times r}} f(\mathbf{U}) = \frac{1}{m} \|\mathbf{y} - \mathcal{A}(\mathbf{U}\mathbf{U}^T)\|_1.$$

By applying the chain rule in Fact 4, we can compute

$$\begin{aligned} & \frac{1}{m} [(\mathcal{A}^*(\text{Sign}(\mathcal{A}(\mathbf{U}\mathbf{U}^T) - \mathbf{y})))^T \mathbf{U} \\ & + \mathcal{A}^*(\text{Sign}(\mathcal{A}(\mathbf{U}\mathbf{U}^T) - \mathbf{y}))\mathbf{U}] \subseteq \partial_C f(\mathbf{U}), \end{aligned}$$

where \mathcal{A}^* is the adjoint of \mathcal{A} ; see [5].

- **(Robust sign retrieval).** Phase retrieval is a classic inverse problem that arises in areas such as crystallography [19], optical imaging [35], and audio signal processing [36]. Here, let us consider a real-valued version of the problem, in which we are interested in recovering a vector $\mathbf{x}^* \in \mathbb{R}^n$ from noisy measurements of the form

$$b_i = (\mathbf{a}_i^T \mathbf{x}^*)^2 + s_i^* \quad \text{for } i = 1, \dots, m, \quad (15)$$

where $\mathbf{a}_1, \dots, \mathbf{a}_m \in \mathbb{R}^n$ are measurement vectors, $\mathbf{s}^* \in \mathbb{R}^m$ is the noise vector, and $b_1, \dots, b_m \in \mathbb{R}$ are the observed values. One approach to tackling this problem is to consider the weakly convex formulation

$$\min_{\mathbf{x} \in \mathbb{R}^n} f(\mathbf{x}) = \frac{1}{m} \sum_{i=1}^m |(\mathbf{a}_i^T \mathbf{x})^2 - b_i|,$$

which aims at handling outliers in the measurements [37], [38]. Using the calculus rules in Fact 4, we have

$$\frac{2}{m} \sum_{i=1}^m (\mathbf{a}_i^T \mathbf{x}) \cdot \text{Sign}((\mathbf{a}_i^T \mathbf{x})^2 - b_i) \cdot \mathbf{a}_i \subseteq \partial_C f(\mathbf{x});$$

see [38].

- **(Robust blind deconvolution).** The blind deconvolution problem, which is found in diverse fields such as astronomy [39] and image processing [40], [41], aims to recover a pair of signals in two low-dimensional structured spaces from observations of their noisy pairwise convolutions. Again, let us focus on a real-valued version of this problem for simplicity. Formally, we consider the task of robustly recovering a pair $(\mathbf{w}^*, \mathbf{x}^*) \in \mathbb{R}^{n_1} \times \mathbb{R}^{n_2}$ from m bilinear measurements:

$$b_i = (\mathbf{a}_i^T \mathbf{w}^*)(\mathbf{c}_i^T \mathbf{x}^*) + s_i^* \quad \text{for } i = 1, \dots, m,$$

where $\mathbf{a}_1, \dots, \mathbf{a}_m \in \mathbb{R}^{n_1}$ and $\mathbf{c}_1, \dots, \mathbf{c}_m \in \mathbb{R}^{n_2}$ are measurement vectors, $b_1, \dots, b_m \in \mathbb{R}$ are the observed values, and $\mathbf{s}^* \in \mathbb{R}^m$ is the noise vector. One non-smooth formulation of the problem reads

$$\min_{\mathbf{w} \in \mathbb{R}^{n_1}, \mathbf{x} \in \mathbb{R}^{n_2}} f(\mathbf{w}, \mathbf{x}) = \frac{1}{m} \sum_{i=1}^m |(\mathbf{a}_i^T \mathbf{w})(\mathbf{c}_i^T \mathbf{x}) - b_i|,$$

in which the ℓ_1 -loss promotes strong recovery and stability guarantees under certain statistical assumptions [42]. By invoking the chain rule in Fact 4, we obtain

$$\begin{aligned} & \frac{1}{m} \sum_{i=1}^m \text{Sign}((\mathbf{a}_i^T \mathbf{w})(\mathbf{c}_i^T \mathbf{x}) - b_i) \cdot \\ & \left((\mathbf{c}_i^T \mathbf{x}) \begin{bmatrix} \mathbf{a}_i \\ \mathbf{0} \end{bmatrix} + (\mathbf{a}_i^T \mathbf{w}) \begin{bmatrix} \mathbf{0} \\ \mathbf{c}_i \end{bmatrix} \right) \subseteq \partial_C f(\mathbf{w}, \mathbf{x}); \end{aligned}$$

see [42].

Another illustrative example is given by the family of weakly convex sparse regularizers [43], [44], such as logarithmic sum penalty [45], smoothly clipped absolute deviation (SCAD) [46], and minimax concave penalty (MCP) [47]. These regularizers take the form

$$\mathbb{R}^n \ni \mathbf{x} \mapsto R(\mathbf{x}) = \sum_{i=1}^n \phi(|x_i|),$$

where $\phi : \mathbb{R} \rightarrow \mathbb{R}$ is a non-decreasing concave but weakly convex function. Although we cannot apply the chain rule in Fact 4 directly, by using the fact that the absolute value function is locally Lipschitz, we can still apply an extended version of the chain rule (see [20, Theorem 10.49]) to compute an element of the subdifferential of R . Let us demonstrate this via the following concrete example:

Let $R : \mathbb{R}^n \rightarrow \mathbb{R}$ be the logarithmic sum penalty function; i.e.,

$$R(\mathbf{x}) = \sum_{i=1}^d \log(|x_i| + \theta),$$

where $\theta > 0$ is a smoothing parameter. Consider the

following regularized least-squares regression problem:

$$\min_{\mathbf{x} \in \mathbb{R}^n} f(\mathbf{x}) = \frac{1}{2m} \sum_{i=1}^m (b_i - \mathbf{a}_i^\top \mathbf{x})^2 + \lambda R(\mathbf{x}),$$

where $\lambda \geq 0$ is a regularization parameter. Observe that f is regular, as the sum of regular functions is regular; see Fact 4. By the extended chain rule, any vector $\mathbf{s} \in \mathbb{R}^n$ with

$$s_i \in \frac{\text{Sign}(x_i)}{|x_i| + \theta} \quad \text{for } i = 1, \dots, n$$

satisfies $\mathbf{s} \in \partial_C R(\mathbf{x})$. It is then straightforward to obtain an element of the subdifferential of f via the sum rule in Fact 4.

In all the above examples, the ability to explicitly calculate the subdifferential of the weakly convex objective function at hand makes it possible to use simple subgradient methods to minimize the function. Moreover, if the objective function satisfies a regularity condition called *sharpness*, then a suitably initialized subgradient method with properly chosen step sizes will converge at a linear rate to an optimal solution to the problem [32] (see also [5], [38]). We also refer the reader to [48], which discusses stochastic methods for tackling optimization problems involving weakly convex objective functions, and to [49], which develops Riemannian subgradient-type methods for weakly convex optimization over the Stiefel manifold.

Although the class of weakly convex functions provides a powerful modeling tool for applications in signal processing and machine learning, there are still other widely-used functions that do not belong to this class. Here are two examples.

- **(Canonical robust sign retrieval).** Besides the squared-amplitude measurement model in (15), another measurement model of interest for phase retrieval problems is

$$b_i = |\mathbf{a}_i^\top \mathbf{x}^*| + s_i^* \quad \text{for } i = 1, \dots, m,$$

where $\mathbf{x}^* \in \mathbb{R}^n$ is the signal to be recovered, $\mathbf{a}_1, \dots, \mathbf{a}_m \in \mathbb{R}^n$ are measurement vectors, $\mathbf{s}^* \in \mathbb{R}^m$ is the noise vector, and $b_1, \dots, b_m \in \mathbb{R}$ are the observed values. Such an amplitude measurement model is used, e.g., in optical wavefront reconstruction; see [50] for details. The corresponding robust phase retrieval problem then takes the form

$$\min_{\mathbf{x} \in \mathbb{R}^n} f(\mathbf{x}) = \frac{1}{m} \sum_{i=1}^m ||\mathbf{a}_i^\top \mathbf{x}| - b_i|.$$

The function f is not weakly convex as it is not even subdifferentially regular [51].

- **(Deep Neural Network).** Deep learning is a powerful paradigm in machine learning that allows one to learn a complicated mapping by decomposing it into a series of nested simple mappings, and it has attracted immense interest in various areas of science and engineering [52]. As an illustration, consider a simple prediction problem, in which one is given N observed feature-label pairs $(\mathbf{x}_i, y_i) \in \mathbb{R}^n \times \mathbb{R}$, where $i = 1, \dots, N$, and the goal is to

learn the feature-label relationship. One can model such a relationship using the one-hidden-layer neural network shown in Figure 3. Here, $\mathbf{W} = [\mathbf{w}_1, \dots, \mathbf{w}_k]$ is the





Fig. 3. Illustration of one-hidden-layer neural network.

matrix of weight parameters, where $\mathbf{w}_j \in \mathbb{R}^n$ denotes the weight with respect to the j -th neuron, and $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ is a (typically non-smooth) activation function (e.g., the ReLU function $x \mapsto \max\{x, 0\}$). Using the square-loss function, the weights that best model the relationship in the given feature-label pairs can be found by solving the following optimization problem:

$$\min_{\mathbf{W} \in \mathbb{R}^{n \times k}} f(\mathbf{W}) = \frac{1}{2N} \sum_{i=1}^N \left(\sum_{j=1}^k \sigma(\mathbf{w}_j^\top \mathbf{x}_i) - y_i \right)^2. \quad (16)$$

Unfortunately, neural networks with non-smooth activation functions typically give rise to objective functions that are not subdifferentially regular [53]. For example, consider the instance of Problem (16) in which $N = n = k = 1$, $x_1 = y_1 = 1$, and $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ is the ReLU function (i.e., $\sigma(x) = \max\{x, 0\}$). Then, the objective function in (16) becomes $f(w) = \frac{1}{2}(\max\{w, 0\} - 1)^2$, whose graph is shown below.



It is a simple exercise to show that $f^\circ(0, 1) > 0 > f'(0, 1)$. Hence, by Definition 2, we see that f is not subdifferentially regular. Roughly speaking, the graph of a subdifferentially regular function cannot have “downward-facing cusps” [53].

In view of the above examples, we are naturally interested in developing other sharper generalized subdifferential concepts that can deal with broader function classes.

V. DIRECTIONALLY DIFFERENTIABLE FUNCTIONS

As we have seen in Section III, the Clarke directional derivative f° does not always yield useful information about the descent directions of a function at a given point. For

instance, for a directionally differentiable locally Lipschitz function f with directional derivative f' , we always have $f' \leq f^\circ$ and hence the Clarke subdifferential is in some sense too large; see Fact 3. We circumvent this problem in Section IV by imposing the assumption $f' = f^\circ$ on the functions we consider, thereby leading us to the class of subdifferentially regular functions. In this section, we present another approach, which begins by constructing subdifferentials that are smaller than the Clarke subdifferential and then trying to refine them so that they possess some of the desirable properties mentioned in Section II. One advantage of such an approach is that it allows us to tackle functions that are not necessarily subdifferentially regular.

To begin, consider a directionally differentiable locally Lipschitz function $f : \mathbb{R}^n \rightarrow \mathbb{R}$; i.e., the directional derivative $f'(\mathbf{x}, \mathbf{d})$ exists for all $\mathbf{x} \in \mathbb{R}^n$ and $\mathbf{d} \in \mathbb{R}^n$. Since $f' \leq f^\circ$ by Fact 3, the following set suggests itself as a natural candidate for a subdifferential of f :

$$\widehat{\partial}f(\mathbf{x}) = \{s \in \mathbb{R}^n : s^T \mathbf{d} \leq f'(\mathbf{x}, \mathbf{d}) \text{ for all } \mathbf{d} \in \mathbb{R}^n\}. \quad (17)$$

The set $\widehat{\partial}f(\mathbf{x})$ is known as the *Fréchet subdifferential* and its elements the *Fréchet subgradients* of f at \mathbf{x} . It is immediate from (17) that $\widehat{\partial}f(\mathbf{x}) \subseteq \partial_C f(\mathbf{x})$ for any $\mathbf{x} \in \mathbb{R}^n$ (see (12)), and that the Fréchet subdifferential coincides with the usual convex subdifferential when f is convex (see (5)). In fact, the Fréchet subdifferential is closely related to the convex subdifferential. Specifically, the former can be obtained by using higher-order minorants in the construction (2) of the convex subdifferential (see [20, Exercises 8.4 and 9.15]):

$$\widehat{\partial}f(\mathbf{x}) = \left\{ s \in \mathbb{R}^n : \begin{aligned} f(\mathbf{y}) &\geq f(\mathbf{x}) + s^T(\mathbf{y} - \mathbf{x}) \\ &+ o(\|\mathbf{y} - \mathbf{x}\|_2) \text{ for all } \mathbf{y} \in \mathbb{R}^n \end{aligned} \right\}.$$

The inequality with the little-oh term in the above expression means that

$$\liminf_{\mathbf{y} \rightarrow \mathbf{x}} \frac{f(\mathbf{y}) - f(\mathbf{x}) - s^T(\mathbf{y} - \mathbf{x})}{\|\mathbf{y} - \mathbf{x}\|_2} \geq 0.$$

Moreover, observe that for any $\mathbf{x} \in \mathbb{R}^n$, we have $f'(\mathbf{x}, t\mathbf{d}) = t \cdot f'(\mathbf{x}, \mathbf{d})$ for any $\mathbf{d} \in \mathbb{R}^n$ and $t > 0$. Hence, by [20, Theorem 8.24], the set $\widehat{\partial}f(\mathbf{x})$ is closed and convex. In addition, since $f'(\mathbf{x}, \mathbf{d}) < \infty$ for all $\mathbf{d} \in \mathbb{R}^n$ due to the Lipschitzian property of f , the support function of $\widehat{\partial}f(\mathbf{x})$ is given by $\text{conv}(f'(\mathbf{x}, \cdot))$; i.e.,

$$\text{conv}(f'(\mathbf{x}, \cdot))(\mathbf{d}) = \sup_{s \in \widehat{\partial}f(\mathbf{x})} s^T \mathbf{d},$$

where $\text{conv}(f'(\mathbf{x}, \cdot))$ is the pointwise supremum of all convex functions g satisfying $g(\mathbf{d}) \leq f'(\mathbf{x}, \mathbf{d})$ for all $\mathbf{d} \in \mathbb{R}^n$. We refer the interested reader to [54], [55] for a detailed treatment of the Fréchet subdifferential.

Although the above discussion suggests that the Fréchet subdifferential possesses many attractive properties, it is still rather limited. Consider, for instance, the directionally differentiable Lipschitz function $\mathbb{R} \ni x \mapsto f(x) = -|x|$. Then, a simple calculation yields $\widehat{\partial}f(0) = \emptyset$. In particular, the Fréchet subdifferential can be empty, even at points that could be of interest (in this case, $\bar{x} = 0$ is the global maximum). Moreover, by taking a sequence $x^k \searrow 0$, we have $-1 \in \widehat{\partial}f(x^k)$ for all

k but $-1 \notin \widehat{\partial}f(0)$; i.e., the mapping $\widehat{\partial}f$ is not closed. This shows that the Fréchet subdifferential is not stable with respect to small perturbations of the point in question, which can cause instabilities in computation. One way of addressing this issue is to “close” the mapping $\widehat{\partial}f$ by defining the following *limiting subdifferential* of f :

$$\partial f(\mathbf{x}) = \left\{ s \in \mathbb{R}^n : \begin{aligned} \exists x^k \rightarrow \mathbf{x} \text{ and } s^k \in \widehat{\partial}f(x^k) \\ \text{such that } s^k \rightarrow s \end{aligned} \right\}. \quad (18)$$

However, such a process can destroy the convexity of the resulting set. Indeed, continuing with the example $f(\cdot) = -|\cdot|$, we have $\partial f(0) = \{-1, 1\}$. Still, the limiting subdifferential possesses nice properties and is very useful in formulating optimality conditions for non-smooth optimization problems [20], [21], [54]. As a first illustration, let us present the following result, which establishes the relationship among the three subdifferentials we have introduced so far, namely the Fréchet subdifferential, the limiting subdifferential, and the Clarke subdifferential.

Fact 5 (cf. [20, Theorem 8.6], [21, Theorem 3.57]) *For any locally Lipschitz function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ and $\mathbf{x} \in \mathbb{R}^n$, we have*

$$\widehat{\partial}f(\mathbf{x}) \subseteq \partial f(\mathbf{x}) \subseteq \partial_C f(\mathbf{x}) \quad (19)$$

and $\partial_C f(\mathbf{x}) = \text{conv}(\partial f(\mathbf{x}))$. Moreover, if f is subdifferentially regular at \mathbf{x} (in particular, f is directionally differentiable at \mathbf{x}), then all the above subdifferentials coincide; i.e., $\widehat{\partial}f(\mathbf{x}) = \partial f(\mathbf{x}) = \partial_C f(\mathbf{x})$.




Fig. 4. Relationship among the various subdifferentials.

Note that each of the inclusions in (19) can be strict. Indeed, in our previous example $f(\cdot) = -|\cdot|$, we have $\widehat{\partial}f(0) = \emptyset \subsetneq \partial f(0) = \{-1, 1\} \subsetneq \partial_C f(0) = [-1, 1]$. Fact 5 reveals that the limiting subdifferential is tighter than the Clarke subdifferential. Moreover, when f is regular, the limiting subdifferential inherits all the properties of the Clarke subdifferential discussed in Sections III and IV. In particular, since a convex function f is regular, there is no danger of confusion as to the meaning of ∂f , as the usual convex subdifferential and the limiting subdifferential coincide in this case.

As a further illustration and in preparation for our discussion of optimality conditions of non-smooth optimization problems, let us consider the Fréchet and limiting subdifferentials of the indicator function associated with a closed but not necessarily convex set $C \subseteq \mathbb{R}^n$. Recall the definition of the indicator \mathbb{I}_C of C in (7). Clearly, the indicator needs not be directionally

differentiable or locally Lipschitz. Nevertheless, a formal calculation using the definition of the Fréchet subdifferential in (17) yields

$$\widehat{\partial}\mathbb{I}_C(\mathbf{x}) = \left\{ \mathbf{s} \in \mathbb{R}^n : \begin{array}{l} \mathbf{s}^T(\mathbf{y} - \mathbf{x}) \leq o(\|\mathbf{y} - \mathbf{x}\|_2) \\ \text{for all } \mathbf{y} \in C \end{array} \right\} \quad (20)$$

if $\mathbf{x} \in C$ and $\widehat{\partial}\mathbb{I}_C(\mathbf{x}) = \emptyset$ otherwise. The defining condition of the set on the right-hand side of (20) can also be written as

$$\limsup_{\mathbf{y} \rightarrow \mathbf{x}, \mathbf{y} \in C} \frac{\mathbf{s}^T(\mathbf{y} - \mathbf{x})}{\|\mathbf{y} - \mathbf{x}\|_2} \leq 0.$$

The formula (20) for $\widehat{\partial}\mathbb{I}_C(\mathbf{x})$ is indeed valid and can be established in a rigorous manner [20, Exercise 8.14]. The set on the right-hand side of (20) is called the *Fréchet normal cone* to C at \mathbf{x} and is denoted by $\widehat{\mathcal{N}}_C(\mathbf{x})$; cf. the discussion following (8). Now, using (18) and (20), we can compute the limiting subdifferential of \mathbb{I}_C as

$$\partial\mathbb{I}_C(\mathbf{x}) = \left\{ \mathbf{s} \in \mathbb{R}^n : \begin{array}{l} \exists \mathbf{x}^k \rightarrow \mathbf{x} \text{ and } \mathbf{s}^k \in \widehat{\mathcal{N}}_C(\mathbf{x}^k) \\ \text{such that } \mathbf{s}^k \rightarrow \mathbf{s} \end{array} \right\}; \quad (21)$$

see [20, Definition 6.3 and Exercise 8.14]. Following the terminology used above, the set on the right-hand side of (21) is called the *limiting normal cone* of C at \mathbf{x} and is denoted by $\mathcal{N}_C(\mathbf{x})$. Figures 5 and 6 show the Fréchet and limiting normal cones of two closed non-convex sets. It is worth noting that the two normal cones do not always coincide; see Figure 6, where $\widehat{\mathcal{N}}_C(\mathbf{x})$ consists of the zero vector only and $\mathcal{N}_C(\mathbf{x})$ consists of the two rays emanating from \mathbf{x} . In general, we always have $\widehat{\mathcal{N}}_C(\mathbf{x}) \subseteq \mathcal{N}_C(\mathbf{x})$ [20, Proposition 6.5].




Fig. 5. A closed non-convex set with $\widehat{\mathcal{N}}_C(\mathbf{x}) = \mathcal{N}_C(\mathbf{x})$.




Fig. 6. A closed non-convex set with $\widehat{\mathcal{N}}_C(\mathbf{x}) \subsetneq \mathcal{N}_C(\mathbf{x})$.

A. Concepts of Stationarity

Armed with the above development, we are now ready to address our primary goal of this paper, which is to introduce

and compare different stationarity concepts for non-convex non-smooth optimization problems. To begin, consider Problem (9), where $g : \mathbb{R}^n \rightarrow \mathbb{R}$ is a directionally differentiable locally Lipschitz function and $C \subseteq \mathbb{R}^n$ is a closed set. We say that $\bar{\mathbf{x}} \in \mathbb{R}^n$ is a *directional stationary* (resp. *limiting stationary* and *Clarke stationary*) point of Problem (9) if $\mathbf{0} \in \widehat{\partial}(g + \mathbb{I}_C)(\bar{\mathbf{x}})$ (resp. $\mathbf{0} \in \partial(g + \mathbb{I}_C)(\bar{\mathbf{x}})$ and $\mathbf{0} \in \partial_C(g + \mathbb{I}_C)(\bar{\mathbf{x}})$). The following result gives a necessary condition for local optimality of a feasible solution to Problem (9):

Fact 6 (cf. [20, Theorems 8.15 and 10.1, Corollary 6.29]) *If $\bar{\mathbf{x}}$ is a local minimum of (9), then $\bar{\mathbf{x}}$ is a directional stationary (d -stationary) point of (9). If in addition g and \mathbb{I}_C are regular at $\bar{\mathbf{x}}$, then*

$$f'(\bar{\mathbf{x}}, \mathbf{d}) \geq 0 \quad \text{for all } \mathbf{d} \in \mathcal{N}_C^\circ(\bar{\mathbf{x}}),$$

where

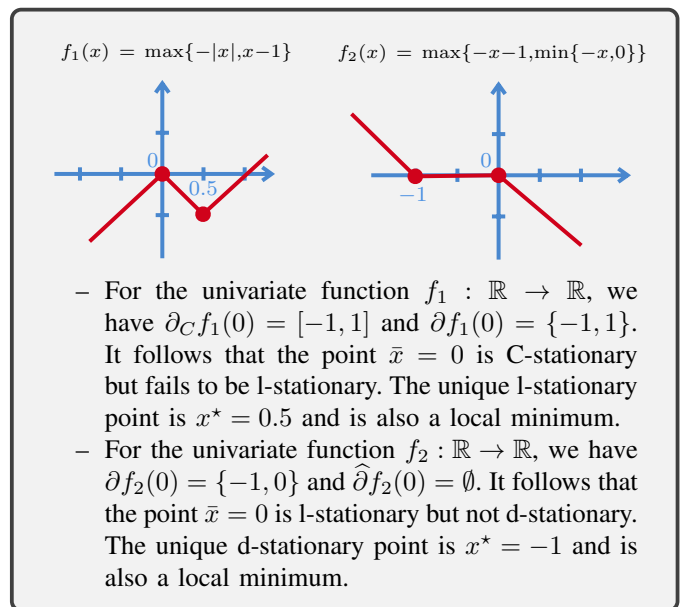
$$\mathcal{N}_C^\circ(\mathbf{x}) = \{ \mathbf{d} \in \mathbb{R}^n : \mathbf{s}^T \mathbf{d} \leq 0 \text{ for all } \mathbf{s} \in \mathcal{N}_C(\mathbf{x}) \}$$

is called the *polar* of $\mathcal{N}_C(\mathbf{x})$.

Note that if $\bar{\mathbf{x}}$ is a d -stationary point of (9), then by Facts 5 and 6 it is also a limiting stationary (l -stationary) and Clarke stationary (C -stationary) point of (9). In particular, we have the following implications:

$$d\text{-stationarity} \implies l\text{-stationarity} \implies C\text{-stationarity}.$$

We now give two examples to show that the reverse implications need not hold in general; see [9].



The above discussion suggests that among the three notions of stationarity, d -stationarity is the sharpest. However, the development of algorithms for computing a d -stationary point of the non-convex non-smooth optimization problem (9) is still in the infancy stage. We will briefly discuss a recent effort in this direction in the next sub-section and refer the reader to [9], [56] for further reading. By contrast, under the assumption that $g + \mathbb{I}_C$ satisfies the so-called *Kurdyka-Łojasiewicz* property, various algorithms will produce iterates

that are provably convergent to a limiting stationary point of (9); see, e.g., [57].

B. Application: Least Squares Piecewise Affine Regression

In this sub-section, we discuss a representative application called *Least Squares Piecewise Affine Regression*, in which the objective function is piecewise linear-quadratic (PLQ) and hence directionally differentiable (see [20, Proposition 10.21]). Specifically, the objective function takes the form

$$\min_{\mathbf{W} \in C} f(\mathbf{W}) = \frac{1}{2N} \sum_{s=1}^N \left(y_s - \max_{1 \leq i \leq k} \mathbf{w}_i^T \mathbf{x}_s \right)^2, \quad (22)$$

where $\mathbf{W} = [\mathbf{w}_1, \dots, \mathbf{w}_k] \in \mathbb{R}^{n \times k}$ is the matrix of decision variables and $C \subseteq \mathbb{R}^{n \times k}$ is the feasible set. By setting $h_s(u) = (y_s - u)^2$ (the square loss) and $g_s(\mathbf{W}) = \max_{1 \leq i \leq k} \mathbf{w}_i^T \mathbf{x}_s$ (a piecewise affine function), we can write the above problem in the following compact form:

$$\min_{\mathbf{W} \in C} f(\mathbf{W}) = \frac{1}{2N} \sum_{s=1}^N h_s(g_s(\boldsymbol{\theta})).$$

The above problem can be used to model the one-layer neural network with the ReLU activation function, in which $k = 1$, $C = \mathbb{R}^n$, and g_s takes the simple form $g_s(\mathbf{w}) = \max\{\mathbf{w}^T \mathbf{x}_s, 0\}$; cf. (16). Our interest in Problem (22) stems from the following:

Fact 7 (cf. [58, Proposition 16]) *The least squares piecewise affine regression problem (22) possesses the following properties:*

- (a) *It attains a finite global minimum value.*
- (b) *The set of d -stationary points is finite.*
- (c) *Every d -stationary point is a local minimizer.*

The above result provides further evidence that the notion of d -stationarity is in some sense the sharpest, as every d -stationary point of Problem (22) is a local minimum. In view of this, it is natural to ask whether we can propose an iterative algorithm to find such points. In [9] the authors proposed a non-monotone majorized-minimization (MM) algorithm with a semi-smooth Newton method as its inner solver to find a d -stationary point of a class of so-called *composite difference-convex-piecewise* optimization problems, of which Problem (22) is an instance. They also showed that the MM algorithm will converge to a d -stationary point of such problems under mild conditions (which are satisfied by (22)). One of the motivations for introducing such an algorithm is that it is not known whether the basic chain rule holds for the objective function in (22). For the purpose of experimentation, let us pretend the basic chain rule holds and use it to compute a pseudo-subgradient (actually back-propagation in deep learning) of the objective function:

$$\frac{\tilde{\partial} f(\mathbf{w}_i)}{\partial \mathbf{w}_i} = \frac{1}{2} \sum_{s=1}^N \left(\max_{1 \leq i \leq k} \mathbf{w}_i^T \mathbf{x}_s - y_s \right) \mathbf{x}_s \mathbb{I}_{\left\{ i \in \arg \max_i \mathbf{w}_i^T \mathbf{x}_s \right\}}.$$

Then, we can try using the subgradient method with such a pseudo-subgradient to tackle Problem (22). However, such an




Fig. 7. Objective values computed by the MM and subgradient algorithms, $N = 10$.

approach does not quite work empirically. Indeed, let us follow the experimental setup in [9] and consider the 2-dimensional convex piecewise linear model

$$y = \max \{x_1 + x_2, x_1 - x_2, -2x_1 + x_2, -2x_1 - x_2\} + \varepsilon$$

with different sample sizes $N = 10, 50, 100$. We test the MM and subgradient algorithms on synthetic data. Using the same initial points for the two algorithms, all the experiment results reported here were collected over 500 independent trials over random seeds. From Figures 7 and 8, we observe




Fig. 8. Number of initial points that lead to the smallest objective values.

that there is an apparent gap between these two algorithms. In particular, the figures show that the subgradient algorithm reaches many limit points that are unsatisfactory. Nevertheless, the MM algorithm can be rather slow. As a future work, it would be interesting to design practically efficient first-order algorithms that can provably return a d -stationary point for this application, and more generally, for other signal processing, machine learning, and statistical applications; see, e.g., [8]–[10], [56], [58] and the references therein.

VI. CONCLUSION

In this article, we elucidated the constructions of various subdifferentials for several important sub-classes of non-

smooth functions and discussed their corresponding stationarity concepts. We also showcased several representative examples and applications to illustrate the differences among various constructions. We hope that this introductory article will serve as a good starting point for readers who would like to utilize the mathematical tools from non-smooth analysis in the design and analysis of iterative methods for non-smooth optimization problems.

REFERENCES

- [1] Y. Koren, R. Bell, and C. Volinsky, "Matrix Factorization Techniques for Recommender Systems," *Computer*, vol. 42, no. 8, pp. 30–37, 2009.
- [2] R. Ge, J. D. Lee, and T. Ma, "Matrix Completion has No Spurious Local Minima," in *Advances in Neural Information Processing Systems 29: Proceedings of the 2016 Conference*, D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, Eds., 2016, pp. 2973–2981.
- [3] S. Tu, R. Boczar, M. Simchowitz, M. Soltanolkotabi, and B. Recht, "Low-Rank Solutions of Linear Matrix Equations via Procrustes Flow," in *Proceedings of the 33rd International Conference on Machine Learning (ICML 2016)*, 2016, pp. 964–973.
- [4] Y. Li, Y. Chi, H. Zhang, and Y. Liang, "Nonconvex Low-Rank Matrix Recovery with Arbitrary Outliers via Median-Truncated Gradient Descent," *Information and Inference: A Journal of the IMA*, p. iaz009, 2019.
- [5] X. Li, Z. Zhu, A. M.-C. So, and R. Vidal, "Nonconvex Robust Low-Rank Matrix Recovery," *SIAM Journal on Optimization*, vol. 30, no. 1, pp. 660–686, 2020.
- [6] P.-L. Loh and M. J. Wainwright, "High-Dimensional Regression with Noisy and Missing Data: Provable Guarantees with Nonconvexity," *The Annals of Statistics*, vol. 40, no. 3, pp. 1637–1664, 2012.
- [7] —, "Regularized M -Estimators with Nonconvexity: Statistical and Algorithmic Theory for Local Optima," *Journal of Machine Learning Research*, vol. 16, no. Mar, pp. 559–616, 2015.
- [8] M. Ahn, J.-S. Pang, and J. Xin, "Difference-of-Convex Learning: Directional Stationarity, Optimality, and Sparsity," *SIAM Journal on Optimization*, vol. 27, no. 3, pp. 1637–1665, 2017.
- [9] Y. Cui, J.-S. Pang, and B. Sen, "Composite Difference-Max Programs for Modern Statistical Estimation Problems," *SIAM Journal on Optimization*, vol. 28, no. 4, pp. 3344–3374, 2018.
- [10] M. Nouiehed, J.-S. Pang, and M. Razaviyayn, "On the Pervasiveness of Difference-Convexity in Optimization and Statistics," *Mathematical Programming, Series B*, vol. 174, no. 1–2, pp. 195–222, 2019.
- [11] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet Classification with Deep Convolutional Neural Networks," in *Advances in Neural Information Processing Systems 25: Proceedings of the 2012 Conference*, P. Bartlett, F. C. N. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, Eds., 2012, pp. 1097–1105.
- [12] D. P. Kingma and J. L. Ba, "Adam: A Method for Stochastic Optimization," in *Proceedings of the 3rd International Conference on Learning Representations (ICLR 2015)*, 2015.
- [13] X. Chen, S. Liu, R. Sun, and M. Hong, "On the Convergence of A Class of Adam-Type Algorithms for Non-Convex Optimization," in *Proceedings of the 7th International Conference on Learning Representations (ICLR 2019)*, 2019.
- [14] J. Sun, Q. Qu, and J. Wright, "Complete Dictionary Recovery Over the Sphere I: Overview and the Geometric Picture," *IEEE Transactions on Information Theory*, vol. 63, no. 2, pp. 853–884, 2017.
- [15] —, "Complete Dictionary Recovery Over the Sphere II: Recovery by Riemannian Trust-Region Method," *IEEE Transactions on Information Theory*, vol. 63, no. 2, pp. 885–914, 2017.
- [16] C.-J. Lin, "Projected Gradient Methods for Nonnegative Matrix Factorization," *Neural Computation*, vol. 19, no. 10, pp. 2756–2779, 2007.
- [17] Y. Li, Y. Liang, and A. Risteski, "Recovery Guarantee of Non-Negative Matrix Factorization via Alternating Updates," in *Advances in Neural Information Processing Systems 29: Proceedings of the 2016 Conference*, D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, Eds., 2016, pp. 4987–4995.
- [18] T. Bendory, Y. C. Eldar, and N. Boumal, "Non-Convex Phase Retrieval from STFT Measurements," *IEEE Transactions on Information Theory*, vol. 64, no. 1, pp. 467–484, 2018.
- [19] V. Elser, T.-Y. Lan, and T. Bendory, "Benchmark Problems for Phase Retrieval," *SIAM Journal on Imaging Sciences*, vol. 11, no. 4, pp. 2429–2455, 2018.
- [20] R. T. Rockafellar and R. J.-B. Wets, *Variational Analysis*, 2nd ed., ser. Grundlehren der mathematischen Wissenschaften. Berlin Heidelberg: Springer-Verlag, 2004, vol. 317.
- [21] B. S. Mordukhovich, *Variational Analysis and Generalized Differentiation I: Basic Theory*, 2nd ed., ser. Grundlehren der mathematischen Wissenschaften. Berlin Heidelberg: Springer-Verlag, 2013, vol. 330.
- [22] J.-B. Hiriart-Urruty and C. Lemaréchal, *Fundamentals of Convex Analysis*, ser. Grundlehren Text Editions. Berlin/Heidelberg: Springer-Verlag, 2001.
- [23] P. Wolfe, "A method of conjugate subgradients for minimizing non-differentiable functions," *Mathematical Programming Study*, vol. 3, pp. 145–173, 1975.
- [24] K. C. Kiwiel, *Methods of Descent for Nondifferentiable Optimization*, ser. Lecture Notes in Mathematics. Berlin Heidelberg: Springer-Verlag, 1985, vol. 1133.
- [25] J. L. Goffin, "On convergence rates of subgradient optimization methods," *Mathematical Programming*, vol. 13, no. 1, pp. 329–347, 1977.
- [26] N. Z. Shor, *Minimization Methods for Non-Differentiable Functions*, ser. Springer Series in Computational Mathematics. Berlin Heidelberg: Springer-Verlag, 1985, vol. 3.
- [27] A. D. Ioffe, "On the theory of subdifferentials," *Advances in Nonlinear Analysis*, vol. 1, no. 1, pp. 47–120, 2012.
- [28] F. H. Clarke, "Generalized Gradients and Applications," *Transactions of the American Mathematical Society*, vol. 205, no. Apr., pp. 247–262, 1975.
- [29] R. T. Rockafellar, *Convex Analysis*, ser. Princeton Landmarks in Mathematics and Physics. Princeton, New Jersey: Princeton University Press, 1997.
- [30] F. H. Clarke, *Optimization and Nonsmooth Analysis*, ser. Classics in Applied Mathematics. Philadelphia, Pennsylvania: Society for Industrial and Applied Mathematics, 1990.
- [31] J.-P. Vial, "Strong and Weak Convexity of Sets and Functions," *Mathematics of Operations Research*, vol. 8, no. 2, pp. 231–259, 1983.
- [32] D. Davis, D. Drusvyatskiy, K. J. MacPhee, and C. Paquette, "Subgradient Methods for Sharp Weakly Convex Functions," *Journal of Optimization Theory and Applications*, vol. 179, no. 3, pp. 962–982, 2018.
- [33] M. A. Davenport and J. Romberg, "An Overview of Low-Rank Matrix Recovery from Incomplete Observations," *IEEE Journal of Selected Topics in Signal Processing*, vol. 10, no. 4, pp. 608–622, 2016.
- [34] N. Srebro, J. D. M. Rennie, and T. S. Jaakkola, "Maximum-Margin Matrix Factorization," in *Advances in Neural Information Processing Systems 17: Proceedings of the 2004 Conference*, L. K. Saul, Y. Weiss, and L. Bottou, Eds., 2004, pp. 1329–1336.
- [35] Y. Shechtman, Y. C. Eldar, O. Cohen, H. N. Chapman, J. Miao, and M. Segev, "Phase Retrieval with Application to Optical Imaging," *IEEE Signal Processing Magazine*, vol. 32, no. 3, pp. 87–109, 2015.
- [36] I. Waldspurger, "Phase Retrieval for Wavelet Transforms," *IEEE Transactions on Information Theory*, vol. 63, no. 5, pp. 2993–3009, 2017.
- [37] J. C. Duchi and F. Ruan, "Solving (Most) of a Set of Quadratic Equalities: Composite Optimization for Robust Phase Retrieval," *Information and Inference: A Journal of the IMA*, vol. 8, no. 3, pp. 471–529, 2019.
- [38] D. Davis, D. Drusvyatskiy, and C. Paquette, "The Nonsmooth Landscape of Phase Retrieval," *IMA Journal of Numerical Analysis*, p. drz031, 2020.
- [39] D. Kundur and D. Hatzinakos, "Blind image deconvolution," *IEEE Signal Processing Magazine*, vol. 13, no. 3, pp. 43–64, 1996.
- [40] T. F. Chan and C.-K. Wong, "Total variation blind deconvolution," *IEEE Transactions on Image Processing*, vol. 7, no. 3, pp. 370–375, 1998.
- [41] A. Levin, Y. Weiss, F. Durand, and W. T. Freeman, "Understanding Blind Deconvolution Algorithms," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 33, no. 12, pp. 2354–2367, 2011.
- [42] V. Charisopoulos, D. Davis, M. Díaz, and D. Drusvyatskiy, "Composite optimization for robust blind deconvolution," 2019, manuscript, available at <https://arxiv.org/abs/1901.01624>.
- [43] X. Shen and Y. Gu, "Nonconvex sparse logistic regression with weakly convex regularization," *IEEE Transactions on Signal Processing*, vol. 66, no. 12, pp. 3199–3211, 2018.
- [44] B. Wen, X. Chen, and T. K. Pong, "A proximal difference-of-convex algorithm with extrapolation," *Computational Optimization and Applications*, vol. 69, no. 2, pp. 297–324, 2018.
- [45] E. J. Candès, M. B. Wakin, and S. P. Boyd, "Enhancing sparsity by reweighted ℓ_1 minimization," *Journal of Fourier Analysis and Applications*, vol. 14, no. 5–6, pp. 877–905, 2008.
- [46] J. Fan and R. Li, "Variable Selection via Nonconcave Penalized Likelihood and Its Oracle Properties," *Journal of the American Statistical Association*, vol. 96, no. 456, pp. 1348–1360, 2001.

- [47] C.-H. Zhang, “Nearly Unbiased Variable Selection under Minimax Concave Penalty,” *The Annals of Statistics*, vol. 38, no. 2, pp. 894–942, 2010.
- [48] D. Davis and D. Drusvyatskiy, “Stochastic model-based minimization of weakly convex functions,” *SIAM Journal on Optimization*, vol. 29, no. 1, pp. 207–239, 2019.
- [49] X. Li, S. Chen, Z. Deng, Q. Qu, Z. Zhu, and A. M.-C. So, “Weakly Convex Optimization over Stiefel Manifold Using Riemannian Subgradient-Type Methods,” 2019, manuscript, available at <https://arxiv.org/abs/1911.05047>.
- [50] D. R. Luke, J. V. Burke, and R. G. Lyon, “Optical wavefront reconstruction: Theory and numerical methods,” *SIAM Review*, vol. 44, no. 2, pp. 169–224, 2002.
- [51] A. Aravkin, J. Burke, and D. He, “On the global minimizers of real robust phase retrieval with sparse noise,” 2019, manuscript, available at <https://arxiv.org/abs/1905.10358>.
- [52] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*, ser. Adaptive Computation and Machine Learning Series. Cambridge, Massachusetts: MIT Press, 2016.
- [53] D. Davis, D. Drusvyatskiy, S. Kakade, and J. D. Lee, “Stochastic subgradient method converges on tame functions,” *Foundations of Computational Mathematics*, vol. 20, no. 1, pp. 119–154, 2020.
- [54] A. Y. Kruger, “On Fréchet Subdifferentials,” *Journal of Mathematical Sciences*, vol. 116, no. 3, pp. 3325–3358, 2003.
- [55] B. S. Mordukhovich, N. M. Nam, and N. D. Yen, “Fréchet subdifferential calculus and optimality conditions in nondifferentiable programming,” *Optimization*, vol. 55, no. 5–6, pp. 685–708, 2006.
- [56] J.-S. Pang, M. Razaviyayn, and A. Alvarado, “Computing B–Stationary Points of Nonsmooth DC Programs,” *Mathematics of Operations Research*, vol. 42, no. 1, pp. 95–118, 2017.
- [57] H. Attouch, J. Bolte, and B. F. Svaiter, “Convergence of Descent Methods for Semi–Algebraic and Tame Problems: Proximal Algorithms, Forward–Backward Splitting, and Regularized Gauss–Seidel Methods,” *Mathematical Programming, Series A*, vol. 137, no. 1–2, pp. 91–129, 2013.
- [58] Y. Cui and J.-S. Pang, “On the Finite Number of Directional Stationary Values of Piecewise Programs,” *arXiv preprint arXiv:1803.00190*, 2018.