Provably Correct Automatic Subdifferentiation for Qualified Programs

Sham Kakade [*]   Jason D. Lee [†]

January 15, 2019

Abstract

The *Cheap Gradient Principle* [Griewank and Walther, 2008] — the computational cost of computing the gradient of a scalar-valued function is nearly the same (often within a factor of 5) as that of simply computing the function itself — is of central importance in optimization; it allows us to quickly obtain (high dimensional) gradients of scalar loss functions which are subsequently used in black box gradient-based optimization procedures. The current state of affairs is markedly different with regards to computing subderivatives: widely used ML libraries, including TensorFlow and PyTorch, do *not* correctly compute (generalized) subderivatives even on simple examples. This work considers the question: is there a *Cheap Subgradient Principle*? Our main result shows that, under certain restrictions on our library of nonsmooth functions (standard in nonlinear programming), provably correct generalized subderivatives can be computed at a computational cost that is within a (dimension-free) factor of 6 of the cost of computing the scalar function itself.

1   Introduction

The widespread implementation of Automatic Differentiation (AD) methods [Baydin et al., 2015] has had a transformative effect on applied machine learning; these methods have eased the difficulty for practitioners, across a range of disciplines, to learn

[*]University of Washington, Seattle, e-mail: sham@cs.washington.edu
[†]University of Southern California, e-mail: jasonlee@marshall.usc.edu

sophisticated machine learning models (including deep neural architectures and richer inferential models). The paradigm is: one simply writes a program to compute the function of interest, say a scalar (loss) function $f(x) : \mathbb{R}^d \to \mathbb{R}$, and then a correctly implemented AD method will return both $f(x)$ and all d of its partial derivatives when provided with x as an input. These partial derivatives are often used in conjunction with some (stochastic) gradient-based optimization approach.

Underlying the effectiveness of this general black-box approach is the *Cheap Gradient Principle* [Griewank and Walther, 2008]: the computational cost of computing the vector of partial derivatives $(\partial f/\partial x_1, \partial f/\partial x_2, \ldots \partial f/\partial x_d)$ is often nearly the same as that of simply computing the scalar function $f(x)$ itself. In fact, for all rational functions, the striking Baur-Strassen theorem [Baur and Strassen, 1983, Griewank, 1989] shows that this increase in computational complexity is a (dimension free) factor of 5.

In many settings, our underlying function $f(x)$ is a nonsmooth function, and we resort to subgradient methods. This work considers the question: is there a *Cheap Subgradient Principle*? Specifically, given a program that computes a (locally Lipschitz) function f and given a point x, can we automatically compute an element of the (Clarke) subdifferential $\partial f(x)$ [Clarke, 1975], and can we do this at a cost which is comparable to computing the function $f(x)$ itself? Informally, the set $\partial f(x)$ is the convex hull of limits of gradients at nearby differentiable points. It can be thought of as generalizing the gradient (for smooth functions) and the subgradient (for convex functions).

Let us briefly consider how current approaches handle nonsmooth functions, which are available to the user as functions in some library. Consider the following three equivalent ways to write the identity function, where $x \in \mathbb{R}$,

$$f_1(x) = x, \quad f_2(x) = \mathrm{ReLU}\,(x) - \mathrm{ReLU}\,(-x), \quad f_3(x) = 10 f_1(x) - 9 f_2(x),$$

where $\mathrm{ReLU}\,(x) = \max\{x, 0\}$, and so $f_1(x) = f_2(x) = f_3(x)$. As these functions are differentiable at 0, the unique derivative is $f_1'(0) = f_2'(0) = f_3'(0) = 1$. However, both TensorFlow [Abadi et al., 2015] and PyTorch [Paszke et al., 2017], claim that $f_1'(0) = 1, f_2'(0) = 0, f_3'(0) = 10$. This particular answer is due to using a subgradient of 0 at $x = 0$. One may ask if a more judicious choice fixes such issues; unfortunately, it is not difficult to see that no such universal choice exists[1].

---

[1] By defining $\mathrm{ReLU}'(0) = 1/2$, the reader may note we obtain the correct derivative on f_2, f_3; however, consider $f_4(x) = \mathrm{ReLU}\,(\mathrm{ReLU}\,(x)) - \mathrm{ReLU}\,(-x)$, which also equals $f_1(x)$. Here, we would need $\mathrm{ReLU}'(0) = \frac{\sqrt{5}-1}{2}$ to obtain the correct answer.

This example should be concerning for a number of reasons. The use of nonsmooth functions in AD go well beyond simple one dimensional nonsmooth functions (such as ReLU $(\cdot)$ or the $|\cdot|$); current methods permit utilizing eigenvalues, SVDs, QR decompositions (there are AD procedures on these nonsmooth linear algebra functions [Maclaurin et al., 2015, Seeger et al., 2017]).

Is correctness important? One option is to disregard these issues — which is the current state of affairs — based on the observation that in most cases these issues are unlikely to harm our optimization method. In numerical linear algebra, one could make the same argument: we never truly encounter degenerate linear systems (or degenerate eigenspaces); nonetheless, in retrospect, numerical issues have made evident the importance of carefully addressing these "corner cases". The situation may be analogous here: numerical issues in these approaches can easily lead to unstable outputs. Note that some numerical instability is certainly to be expected due to nonsmoothness (a point we return to in the Discussion under the notion of *mixed stability*); yet we would still hope to have nontrivial stability guarantees in our widely used AD libraries, much in the manner we have for our established numerical linear algebra libraries [Trefethen and Bau III, 1997, Demmel, 1997].

Ultimately, the importance of correctness in these methods is a decision that must be made by the broader ML community. Here, it is worthwhile to consider that AD software has a range of applications: from physical simulators to health care/social science applications to deployed online learning systems to differentiable programming. For example, when using physical simulators (say in robotics or in the sciences), a strong notion of stability may be critical when doing AD through nonsmooth system dynamics. In safety-critical settings, we may seek to have deployed online learning methods which are not susceptible to errors due to misspecified input-output behavior in our programs. Perhaps the most compelling reason for provably correct software implementations is to avoid costly failure modes due to the utilization of the methods in novel and unforeseen manners.

Related Work: These issues are in fact known in the mathematical AD literature (see Griewank and Walther [2008, Chapter 14]). Once we include either nonsmooth primitive functions or permit branching in a program, the usual chain rule fails to hold and incorrect input-out behavior is easy to observe. Due to that established calculus properties of nonsmooth functions [Klatte and Kummer, 2002, Mordukhovich, 2006] do not seem amenable to AD approaches, the current prov-

able methods do not have general purpose, computationally efficient AD methods for subdifferentials.

One influential and powerful idea is that of lexicographic differentiation [Nesterov, 2005]; it is a property of a subclass of nonsmooth functions which allow these function to inherit a generalized notion of a chain rule. This idea has been utilized for obtaining correct generalized derivatives in Khan and Barton [2013], Griewank [2013]. The difficulty is that lexicographic approach often is expensive in that it involves a dimensional factor in the computational cost increase.

The other relatively few works that do focus on automatic generalized differentiation go through some notion of algorithmic linearization [A.Griewank, 1995, Nesterov, 2005, Khan and Barton, 2013, 2015, Fiege et al., 2017], where often piecewise smooth functions are considered, and the approach attempts at correct AD through probing the pieces through some linearization (see Griewank [2014] for review). The difficulties are due to understanding what information we can extract through linear "probes" into the function.

One of the first ideas along this line of thought is due to [A.Griewank, 1995], which shows how to compute directional derivatives of nonsmooth functions through following a "branch" the program would take on an input (where the branch corresponds to the approach direction in the directional derivative). In fact, our work uses this basic idea, as does the "branch locking" approach in Khan [2017], Griewank [2013]. The difficulty in these approaches is in finding a means to relate this linearization to properties of the (nonsmooth) functions, which will allow the algorithm to succeed; naively, we can tell when a method might have failed though it is difficult to guarantee if it will succeed.

As such, the extant body of work does not contain methods which contain only a constant factor blow up in the computational cost. Notable differences in this work is that our assumptions make strong connections to nonlinear programming [Abadie, 1967, Peterson, 1973, Gould and Tolle, 1971], which help in characterizing when the linearization approach is informative, and we provide a key technical result showing a certain chain rule holds for randomized algorithms. Furthermore, our focus is on generalizing the reverse mode for scalar functions (as opposed to focusing on multivariate functions where there is no known Cheap Gradient Principle).

Our contributions: Our main result provides — under a natural set of assumptions widely used in nonlinear programming — a provably correct Automatic Subdifferentiation procedure, which given some x, computes *both* the functional value

---

Algorithm 1: Straight Line Program for $f(x)$

Input: $x = (x_1, \ldots x_d)$
  1: **for** $k = d+1, d+2, \ldots T$ **do**
  2:     Compute:
$$x_k = g_k(x_{\text{parents}(k)})$$
     where parents(k) is the index set of the "parent" variables of k.
  3: **end for**
Return: x_T.

Algorithm 2: The Reverse Mode of AD

Input: variables $(x_1, \ldots x_T)$; a computational graph $\{\text{children}(t)\}_{t \in \{1, \ldots T\}}$; the associated derivatives
  1: Initialize: $\frac{\partial x_T}{\partial x_T} = 1$
  2: **for** $t = T, T-1, \ldots 1$ **do**
  3:     Compute:
$$\frac{\partial x_T}{\partial x_t} = \sum_{i \in \text{children}(t)} \frac{\partial x_T}{\partial x_i} \frac{\partial x_i}{\partial x_t}$$
  4: **end for**
Return: $\frac{\partial x_T}{\partial x} = \left(\frac{\partial x_T}{\partial x_1}, \frac{\partial x_T}{\partial x_2}, \ldots \frac{\partial x_T}{\partial x_d} \right)$.

$f(x)$ and a d dimensional subdifferential $(u_1, \ldots u_d) \in \partial f(x)$, with a computational cost that is a factor of at most 6 times that of computing the scalar function $f(x)$ itself. Our assumption is that our library of functions be implemented in a manner consistent with the standard *constraint qualification* assumptions in nonlinear programming [Abadie, 1967]. In short, this work shows that in fact there is a *Cheap Subgradient Principle*.

2   Preliminaries

Assume $f: \mathbb{R}^d \to \mathbb{R}$ is a locally Lipschitz function, and recall, that by Rademacher's theorem, this implies that f is differentiable almost everywhere. The *Clarke subdifferential* of f at any point x is the set [Clarke et al., 2008, Theorem 8.1]

$$\partial f(x) := \text{conv} \left\{ \lim_{i \to \infty} \nabla f(x_i) : x_i \xrightarrow{\Omega} x \right\}, \tag{1}$$

5

where Ω is any full-measure subset of $\mathbb{R}^d$ such that f is differentiable at each of its points. Here, the limit is taken to be the set of all limit points. In classical circumstances, the subdifferential reduces to more familiar objects. Namely, when f is C^1-smooth at x, the subdifferential $\partial f(x)$ consists only of the gradient $\nabla f(x)$, while for convex functions, it reduces to the subdifferential in the sense of convex analysis.

2.1 AD Review and The Baur-Strassen Theorem

A *straight line program* for computing $f(x) : \mathbb{R}^d \to \mathbb{R}$ is specified by a program of the form shown in Algorithm 1. Here the functions $g_1, g_2, \ldots$ are assumed to be some function from a library of functions. In the algebraic circuit complexity model, these functions are either monomials or affine functions of its inputs.

More generally, we will be interested in utilizing a richer class of functions where $g \in \mathcal{L}$, a library of functions, e.g. we may desire functions like the $|\cdot|$, ReLU (x), or ever richer nonsmooth functions like eigenvalues.

Define $\mathrm{Runtime}(f; x)$ to be the time it takes to compute $f(x)$ under a given program for f.

Theorem 2.1. *[Baur and Strassen, 1983, Griewank, 1989] Assume all multiplications and additions have unit runtime cost. If we restrict to the algebraic circuit complexity model (where the functions g_k are either monomials or affine functions), then it is possible to compute both $f(x)$ and all its partial derivatives $\nabla f(x)$ in time that is at most $5 * \mathrm{Runtime}(f; x)$.*

An algorithm achieving this guarantee is to first compute $f(x)$ and then use the *reverse mode* of AD, in Algorithm 2. To see the specific counting argument, see [Morgenstern, 1985]. This theorem is often more general: the reverse mode also correctly returns the derivatives even with a richer family of smooth functions in our library $\mathcal{L}$, often with a constant factor cost increase as well [Griewank, 1989]. The reverse mode itself has been rediscovered many times [Griewank, 2012]; the well known back-propagation algorithm [Rumelhart et al., 1986] is one example of the reverse mode of AD. The reverse mode (and the back-propagation algorithm) is not a direct application of the chain rule; the direct application of the chain rule is referred to as the *forward mode* of AD (see Griewank and Walther [2008]), which is d times more expensive procedure to compute the gradient. The reverse mode can be viewed as a form of dynamic programming. To compare the two, in the reverse mode of AD, we compute the derivatives $\frac{\partial x_T}{\partial x_t}$, referred to as the

6

adjoints[2], while in the forward mode of AD we would compute (d-dimensional) derivatives of the form $\frac{\partial x_t}{\partial x}$ (referred to as dual numbers).

2.2 Nonsmooth functions and our computational model

To specify how our nonsmooth functions are implemented, we extend the computational model to allow for branching, using (a restricted version[3] of) the Blum-Shub-Smale model of computation [Blum et al., 1988].

Definition 2.1 (Computation Model)**.** *The computational model for computing any $g(x) : \mathbb{R}^d \to \mathbb{R}$ in our library (d may be different for each function) is specified by a program of the form shown in Algorithm 3. We assume that the function $g_{k,z}$ is either a monomial or an affine function of its inputs. Furthermore, for every g, we assume that there exists a time T, where the program terminates in at most this amount of time.*

Throughout, we make the following assumption:

Assumption 2.1. *(Computational Cost) Assume all multiplications and additions have unit runtime cost and that an execution of an "If" statement is also unit cost. For example, the cost of computing a monomial is the number of multiplications.*

The program implicitly encodes a function that has the following representation:

$$f(x) = \sum_{z \in \{-1,1\}^T} \mathbb{I}_{S_z}(x) p_z(x), \tag{2}$$

where each p_z is a polynomial; $\mathbb{I}_{S_z}$ is the indicator function on the set S_z; and S_z consists of all x where the program executes branch z when given x as input. The set S_z can be explicitly defined as follows: for steps k where the programs branches on z, define $h_{k,z}(x) = x_k$; on non-branching k, define $h_{k,z}(x) = -1$; define the vector valued function $h_z(x) = (h_1(z), \dots h_T(x))$;

$$S_z := \{x \,|\, \mathbb{I}(\text{sign}(h_z(x)) = z)\} \tag{3}$$

---

[2]For a variable $x_T = g(x_{\text{parents}})$, the notation $\frac{\partial x_T}{\partial x_t}$ refers to the derivative with respect to x_t, but holding all parent variables of x_t as fixed. If x_t is an input variable, then this is the usual partial derivative.

[3]We avoid halting concerns by assuming our programs halt in a bounded amount of time. We also explicitly avoid discussing tapes and registers in our computational cost model.

Algorithm 3: Program for a Nonsmooth function $g(x)$

Input: $x = (x_1, \ldots x_d)$

1: Initialize a vector z to be all -1's. z is for notational convenience to keep track of the branch.

2: **for** $k = d + 1, d + 2, \ldots T$ **do**

3:    Compute:

$$x_k = g_{k,z}(x_{\text{parents}(k,z)})$$

4:    If the program branches at (k, z), then

- If: $x_k \geqslant 0$, $z_k = 1$.

- Else: $z_k = -1$.

5:    If the program halts at (k, z), then terminate the **for** loop.

6: **end for**

Return: x_k.

---

where the $\text{sign}(\cdot)$ is the usual sign function (applied componentwise) taking values in $\{-1, 1\}$ (where we take $\text{sign}(0) = 1$). Note that S_z is specified by a set of polynomial inequalities as defined by the functions $h_{k,z}(x)$.

3 Provable Automatic Subdifferentiation

In the algebraic circuit complexity model, where AD is provably correct, branching is not permitted. The inclusion of branching into our program leads to a number of subtle issues. Branching allows us to implement the same nonsmooth function in different manners, which have important consequences in linearization approaches. Consider two different programs (with the same input-output behavior) for the ReLU (x) function in Figure 1. The left program returns x on the constraint set that is encoded as $S_1 = \{x | x \geqslant 0\}$, while the right program returns x on the constraint set that is encoded as $S_1 = \{x | x^3 \geqslant 0\}$. In nonlinear programming, the importance of avoiding encoding constraints in the latter manner is well-known [Abadie, 1967, Peterson, 1973, Gould and Tolle, 1971].

This example motivates our restriction to only consider library functions that are encoded like the former set. We will make the standard constraint qualification

Algorithm 4: ReLU (x)	**Algorithm 5:** ReLU (x)
Input: $x = x_1$	**Input:** $x = x_1$
1: Branch:	1: Branch:
• If: $x_1 \geqslant 0$, set $x_2 = x_1$.	• If: $x_1^3 \geqslant 0$, set $x_2 = x_1$.
• Else: set $x_2 = 0$.	• Else: set $x_2 = 0$.
Return: x_2.	**Return:** x_2.

Figure 1: Two programs that implement ReLU (x): Both programs are correct and return the same value. However, the program on the right violates Assumption 3.1 since the gradient of the constraint function at $x = 0$, $\nabla(x_1^3) = 3x_1^2 = 0$.

assumption[4]. Roughly speaking, the assumption states that first order information characterizes the set of feasible perturbations. We state this assumption in a manner more directly applicable to our setting (see [Abadie, 1967, Peterson, 1973, Gould and Tolle, 1971]).

Assumption 3.1. *(Constraint Qualification on our Library) Assume for all $g \in \mathcal{L}$ that g is locally Lipschitz and our program for g (in our computational model) satisfies the constraint qualification condition on all sets S_z in the following sense: suppose $\{h_z\}$ (for binary z) are the corresponding constraint functions in our program. For any x, v (of the same input dimensionality of g), assume that for all z:*

$$\lim_{\delta \downarrow 0}(\text{sign}(h_z(x + \delta v))) = \lim_{\delta \downarrow 0}(\text{sign}(h_z(x) + \delta \nabla h_z(x) \cdot v)).$$

Roughly, this states that the set approached along the limiting direction $x + \delta v$, when $\delta \downarrow 0$, can be determined with first order information.

Before we state our main theorem, one more definition is in order, due to that $\text{Runtime}(f; x)$ may not be continuous. Define the limiting runtime $\text{Runtime}^*(f; x)$ of f at x as the (supremum) runtime to compute $f(x)$, as x is approached from nearby points. Precisely,

$$\text{Runtime}^*(f; x) := \sup \left\{ \lim_{i \to \infty} \text{Runtime}(f; x_i) : x_i \to x \right\},$$

(where the limit is taken to be the set of all limit points).

---

[4]The standard constraint qualification assumption on a constraint set is that the tangent cone of the constraint set equals the linearized cone (of the functions which define the constraints).

Theorem 3.1. *(A Cheap Subgradient Principle) Assume that our program for $f(x)$, in Algorithm 1, is allowed to use nonsmooth functions from our library $\mathcal{L}$ (in addition to affine functions and monomials). Suppose assumptions 2.1 and 3.1 hold. There exists a (randomized) algorithm, which upon input x, terminates in time that is at most $6 * \mathrm{Runtime}^*(f; x)$, and, almost surely, returns both $f(x)$ and an element $u \in \partial f(x)$.*

The following example shows one subtle issue with regards to constraint qualification.

Example 3.1. *(Constraint qualification on programs do not compose) Consider the function $f(x) = ReLU\left(x^2\right)$ (which is equivalent to the smooth function x^2). It is straight forward to see that the induced program for $f(x) = ReLU\left(x^2\right)$ (when we unravel it) does not satisfy the constraint qualification assumption, even if we do use an implementation of $ReLU\left(\cdot\right)$ that does satisfy this assumption. Regardless, in Example 3.4, we show that our algorithm does indeed correctly compute the gradient on this (continuous) function.*

Before we present the construction, we first provide a chain rule for nonsmooth functions.

3.1 A Chain Rule for Nonsmooth Functions

Let $D[g; v](x)$ denote the one-sided (Dini) directional derivative:

$$D[g; v](x) := \lim_{\delta \downarrow 0} \frac{g(x + \delta v) - g(x)}{\delta}.$$

(note that we are not assuming that v is a unit vector). This derivative exists for all piecewise polynomials and semialgebraic functions [Coste, 2000, Lemma 6.2].

Assumption 3.2. *(Overloading the library with ASD subroutines) Assume we have a library of (locally Lipschitz) functions $\mathcal{L}$ computable in our computational model. For any $g \in \mathcal{L}$, with the representation $g(x) = \sum_{z \in \{-1,1\}^T} \mathbb{I}_{S_z}(x)p_z(x)$, assume we have the following associated automatic subdifferentiation subroutine $ASD[g]$ with the following behavior: upon input $(x; v)$, the output $[a, d, u] = ASD[g](x; v)$ satisfies*

$$a = g(x), \ d = D[g; v](x), \ u = \nabla p_z(x)$$

where z is such that:

$$\lim_{\delta\downarrow 0}(\mathbb{I}_{S_z}(x + \delta v)) = 1\,.$$

Roughly speaking, u is the derivative determined by the set S_z which is approached along the limiting direction $x + \delta v$, when $\delta \downarrow 0$.

For any locally Lipschitz function h, define the limiting total derivate as:

$$\partial[h; v](x) := \lim_{\delta\downarrow 0} \nabla h(x + \delta v)$$

if the limit exists. For almost all v, the limit exists, and $\partial[h; v](x)$ is a subdifferential of h.

Theorem 3.2. *(A Chain Rule for Nonsmooth Functions) Assume $h : \mathbb{R}^m \to \mathbb{R}$ and $g_1, \ldots g_m$ (where $g_i : \mathbb{R}^d \to \mathbb{R}$) are locally Lipschitz functions computable in our computational model and that the function h is overloaded with an ASD subroutine as specified in Assumption 3.2. Define:*

$$f(x) := h(g_1(x), \ldots g_m(x)) = h(g(x))\,,$$

where $g(x)$ is the vector valued function $(g_1(x), \ldots g_m(x))^\top$. Denote the $m \times 1$ vector of (one-sided) directional derivatives as $D[g; v](x)$. If it exists, let $\partial[g; v](x)$ denote $m \times d$ limiting Jacobian matrix (whose rows are given by the vectors $\partial[g_i; v](x)$'s). Set:

$$[a, d, u] = ASD[h](g(x); D[g; v](x))$$

For all but a measure 0 set of v, we have that $\partial[f; v](x)$ and $\partial[g; v](x)$ exist and that:

$$\partial[f; v](x) = \partial[g; v](x)^\top u\,. \tag{4}$$

Example 3.2. *Consider the example $x = f_2(x) = ReLU(x) - ReLU(-x)$. We define $h(y_1, y_2) = y_1 - y_2$, $g_1(x) = ReLU(x)$, and $g_2(x) = ReLU(-x)$, so that $f_2 = h(g_1(x), g_2(x))$. By applying the ASD subroutine to h, starting at $x = 0$ with $v = 1$ which leads to running $ASD[h]((0, 0); (1, 0)) = [a, d, u]$ (where it is straightforward to verify that $u = [1, -1]^\top$), we obtain*

$$\partial[f_2; v](0) = \partial[g; v](0)^T u$$

$$= \begin{bmatrix} 1 \\ 0 \end{bmatrix}^\top \begin{bmatrix} 1 \\ -1 \end{bmatrix}$$

$$= 1,$$

which is correct. Furthermore, note a correct answer is obtained for any $v \neq 0$.

Algorithm 6: Automatic Subdifferentiation

Input: $x = (x_1, \ldots x_d)$, $v \in R^d$.
Initialize: Set $\dot{x}_1 = v_1$, $\dot{x}_2 = v_2, \ldots \dot{x}_d = v_d$.
  1: **for** $k = d + 1, d + 2, \ldots T$ **do**
  2:     Compute $[a, d, u] = \text{ASD}[g_k](x_{\text{parents}(k)}; \dot{x}_{\text{parents}(k)})$ and set:

$$x_k = a, \ \dot{x}_k = d, \quad \frac{\partial x_k}{\partial x_{\text{parents}(k)}} = u$$

  3: **end for**
  4: Compute $\frac{\partial x_T}{\partial x}$ using the Reverse Mode on these precomputed variables.
Return: x_T, and $\frac{\partial x_T}{\partial x}$.

Example 3.3. *We return to $f(x) = ReLU\,(x^2)$ from Example 3.1. Define $h(y) = ReLU\,(y)$, $g(x) = x^2$, and so $f(x) = h(g(x))$. By applying the chain rule lemma at $x = 0$ with $v = 1$,*

$$\partial[f; v](0) = \partial[g; v](0)u = 0 \cdot u = 0$$

Subtly, note that $[a, d, u] = ASD[h](0; 0)$, so we are feeding a degenerate direction $d = 0$ into our subroutine. Regardless, the chain rule lemma still applies (for any v in this case).

3.2 The algorithm

We first present the algorithm that utilizes an overloaded library. We then provide a provably correct construction of this overloaded library. All proofs are provided in the appendix.

Subdifferentiation with the overloaded library

Algorithm 6 is the Automatic Subdifferentiation procedure. Correctness follows from Lemma 3.1.

Lemma 3.1. *Suppose Assumptions 2.1 and 3.2 hold. Upon input of an arbitrary x, and if v is sampled uniformly at random from the unit sphere, then, almost surely, Algorithm 6 returns both $f(x)$ and an element $u \in \partial f(x)$.*

Algorithm 7: Overloading the function $g(x)$

Input: $x = (x_1, \ldots x_d)$, $v \in R^d$.
Initialize: Set $\dot{x}_1 = v_1$, $\dot{x}_2 = v_2, \ldots \dot{x}_d = v_d$.
1: **for** $k = d + 1, d + 2, \ldots T$ **do**
2:    Compute x_k, its partial derivatives, and the directional derivative:

$$x_k = g_{k,z}(x_{\text{parents}(k,z)}), \quad \left\{ \frac{\partial x_k}{\partial x_j} \bigg| j \in \text{parents}(k, z) \right\},$$

$$\dot{x}_k = \sum_{j \in \text{parents}(k,z)} \frac{\partial x_k}{\partial x_j} \dot{x}_j$$

3:    If the program branches at (k, z), then:

   •   If: $x_k > 0$, then $z_k = 1$.

   •   Elseif: $x_k = 0$ *and* $\dot{x}_k \geqslant 0$, then $z_k = 1$.

   •   Else: $z_k = -1$

4:    If the program halts at (k, z), then terminate the **for** loop.
5: **end for**
6: Compute $\frac{\partial x_k}{\partial x}$ using the Reverse Mode on these pre-computed variables.
Return: $[a, d, u] = [x_k, \dot{x}_k, \frac{\partial x_k}{\partial x}]$.

Proof. Fix $k \in [d + 1, \ldots, T]$. Every parent variable $j \in \text{parent}(k)$ can be expressed as $x_j = \tilde{g}_j(x)$, where g_j is a piecewise polynomial on the d dimensional input x. Thus

$$x_k = g_k(\tilde{g}_1(x), \ldots, \tilde{g}_{k-1}(x)).$$

Now the usual chain rule holds for directional derivatives [Shapiro, 1990]. As the forward mode of AD implements the usual chain rule of directional derivatives, then we have $\dot{x}_j = D[\tilde{g}_j; v]$.

By Assumption 3.2 and Theorem 3.2, $\text{ASD}[g_k](x_{\text{parents}(k)}, \dot{x}_{\text{parents}(k)})$ returns $u = \frac{\partial x_k}{\partial x_{\text{parents}(k)}} = \partial[g_k; \dot{x}_{\text{parents}(k)}]$ and this limiting total derivate satisfies the chain rule $\partial[x_k; v](x) = \partial[\tilde{g}; v](x)^\top u$. Since the limiting total derivates satisfies the chain rule and the validity of reverse mode AD algorithm relies only on the chain rule, Algorithm 6 correctly computes $\partial[f(x); v]$.

13

By Rademacher's theorem and the definition of Clarke subgradient in Equation (1), $\partial[f(x); v] \in \partial f(x)$, for almost all v. $\qquad\square$

Overloading the Library Functions

The following lemma shows that we can provide a method to correctly overload the library, which we use in Algorithm 6.

Lemma 3.2. *(Correct Library Overloading) Assume g satisfies the constraint qualification conditions in Assumption 3.1, Suppose the corresponding representation is $g(x) = \sum_{z \in \{-1,1\}^T} \mathbb{I}_{S_z}(x) p_z(x)$, On an arbitrary input x and v, Algorithm 7 returns $g(x)$, $D[g; v](x)$, and an element $u = \nabla p_z(x)$ where z is such that: $\lim_{\delta \downarrow 0}(\mathbb{I}_{S_z}(x + \delta v)) = 1$.*

Example 3.4. *We again return to ReLU (x^2) from Example 3.1. Here we examine how h is overloaded based on the implementation in Algorithm 7. When $(x, v) = (0, 1)$, we are running $ASD[h](0; 0)$ and this may not follow the same branch had we run on the (infinitesimal) input $x = \epsilon v$ which leads to running $h(\epsilon^2 v^2)$. However, the gradient is correctly computed, $\partial ReLU(x^2) = 0$, regardless of the branch taken.*

4 Discussion and Open Questions

Overloading the Library Functions: It is not difficult to see that piecewise univariate functions can be implemented in our library.

Example 4.1. *Univariate Piecewise Polynomial (Algorithm 8). Let $\sigma : \mathbb{R} \to \mathbb{R}$ be a univariate piecewise polynomial, meaning that the domain $\mathbb{R}$ is partitioned into a set of k intervals $(-\infty, b_1), (b_1, b_2), \ldots, (b_{k-1}, \infty)$. On each interval, the function is equal to a polynomial $p_1, \ldots, p_k$.*

Algorithm 8 provides a constraint qualified program for the function $\sigma(\cdot)$, which can be used as a library function.

An important step would be in extending our computational model to allow the incorporation of provably correct automatic subdifferentiation libraries for linear algebra libraries. AutoGrad [Maclaurin et al., 2015] does do AD through linear algebra methods though it can not be used to obtain correct subdifferentials in programs (at nondifferentiable points); obtaining correct generalized derivatives may be particularly important in cases where we deal with low rank methods. We

Algorithm 8: $\sigma(x)$

Input: $x = x_1$

1: Branch:

- If: $x_1 \leqslant b_1$, set $x_2 = p_1(x)$.

- Elseif: $x_1 \leqslant b_2$, set $x_2 = p_2(x)$.
  $\vdots$

- Elseif: $x_1 \leqslant b_{k-1}$, set $x_2 = p_{k-1}(x)$.

- Else: set $x_2 = p_k(x)$.

Return: x_2.

---

conjecture our results can be extended, by extending the computational model, to handle these cases (there is already much known about the first order structure of these methods [Seeger et al., 2017]); technically, SVDs are not exactly computable in either the algebraic circuit complexity model or the Blum-Shub-Smale model.

Numerical Analysis: The most important open question is how to obtain numerically stable and accurate solutions [Trefethen and Bau III, 1997, Demmel, 1997]. We conjecture the techniques developed here will help in characterizing these issues. In particular, the most natural question is how to develop algorithms that satisfy the *mixed stability* criterion: the algorithm should give "nearly the right answer to nearly the right problem" (as in [Trefethen and Bau III, 1997]). For example, for the abs($\cdot$) function, it should be acceptable for an AD method to provide a subgradient near to -1 for a small input $\epsilon > 0$ due to roundoff error; however, it would undesirable for numerical error to lead vastly different gradients than those that arise from any nearby problem. This may be particularly important when doing AD in physical simulators.

References

M. Abadi, A. Agarwal, P. Barham, E. Brevdo, et al. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. URL https://www.tensorflow.org/. Software available from tensorflow.org.

J. Abadie. On the kuhn tucker theorem. *Nonlinear Programming*, pages 19–36, 1967.

A.Griewank. Automatic directional differentiation of nonsmooth composite functions. In R.Durier, editor, *Recent developments in Optimization / Seventh French-German Conference on Optimization, Dijon 1994*, pages 155–169. Springer Verlag, 1995.

Walter Baur and Volker Strassen. The complexity of partial derivatives. *Theoretical Computer Science*, 22:317–330, 1983.

Atilim Gunes Baydin, Barak A. Pearlmutter, and Alexey Radul. Automatic differentiation in machine learning: a survey. *CoRR*, abs/1502.05767, 2015.

Lenore Blum, Mike Shub, and Steve Smale. On a theory of computation over the real numbers; NP completeness, recursive functions and universal machines (extended abstract). In *FOCS*, pages 387–397. IEEE Computer Society, 1988.

F.H. Clarke. Generalized gradients and applications. *Trans. Amer. Math. Soc.*, 205:247–262, Apr. 1975.

F.H. Clarke, Y.S. Ledyaev, R.J. Stern, and P.R. Wolenski. *Nonsmooth analysis and control theory*, volume 178. Springer Science & Business Media, 2008.

Michel Coste. *An introduction to o-minimal geometry*. Istituti editoriali e poligrafici internazionali Pisa, 2000.

J. W. Demmel. *Applied numerical linear algebra*. Society for Industrial Mathematics, 1997.

Sabrina Fiege, Andrea Walther, Kshitij Kulshreshtha, and Andreas Griewank. Algorithmic differentiation for piecewise smooth functions: a case study for robust optimization. pages 1–16, 06 2017.

F.J. Gould and J.W. Tolle. A necessary and sufficient qualification for constrained optimization. 20, 03 1971.

Andreas Griewank. On automatic differentiation. In *Mathematical Programming: Recent Developments and Applications*, pages 83–108. Kluwer Academic Publishers, 1989.

Andreas Griewank. Who invented the reverse mode of differentiation? *Optimization Stories*, Documenta Matematica, Extra Volume ISMP (2012):389–400, 2012.

Andreas Griewank. On stable picewise linearization and generalized differentiation. *Optimization Methods and Software*, 28(6):1139–1178, 2013.

Andreas Griewank. On Automatic Differentiation and Algorithmic Linearization. *Pesquisa Operacional*, 34:621 – 645, 12 2014. ISSN 0101-7438.

Andreas Griewank and Andrea Walther. *Evaluating Derivatives: Principles and Techniques of Algorithmic Differentiation*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, second edition, 2008.

K. A. Khan and P.I. Barton. Evaluating an element of the clarke generalized jacobian of a composite piecewise differentiable function. *ACM Trans. Math. Softw.*, 39:23:1, 2013.

K. A. Khan and P.I. Barton. A vector forward mode of automatic differentiation for generalized derivative evaluation. *Optim. Method Softw.*, 30:1185, 2015.

Kamil A. Khan. Branch-locking ad techniques for nonsmooth composite functions and nonsmooth implicit functions. *Optimization Methods and Software*, 0 (0):1–29, 2017.

D. Klatte and B. Kummer. *Nonsmooth equations in optimization*, volume 60 of *Nonconvex Optimization and its Applications*. Kluwer Academic Publishers, Dordrecht, 2002. ISBN 1-4020-0550-4. Regularity, calculus, methods and applications.

Dougal Maclaurin, David Duvenaud, Matthew Johnson, and Ryan P. Adams. Autograd: Reverse-mode differentiation of native Python, 2015. URL http://github.com/HIPS/autograd.

B. S. Mordukhovich. *Variational Analysis and Generalized Differentiation II: Applications*. Springer Berlin Heidelberg, 2006. ISBN 9783540312468. URL https://books.google.com/books?id=lmdmY75lrokC.

Jacques Morgenstern. How to compute fast a function and all its derivatives: a variation on the theorem of baur-strassen. In *SIGA*, 1985.

Yurii Nesterov. Lexicographic differentiation of nonsmooth functions. *Math. Program.*, 104(2-3):669–700, 2005.

A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer. Automatic differentiation in pytorch. In *NIPS-W*, 2017.

David W. Peterson. A review of constraint qualifications in finite-dimensional spaces. *SIAM Review*, 15(3):639–654, 1973.

D. E. Rumelhart, G. E. Hinton, and R. J. Williams. Parallel distributed processing: Explorations in the microstructure of cognition, vol. 1. chapter Learning Internal Representations by Error Propagation, pages 318–362. MIT Press, Cambridge, MA, USA, 1986.

Matthias W. Seeger, Asmus Hetzel, Zhenwen Dai, and Neil D. Lawrence. Auto-differentiating linear algebra. *CoRR*, abs/1710.08717, 2017. URL http://arxiv.org/abs/1710.08717.

A. Shapiro. On concepts of directional differentiability. *Journal of Optimization Theory and Applications*, 66(3):477–487, Sep 1990.

Lloyd N Trefethen and David Bau III. *Numerical linear algebra*, volume 50. SIAM, 1997.

A Proofs

First, we provide a helpful technical lemma.

Lemma A.1. *Let $z \in \{-1, 1\}^T$, and suppose the functions $p_1(x), \ldots p_T(x)$ are analytic functions from $\mathbb{R}^n \to \mathbb{R}$. Let $S_z \subset \mathbb{R}^n$ be a set defined as follows:*

$$S_z = \{x \mid \text{sign}(p_1(x)) = z_1, \ldots \text{sign}(p_T(x)) = z_T\}.$$

Fix any $x \in \mathbb{R}^n$. For all $z \in \{-1, 1\}^T$ and for almost all u, there is an open neighborhood $B(u, r)$ such that for all $u' \in B(u, r)$,

$$\lim_{\delta \downarrow 0} \mathbb{I}_{S_z}(x + \delta u) = \lim_{\delta \downarrow 0} \mathbb{I}_{S_z}(x + \delta u').$$

Proof. We need only consider one such z for the proof, as the set of possible z is finite. Without loss of generality, we can say $S_z = S$ is a set of the form:

$$S = \{p_1(x) \geqslant 0, \ldots, p_k(x) \geqslant 0, p_{k+1}(x) < 0, \ldots, p_T(x) < 0\}.$$

For the proof, it suffices to show that for all j and almost all u, there is an open neighborhood $B(u, r)$ such that for all $u' \in B(u, r)$,

$$\lim_{\delta \downarrow 0} \text{sign}(p_j(x + \delta u)) = \lim_{\delta \downarrow 0} \text{sign}(p_j(x + \delta u')).$$

Let us split the constraints into a set of active constraints $p_j(x) = 0$ for $j \in A$ and inactive constraints $p_j(x) \neq 0$ for $j \in I$. For inactive constraints, the above holds due to continuity. It remains to show the above for active constraints. Let us also assume that the functions p_j in the active set are not identically equal to 0 (the claim is true for any zero function).

For every active constraint $j \in A$, testing $\lim_{\delta \downarrow 0} \text{sign}(p_j(x + \delta u)) = 1$ can be done by the Taylor expansion,

$$p_j(x + \delta u) = p_j(x) + \sum_{i=1}^{\infty} \frac{\delta^i \|u\|^i}{i!} \nabla^i p_j(x)[u^{\otimes i}],$$

which exists since p_j is analytic. For any i, note that $\nabla^i p_j(x)[u^{\otimes i}]$ is a polynomial function in u. As it is a polynomial, this function will either be equal to 0 (for all u) or it will be nonzero for almost all u[5]. Let i_j be the first index that the function $\nabla^i p_j(x)[u^{\otimes i}]$ is nonzero for almost all u (i_j is finite as p_j is not the zero function). Then for almost every direction u,

$$\lim_{\delta \downarrow 0} \text{sign}(p_j(x + \delta u)) = \text{sign}\left(\nabla^{i_j} p_j(x)[u^{\otimes i_j}]\right).$$

---

[5]This can proven by induction on the dimension of u.

For almost all u, $\nabla^{i_j} p_j(x)[u^{\otimes i_j}] \neq 0$. Since $\nabla^{i_j} p_j(x)[u^{\otimes i_j}]$ is strictly non-zero, then, by continuity,

$$\text{sign}(\nabla^{i_j} p_j(x)[u^{\otimes i_j}]) = \text{sign}(\nabla^{i_j} p_j(x)[u'^{\otimes i_j}])$$

for $\|u' - u\| < r$ (for sufficiently small r). This implies that for $\|u' - u\| < r$

$$\lim_{\delta \downarrow 0} \text{sign}(p_j(x + \delta u)) = \lim_{\delta \downarrow 0} \text{sign}(p_j(x + \delta u')),$$

which completes the proof. □

A.1 Proof of Theorem 3.2

Proof. Suppose that a program for $h(y) = h(y_1, \ldots, y_m)$ has the representation:

$$h(y) = \sum_{z \in \{-1,1\}^T} \mathbb{I}_{S_z}(y) p_z(y),$$

where S_z is specified as in Equation 2. Since f is the composition of two programs, it also has a representation of the form:

$$f(x) = \sum_{z \in \{-1,1\}^{T'}} \mathbb{I}_{\tilde{S}_z}(x) \tilde{p}_z(x),$$

for different polynomials $\tilde{p}_z$ and sets $\tilde{S}_z$.

By Rademacher's theorem, $\partial[f; v]$ exists for almost all v. Since f itself is a piecewise polynomial, $\partial[f; v](x) = \nabla \tilde{p}_z(x)$ where z is such that $\lim_{\delta \downarrow 0} \mathbb{I}_{\tilde{S}_z}(x + \delta v) = 1$. Using Lemma A.1, for all $z' \in \{-1, 1\}^{T'}$ and almost all v, there is a (full dimensional) neighborhood $\mathcal{B}$ around v, in which $\lim_{\delta \downarrow 0} \mathbb{I}_{\tilde{S}_{z'}}(x + \delta v') = 1$, which implies $\partial[f; v'](x) = \nabla \tilde{p}_z(x)$ (using our choice of z). Hence,

$$\partial[f; v](x) = \partial[f; v'](x)$$

for all $v' \in \mathcal{B}$.

Furthermore, using this and the definition of the directional derivative, we have for all $v' \in \mathcal{B}$,

$$D[f; v'](x) = \partial[f; v'](x)^\top v' = \partial[f; v](x)^\top v'. \tag{5}$$

The remainder of the proof seeks to show this directional derivative can be written as:

$$D[f; v'](x) = u^\top \partial[g; v](x) \, v'$$

(for v' in a sufficiently small full dimensional ball). Note this would imply that for all v' in a full dimensional ball, $u^\top \partial[g; v](x) \, v' = \partial[f; v](x)^\top v'$, which would complete the proof by choosing d linearly independent v''s.

Define $M = \partial[g; v](x) \in \mathbb{R}^{m \times d}$ be a matrix whose rows are $\partial[g_i; v](x)$. Observe $D[g; v](x) = Mv$. Now let us show that:

$$D[g; v'](x) = Mv'$$

for all $v' \in \mathcal{B}_1$, where $\mathcal{B}_1$ is a sufficiently small (full dimensional) neighborhood around v. To see this, note the function $g(x)$ itself is a piecewise (vector) valued function, and, by Lemma A.1, for almost all v, there is a neighborhood $\mathcal{B}_1$ around v, in which for all $v' \in \mathcal{B}_1$, v' selects the same piece as v, which implies $\partial[g; v'](x) = M$ as claimed.

Now due to the assumed properties of ASD, we know

$$u = \nabla p_z(y)|_{y = g(x)} \, .$$

for some z which satisfies:

$$\lim_{\delta \downarrow 0} \mathbb{I}_{S_z}(g(x) + \delta D[g; v](x)) = 1 \, .$$

Since $D[g; v](x) = Mv$, we have:

$$\lim_{\delta \downarrow 0} \mathbb{I}_{S_z}(g(x) + \delta Mv) = 1 \, .$$

Also, note

$$D[h; Mv](y) = \nabla p_z(y)^\top Mv$$

by the definition of the directional derivative and where $y = g(x)$.

Using Lemma A.1 again, we will show that

$$D[h; Mv'](y) = \nabla p_z(y)^\top Mv' \tag{6}$$

for all $v' \in \mathcal{B}_2$ (for a sufficiently small ball $\mathcal{B}_2$). Note that the above holds if:

$$\lim_{\delta \downarrow 0} \mathbb{I}_{S_z}(g(x) + \delta Mv') = 1 \, , \tag{7}$$

21

for all $v' \in \mathcal{B}_2$. Now consider the set $\bar{S}_z$ defined by

$$\bar{S}_z = \{\tilde{v} \,|\, \mathrm{sign}(p_1(g(x) + M\tilde{v})) = z_1, \ldots \mathrm{sign}(p_T(g(x) + M\tilde{v})) = z_T\}.$$

Lemma A.1 guarantees that there exists a ball $\mathcal{B}_2$ (centered on v) such that:

$$\lim_{\delta \downarrow 0} \mathbb{I}_{\bar{S}_z}(\delta v') = 1\,,$$

for all $v' \in \mathcal{B}_2$, which is equivalent to Equation (7). Thus we have proven the Equation (6), for all $v' \in \mathcal{B}_2$.

As the chain rule holds for directional derivatives (of locally Lipschitz piecewise polynomial/semialgebraic functions) [Shapiro, 1990], we have that:

$$D[f; v'](x) = D[h; Mv'](y) = \nabla p_z(y)^\top M v' = u^\top \partial[g; v](x)v'$$

for v' that is in some sufficiently small full dimensional neighborhood $\mathcal{B}_3$ (that is contained in $\mathcal{B}_2$ and $\mathcal{B}_1$).

Now define the full dimensional neighborhood $\mathcal{B}_4$ so that it is contained in both $\mathcal{B}_3$ and $\mathcal{B}$. We have shown, for $v' \in \mathcal{B}_4$,

$$u^\top \partial[g; v](x)\, v' = \partial[f; v](x)^\top v'\,,$$

using Equation 5. The proof is completed by choosing d linearly independent v''s, which implies that $u^\top \partial[g; v](x) = \partial[f; v](x)^\top$. $\qquad\square$

A.2 Proof of Lemma 3.2

Proof of Lemma 3.2. Let $h_{t,z}(x)$ be polynomials that represent $g_{t,z}$ by expressing in terms of the input variables x only. By the Constraint Qualification Assumption 3.1, the branch taken by Algorithm 7 is the same branch that would be taken by running on the perturbed input $x + \delta v$ for all $\delta < \delta_0(v)$.

To finish the proof, we simply note that the reverse mode is performed on a straight-line program that computes $\nabla p_z(x)$, where z is such that $\lim_{\delta \downarrow 0} \mathbb{I}_{S_z}(x + \delta v) = 1$. $\qquad\square$

A.3 Completing the proof of Theorem 3.1

Lemma 3.2 shows that every library function g that satisfies Assumption 3.1 can be implemented via Algorithm 7 to satisfy Assumption 3.2. Then Lemma 3.1 shows that we compute an element $\partial[f; v](x) \in \partial f(x)$.

The second part of the proof counts the number of unit operations. See Morgenstern [1985], Griewank [1989] for the counting argument for a factor of 5, where the reverse mode (after we have computed the function and stored the intermediate variables) is a factor of 4 more than computing the function itself.

To obtain a factor of 6 in the existence claim in Theorem 3.1, the construction uses a minor modification of the algorithm presented, where the algorithm is modified to only run one (global) reverse mode (as opposed to using a reverse mode computation associated with each library function). Algorithm 6, the one presented, has a runtime that is within a factor of 10 of the cost of computing just $f(x)$, due to that it runs the reverse mode when calling every library function (which costs a factor of 4 more than just computing $g(x)$). This algorithm presented is what we actually suggest to use, as the algorithm that achieves a factor of 6 may require more memory in order to build a larger computational graph (and the factor of 10 is rather pessimistic). To see the factor of 10, the reverse mode (associated with each library function) gives an additional factor of 4. The directional derivative costs one additional factor (this is referred to as the forward mode of AD [Griewank, 1989]). After all T (overloaded) library functions are completed, one additional reverse mode is called at the end of Algorithm 6, which incurs (at most) an additional factor of 4.

To obtain the factor of 6, it is straightforward to modify the algorithm, as the following proof shows.

Proof of Theorem 3.1. First, observe that the algorithm correctly computes the function value on some (limiting branch), so the time it takes to compute $f(x)$ on this branch is some T where $T \leqslant \mathrm{Runtime}^*(f; x)$. Now the directional derivative computation can be computed when doing the forward pass of the algorithm, which incurs at most one additional factor of T (this is also referred to as the forward mode of AD). Instead of computing the intermediate derivatives $\frac{\partial x_k}{\partial x_j}$ (as in Algorithm 7), we will unravel the entire computational graph of $f(x)$, where we replace each function $g_k(\cdot)$, by inserting the corresponding computational graph of $g_k(\cdot)$ into a global computational graph. We then run the reverse mode of AD on this global computational graph, giving only one additional factor of 4. $\qquad\square$