

On Correctness of Automatic Differentiation for Non-Differentiable Functions

Wonyeol Lee, Hangeol Yu, Xavier Rival, Hongseok Yang

► **To cite this version:**

Wonyeol Lee, Hangeol Yu, Xavier Rival, Hongseok Yang. On Correctness of Automatic Differentiation for Non-Differentiable Functions. NeurIPS 2020 - 34th Conference on Neural Information Processing Systems, Dec 2020, Vancouver / Virtual, Canada. hal-03081582

HAL Id: hal-03081582

<https://hal.inria.fr/hal-03081582>

Submitted on 18 Dec 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

On Correctness of Automatic Differentiation for Non-Differentiable Functions

Wonyeol Lee[†] Hangyeol Yu[†] Xavier Rival[‡] Hongseok Yang[†]

[†]School of Computer Science, KAIST, South Korea

[‡]INRIA Paris, Département d'Informatique of ENS, and CNRS/PSL University, France

{wonyeol.lee.cs, yhk1344}@gmail.com rival@di.ens.fr hongseok.yang@kaist.ac.kr

Abstract

Differentiation lies at the core of many machine-learning algorithms, and is well-supported by popular autodiff systems, such as TensorFlow and PyTorch. Originally, these systems have been developed to compute derivatives of differentiable functions, but in practice, they are commonly applied to functions with non-differentiabilities. For instance, neural networks using ReLU define non-differentiable functions in general, but the gradients of losses involving those functions are computed using autodiff systems in practice. This status quo raises a natural question: are autodiff systems correct in any formal sense when they are applied to such non-differentiable functions? In this paper, we provide a positive answer to this question. Using counterexamples, we first point out flaws in often-used informal arguments, such as: non-differentiabilities arising in deep learning do not cause any issues because they form a measure-zero set. We then investigate a class of functions, called PAP functions, that includes nearly all (possibly non-differentiable) functions in deep learning nowadays. For these PAP functions, we propose a new type of derivatives, called intensional derivatives, and prove that these derivatives always exist and coincide with standard derivatives for almost all inputs. We also show that these intensional derivatives are what most autodiff systems compute or try to compute essentially. In this way, we formally establish the correctness of autodiff systems applied to non-differentiable functions.

1 Introduction

Automatic differentiation or autodiff is one of the key technologies behind the dramatic progress of deep learning in recent years [4, 26, 38, 39]. It refers to the idea of developing and using a generic tool that can differentiate any function expressed as a program in a general-purpose programming language [20, 34]. Effective autodiff systems have been developed for popular programming languages [2, 3, 5, 9, 22, 28, 33, 35, 42, 44, 45]. They have enabled the development of sophisticated models and algorithms in machine learning that, in particular, involve deep neural networks [19].

This paper is concerned with one seeming contradiction of these autodiff systems: the systems have originally been developed to compute derivatives of differentiable functions, but in practice, they are commonly applied to functions with non-differentiabilities. For instance, neural networks using ReLU define non-differentiable functions in general, but the derivatives of losses involving those functions are computed using autodiff systems in practice. This status quo raises a natural question: are autodiff systems correct in any formal sense when applied to such non-differentiable functions?

A common reaction to the question is: non-differentiabilities arising in deep learning (e.g., from ReLU) do not cause any issues because they occur rarely (i.e., they form a Lebesgue-measure-zero set). In the paper, we first show that this reaction needs to be carefully re-examined at least. Using counterexamples, we point out flaws in three often-used arguments derived from this reaction. We

then present our answer. It is also positive, but based on a class of functions that satisfy a condition called *piecewise analyticity under analytic partition* (in short, PAP). These PAP functions include nearly all (possibly non-differentiable) functions in deep learning nowadays. For these PAP functions, we propose a new type of derivatives, called *intensional derivatives*, and prove that these derivatives always exist and coincide with standard derivatives for almost all inputs. These intensional derivatives behave almost as well as, and sometimes even better than, usual derivatives for differentiable functions. For instance, they always satisfy a chain rule even if functions are non-differentiable. Using these properties of intensional derivatives, we show that the intensional derivatives are what most autodiff systems compute or try to compute essentially. In this way, we formally establish the correctness of autodiff systems that compute derivatives of non-differentiable functions.

We use (a, b) , $(a, b]$, $[a, b)$, and $[a, b]$ to denote intervals in \mathbb{R} , and $\langle a_1, \dots, a_n \rangle$ to denote tuples. For $n \in (\mathbb{Z}_{>0} \cup \{\infty\})$, $[n]$ means the set $\{1, 2, \dots, n\}$. We call Lebesgue measure simply by measure. The detailed statements and missing proofs of our results can be found in the appendix.

2 Challenges

As mentioned in the introduction, practitioners frequently apply autodiff systems to functions with non-differentiabilities, and justify these out-of-scope use cases with plausible yet heuristic arguments. In this section, we analyse these arguments. We go through three claims that are often used in the arguments implicitly, and show that although looking innocent at the outset, the claims have serious flaws; they are wrong, and we provide counterexamples.

Recall a notion of correctness for an autodiff system covering non-differentiable functions [6, 20, 24]:

Definition 1 (Correctness of Autodiff). *We say that an autodiff system is **correct** if the following condition holds: for every measurable function $f : \mathcal{X} \rightarrow \mathbb{R}^m$ defined on an open $\mathcal{X} \subseteq \mathbb{R}^n$ and implemented by a program, if f is differentiable almost everywhere (i.e., the set of inputs making f non-differentiable is contained in a measure-zero set), then for almost every $x \in \mathcal{X}$, the autodiff system applied to (the program of) f and the input x computes the derivative of f at x .*

The definition permits a non-differentiable function as an input to an autodiff system, as long as its non-differentiability occurs rarely (i.e., at a measure-zero subset of the input domain). For such a function, it may be impossible to compute the correct derivative for all inputs, simply because the derivative does not exist for inputs where the function is non-differentiable. Thus, the definition just requires that the system should compute the correct derivative for *most* inputs instead (i.e., for a subset of the input domain whose complement inside the domain is contained in a measure-zero set).

Proving the correctness of an autodiff system is surprisingly difficult. Nearly every autodiff system is based on a chain rule for computing the derivative of function composition, but when the component functions are non-differentiable, designing a correct version of the rule is challenging. To help the reader see this challenge, we analyse three plausible yet flawed claims about the derivative of function composition, which are sometimes used implicitly in heuristic justifications of autodiff systems.

Let $f : \mathcal{X} \rightarrow \mathcal{Y}$ and $g : \mathcal{Y} \rightarrow \mathbb{R}^l$ be measurable functions defined over open sets $\mathcal{X} \subseteq \mathbb{R}^n$ and $\mathcal{Y} \subseteq \mathbb{R}^m$. Here is the first claim:

Claim 1 If f and g are differentiable almost everywhere and continuous, then $g \circ f$ should be differentiable almost everywhere.

A rationale for the claim goes as follows. In order for $g \circ f$ to be non-differentiable at x_0 , the function f has to be non-differentiable at x_0 , or it should map x_0 to a non-differentiable input to g and be able to vary enough in a neighbourhood of x_0 . The claim says that such an x_0 is rare (from the perspective of measure). Of course, the first case that f is non-differentiable at x_0 occurs rarely by assumption. The second case seems to happen rarely as well, because the non-differentiable inputs to g are rare and f is continuous: because of continuity, if f maps many x_0 's (i.e., all the x_0 in some non-measure-zero set) to those rare non-differentiable inputs of g , it should behave as a constant function in the neighbourhoods of most of those x_0 's.

The rationale has a flaw, and the claim is false. The inputs x_0 falling into the second case are not necessarily rare. Although f is continuous, it is possible that f maps many x_0 's to some of those rare non-differentiable inputs of g without acting as a constant function in a neighbourhood of each of those x_0 's. The precise result is summarised in the following proposition:




Figure 1: The graphs of ϕ_1 , f , and g that are described in the proof of Proposition 1.

Proposition 1. *There exist functions $f : (0, 1) \rightarrow (0, 1)$ and $g : (0, 1) \rightarrow [0, 1]$ such that f and g are differentiable almost everywhere and continuous, but $g \circ f$ fails to be almost-everywhere differentiable.*

Proof. Before our proof, we review a generalised version of the Cantor set and Cantor function [11, 14] defined over the open interval $(0, 1)$: the λ -Cantor set C_λ and the λ -Cantor function ϕ_λ for $\lambda \in (0, 1]$, which reduce to the original Cantor set and function when $\lambda = 1$. The C_λ and ϕ_λ will serve as building blocks of counterexamples presented in this paper. The set $C_\lambda \subset (0, 1)$ consists of all the real numbers not removed during the following procedure: at step $k = 0$, start with a closed interval $[0, 1] \subset \mathbb{R}$; at each step $k > 0$, remove every open interval of length $\lambda/3^k$ which locates exactly at the middle of each remaining closed interval; and after all these steps, remove 0 and 1. Note that what remains after the step k are 2^k closed intervals. The function $\phi_\lambda : (0, 1) \rightarrow (0, 1)$ is defined as follows, based on the construction of C_λ . At each step $k > 0$, define ϕ_λ over the i -th open interval to be removed as $\phi_\lambda(x) = (2i - 1)/2^k$ ($i \in [2^{k-1}]$). After all these steps, ϕ_λ is defined over some dense open subset S of $[0, 1]$ (in fact, $S = (0, 1) \setminus C_1$). Since $\phi_\lambda : S \rightarrow [0, 1]$ is uniformly continuous with S being dense in $[0, 1]$, it has the unique continuous extension $\phi_\lambda : [0, 1] \rightarrow [0, 1]$ [37, Exercise 4.13]. The λ -Cantor function refers to its restricted version $\phi_\lambda : (0, 1) \rightarrow (0, 1)$.

Returning to the proof, let f be the inverse of the homeomorphism $F : (0, 1) \rightarrow (0, 1)$ defined by $F(x) = \frac{1}{2}(\phi_1(x) + x)$ [8, Example 2.3.1]. It is known that $f(C_{1/2}) = C_1$ [8, Example 2.3.2]. Construct $g : (0, 1) \rightarrow [0, 1]$ based on the construction of C_1 as follows, similarly to [18, Example 8.18]: at each step $k > 0$, define g over the i -th open interval (l_i, r_i) to be removed, as $g(y) = 2^{-k} \cdot [1 - (y - (r_i + l_i)/2)^2 / ((r_i - l_i)/2)^2]$ ($i \in [2^{k-1}]$); and define g over C_1 as $g(y) = 0$. See Figure 1 for the graphs of ϕ_1 , f , and g constructed so far. Clearly, f is continuous. Also, g is continuous, since the height of the parabolas defined at the k -th step of g 's construction converges to 0 as $k \rightarrow \infty$, and $g(C_1) = \{0\}$ (see Appendix A for the details). Hence, $g \circ f$ is continuous. Note that f is even Lipschitz continuous. To prove this, observe that f can be constructed similarly to $\phi_{1/2}$, and repeat the proof of the Lipschitz continuity of $\phi_{1/2}$ [14].

We now show that f and g are differentiable almost everywhere, but $g \circ f$ is not. First, since f is Lipschitz continuous, it is differentiable almost everywhere by Rademacher's theorem [41, Theorem 2.2.4]. Next, since g is differentiable on $(0, 1) \setminus C_1$ by its construction, it is differentiable almost everywhere (as C_1 has measure 0). Lastly, $g \circ f$ is non-differentiable on $C_{1/2}$, which has measure $1/2$, due to that: the parabolas defined at the k -th step of g 's construction get sharper as $k \rightarrow \infty$; $g(C_1) = \{0\}$; and f is a homeomorphism with $f(C_{1/2}) = C_1$ (see Appendix A for the details). \square

The second claim is about a chain rule.

Claim 2 If f , g , and $g \circ f$ are differentiable almost everywhere and continuous, then the standard chain rule for $g \circ f$ should hold almost everywhere. In particular, if $f, g : (0, 1) \rightarrow (0, 1)$, then $(g \circ f)'(x_0) = g'(f(x_0)) \cdot f'(x_0)$ for almost all $x_0 \in (0, 1)$.

Note that all of f , g , and $g \circ f$ in the claim are assumed to be differentiable almost everywhere. The claim comes from heuristic reasoning that if we just avoid those rare non-differentiable inputs of $g \circ f$, we should be able to use the standard result for differentiation, including the chain rule.

The second claim is also wrong. The flaw in the heuristic reasoning from above is that the almost-everywhere differentiability of f , g , and $g \circ f$ does not stop $g'(f(x_0))$ from being undefined for many x_0 's in $(0, 1)$. This is related to the flaw in the justification for the first claim that we explained. The next proposition and its proof provide a concrete example for this phenomenon:

Proposition 2. *There exist functions $f, g : (0, 1) \rightarrow (0, 1)$ such that f , g , and $g \circ f$ are differentiable almost everywhere and continuous, but it is not that $g'(f(x))$ is defined for almost all $x \in (0, 1)$.*

Proof. Let $f(x) = 1/2$ and $g(y) = \text{ReLU}(y - 1/2) + 1/2$. Then, $(g \circ f)(x) = 1/2$. Certainly, f , g , and $g \circ f$ are differentiable almost everywhere and Lipschitz continuous. But, g is not differentiable at $f(x) = 1/2$ for all $x \in (0, 1)$. So, it is not that $g'(f(x))$ is defined for almost all $x \in (0, 1)$. \square

The third claim is a natural reaction to the failure of the second claim. It implements the strategy of making the chain rule in the second claim more permissive such that the counter argument of Proposition 2 no longer applies. The claim expresses a weaker version of the rule that allows one to set the derivatives of f and g to arbitrary values wherever f and g are not differentiable.

Claim 3 Let $f, g : (0, 1) \rightarrow (0, 1)$. If f , g , and $g \circ f$ are differentiable almost everywhere and continuous, then there should exist $df, dg : (0, 1) \rightarrow \mathbb{R}$ such that $df(x_0) = f'(x_0)$, $dg(y_0) = g'(y_0)$, and $(g \circ f)'(x_0) = dg(f(x_0)) \cdot df(x_0)$ for almost all $x_0, y_0 \in (0, 1)$.

The functions df and dg in the claim are the extensions of f' and g' that set $df(x)$ and $dg(y)$ to arbitrary values whenever $f'(x)$ and $g'(y)$ are undefined. The chain rule in the claim is phrased in terms of these extensions df and dg , so that it does not suffer from the problem pointed out in Proposition 2. However, this new rule is still flawed as shown in the next proposition:

Proposition 3. *There exist functions $f, g : (0, 1) \rightarrow (0, 1)$ such that f , g , and $g \circ f$ are differentiable almost everywhere and continuous, but for some measurable subset $A \subseteq (0, 1)$ with non-zero measure, they satisfy the following property: $f'(x) = 0$ and $(g \circ f)'(x) \neq 0$ for all $x \in A$.*

Proof. Consider the function f in the proof of Proposition 1. Let g be the 1-Cantor function ϕ_1 . Then, $g \circ f$ is the $(1/2)$ -Cantor function $\phi_{1/2}$. We already showed f is differentiable almost everywhere and even Lipschitz continuous. Since g and $g \circ f$ are monotone on $(0, 1)$, they are differentiable almost everywhere by the monotone differentiation theorem [41, Theorem 1.6.25]; and they are clearly continuous. We now show there exists $A \subseteq C_{1/2}$ with the desired properties. Since $C_{1/2}$ has measure $1/2$, it suffices to prove $f'(x) = 0$ and $(g \circ f)'(x) = 2$ for almost all $x \in C_{1/2}$. The claim indeed holds due to the following: f and $g \circ f$ are Lipschitz, so absolutely continuous; $f'(x) = 2$ and $(g \circ f)'(x) = 0$ for all $x \notin C_{1/2}$; $f'(x) \geq 0$ and $(g \circ f)'(x) \leq 2$ for all $x \in C_{1/2}$ whenever these derivatives exist; and $C_{1/2}$ has measure $1/2$. For the details, see [8, Example 2.3.2] and [14]. \square

The proposition implies the third claim is doomed. The claim says that $df(x_0) = f'(x_0)$ and $(g \circ f)'(x_0) = dg(f(x_0)) \cdot df(x_0)$ for almost all $x_0 \in A$. But both equalities cannot hold simultaneously: if they do, by Proposition 3, $(g \circ f)'(x_0) = dg(f(x_0)) \cdot df(x_0) = dg(f(x_0)) \cdot f'(x_0) = dg(f(x_0)) \cdot 0 = 0$, but the same proposition also entails $(g \circ f)'(x_0) \neq 0$, leading to a contradiction.

A lesson from these flawed claims is that although the notion of correctness in Definition 1 only refers to almost-everywhere differentiability, we need a condition stronger than it, which behaves better in handling function composition and gives rise to a chain rule. We describe such a condition next.

3 PAP Function and Intensional Derivative

Our justification of autodiff systems relies on two key concepts: *piecewise analyticity under analytic partition*, and *intensional derivative*. The first is a (strictly) stronger property about functions than almost-everywhere differentiability, and yet it is satisfied by practically all programs targeted at by existing autodiff systems, as we will show in §4. Functions with this new property, called PAP functions, have an unusual type of derivatives, called intensional derivatives, which form the second concept. Intensional derivatives of PAP functions are defined everywhere and satisfy a chain rule, while still agreeing with standard derivatives for almost all inputs. In fact, the PAP functions have not just first-order but also all higher-order intensional derivatives. In §4, we will show that most autodiff systems compute intensional derivatives when applied to functions with non-differentiabilities.

To expand the overview of the two concepts just given, we need a notion of piecewise representation:

Definition 2 (Piecewise Representation). *Let $\mathcal{X} \subseteq \mathbb{R}^n$ and $\mathcal{Y} \subseteq \mathbb{R}^m$. A set of pairs $\gamma = \{A^i, f^i\}_{i \in [I]}$ is a **piecewise representation** of a function from \mathcal{X} to \mathcal{Y} or simply **representation** if and only if (i) $I \in \mathbb{Z}_{>0} \cup \{\infty\}$; (ii) $\{A^i\}_{i \in [I]}$ is a partition of \mathcal{X} ; and (iii) each $f^i : \mathcal{X}^i \rightarrow \mathcal{Y}$ is a function on an open domain $\mathcal{X}^i \subseteq \mathbb{R}^n$ with $A^i \subseteq \mathcal{X}^i$. The **evaluation** of a representation $\gamma = \{A^i, f^i\}_{i \in [I]}$ is the map $\langle\langle \gamma \rangle\rangle : \mathcal{X} \rightarrow \mathcal{Y}$ defined by: $\langle\langle \gamma \rangle\rangle(x) = f^i(x)$ for all $i \in [I]$ and $x \in A^i$.*

A representation $\gamma = \{\langle A^i, f^i \rangle\}_{i \in [I]}$ describes a function f in terms of a finite or countable number of component functions f^i and their scopes A^i of application. For each input x , the value of f at x is determined by some component f^i , chosen based on which piece of the input partition $\{A^i\}_{i \in [I]}$, the x belongs to. Our notation for the described f is $\langle\langle \gamma \rangle\rangle$. Note that the domain \mathcal{X}^i of each f^i has to be open, and it may be larger than A^i , the set of inputs where f^i is used. As a result, a function can be described by an infinite number of representations. For instance, ReLU has at least two representations: $\{\langle \mathbb{R}_{<0}, x \mapsto 0 \rangle, \langle \mathbb{R}_{\geq 0}, x \mapsto x \rangle\}$ and $\{\langle \mathbb{R}_{<0}, x \mapsto 0 \rangle, \langle \{0\}, x \mapsto 2x \rangle, \langle \mathbb{R}_{>0}, x \mapsto x \rangle\}$.

Having a piecewise representation of a function, instead of just the function, has several advantages. One of them is that for many properties of functions, we can derive their piecewise variants using those representations. We apply this advantage to analyticity. Recall that a real-valued function is analytic if and only if it is infinitely differentiable and is equal to its Taylor expansion. For \mathbb{R}^m -valued function g , the analyticity means the coordinate-wise analyticity: for all $j \in [m]$, the composition $\pi_j \circ g$ of the j -th projection π_j and g is analytic. Analytic functions are favoured by autodiff systems, because they are infinitely differentiable and become zero only at a measure-zero set of inputs (unless they are zero everywhere or their domains are disconnected). Let $\mathcal{X} \subseteq \mathbb{R}^n$ and $\mathcal{Y} \subseteq \mathbb{R}^m$ be any sets.

Definition 3 (Analytic Partition). *A set $A \subseteq \mathbb{R}^n$ is **analytic** if and only if for some $J, L \in \mathbb{Z}_{>0}$, there are analytic functions $g_j^+ : \mathcal{X}_j^+ \rightarrow \mathbb{R}$ and $g_l^- : \mathcal{X}_l^- \rightarrow \mathbb{R}$ over open domains $\mathcal{X}_j^+, \mathcal{X}_l^- \subseteq \mathbb{R}^n$ ($j \in [J]$, $l \in [L]$) such that $A = \{x \in \mathbb{R}^n \mid (\bigwedge_{j \in [J]} x \in \mathcal{X}_j^+ \wedge g_j^+(x) > 0) \wedge (\bigwedge_{l \in [L]} x \in \mathcal{X}_l^- \wedge g_l^-(x) \leq 0)\}$. A partition $\{A^i\}_{i \in [I]}$ of \mathcal{X} is **analytic** if and only if all A^i are analytic.*

Definition 4 (PAP Representation). *A representation $\gamma = \{\langle A^i, f^i \rangle\}_{i \in [I]}$ from \mathcal{X} to \mathcal{Y} is **piecewise analytic under analytic partition** (in short, **PAP**) if and only if $\{A^i\}_{i \in [I]}$ is an analytic partition of \mathcal{X} and f^i is analytic over its domain \mathcal{X}^i for all $i \in [I]$.*

Definition 5 (PAP Function). *A function $f : \mathcal{X} \rightarrow \mathcal{Y}$ is **piecewise analytic under analytic partition** (in short, **PAP**) if $f = \langle\langle \gamma \rangle\rangle$ for some PAP representation γ .*

The definitions identify PAP representations and PAP functions as those built by the two-step process: we first split the input domain such that boundaries of the split regions are expressed by the zero sets of analytic functions, and next choose an appropriate analytic function for each piece of the split. Note the use of analytic functions in both steps. Thus, just like the standard analyticity, the PAP property implies almost-everywhere differentiability (Proposition 4), but not vice versa (Proposition 5).

Proposition 4 ([47]). *All PAP functions are differentiable almost everywhere.*

Proof. The proof extends the one for a similar result in [47]. The key idea is to use the fact that the zero set of a non-constant analytic function over a connected open domain has measure zero [32]. To prove the proposition, we show that for each PAP function f , there exist countably many non-constant analytic functions $\{g_j\}_j$ over connected open domains such that if f is non-differentiable at $x \in \mathcal{X}$, then x belongs to the zero set of some g_j . For the details, see Appendix B.3. \square

Proposition 5. *There is a continuous almost-everywhere differentiable yet non-PAP function.*

Proof. We have two sufficient conditions for a function f to be non-PAP: (i) the k -th order derivative of f is undefined on a set of positive measure for some $k \geq 1$ (Proposition 8); and (ii) the k -th order derivative of f is undefined on an uncountable set for some $k \geq 1$, and f is defined on a subset of \mathbb{R} (Appendix B.5). The following functions satisfy (i) with $k = 2$: the λ -Cantor function for all $\lambda \in (0, 1)$, f in the proof of Proposition 1, and Volterra's function $V : (0, 1) \rightarrow \mathbb{R}$ [18, Example 8.35]. The following functions satisfy (ii) with $k = 1$: the 1-Cantor function and g in the proof of Proposition 1. All these functions are known (or already proven above) to be continuous and almost-everywhere differentiable. Hence, all of them are desired functions. \square

We now define intensional derivatives of PAP representations and PAP functions. Let $\gamma = \{\langle A^i, f^i \rangle\}_{i \in [I]}$ be a PAP representation from \mathcal{X} to \mathcal{Y} for some $\mathcal{X} \subseteq \mathbb{R}^n$ and $\mathcal{Y} \subseteq \mathbb{R}^m$.

Definition 6 (Intensional Derivative of PAP Representation). *An **intensional derivative** of γ is the set $D\gamma = \{\langle A^i, Df^i \rangle\}_{i \in [I]}$, where Df^i is the standard derivative of f^i viewed as a function from the domain \mathcal{X}^i of f^i to $\mathbb{R}^{m \times n}$.*

Proposition 6. *The intensional derivative $D\gamma$ is a PAP representation from \mathcal{X} to $\mathbb{R}^{m \times n}$.*

Proof. Nearly all the requirements for $D\gamma$ to be a PAP representation directly follow from the fact that γ is PAP. The only exception is the analyticity of Df^i . There we use the fact that the operation of taking a (standard) partial derivative of a function preserves analyticity [25, Proposition 2.2.3]. \square

Definition 7 (First-Order Intensional Derivative). *Let $f : \mathcal{X} \rightarrow \mathcal{Y}$ be a PAP function. Define $\partial_{\bullet} f$ to be the following set of functions: $\partial_{\bullet} f = \{\langle\langle D\gamma \rangle\rangle \mid \gamma \text{ is a PAP representation of } f\}$. Each $df \in \partial_{\bullet} f$ is called a **(first-order) intensional derivative** of f .*

By Proposition 6, only PAP functions live in $\partial_{\bullet} f$. Thus, we can also take intensional derivatives of functions in $\partial_{\bullet} f$. We push this observation further and define higher-order intensional derivatives:

Definition 8 (Higher-Order Intensional Derivative). *Let $f : \mathcal{X} \rightarrow \mathcal{Y}$ be a PAP function. For each $k \in \mathbb{Z}_{\geq 0}$, inductively define the set $\partial_{\bullet}^k f$ of functions by: $\partial_{\bullet}^0 f = \{f\}$, and $\partial_{\bullet}^k f = \{df^k \in \partial_{\bullet}(df^{k-1}) \mid df^{k-1} \in \partial_{\bullet}^{k-1} f\}$ for $k \geq 1$. Each $df^k \in \partial_{\bullet}^k f$ is called a **k -th order intensional derivative** of f .*

A function $f : \mathcal{X} \rightarrow \mathcal{Y}$ may have zero, one, or more than one intensional derivatives. Having at least one intensional derivative corresponds to f being differentiable in the standard sense. The next propositions show that every PAP function is infinitely differentiable in the new sense (Proposition 7), and that these standard and new notions of differentiability coincide if we ignore a measure-zero set of inputs (Proposition 8). Let $f : \mathcal{X} \rightarrow \mathcal{Y}$ be a PAP function for some $\mathcal{X} \subseteq \mathbb{R}^n$ and $\mathcal{Y} \subseteq \mathbb{R}^m$.

Proposition 7. *For all $k \in \mathbb{Z}_{\geq 0}$, the set $\partial_{\bullet}^k f$ of the k -th order intensional derivatives is not empty. Furthermore, its elements are again PAP functions from \mathcal{X} to $\mathbb{R}^{m \times n^k}$.*

Proof. The proof is by induction on k . First, for the case $k = 0$, $\partial_{\bullet}^0 f$ is a singleton set and its unique element f is PAP. Hence, the proposition holds. Next, consider the case $k > 0$. By induction hypothesis, $\partial_{\bullet}^{k-1} f$ is not empty and consists of PAP functions only. Thus, the set $\{\gamma \mid \gamma \text{ is a PAP representation of } g \text{ for some } g \in \partial_{\bullet}^{k-1} f\}$ is not empty. This and Proposition 6 imply that $\partial_{\bullet}^k f$ (obtained by applying $\langle\langle D- \rangle\rangle$ on the above set) is nonempty and contains only PAP functions. \square

Proposition 8. *For all $k \in \mathbb{Z}_{\geq 0}$, f has the k -th order standard derivative almost everywhere, and this derivative agrees with any k -th order intensional derivative $df^k \in \partial_{\bullet}^k f$ almost everywhere.*

In the proposition, we view the k -th order standard derivative of f as a function of type $\mathcal{X} \rightarrow \mathbb{R}^{m \times n^k}$. For instance, when $k = 1$, the derivative maps an input to the Jacobian matrix of f .

Proof of Proposition 8. The first claim is proven similarly to Proposition 4, except that we additionally use the following: an analytic function is infinitely differentiable. As in Proposition 4, we prove a stronger statement: there exist countably many non-constant analytic functions $\{g_j\}_j$ over connected open domains such that for all k , if the k -th order standard derivative of f is not defined at $x \in \mathcal{X}$, then x is in the zero set of some g_j . Next, consider the second claim. Its current form is not strong enough to enable inductive proofs. We instead prove a stronger statement by induction on k : for each $df^k \in \partial_{\bullet}^k f$, there exist countably many non-constant analytic functions $\{h_l\}_l$ over connected open domains such that the k -th order standard derivative of f is well-defined, and agrees with df^k , at all those inputs not in the zero sets of $\{h_l\}_l$. For the details, see Appendices B.3 and B.4. \square

Intensional derivatives behave better than standard derivatives. First, the intensional derivatives of a PAP function f are total functions, i.e., functions defined for all inputs (Proposition 7). Contrast this with the fact that the standard derivative of f is a partial function in general. The intensional derivatives can be understood as extensions of this standard derivative of f to those problematic inputs that make f non-differentiable (Proposition 8). The totality simplifies the reasoning about intensional derivatives. Second, the intensional derivatives satisfy a chain rule for all PAP functions (Proposition 10). Let $f : \mathcal{X} \rightarrow \mathcal{Y}$ and $g : \mathcal{Y} \rightarrow \mathbb{R}^l$ be PAP functions for some $\mathcal{X} \subseteq \mathbb{R}^n$ and $\mathcal{Y} \subseteq \mathbb{R}^m$.

Proposition 9. *$g \circ f$ is a PAP function.*

Proof. Since f and g are PAP, they have PAP representations $\gamma_f = \{\langle A^i, f^i \rangle\}_{i \in [I]}$ and $\gamma_g = \{\langle B^j, g^j \rangle\}_{j \in [J]}$. Define their composition as follows: $\gamma_g \circ \gamma_f = \{\langle C^{(i,j)}, g^j \circ f^i \rangle\}_{(i,j) \in [I] \times [J]}$ where $C^{(i,j)} = \{x \in \mathcal{X} \mid x \in A^i \wedge f^i(x) \in B^j\}$. Then, $\gamma_g \circ \gamma_f$ is a representation of $g \circ f$. Also, it is PAP as the composition of analytic functions is analytic [25, Proposition 2.2.8]. Thus, $g \circ f$ is PAP. \square

Proposition 10 (Chain Rule for Intensional Derivatives). *Let $df : \mathcal{X} \rightarrow \mathbb{R}^{m \times n}$ and $dg : \mathcal{Y} \rightarrow \mathbb{R}^{l \times m}$ be intensional derivatives of f and g (i.e., $df \in \partial_{\bullet} f$ and $dg \in \partial_{\bullet} g$). Let $h = g \circ f$. Then, the following function $dh : \mathcal{X} \rightarrow \mathbb{R}^{l \times n}$ is an intensional derivative of h (i.e., $dh \in \partial_{\bullet} h$): $dh(x) = dg(f(x)) \cdot df(x)$ for all $x \in \mathcal{X}$, where $dg(f(x)) \cdot df(x)$ is the multiplication of matrices $dg(f(x))$ and $df(x)$.*

Proof. By the definition of intensional derivative, $df = \langle\langle D\gamma_f \rangle\rangle$ and $dg = \langle\langle D\gamma_g \rangle\rangle$ for some PAP representations $\gamma_f = \{\langle A^i, f^i \rangle\}_{i \in [I]}$ and $\gamma_g = \{\langle B^j, g^j \rangle\}_{j \in [J]}$ of f and g . Let $\gamma_g \circ \gamma_f$ be the composed representation defined in the proof of Proposition 9. Then, $\gamma_g \circ \gamma_f$ is a PAP representation

of $g \circ f$. So, $\langle\langle D(\gamma_g \circ \gamma_f) \rangle\rangle$ is an intensional derivative of $g \circ f$. Also, for all $x \in \mathcal{X}$, if we let $\langle i, j \rangle \in [I] \times [J]$ with $x \in A^i$ and $f^i(x) \in B^j$, then $\langle\langle D(\gamma_g \circ \gamma_f) \rangle\rangle(x) = D(g^j \circ f^i)(x) = D(g^j)(f^i(x)) \cdot D(f^i)(x) = \langle\langle D\gamma_g \rangle\rangle(f(x)) \cdot \langle\langle D\gamma_f \rangle\rangle(x)$. The first and last equalities follow from the definition of intensional derivative. The second equality uses the chain rule for standard derivatives: the rule holds here because f^i and g^j are analytic in neighbourhoods of x and $f^i(x)$, respectively. \square

We next use these properties to show that existing autodiff systems compute intensional derivatives.

4 Correctness of Autodiff Systems

Consider a simple programming language that assumes real-valued input variables x_1, \dots, x_N and has the following syntax for programs: $e ::= c \mid x_i \mid \bar{\mathbf{f}}(e_1, \dots, e_n) \mid \text{if } (e_1 > 0) e_2 e_3$.

A program e in the language describes a real-valued computation. It is a real number c , an input variable x_i , the application of a primitive function $\bar{\mathbf{f}}$ to arguments e_1, \dots, e_n , or a conditional expression. In the third case, the applied function is required to be a PAP function of the right type (i.e., $\mathbb{R}^n \rightarrow \mathbb{R}$ in this case). We remark that practically all primitive functions supported by autodiff systems are indeed PAP; we are unaware of any non-PAP primitive function used in practice. If this requirement is met, all programs e mean PAP functions of type $\mathbb{R}^N \rightarrow \mathbb{R}$. More precisely, we interpret each program e as a function $\llbracket e \rrbracket : \mathbb{R}^N \rightarrow \mathbb{R}$ inductively, as shown below: for all $v \in \mathbb{R}^N$,

$$\begin{aligned} \llbracket c \rrbracket v &= c, & \llbracket x_i \rrbracket v &= v_i, & \llbracket \bar{\mathbf{f}}(e_1, \dots, e_n) \rrbracket v &= \mathbf{f}(\llbracket e_1 \rrbracket v, \dots, \llbracket e_n \rrbracket v), \\ \llbracket \text{if } (e_1 > 0) e_2 e_3 \rrbracket v &= \text{if } (\llbracket e_1 \rrbracket v > 0) \text{ then } \llbracket e_2 \rrbracket v \text{ else } \llbracket e_3 \rrbracket v. \end{aligned}$$

where v_i is the i -th component of the vector v , and $\mathbf{f} : \mathbb{R}^n \rightarrow \mathbb{R}$ is the PAP function denoted by the function symbol $\bar{\mathbf{f}}$ in the language. Then, the defined $\llbracket e \rrbracket$ is always PAP.

Proposition 11. *For every program e , its denotation $\llbracket e \rrbracket$ is a PAP function from \mathbb{R}^N to \mathbb{R} .*

Proof. The proof is by induction on the structure of e . The cases of $e \equiv c$ and $e \equiv x_i$ follow from the fact that both constant functions and projections are PAP. For the case of $e \equiv \bar{\mathbf{f}}(e_1, \dots, e_n)$, we note two facts about PAP functions: the composition of two PAP functions is PAP (Proposition 9); and for PAP functions $g_1, \dots, g_n : \mathbb{R}^N \rightarrow \mathbb{R}$, the function $v \mapsto \langle g_1(v), \dots, g_n(v) \rangle$ of type $\mathbb{R}^N \rightarrow \mathbb{R}^n$ is PAP again, mainly because any finite intersection of open sets is again open. By these two facts, the claimed property of the proposition holds in this case.

The only remaining case is $e \equiv (\text{if } (e_1 > 0) e_2 e_3)$. By induction hypothesis, all of $\llbracket e_1 \rrbracket$, $\llbracket e_2 \rrbracket$, and $\llbracket e_3 \rrbracket$ are PAP. Let $\gamma_1 = \{\langle A^i, f^i \rangle\}_{i \in [I]}$, $\gamma_2 = \{\langle B^j, g^j \rangle\}_{j \in [J]}$, and $\gamma_3 = \{\langle C^k, h^k \rangle\}_{k \in [K]}$ be their PAP representations, and define their conditional composition $\text{cond}(\gamma_1, \gamma_2, \gamma_3)$ as follows: $\text{cond}(\gamma_1, \gamma_2, \gamma_3) = \{\langle E^{(i,j,k,l)}, t^{(i,j,k,l)} \rangle\}_{\langle i,j,k,l \rangle \in ([I] \times [J] \times [K] \times \{0,1\})}$, $E^{(i,j,k,l)} = A_i^i \cap B^j \cap C^k$, and $t^{(i,j,k,l)} = \text{if } (l = 1) \text{ then } g^j \text{ else } h^k$, where $A_i^i = \{v \in A^i \mid f^i(v) > 0\}$ and $A_0^i = \{v \in A^i \mid f^i(v) \leq 0\}$. Then, $\{E^{(i,j,k,l)}\}_{\langle i,j,k,l \rangle}$ is an analytic partition of \mathbb{R}^N , every $t^{(i,j,k,l)}$ is an analytic function, and its domain is an open set containing $E^{(i,j,k,l)}$. Thus, $\text{cond}(\gamma_1, \gamma_2, \gamma_3)$ is a PAP representation. Furthermore, its evaluation $\langle\langle \text{cond}(\gamma_1, \gamma_2, \gamma_3) \rangle\rangle$ is equal to $\llbracket \text{if } (e_1 > 0) e_2 e_3 \rrbracket$. Hence, the proposition holds in this case. \square

We show that if an autodiff system for the language satisfies two requirements to be described shortly, it essentially computes an intensional derivative, even when the input function is not differentiable. The first requirement is that for every primitive operation $\bar{\mathbf{f}}$ of type $\mathbb{R}^n \rightarrow \mathbb{R}$, the system should come with a function $\tilde{D}\mathbf{f} : \mathbb{R}^n \rightarrow \mathbb{R}^{1 \times n}$ that satisfies $\tilde{D}\mathbf{f} \in \partial_\bullet \mathbf{f}$ and serves as a “derivative” of \mathbf{f} . To describe the second requirement, we inductively define the function $\llbracket e \rrbracket^\nabla : \mathbb{R}^N \rightarrow \mathbb{R}^{1 \times N}$, formalising symbolic differentiation of a program e , as follows: for all $v \in \mathbb{R}^n$,

$$\begin{aligned} \llbracket c \rrbracket^\nabla v &= \vec{0}_{1 \times N}, & \llbracket \bar{\mathbf{f}}(e_1, \dots, e_n) \rrbracket^\nabla v &= (\tilde{D}\mathbf{f})(\llbracket e_1 \rrbracket v, \dots, \llbracket e_n \rrbracket v) \cdot [\llbracket e_1 \rrbracket^\nabla v; \dots; \llbracket e_n \rrbracket^\nabla v], \\ \llbracket x_i \rrbracket^\nabla v &= [\vec{0}_{(i-1) \times 1}; \vec{1}_{1 \times 1}; \vec{0}_{(N-i) \times 1}]^\top, & \llbracket \text{if } (e_1 > 0) e_2 e_3 \rrbracket^\nabla v &= \text{if } (\llbracket e_1 \rrbracket v > 0) \text{ then } \llbracket e_2 \rrbracket^\nabla v \text{ else } \llbracket e_3 \rrbracket^\nabla v. \end{aligned}$$

Here $\vec{c}_{i \times j}$ denotes the $i \times j$ matrix containing only the constant c , and the RHS of the second equation means the multiplication of the $1 \times n$ matrix $(\tilde{D}\mathbf{f})(\llbracket e_1 \rrbracket v, \dots, \llbracket e_n \rrbracket v)$ and the $n \times N$ matrix $[\llbracket e_1 \rrbracket^\nabla v; \dots; \llbracket e_n \rrbracket^\nabla v]$ that is constructed by concatenating n matrices $\llbracket e_1 \rrbracket^\nabla v, \dots, \llbracket e_n \rrbracket^\nabla v$ of size $1 \times N$. The definition of $\llbracket e \rrbracket^\nabla v$ computes a “Jacobian” matrix of $\llbracket e \rrbracket$ (which is a $1 \times N$ matrix in this case) in the usual way of symbolic differentiation: run the program e under the input v , and calculate the “Jacobian” matrix of the encountered operations using the chain rule with the “derivatives” $\tilde{D}\mathbf{f}$

of those operations \mathbf{f} [1, 4, 23, 46]. Our second requirement is that given a program e and an input $v \in \mathbb{R}^N$, if the system performs forward-mode (or reverse-mode) autodiff with a tangent vector $w \in \mathbb{R}^N$ (or a cotangent vector $u \in \mathbb{R}$), it should output the Jacobian-vector product $(\llbracket e \rrbracket^\nabla v) \cdot w \in \mathbb{R}$ (or the vector-Jacobian product $u^\top \cdot (\llbracket e \rrbracket^\nabla v) \in \mathbb{R}^N$) [20, 27, 31]. Intuitively, the second requirement says that the output of an autodiff system should coincide with that of symbolic differentiation. Note that the requirement does not fix the choice of an autodiff algorithm, and has two separate conditions for forward-mode and reverse-mode algorithms. Our next results show that these requirements ensure the correctness of an autodiff system.

Theorem 12. *If $\tilde{D}\mathbf{f} \in \partial_\bullet \mathbf{f}$ for all primitive functions $\bar{\mathbf{f}}$, then $\llbracket e \rrbracket^\nabla \in \partial_\bullet \llbracket e \rrbracket$ for all programs e .*

Proof. The proof is by induction on the structure of e . When $e \equiv c$, the trivial partition $\{\mathbb{R}^N\}$ and the constant function $v \mapsto c$ form a PAP representation of $\llbracket e \rrbracket$. The intensional derivative of this representation is $\{\langle \mathbb{R}^N, v \mapsto \vec{0}_{1 \times N} \rangle\}$, and its evaluation is $\llbracket c \rrbracket^\nabla$, as claimed by the theorem. The other base case is $e \equiv x_i$. We use the trivial partition again with the projection function $v \mapsto v_i$, and form a PAP representation of $\llbracket e \rrbracket$. The intensional derivative of this representation is $\{\langle \mathbb{R}^N, v \mapsto [\vec{0}_{(i-1) \times 1}; \vec{1}_{1 \times 1}; \vec{0}_{(N-i) \times 1}]^\top \rangle\}$, and its evaluation is $\llbracket x_i \rrbracket^\nabla$.

The next case is $e \equiv \bar{\mathbf{f}}(e_1, \dots, e_n)$. By induction hypothesis, $\llbracket e_i \rrbracket^\nabla$ is an intensional derivative of $\llbracket e_i \rrbracket$ for every $i \in [n]$. Let $g : \mathbb{R}^N \rightarrow \mathbb{R}^n$ and $dg : \mathbb{R}^N \rightarrow \mathbb{R}^{n \times N}$ be functions defined by $g(v) = \langle \llbracket e_1 \rrbracket v, \dots, \llbracket e_n \rrbracket v \rangle$ and $dg(v) = \llbracket \llbracket e_1 \rrbracket^\nabla v; \dots; \llbracket e_n \rrbracket^\nabla v \rangle$ for all v . Then, dg is an intensional derivative of g . Also, $\llbracket \bar{\mathbf{f}}(e_1, \dots, e_n) \rrbracket = \mathbf{f} \circ g$. Therefore, by the chain rule for intensional derivative (Proposition 10), the function $v \mapsto (\tilde{D}\mathbf{f})(g(v)) \cdot dg(v)$ is an intensional derivative of $\llbracket \bar{\mathbf{f}}(e_1, \dots, e_n) \rrbracket = \mathbf{f} \circ g$. Here we use the assumption that $\tilde{D}\mathbf{f}$ is an intensional derivative of \mathbf{f} . Note that the function is equal to $\llbracket \bar{\mathbf{f}}(e_1, \dots, e_n) \rrbracket^\nabla$. So, the theorem holds.

The last case is $e \equiv (\text{if } (e_1 > 0) e_2 e_3)$. By induction hypothesis, $\llbracket e_i \rrbracket^\nabla$ is an intensional derivative of $\llbracket e_i \rrbracket$ for all $i \in [3]$. Let γ_1, γ_2 , and γ_3 be the PAP representations of $\llbracket e_1 \rrbracket, \llbracket e_2 \rrbracket$, and $\llbracket e_3 \rrbracket$ such that $\langle\langle D\gamma_2 \rangle\rangle = \llbracket e_2 \rrbracket^\nabla$ and $\langle\langle D\gamma_3 \rangle\rangle = \llbracket e_3 \rrbracket^\nabla$. Let $\text{cond}(\gamma_1, \gamma_2, \gamma_3)$ be the conditionally composed representation defined in the proof of Proposition 11. Then, it is a PAP representation of $\llbracket \text{if } (e_1 > 0) e_2 e_3 \rrbracket$. But by the definitions of $\gamma_2, \gamma_3, \llbracket - \rrbracket^\nabla, \text{cond}(-, -, -)$, and $\langle\langle D- \rangle\rangle$, we have $\langle\langle D\text{cond}(\gamma_1, \gamma_2, \gamma_3) \rangle\rangle(v) = \llbracket \text{if } (e_1 > 0) e_2 e_3 \rrbracket^\nabla v$ for all $v \in \mathbb{R}^N$. Hence, the theorem holds in this case as well. \square

Corollary 13. *Assume that an autodiff system for the language in this section satisfies the two requirements. Then, for each program e , there exists an intensional derivative df of $\llbracket e \rrbracket$ such that if the system performs forward-mode (or reverse-mode) autodiff with a tangent vector $w \in \mathbb{R}^N$ (or a cotangent vector $u \in \mathbb{R}$), it computes the Jacobian-vector product $df(v) \cdot w \in \mathbb{R}$ (or the vector-Jacobian product $u^\top \cdot df(v) \in \mathbb{R}^N$) for every input $v \in \mathbb{R}^N$. Furthermore, the computed entity is the corresponding Jacobian-vector product (or vector-Jacobian product) with the standard derivative of $\llbracket e \rrbracket$ for almost all inputs $v \in \mathbb{R}^N$.*

Proof. Suppose that the two requirements are met, and let e be a program. Consider the case when the system performs forward-mode autodiff. Let $df = \llbracket e \rrbracket^\nabla$. Then, by Theorem 12, df is an intensional derivative in $\partial_\bullet \llbracket e \rrbracket$. Moreover, by the second requirement, the output of the system for e with a tangent vector $w \in \mathbb{R}^N$ is $(\llbracket e \rrbracket^\nabla v) \cdot w = df(v) \cdot w \in \mathbb{R}$ for all inputs $v \in \mathbb{R}^N$. This proves the first part of the corollary. The other part of the corollary follows immediately from Proposition 8. The proof for the case when reverse-mode autodiff is performed is essentially the same, so we omit it. \square

Remark 1 (Intensional Derivatives in Practice). We briefly discuss whether the first requirement mentioned earlier is indeed met by autodiff systems. TensorFlow [2] and PyTorch [33], two popular autodiff systems, support a wide range of primitive functions that are not differentiable for some inputs. A well-known example is `relu` : $\mathbb{R} \rightarrow \mathbb{R}$ (i.e., the map $x \mapsto \max(0, x)$). This function is not differentiable at 0, but when these autodiff systems are applied to differentiate `relu` at 0, they return 0 instead of an error (i.e., $(\tilde{D}\text{relu})(0) = 0$). It means that the systems compute the intensional derivative $\langle\langle D\gamma \rangle\rangle$ of `relu` for $\gamma = \{\langle \mathbb{R}_{>0}, x \mapsto x \rangle, \langle \mathbb{R}_{\leq 0}, x \mapsto 0 \rangle\}$. Thus, they fulfil the first requirement on `relu`. Note that setting the derivative at the input 0 to any $c \in \mathbb{R}$ is an acceptable option; it can be justified by a different PAP representation γ' of `relu` where $\gamma' = \{\langle \mathbb{R}_{>0}, x \mapsto x \rangle, \langle \{0\}, x \mapsto c \cdot x \rangle, \langle \mathbb{R}_{<0}, x \mapsto 0 \rangle\}$. Another example is `reciprocal_no_nan` from TensorFlow. It means the map $f(x) = 1/x$ if $x \neq 0$, and 0 if $x = 0$. Differentiating this map by TensorFlow gives the function $(\tilde{D}f)(x) = -1/x^2$ if $x \neq 0$, and 0 if $x = 0$. The $\tilde{D}f$ is an intensional derivative of f , because $\gamma = \{\langle \mathbb{R} \setminus \{0\}, x \mapsto 1/x \rangle, \langle \{0\}, x \mapsto 0 \rangle\}$ is a PAP representation of f and $\langle\langle D\gamma \rangle\rangle$ coincides with $\tilde{D}f$. Hence, TensorFlow meets the first requirement on `reciprocal_no_nan`.

However, for some functions, TensorFlow and PyTorch behave more conservatively than our theory of PAP functions permits. A good example is `sqrt`. The domain of `sqrt` is $\mathbb{R}_{\geq 0}$, but differentiating `sqrt` at 0 using these autodiff systems returns `+inf`, which means ∞ (i.e., $(\bar{D}\text{sqrt})(0) = \infty$). Thus, the first requirement on `sqrt` is violated by these systems. One consequence of this is that when the systems are applied to differentiate `sqrt(mult(x, 0))`, they end up with computing `0*(+inf)` that evaluates to `NaN` (meaning an undefined number in floating point arithmetic) for all $x \in \mathbb{R}$. Note that `sqrt(mult(x, 0))` is constantly zero and has zero as its derivative. This conservative behaviour could have been avoided if TensorFlow and PyTorch had recognised that `sqrt` is a PAP function and had used its intensional derivative instead. For instance, the function $(d\text{sqrt})(x) = 1/(2\sqrt{x})$ if $x > 0$, and 0 if $x = 0$, is an intensional derivative of `sqrt`, since $\gamma = \{(\mathbb{R}_{>0}, x \mapsto \sqrt{x}), (\{0\}, x \mapsto 0)\}$ is a PAP representation of `sqrt`. If TensorFlow and PyTorch used `dsqrt` as a derivative of `sqrt`, they would be able to differentiate `sqrt(mult(x, 0))` correctly for all $x \in \mathbb{R}$.

5 Related Work and Discussion

Autodiff has a long history with a large body of literature [4, 20, 34]. Its community has been aware of some issues with non-differentiable functions [20, Chapter 14]. These issues have become ever more important, as autodiff has been increasingly applied to a variety of non-differentiable functions, including sophisticated linear algebra functions [24, 27, 40]. In this paper, we investigate the issues in a more systematic and rigorous way, by presenting non-trivial concrete counterexamples that illuminate subtleties of non-differentiable functions in autodiff, and also proposing intensional derivatives, a new notion of derivatives, that enable us to formally prove the correctness of, and better understand the behaviour of, autodiff systems applied to non-differentiable functions.

Clarke subdifferential [12, 13] and related notions have been favourite tools for studying the use of non-differentiable functions in the context of autodiff and stochastic gradient descent [6, 15, 21, 24, 29]. Let f be an input program (or function) to be differentiated. Assume that f is locally Lipschitz continuous but possibly non-differentiable at some inputs. Kakade and Lee [24] proposed a linear-time algorithm that correctly computes a *subdifferential* of f at x *almost surely*, for *all* inputs x , when f satisfies some qualification condition. Their work, however, could not provide the correctness of the current autodiff systems, because those systems do not use their algorithm. Bolte and Pauwels [6] presented a result closest to ours: standard autodiff algorithms correctly compute the *standard derivative* of f at x *always*, for *almost all* inputs x , if f is definable [43]. To do so, they extended Clarke subdifferential to define a new notion of derivatives, called conservative fields, and showed that autodiff algorithms compute conservative fields and the fields coincide with the standard derivatives almost everywhere. We point out three key differences between these works and ours. First, they use highly advanced mathematical tools (e.g., variational analysis [36]). Thanks to those powerful tools, they could prove more results than ours, such as the almost-sure convergence of the stochastic gradient descent with conservative fields [6], but they are less accessible than our work, which requires much simpler mathematics (probably at an undergraduate level) but still gives the correctness proof of autodiff systems. Second, subdifferentials and conservative fields are difficult to be generalised to higher-order derivatives, because they are defined in terms of a *single set-valued* function. In contrast, our intensional derivatives are defined in terms of a *set of standard* (i.e., singleton-valued) functions, and thus have a natural generalisation to higher-order derivatives. Lastly, different function classes are considered: those works apply to locally Lipschitz continuous functions with additional properties (e.g., definability), whereas ours applies to PAP functions. We point out that PAP functions include many non-locally-Lipschitz functions, such as discontinuous $x \in \mathbb{R} \mapsto \mathbb{1}[x > 0]$ and continuous $x \in \mathbb{R} \mapsto \sqrt{|x|}$; the opposite direction is also true (e.g., the λ -Cantor functions for $\lambda \in (0, 1)$).

The PAP property is similar to some other piecewise properties, such as piecewise smoothness [47] which inspired PAP, and piecewise analyticity [20]. The PAP, however, differs from the others in a few ways: unlike piecewise smoothness, it allows piecewise functions to be defined on open subsets of \mathbb{R}^n (not just on \mathbb{R}^n), but restricts them to be analytic (rather than smooth); unlike piecewise analyticity, it allows piecewise domains to be countably many (not just finitely many), but restricts them to be analytic (rather than unrestricted). Unlike prior works, we also give an in-depth study of the PAP property, such as its closedness under operations (e.g., composition) and its relationship to other properties (e.g., almost-everywhere differentiability). The property of almost-everywhere differentiability has been studied as well for programs and functions (that need not be related to programs). E.g., it is shown to hold for all programs in some programming languages [20, 30], and

for all piecewise smooth functions [47]. We prove the property for all PAP functions and utilise it in proving that higher-order intensional derivatives agree with standard derivatives almost everywhere.

Recently and concurrently with this work, Bolte and Pauwels [7] studied some concepts and results similar to ours. They proposed a new class of functions and a new notion of derivatives, called elementary selections and selection derivatives, which roughly correspond to our PAP functions and intensional derivatives; and proved properties of those new functions and derivatives, which roughly correspond to our Propositions 8 to 11 and Theorem 12. Although having some similarities, our work and their work have three key differences, complementing each other. First, their work is applicable to a strictly smaller class of functions than ours, as any elementary selection is PAP (and locally Lipschitz) but not vice versa. Second, it considers selection derivatives of first order only, whereas our work considers intensional derivatives of higher orders as well. Third, their work provides some results not in our work (e.g., convergence of stochastic gradient descent with selection derivatives), and vice versa (e.g., the results in §2).

We conclude the section by discussing a few remarks on a notion of correctness for autodiff systems. First, there can be multiple different correctness conditions of autodiff systems and we do not claim that our correctness condition (Definition 1 or Corollary 13) is “the” gold standard. Rather we are just suggesting “a” correctness condition that can serve as a reasonable (possibly minimal) requirement for existing and future autodiff systems. Second, as important as the correctness of autodiff systems is the correctness of the applications built upon autodiff systems (e.g., gradient descent, variational inference, and Hamiltonian Monte Carlo), but the latter might not necessarily follow from the former. For example, consider an autodiff system that is correct in the sense of Corollary 13 (i.e., computes an intensional derivative), and a gradient descent algorithm that follows the “derivative”, computed by the system, of an objective function $f : \mathbb{R} \rightarrow \mathbb{R}$ to minimise f . Suppose that f is given by $f(x) = x^2$ if $x \neq 1$, and 1 if $x = 1$; the autodiff system computes $df(x) = 2x$ if $x \neq 1$, and 0 if $x = 1$, as the “derivative” of f ; and the gradient descent starts at $x = 1$.¹ Then, the gradient descent gets stuck at $x = 1$ forever since $df(1) = 0$, even though 1 is never a critical point of f . This example illustrates that the correctness of autodiff systems does not necessarily imply the correctness of gradient descents built upon those systems (where the latter correctness is given by that gradient descents should converge to Clarke critical points).² Nevertheless, the two notions of correctness discussed so far, one for autodiff systems and the other for applications using autodiff systems, address two separate issues and this paper is mainly about the former notion of correctness.

Broader Impact

This work focuses mainly on theoretical aspects of autodiff systems. In particular, we formally prove that the systems, though developed to handle differentiable functions, remain correct even when applied to non-differentiable functions. Our result justifies, at least in part, the current situation in machine learning, in which the systems are frequently applied to non-differentiable functions without much consideration to their correctness under such out-of-scope use cases. Other than the justification, this work does not present any other foreseeable societal consequence due to its theoretical nature.

Acknowledgments and Disclosure of Funding

We thank anonymous reviewers for their insightful and constructive comments. Lee, Yang, and Yu were supported by the Engineering Research Center Program through the National Research Foundation of Korea (NRF) funded by the Korean Government MSIT (NRF-2018R1A5A1059921), and also by Next-Generation Information Computing Development Program through the National Research Foundation of Korea (NRF) funded by the Ministry of Science, ICT (2017M3C4A7068177). Rival was supported by a Facebook gift and by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No 825492).

¹We thank an anonymous reviewer for coming up with this example.

²On the other hand, we conjecture that for any PAP and locally Lipschitz $f : \mathbb{R}^n \rightarrow \mathbb{R}$, if a gradient descent follows an intensional derivative of f and starts at *randomly chosen* initial points, then it would be correct *almost surely*. A partial result confirming the conjecture was proven recently in [7], concurrently with this work.

References

- [1] M. Abadi and G. D. Plotkin. A simple differentiable programming language. *Proceedings of the ACM on Programming Languages*, 4(POPL):38:1–38:28, 2020.
- [2] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, M. Kudlur, J. Levenberg, R. Monga, S. Moore, D. G. Murray, B. Steiner, P. A. Tucker, V. Vasudevan, P. Warden, M. Wicke, Y. Yu, and X. Zheng. TensorFlow: A system for large-scale machine learning. In *Operating Systems Design and Implementation (OSDI)*, pages 265–283, 2016.
- [3] A. G. Baydin, B. A. Pearlmutter, and J. M. Siskind. Diffsharp: An AD library for .NET languages. In *International Conference on Algorithmic Differentiation (AD)*, 2016. Also arXiv preprint arXiv:1611.03423.
- [4] A. G. Baydin, B. A. Pearlmutter, A. A. Radul, and J. M. Siskind. Automatic differentiation in machine learning: A survey. *Journal of Machine Learning Research*, 18:153:1–153:43, 2017.
- [5] J. Bergstra, O. Breuleux, F. Bastien, P. Lamblin, R. Pascanu, G. Desjardins, J. Turian, D. Warde-Farley, and Y. Bengio. Theano: A CPU and GPU math compiler in Python. In *Python in Science Conference (SciPy)*, pages 18–24, 2010.
- [6] J. Bolte and E. Pauwels. Conservative set valued fields, automatic differentiation, stochastic gradient method and deep learning. *Mathematical Programming*, 2020.
- [7] J. Bolte and E. Pauwels. A mathematical model for automatic differentiation in machine learning. *arXiv preprint arXiv:2006.02080*, 2020.
- [8] A. Boyarsky and P. Gora. *Laws of chaos: Invariant measures and dynamical systems in one dimension*. Springer Science & Business Media, 1997.
- [9] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, and S. Wanderman-Milne. JAX: Composable transformations of Python+NumPy programs, 2018. Refer to <http://github.com/google/jax> and [17].
- [10] A. M. Bruckner, J. B. Bruckner, and B. S. Thomson. *Real analysis*. Prentice-Hall, first edition, 1997.
- [11] G. Cantor. De la puissance des ensembles parfaits de points: Extrait d’une lettre adressée à l’éditeur [On the power of perfect sets of points]. *Acta Mathematica*, 4:381–392, 1884. Translated in English, e.g., in [16, Section 2].
- [12] F. H. Clarke. *Optimization and nonsmooth analysis*. Classics in Applied Mathematics: Volume 5. SIAM, 1990.
- [13] F. H. Clarke, Y. S. Ledyae, R. J. Stern, and P. R. Wolenski. *Nonsmooth analysis and control theory*. Graduate Texts in Mathematics: Volume 178. Springer Science & Business Media, 1998.
- [14] R. B. Darst. Some Cantor sets and Cantor functions. *Mathematics Magazine*, 45(1):2–7, 1972.
- [15] D. Davis, D. Drusvyatskiy, S. M. Kakade, and J. D. Lee. Stochastic subgradient method converges on tame functions. *Foundations of Computational Mathematics*, 20(1):119–154, 2020.
- [16] G. A. Edgar. *Classics on fractals*. CRC Press, 2019.
- [17] R. Frostig, M. Johnson, and C. Leary. Compiling machine learning programs via high-level tracing. In *SysML Conference*, 2018.
- [18] B. R. Gelbaum and J. M. H. Olmsted. *Counterexamples in analysis*. Dover Publications, 2003.
- [19] I. Goodfellow, Y. Bengio, and A. Courville. *Deep learning*. MIT Press, 2016. <http://www.deeplearningbook.org>.

- [20] A. Griewank and A. Walther. *Evaluating derivatives: Principles and techniques of algorithmic differentiation*. SIAM, second edition, 2008.
- [21] A. Griewank, A. Walther, S. Fiege, and T. Bosse. On Lipschitz optimization based on gray-box piecewise linearization. *Mathematical Programming*, 158(1-2):383–415, 2016.
- [22] L. Hascoët and V. Pascual. The Taped automatic differentiation tool: Principles, model, and specification. *ACM Transactions on Mathematical Software*, 39(3):20:1–20:43, 2013.
- [23] M. Huot, S. Staton, and M. Vákár. Correctness of automatic differentiation via diffeologies and categorical gluing. In *Foundations of Software Science and Computation Structures (FoSSaCS)*, pages 319–338, 2020.
- [24] S. M. Kakade and J. D. Lee. Provably correct automatic sub-differentiation for qualified programs. In *Neural Information Processing Systems (NeurIPS)*, pages 7125–7135, 2018.
- [25] S. G. Krantz and H. R. Parks. *A primer of real analytic functions*. Springer Science & Business Media, second edition, 2002.
- [26] Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015.
- [27] D. Maclaurin. *Modeling, inference and optimization with composable differentiable procedures*. PhD thesis, Harvard University, 2016.
- [28] D. Maclaurin, D. Duvenaud, and R. P. Adams. Autograd: Effortless gradients in Numpy. In *ICML AutoML Workshop*, 2015.
- [29] S. Majewski, B. Miasojedow, and E. Moulines. Analysis of nonsmooth stochastic approximation: The differential inclusion approach. *arXiv preprint arXiv:1805.01916*, 2018.
- [30] C. Mak, C. L. Ong, H. Paquet, and D. Wagner. Densities of almost-surely terminating probabilistic programs are differentiable almost everywhere. *arXiv preprint arXiv:2004.03924*, 2020.
- [31] C. C. Margossian. A review of automatic differentiation and its efficient implementation. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 9(4), 2019.
- [32] B. Mityagin. The zero set of a real analytic function. *arXiv preprint arXiv:1512.07276*, 2015.
- [33] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Köpf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala. PyTorch: An imperative style, high-performance deep learning library. In *Neural Information Processing Systems (NeurIPS)*, pages 8024–8035, 2019.
- [34] B. A. Pearlmutter and J. M. Siskind. Reverse-mode AD in a functional framework: Lambda the ultimate backpropagator. *ACM Transactions on Programming Languages and Systems*, 30(2):7:1–7:36, 2008.
- [35] J. Revels, M. Lubin, and T. Papamarkou. Forward-mode automatic differentiation in Julia. *arXiv preprint arXiv:1607.07892*, 2016.
- [36] R. T. Rockafellar and R. J.-B. Wets. *Variational analysis*. A Series of Comprehensive Studies in Mathematics: Volume 317. Springer Science & Business Media, 1998.
- [37] W. Rudin. *Principles of mathematical analysis*. McGraw-Hill, third edition, 1976.
- [38] D. E. Rumelhart, G. E. Hinton, and R. J. Williams. Learning representations by back-propagating errors. *Nature*, 323(6088):533–536, 1986.
- [39] J. Schmidhuber. Deep learning in neural networks: An overview. *Neural Networks*, 61:85–117, 2015.
- [40] M. W. Seeger, A. Hetzel, Z. Dai, and N. D. Lawrence. Auto-differentiating linear algebra. *arXiv preprint arXiv:1710.08717*, 2017.

- [41] T. Tao. *An introduction to measure theory*. Graduate Studies in Mathematics: Volume 126. American Mathematical Society, 2011.
- [42] S. Tokui, R. Okuta, T. Akiba, Y. Niitani, T. Ogawa, S. Saito, S. Suzuki, K. Uenishi, B. Vogel, and H. Y. Vincent. Chainer: A deep learning framework for accelerating the research cycle. In *International Conference on Knowledge Discovery & Data Mining (KDD)*, pages 2002–2011, 2019.
- [43] L. van den Dries and C. Miller. Geometric categories and o-minimal structures. *Duke Mathematical Journal*, 84(2):497–540, 1996.
- [44] B. van Merriënboer, D. Moldovan, and A. B. Wiltschko. Tangent: Automatic differentiation using source-code transformation for dynamically typed array programming. In *Neural Information Processing Systems (NeurIPS)*, pages 6259–6268, 2018.
- [45] A. Walther and A. Griewank. Getting started with ADOL-C. In *Combinatorial Scientific Computing*, chapter 7, pages 181–202. Chapman & Hall/CRC Computational Science, 2012.
- [46] F. Wang, D. Zheng, J. M. Decker, X. Wu, G. M. Essertel, and T. Rompf. Demystifying differentiable programming: Shift/reset the penultimate backpropagator. *Proceedings of the ACM on Programming Languages*, 3(ICFP):96:1–96:31, 2019.
- [47] Y. Zhou, B. J. Gram-Hansen, T. Kohn, T. Rainforth, H. Yang, and F. Wood. LF-PPL: A low-level first order probabilistic programming language for non-differentiable models. In *Artificial Intelligence and Statistics (AISTATS)*, pages 148–157, 2019.

A Comments on Results in §2

A.1 Comments on the proof of Proposition 1

First, we elaborate on our proof in Proposition 1 that g is continuous on $(0, 1)$. Since g is continuous on $(0, 1) \setminus C_1$ by its construction, we only need to show that g is continuous on C_1 . Consider any $x \in C_1$ and $\epsilon > 0$. It suffices to show that there is $\delta \in (0, x)$ such that

$$0 < |x - x'| < \delta \implies |g(x) - g(x')| = |g(x')| = g(x') < \epsilon. \quad (1)$$

Let $k > 0$ be an integer with $2^{-k} < \epsilon$. Consider the set

$$S = \{x' \in (0, 1) \setminus C_1 \mid g(x') \geq 2^{-k}\}.$$

By the construction of g , S is the union of some finitely many closed intervals in $(0, 1)$ that do not contain x . (Note that each of those closed intervals is contained in an open interval removed at some $k'(\leq k)$ -th step of g 's construction.) Hence,

$$\delta = \min \left(\inf_{x' \in S} |x - x'|, x \right)$$

is positive. We now show that δ satisfies (1). Consider any x' with $0 < |x - x'| < \delta$. If $x' \in C_1$, then $g(x') = 0$. If $x' \notin C_1$, then $x' \in (0, 1) \setminus C_1$ and $x' \notin S$ by the definition of δ , and thus $g(x') < 2^{-k} < \epsilon$ by the definition of S and k . Hence, (1) holds and this completes the proof. \square

Second, we elaborate on our proof in Proposition 1 that $g \circ f$ is not differentiable on $C_{1/2}$. Consider any $x \in C_{1/2}$. It suffices to show that for any $\delta \in (0, x)$, there exist $x_1, x_2 \in (x - \delta, x + \delta) \setminus \{x\}$ such that

$$\left| \frac{(g \circ f)(x) - (g \circ f)(x_1)}{x - x_1} \right| = 0 \quad \text{and} \quad \left| \frac{(g \circ f)(x) - (g \circ f)(x_2)}{x - x_2} \right| > 1. \quad (2)$$

Consider any $\delta \in (0, x)$. Since $x \in C_{1/2}$ is a limit point of $C_{1/2}$, there exists $x_1 \in (x - \delta, x + \delta) \setminus \{x\}$ with $x_1 \in C_{1/2}$. For this x_1 , the first equality in (2) holds, since $(g \circ f)(C_{1/2}) = g(C_1) = \{0\}$. To find x_2 , let $k > 0$ be an integer such that

$$d(k) = \frac{1}{2} \cdot \frac{1}{2^k} + \frac{3}{4} \cdot \frac{1}{3^k} < \delta \quad \text{and} \quad \frac{3}{4} \cdot \left(\frac{2}{3}\right)^k < \frac{1}{2}. \quad (3)$$

We claim that there exists $x_2 \in (0, 1)$ such that

$$0 < |x - x_2| \leq d(k) \quad \text{and} \quad (g \circ f)(x_2) = 2^{-k}. \quad (4)$$

If the claim holds, then

$$\begin{aligned} \left| \frac{(g \circ f)(x) - (g \circ f)(x_2)}{x - x_2} \right| &\geq 2^{-k}/d(k) && \text{(by (4) and } (g \circ f)(x) = 0) \\ &= 2^{-k} / \left(\frac{1}{2} \cdot \frac{1}{2^k} + \frac{3}{4} \cdot \frac{1}{3^k} \right) && \text{(by the equality in (3))} \\ &= 1 / \left(\frac{1}{2} + \frac{3}{4} \cdot \left(\frac{2}{3}\right)^k \right) \\ &> 1 / \left(\frac{1}{2} + \frac{1}{2} \right) = 1 && \text{(by the inequality in (3)),} \end{aligned}$$

and thus the second equality in (2) holds. Hence, finding $x_2 \in (0, 1)$ satisfying (4) completes the proof. We now show that such x_2 exists. Consider the situation right after the k -th step of $C_{1/2}$'s construction is performed. Then, the total length of the closed intervals that still remain is

$$1 - \frac{1}{2} \cdot \left(\frac{1}{3^1} \cdot 2^0 + \frac{1}{3^2} \cdot 2^1 + \cdots + \frac{1}{3^k} \cdot 2^{k-1} \right) = \frac{1}{2} \cdot \left(1 + \left(\frac{2}{3}\right)^k \right),$$

so the length of each of those closed intervals is $\frac{1}{2}(2^{-k} + 3^{-k})$, since those closed intervals have the same length and there are 2^k such intervals. Due to this, and by the construction of $C_{1/2}$, there is some open interval I that is removed exactly at the k -th step of $C_{1/2}$'s construction and satisfies

$$\text{dist}(x, I) \leq \frac{1}{2} \cdot \left(\frac{1}{2^k} + \frac{1}{3^k} \right).$$

Let $x_2 \in (0, 1)$ be the midpoint of I . By the construction of f and g , we have $(g \circ f)(x_2) = 2^{-k}$. Furthermore, since the length of I is $3^{-k}/2$, we have

$$|x - x_2| \leq \text{dist}(x, I) + \frac{1}{2} \cdot \text{len}(I) \leq \frac{1}{2} \cdot \left(\frac{1}{2^k} + \frac{1}{3^k} \right) + \frac{1}{4} \cdot \frac{1}{3^k} = d(k).$$

Hence, $x_2 \in (0, 1)$ satisfies (4), and this concludes the proof. \square

Next, we make a remark on non-differentiable inputs of f , g , and $g \circ f$ in the proof. One might guess that f should be non-differentiable exactly on $C_{1/2}$, given that f maps $(0, 1) \setminus C_{1/2}$ onto $(0, 1) \setminus C_1$ in a linear way and maps $C_{1/2}$ onto C_1 in a non-smooth-looking way. Surprisingly, the guess is wrong: f is in fact non-differentiable only on a measure-zero subset of $C_{1/2}$. On the other hand, g and $g \circ f$ are non-differentiable exactly on C_1 and $C_{1/2}$, respectively. The proof that g is non-differentiable on C_1 is similar to the above proof that $g \circ f$ is non-differentiable on $C_{1/2}$, and thus we omit it.

Finally, we connect the examples in the proof with our results in §3. Both f and g are shown to be non-PAP (Proposition 5). Hence, our results do not guarantee that $g \circ f$ is PAP and so almost-everywhere differentiable. In fact, $g \circ f$ is non-PAP, since $g \circ f$ is not almost-everywhere differentiable.

A.2 Comments on the proof of Proposition 2

We explain how the counterexample in the proof does not contradict to our results in §3. The functions f and g are PAP (and thus $g \circ f$ is so). Although g' is undefined at 0, we can extend it to an intensional derivative $dg \in \partial_\bullet g$ such that dg is defined everywhere (even at 0) and coincides with g' at all but countably many inputs. With such dg , the following version of the chain rule holds almost everywhere:

$$(g \circ f)'(x) = dg(f(x)) \cdot f'(x) \quad \text{for almost all } x \in (0, 1).$$

This is because we have the chain rule for intensional derivatives and these intensional derivatives and standard derivatives coincide almost everywhere (Propositions 8 and 10).

A.3 Comments on the proof of Proposition 3

The functions f , g , and $g \circ f$ in the proof do not contradict to our results. Neither f nor g is a PAP function (Proposition 5). Hence, our results do not guarantee the validity of our version of the chain rule for $g \circ f$.

B Comments on and Proofs for Results in §3

Let $\mathcal{X}, \mathcal{X}_f, \mathcal{X}_g \subseteq \mathbb{R}^n$ and $\mathcal{Y} \subseteq \mathbb{R}^m$ be arbitrary sets.

B.1 Comments on the proof of Proposition 5

We prove the following argument used in the proof of Proposition 5: the functions listed in the proof satisfy the sufficient condition (i) or (ii) mentioned in the proof.

First, consider (i). To show that a function $h : \mathcal{X} \rightarrow \mathbb{R}^m$ satisfies (i) with $k = 2$, it suffices to show the claim that the set

$$S = \{x \in \mathcal{X} \mid h' \text{ is undefined or discontinuous at } x\}$$

has positive measure. For the λ -Cantor function ϕ_λ with $\lambda \in (0, 1)$, S is a full measure subset of C_λ due to the following: $\phi'_\lambda(x) = 1/(1 - \lambda) \neq 0$ for almost all $x \in C_\lambda$; $\phi'_\lambda(x) = 0$ for all $x \notin C_\lambda$; and $C_\lambda \subset (0, 1)$ has no interior. Since C_λ has measure $1 - \lambda > 0$, the claim holds for ϕ_λ . For f in the

proof of Proposition 1, S is a full measure subset of $C_{1/2}$ due to similar reasons. So the claim holds for f . For Volterra's function, S is known to have positive measure [18, Example 8.35]. So the claim holds for Volterra's function.

Next, consider (ii). Observe that the 1-Cantor function and g in the proof of Proposition 1 are both defined on $(0, 1) \subseteq \mathbb{R}$ and non-differentiable exactly on C_1 (for the non-differentiability of g , see §A.1). Since C_1 is uncountable, the two functions satisfy (ii) with $k = 1$. This completes the proof of the argument. \square

As a side note, we remark that Volterra's function V is more pathological than the λ -Cantor function ϕ_λ for $\lambda \in (0, 1)$, in that V' is discontinuous on a set of positive measure even though V is differentiable everywhere and V' is bounded. Contrast this with the fact that ϕ'_λ is also discontinuous on a set of positive measure, but ϕ_λ is differentiable just almost everywhere, not everywhere. In fact, there even exists a more pathological function $W : (0, 1) \rightarrow \mathbb{R}$ such that W is differentiable everywhere and W' is bounded, but W' is discontinuous almost everywhere [10, Exercise 5:5.5]. Certainly, W is another example for Proposition 5: it is continuous and differentiable almost everywhere, but not PAP.

B.2 Interior and subinterior of analytic partition

Definition 9 (Interior of Analytic Partition). *Let $A = \{A^i\}_{i \in [I]}$ be an analytic partition of \mathcal{X} . The interior of A , $\text{int}(A)$, is defined by the following open subset of \mathbb{R}^n :*

$$\text{int}(A) = \bigcup_{i \in [I]} \text{int}(A^i; \mathbb{R}^n),$$

where $\text{int}(A^i; \mathbb{R}^n)$ denotes the largest open subset of \mathbb{R}^n that is included in A^i .

Definition 10 (Subinterior of Analytic Partition). *Let $A = \{A^i\}_{i \in [I]}$ be an analytic partition of \mathcal{X} with $I \in (\mathbb{Z}_{>0} \cup \{\infty\})$. Suppose that a subset $\mathcal{X}' \subseteq \mathcal{X}$ and a partition $B = \{B^t\}_{t \in [T]}$ of \mathcal{X}' satisfy the following conditions:*

- (i) $T \in (\mathbb{Z}_{>0} \cup \{\infty\})$.
- (ii) For all $t \in [T]$, B^t is **subanalytic**. That is, for all $t \in [T]$, there exist $J_t, L_t \in \mathbb{Z}_{\geq 0}$ and analytic functions $g_{t,j}^+ : \mathcal{X}_{t,j}^+ \rightarrow \mathbb{R}$ and $g_{t,l}^- : \mathcal{X}_{t,l}^- \rightarrow \mathbb{R}$ over open domains $\mathcal{X}_{t,j}^+, \mathcal{X}_{t,l}^- \subseteq \mathbb{R}^n$ ($j \in [J_t]$, $l \in [L_t]$) such that: $\mathcal{X}_{t,l}^-$ is connected in \mathbb{R}^n and $(g_{t,l}^-)^{-1}(\{0\}) \neq \mathcal{X}_{t,l}^-$ for all $l \in [L_t]$; and

$$B^t = \left(\bigcap_{j \in [J_t]} (g_{t,j}^+)^{-1}(\mathbb{R}_{>0}) \right) \cap \left(\bigcap_{l \in [L_t]} (g_{t,l}^-)^{-1}(\mathbb{R}_{<0}) \right). \quad (5)$$

- (iii) Let $C^t \subseteq \mathbb{R}^n$ be the set defined as follows:

$$C^t = \left(\bigcap_{j \in [J_t]} (g_{t,j}^+)^{-1}(\mathbb{R}_{>0}) \right) \cap \left(\bigcap_{l \in [L_t]} (g_{t,l}^-)^{-1}(\mathbb{R}_{\leq 0}) \right).$$

Then, $C = \{C^t\}_{t \in [T]}$ is a finer partition of \mathcal{X} than $\{A^i\}_{i \in [I]}$. That is, C is a partition of \mathcal{X} , and for all $t \in [T]$, $C^t \subseteq A^i$ for some $i \in [I]$.

We call the set \mathcal{X}' a **subinterior** of A , and the partition B a **subanalytic** partition of \mathcal{X}' . We use $\text{subint}(A)$ to denote the set of all subinteriors of A .

Lemma 14. *Let $A = \{A^i\}_{i \in [I]}$ be an analytic partition of \mathcal{X} with $I \in (\mathbb{Z}_{>0} \cup \{\infty\})$. Then, $\text{subint}(A) \neq \emptyset$. Furthermore, for any $\mathcal{X}' \in \text{subint}(A)$, the following hold:*

- (a) There is an analytic partition $B = \{B^t\}_{t \in [T]}$ of \mathcal{X}' with $T \in (\mathbb{Z}_{>0} \cup \{\infty\})$ and $\text{int}(B) = \mathcal{X}'$.
- (b) $\mathcal{X}' \subseteq \text{int}(A)$ and \mathcal{X}' is open in \mathbb{R}^n .
- (c) $\mathcal{X} \setminus \mathcal{X}'$ is contained in some measure-zero set.
- (d) $\mathcal{X} \setminus \text{int}(A)$ is contained in some measure-zero set.

Proof. We first prove that $\text{subint}(A) \neq \emptyset$. Consider any $i \in [I]$. Since A^i is analytic, there exist $J_i, L_i \in \mathbb{Z}_{>0}$ and analytic functions $g_{i,j}^+ : \mathcal{X}_{i,j}^+ \rightarrow \mathbb{R}$ and $g_{i,l}^- : \mathcal{X}_{i,l}^- \rightarrow \mathbb{R}$ over open domains $\mathcal{X}_{i,j}^+, \mathcal{X}_{i,l}^- \subseteq \mathbb{R}^n$ ($j \in [J_i]$, $l \in [L_i]$) such that

$$A^i = \left(\bigcap_{j \in [J_i]} (g_{i,j}^+)^{-1}(\mathbb{R}_{>0}) \right) \cap \left(\bigcap_{l \in [L_i]} (g_{i,l}^-)^{-1}(\mathbb{R}_{\leq 0}) \right).$$

We use the following fact: any open set in \mathbb{R}^n is a union of countably many *open balls* in \mathbb{R}^n , thereby a union of countably many *disjoint connected* open sets in \mathbb{R}^n . For every $l \in [L_i]$, since $\mathcal{X}_{i,l}^-$ is open in \mathbb{R}^n , there exists a partition $\{\mathcal{X}_{i,\langle l,t_l \rangle}^-\}_{t_l \in [\infty]}$ of $\mathcal{X}_{i,l}^-$ such that $\mathcal{X}_{i,\langle l,t_l \rangle}^-$ is connected and open in \mathbb{R}^n . For each $\langle t_1, \dots, t_{L_i} \rangle \in [\infty]^{L_i}$, let

$$B^{i,\langle t_1, \dots, t_{L_i} \rangle} = \left(\bigcap_{j \in [J_i]} (g_{i,j}^+)^{-1}(\mathbb{R}_{>0}) \right) \cap \left(\bigcap_{l \in [L_i]} (g_{i,\langle l,t_l \rangle}^-)^{-1}(\mathbb{R}_{<0}) \right)$$

$$C^{i,\langle t_1, \dots, t_{L_i} \rangle} = \left(\bigcap_{j \in [J_i]} (g_{i,j}^+)^{-1}(\mathbb{R}_{>0}) \right) \cap \left(\bigcap_{l \in [L_i]} (g_{i,\langle l,t_l \rangle}^-)^{-1}(\mathbb{R}_{\leq 0}) \right)$$

where $g_{i,\langle l,t_l \rangle}^- : \mathcal{X}_{i,\langle l,t_l \rangle}^- \rightarrow \mathbb{R}$ denotes the restriction of $g_{i,l}^-$ to $\mathcal{X}_{i,\langle l,t_l \rangle}^-$. Here, if $(g_{i,\langle l,t_l \rangle}^-)^{-1}(\{0\}) = \mathcal{X}_{i,\langle l,t_l \rangle}^-$ for some i, l, t_l , then we set $g_{i,\langle l,t_l \rangle}^-$ to the constant function -1 on the domain $\mathcal{X}_{i,\langle l,t_l \rangle}^-$. Then, every $g_{i,\langle l,t_l \rangle}^-$ is analytic on its connected open domain $\mathcal{X}_{i,\langle l,t_l \rangle}^-$, and $(g_{i,\langle l,t_l \rangle}^-)^{-1}(\{0\}) \neq \mathcal{X}_{i,\langle l,t_l \rangle}^-$. Finally, let

$$\mathcal{X}' = \bigcup_{i \in [I], \langle t_1, \dots, t_{L_i} \rangle \in [\infty]^{L_i}} B^{i,\langle t_1, \dots, t_{L_i} \rangle}$$

$$B = \{B^{i,\langle t_1, \dots, t_{L_i} \rangle} \mid i \in [I], \langle t_1, \dots, t_{L_i} \rangle \in [\infty]^{L_i}\}$$

$$C = \{C^{i,\langle t_1, \dots, t_{L_i} \rangle} \mid i \in [I], \langle t_1, \dots, t_{L_i} \rangle \in [\infty]^{L_i}\}.$$

Then, \mathcal{X}' is a subinterior of A , and B is a subanalytic partition of \mathcal{X}' , because of the following:

- B is a partition of \mathcal{X}' .
- $\{i, \langle t_1, \dots, t_{L_i} \rangle \mid i \in [I], \langle t_1, \dots, t_{L_i} \rangle \in [\infty]^{L_i}\}$ is a countable set.
- For all $i \in [I]$ and $\langle t_1, \dots, t_{L_i} \rangle \in [\infty]^{L_i}$, $B^{i,\langle t_1, \dots, t_{L_i} \rangle}$ is subanalytic.
- C is a finer partition of \mathcal{X} than A . This holds because $\{A^i\}_{i \in [I]}$ is a partition of \mathcal{X} , and $\{C^{i,\langle t_1, \dots, t_{L_i} \rangle}\}_{\langle t_1, \dots, t_{L_i} \rangle \in [\infty]^{L_i}}$ is a partition of A^i for all $i \in [I]$, by its construction.

This completes the proof that $\text{subint}(A) \neq \emptyset$.

We now prove the remaining claims. Let $\mathcal{X}' \in \text{subint}(A)$ and $B = \{B^t\}_{t \in [T]}$ be a subanalytic partition of \mathcal{X}' that satisfies the equations in Definition 10.

Proof of (a). By the definition of subanalytic partition, and since $h(x) < 0 \iff -h(x) > 0$ for any function h and input x , B is an analytic partition of \mathcal{X}' with $T \in (\mathbb{Z}_{>0} \cup \{\infty\})$. We argue that for any $t \in [T]$, B^t is open in \mathbb{R}^n . Recall the equation (5):

$$B^t = \left(\bigcap_{j \in [J_t]} (g_{t,j}^+)^{-1}(\mathbb{R}_{>0}) \right) \cap \left(\bigcap_{l \in [L_t]} (g_{t,l}^-)^{-1}(\mathbb{R}_{<0}) \right).$$

Since $g_{t,j}^+ : \mathcal{X}_{t,j}^+ \rightarrow \mathbb{R}$ and $g_{t,l}^- : \mathcal{X}_{t,l}^- \rightarrow \mathbb{R}$ are continuous, and $\mathbb{R}_{>0}$ and $\mathbb{R}_{<0}$ are open in \mathbb{R} , we have that $(g_{t,j}^+)^{-1}(\mathbb{R}_{>0})$ and $(g_{t,l}^-)^{-1}(\mathbb{R}_{<0})$ are open in $\mathcal{X}_{t,j}^+$ and $\mathcal{X}_{t,l}^-$, respectively, by the definition of continuity. Furthermore, since $\mathcal{X}_{t,j}^+$ and $\mathcal{X}_{t,l}^-$ are open in \mathbb{R}^n , we have that $(g_{t,j}^+)^{-1}(\mathbb{R}_{>0})$ and $(g_{t,l}^-)^{-1}(\mathbb{R}_{<0})$ are open in \mathbb{R}^n as well. Since any finite intersection of open subsets is again open, B^t is open in \mathbb{R}^n . Hence,

$$\text{int}(B) = \bigcup_{t \in [T]} \text{int}(B^t; \mathbb{R}^n) = \bigcup_{t \in [T]} B^t = \mathcal{X}'.$$

Proof of (b). We continue the proof from (a). Since B^t is open in \mathbb{R}^n , and $B^t \subseteq A^i$ for some $i \in [I]$ (by the definition of subanalytic partition), we have $B^t \subseteq \text{int}(A^i; \mathbb{R}^n)$. From this, we obtain

$$\mathcal{X}' = \bigcup_{t \in [T]} B^t \subseteq \bigcup_{i \in [I]} \text{int}(A^i; \mathbb{R}^n) = \text{int}(A).$$

Moreover, since any union of open sets is again open, $\mathcal{X}' = \bigcup_{t \in [T]} B^t$ is open in \mathbb{R}^n .

Proof of (c). For this, we use the following theorem [32]: for any open connected $U \subseteq \mathbb{R}^n$ and analytic function $f : U \rightarrow \mathbb{R}$, the zero set $\{x \in U \mid f(x) = 0\}$ of f is either U or contained in some measure-zero set. Observe that

$$\begin{aligned}
\mathcal{X} \setminus \mathcal{X}' &= \left(\bigcup_{t \in [T]} C^t \right) \setminus \left(\bigcup_{t \in [T]} B^t \right) \\
&= \bigcup_{t \in [T]} \left(C^t \setminus B^t \right) \\
&\subseteq \bigcup_{t \in [T]} \left(\bigcap_{l \in [L_t]} (g_{t,l}^-)^{-1}(\mathbb{R}_{\leq 0}) \setminus \bigcap_{l \in [L_t]} (g_{t,l}^-)^{-1}(\mathbb{R}_{< 0}) \right) \\
&\subseteq \bigcup_{t \in [T]} \bigcup_{l \in [L_t]} \left((g_{t,l}^-)^{-1}(\mathbb{R}_{\leq 0}) \setminus (g_{t,l}^-)^{-1}(\mathbb{R}_{< 0}) \right) \\
&= \bigcup_{t \in [T]} \bigcup_{l \in [L_t]} (g_{t,l}^-)^{-1}(\{0\}).
\end{aligned}$$

Since each $g_{t,l}^-$ is analytic, and not everywhere-zero, on its connected open domain $\mathcal{X}_{t,l}^-$ (by the definition of subanalytic partition), the above theorem and equation imply that $(g_{t,l}^-)^{-1}(\{0\})$ is contained in some measure-zero set. Since any countable union of measure-zero sets has measure zero, $\mathcal{X} \setminus \mathcal{X}'$ is contained in some measure-zero set.

Proof of (d). This follows immediately from (b) and (c). \square

B.3 Proofs of Proposition 4 and Proposition 8 (part I)

We remind the reader that the notation $D(f)(x)$ means the standard derivative of f at x .

Definition 11 (Interior of PAP Representation). *Let $\gamma = \{A^i, f^i\}_{i \in [I]}$ be a PAP representation from \mathcal{X} to \mathcal{Y} . The **interior** and **subinterior** of γ are defined by:*

$$\text{int}(\gamma) = \text{int}(\{A^i\}_{i \in [I]}), \quad \text{subint}(\gamma) = \text{subint}(\{A^i\}_{i \in [I]}).$$

Lemma 15. *Let $f : \mathcal{X} \rightarrow \mathcal{Y}$ be a PAP function, γ be a PAP representation of f , and $k \in \mathbb{Z}_{\geq 0}$. Then, for all $x \in \text{int}(\gamma)$, f has the k -th order standard derivative at x . Furthermore, the derivative agrees with the k -th order intensional derivative of γ at x :*

$$D^{(k)}(f)(x) = \langle\langle D^{(k)}(\gamma) \rangle\rangle(x) \quad \text{for all } x \in \text{int}(\gamma),$$

where $F^{(k)}$ denotes the k -time composition of the operator F .

Proof. Consider any $x \in \text{int}(\gamma)$. By the definition of $\text{int}(\gamma)$, we have $x \in \text{int}(A^i; \mathbb{R}^n)$ for some $i \in [I]$. So there exists an open neighbourhood $U \subseteq \mathbb{R}^n$ of x such that $U \subseteq A^i$. Since γ is a representation of f and $U \subseteq A^i$, we have $f = f^i$ on U for the i -th component function of γ . Now focus on f^i . Since f^i is analytic on U (due to γ being PAP), f^i is infinitely differentiable on U and, in particular, has the k -th order standard derivative at x , namely $D^{(k)}(f^i)(x)$. By the definition of intensional derivative, $D^{(k)}(f^i)(x) = \langle\langle D^{(k)}(\gamma) \rangle\rangle(x)$. Since $x \in U$, U is open in \mathbb{R}^n , and $f = f^i$ on U , we obtain that $D^{(k)}(f)(x) = D^{(k)}(f^i)(x) = \langle\langle D^{(k)}(\gamma) \rangle\rangle(x)$. \square

Proposition 16. *Let $f : \mathcal{X} \rightarrow \mathcal{Y}$ be a PAP function and $k \in \mathbb{Z}_{\geq 0}$. Then, f has the k -th order standard derivative almost everywhere. Moreover, the first-order standard derivative Df agrees with any first-order intensional derivative $df \in \partial_{\bullet} f$ almost everywhere.*

Proof. By Lemma 14(d), the interior of any PAP representation of f has the full measure in \mathcal{X} . Hence, Lemma 15 implies the first claim. For the second claim, let $df \in \partial_{\bullet} f$. By the definition of $\partial_{\bullet} f$, there exists a PAP representation γ of f such that $df = \langle\langle D\gamma \rangle\rangle$. Applying Lemma 15 to f , γ , and $k = 1$ gives the second claim, since $\text{int}(\gamma)$ has the full measure in \mathcal{X} . \square

B.4 Proof of Proposition 8 (part II)

Lemma 17. *Let $f : \mathcal{X}_f \rightarrow \mathcal{Y}$ and $g : \mathcal{X}_g \rightarrow \mathcal{Y}$ be PAP functions, and γ_f and γ_g be their PAP representations. If $f(x) = g(x)$ for all $x \in \text{int}(\gamma_f) \cap \text{int}(\gamma_g)$, then*

$$D(f)(x) = D(g)(x) \quad \text{for all } x \in \text{int}(\gamma_f) \cap \text{int}(\gamma_g),$$

where both sides are well-defined for each x .

Proof. Let $U = \text{int}(\gamma_f) \cap \text{int}(\gamma_g)$, and consider any $x \in U$. Since $x \in U$ and U is open in \mathbb{R}^n , we have $D(f)(x) = D(g)(x)$ if both sides are well-defined. Indeed, they are well-defined by Lemma 15 with $k = 1$, since γ_f and γ_g are PAP representations of f and g , respectively. \square

Lemma 18. *Let $f : \mathcal{X} \rightarrow \mathcal{Y}$ be a PAP function. Then, for any intensional derivative $df \in \partial_\bullet f$, there exists a PAP representation γ_{df} of df such that*

$$df(x) = D(f)(x) \quad \text{for all } x \in \text{int}(\gamma_{df}).$$

Proof. By the definition of $\partial_\bullet f$, there exists a PAP representation γ_f of f such that $df = \langle\langle D\gamma_f \rangle\rangle$. By Lemma 15 with $k = 1$, we have $D(f)(x) = \langle\langle D\gamma_f \rangle\rangle(x) = df(x)$ for all $x \in \text{int}(\gamma_f)$. Let $\gamma_{df} = D\gamma_f$. Since $df = \langle\langle D\gamma_f \rangle\rangle = \langle\langle \gamma_{df} \rangle\rangle$ and γ_f is PAP, γ_{df} is a PAP representation of df . Moreover, by the definition of $D\gamma_f$, we have $\text{int}(\gamma_{df}) = \text{int}(\gamma_f)$. Hence, the claim holds with γ_{df} . \square

Definition 12 (Refinement of Representation). *Let $\gamma_f = \{\langle A^i, f^i \rangle\}_{i \in [I]}$ be a representation of a function from \mathcal{X}_f to \mathcal{Y} , and $B = \{B^j\}_{j \in [J]}$ be a partition of \mathcal{X}_g . The **refinement** of γ_f with B is defined by:*

$$\text{refine}(\gamma_f; B) = \{\langle A^i \cap B^j, f^i \rangle\}_{(i,j) \in [I] \times [J]}.$$

*Moreover, for any representation $\gamma_g = \{\langle C^l, g^l \rangle\}_{l \in [L]}$ of a function from \mathcal{X}_g to \mathcal{Z} , the **refinement** of γ_f with γ_g is defined by:*

$$\text{refine}(\gamma_f; \gamma_g) = \text{refine}(\gamma_f; \{C^l\}_{l \in [L]}).$$

Lemma 19. *Let $f : \mathcal{X}_f \rightarrow \mathcal{Y}$ be a PAP function, γ be a PAP representation of f , and $B = \{B^j\}_{j \in [J]}$ be an analytic partition of \mathcal{X}_g with $J \in (\mathbb{Z}_{>0} \cup \{\infty\})$. Let $\gamma' = \text{refine}(\gamma; B)$. Then, γ' is a PAP representation of $f|_{\mathcal{X}_f \cap \mathcal{X}_g}$ with*

$$\text{int}(\gamma') = \text{int}(\gamma) \cap \text{int}(B).$$

Proof. Let $\gamma = \{\langle A^i, f^i \rangle\}_{i \in [I]}$. Since γ is PAP and B is an analytic partition, $\{A^i \cap B^j\}_{(i,j) \in [I] \times [J]}$ is an analytic partition. Also, since $[J]$ is countable, $[I] \times [J]$ is also countable. Thus, γ' is PAP. Since

$$\bigcup_{(i,j) \in [I] \times [J]} (A^i \cap B^j) = \left(\bigcup_{i \in [I]} A^i \right) \cap \left(\bigcup_{j \in [J]} B^j \right) = \mathcal{X}_f \cap \mathcal{X}_g,$$

γ' is a representation of $f|_{\mathcal{X}_f \cap \mathcal{X}_g}$. Finally, we obtain the last claim as follows:

$$\begin{aligned} \text{int}(\gamma') &= \bigcup_{(i,j) \in [I] \times [J]} \text{int}(A^i \cap B^j; \mathbb{R}^n) \\ &= \bigcup_{(i,j) \in [I] \times [J]} \text{int}(A^i; \mathbb{R}^n) \cap \text{int}(B^j; \mathbb{R}^n) \\ &= \left(\bigcup_{i \in [I]} \text{int}(A^i; \mathbb{R}^n) \right) \cap \left(\bigcup_{j \in [J]} \text{int}(B^j; \mathbb{R}^n) \right) \\ &= \text{int}(\gamma) \cap \text{int}(B). \end{aligned}$$

For the second equality, we use the following fact: $\text{int}(S_1 \cap S_2; X) = \text{int}(S_1; X) \cap \text{int}(S_2; X)$ for any $S_1, S_2 \subseteq X$. \square

Lemma 20. *Let $f : \mathcal{X} \rightarrow \mathcal{Y}$ be a PAP function. Consider any PAP representation γ of f , and any subinterior $\mathcal{X}' \in \text{subint}(\gamma)$. Then, $\mathcal{X}' \subseteq \mathcal{X}$ is open in \mathbb{R}^n and $\mathcal{X} \setminus \mathcal{X}'$ is contained in a measure-zero set. Moreover, for all $k \in \mathbb{Z}_{\geq 0}$, $D^{(k)}(f|_{\mathcal{X}'})$ is a total function on \mathcal{X}' , and there exists a PAP representation γ_D^k of $D^{(k)}(f|_{\mathcal{X}'})$ such that $\text{int}(\gamma_D^k) = \mathcal{X}'$.*

Proof. Let $k \in \mathbb{Z}_{\geq 0}$. By Lemma 14(b) and 14(c), $\mathcal{X}' \subseteq \text{int}(\gamma) \subseteq \mathcal{X}$ is open in \mathbb{R}^n and $\mathcal{X} \setminus \mathcal{X}'$ is contained in a measure-zero set. This proves the first claim. Since $\mathcal{X}' \subseteq \text{int}(\gamma)$, Lemma 15 implies that $D^{(k)}(f)(x)$ exists for all $x \in \mathcal{X}'$. Since \mathcal{X}' is open in \mathbb{R}^n , $D^{(k)}(f|_{\mathcal{X}'})(x) = D^{(k)}(f)(x)$ for all $x \in \mathcal{X}'$. This proves the second claim that $D^{(k)}(f|_{\mathcal{X}'})$ is a total function on \mathcal{X}' .

We now prove the last claim. By Lemma 14(a), there exists an analytic partition $B = \{B^j\}_{j \in [J]}$ of \mathcal{X}' such that $J \in (\mathbb{Z}_{>0} \cup \{\infty\})$ and $\text{int}(B) = \mathcal{X}'$. Let $\gamma' = \text{refine}(\gamma; B)$. By Lemma 19, γ' is a PAP representation of $f|_{\mathcal{X}'}$ with $\text{int}(\gamma') = \text{int}(\gamma) \cap \text{int}(B) = \mathcal{X}'$. Consider

$$\gamma_D^k = D^{(k)}(\gamma').$$

We show that it satisfies the last claim. Since $f|_{\mathcal{X}'} : \mathcal{X}' \rightarrow \mathcal{Y}$ is a PAP function with a PAP representation γ' , Lemma 15 implies that $D^{(k)}(f|_{\mathcal{X}'}) (x) = \langle\langle D^{(k)}(\gamma') \rangle\rangle(x) = \langle\langle \gamma_D^k \rangle\rangle(x)$ for all $x \in \text{int}(\gamma') = \mathcal{X}'$. Hence, γ_D^k is a representation of $D^{(k)}(f|_{\mathcal{X}'})$. The rest of the claim also holds as follows: γ_D^k is PAP since γ' is PAP; and $\text{int}(\gamma_D^k) = \text{int}(\gamma') = \mathcal{X}'$ by the definition of $D(\gamma')$. \square

Lemma 21. *Let $f : \mathcal{X} \rightarrow \mathcal{Y}$ be a PAP function and $\mathcal{X}' \subseteq \mathcal{X}$ be a set described in Lemma 20, which has full measure in \mathcal{X} . Consider any $k \in \mathbb{Z}_{\geq 0}$. Then, for any $df^k \in \partial_{\bullet}^k f$, there exists a PAP representation γ_d^k of df^k such that*

$$df^k(x) = D^{(k)}(f)(x) \quad \text{for all } x \in \text{int}(\gamma_d^k) \cap \mathcal{X}'.$$

Proof. The proof proceeds by induction on k . For $k = 0$, we have $df^k = D^{(k)}(f) = f$. So any PAP representation γ_d^k of $df^k = f$ satisfies the claim. Now suppose $k > 0$. By the definition of $\partial_{\bullet}^k f$, there exists $df^{k-1} \in \partial_{\bullet}^{k-1} f$ such that $df^k \in \partial_{\bullet}(df^{k-1})$. We construct the desired PAP representation γ_d^k as follows. First, focus on $df^{k-1} \in \partial_{\bullet}^{k-1} f$. By the induction hypothesis on $k - 1$ for df^{k-1} , there exists a PAP representation γ_d^{k-1} of df^{k-1} such that

$$df^{k-1}(x) = D^{(k-1)}(f)(x) \quad \text{for all } x \in \text{int}(\gamma_d^{k-1}) \cap \mathcal{X}'.$$

By Lemma 20, $D^{(k-1)}(f|_{\mathcal{X}'})$ is a total function on \mathcal{X}' and there exists a PAP representation γ_D^{k-1} of $D^{(k-1)}(f|_{\mathcal{X}'})$ such that $\text{int}(\gamma_D^{k-1}) = \mathcal{X}'$. Since \mathcal{X}' is open in \mathbb{R}^n by Lemma 20, $D^{(k-1)}(f) = D^{(k-1)}(f|_{\mathcal{X}'})$ over \mathcal{X}' . Combining the two results gives that

$$df^{k-1}(x) = D^{(k-1)}(f|_{\mathcal{X}'}) (x) \quad \text{for all } x \in \text{int}(\gamma_d^{k-1}) \cap \text{int}(\gamma_D^{k-1}),$$

where γ_d^{k-1} and γ_D^{k-1} are PAP representations of df^{k-1} and $D^{(k-1)}(f|_{\mathcal{X}'})$, respectively. By Lemma 17 applied to this result, we obtain

$$D(df^{k-1})(x) = D^{(k)}(f|_{\mathcal{X}'}) (x) \quad \text{for all } x \in \text{int}(\gamma_d^{k-1}) \cap \text{int}(\gamma_D^{k-1}).$$

Since $D^{(k)}(f|_{\mathcal{X}'}) = D^{(k)}(f)$ over \mathcal{X}' (as \mathcal{X}' is open in \mathbb{R}^n), and $\text{int}(\gamma_D^{k-1}) = \mathcal{X}'$, we have

$$D(df^{k-1})(x) = D^{(k)}(f)(x) \quad \text{for all } x \in \text{int}(\gamma_d^{k-1}) \cap \mathcal{X}'.$$

Next, focus on $df^k \in \partial_{\bullet}(df^{k-1})$. By Lemma 18 applied to df^{k-1} , there exists a PAP representation $\gamma_d'^k$ of df^k such that

$$df^k(x) = D(df^{k-1})(x) \quad \text{for all } x \in \text{int}(\gamma_d'^k).$$

Combining the last two equations, we obtain

$$df^k(x) = D^{(k)}(f)(x) \quad \text{for all } x \in \text{int}(\gamma_d'^k) \cap \text{int}(\gamma_d^{k-1}) \cap \mathcal{X}'.$$

Now let $\gamma_d^k = \text{refine}(\gamma_d'^k; \gamma_d^{k-1})$. By Lemma 19 applied to $\gamma_d'^k$, we have that γ_d^k is a PAP representation of df^k with $\text{int}(\gamma_d^k) = \text{int}(\gamma_d'^k) \cap \text{int}(\gamma_d^{k-1})$. From this, we obtain the desired claim:

$$df^k(x) = D^{(k)}(f)(x) \quad \text{for all } x \in \text{int}(\gamma_d^k) \cap \mathcal{X}'.$$

\square

Proposition 22. *Let $f : \mathcal{X} \rightarrow \mathcal{Y}$ be a PAP function and $k \in \mathbb{Z}_{\geq 0}$. Then, any k -th order intensional derivative $df^k \in \partial_{\bullet}^k f$ satisfies the following:*

$$df^k(x) = D^{(k)}(f)(x) \quad \text{for almost all } x \in \mathcal{X}.$$

Proof. The claim follows from Lemma 21 and the following: \mathcal{X}' and $\text{int}(\gamma_d^k)$ described in Lemma 21 have the full measure in \mathcal{X} , by Lemma 21 and Lemma 14(d). \square

B.5 Additional property on PAP functions

Proposition 23. *Let $f : \mathcal{X} \rightarrow \mathcal{Y}$ be a PAP function, and $k \in \mathbb{Z}_{>0}$. Let $N_{f,k} \subseteq \mathcal{X}$ be the set defined by $\{x \in \mathcal{X} \mid D^{(k)}(f)(x) \text{ is undefined}\}$. If $n = 1$, then $N_{f,k}$ is countable. But if $n > 1$, then it could be uncountable.*

Proof. Suppose $n = 1$. Recall the following two well-known results: (i) if $g : U \rightarrow \mathbb{R}$ is an analytic function on an open interval $U \subseteq \mathbb{R}$ and is not everywhere-zero, then its zero set $Z_g = \{x \in U \mid g(x) = 0\}$ contains none of the limit points of Z_g [25, Corollary 1.2.7]; (ii) if X is a second-countable space and $S \subseteq X$ contains none of the limit points of S , then S is countable. Since \mathbb{R} is second-countable, every g satisfying the assumption of (i) has at most countably many zeros.

Now return to our claim. By the statement and proof of Lemma 14(c) and Lemma 15, there are countably many analytic functions $\{g_j : U_j \rightarrow \mathbb{R}\}_j$ defined over connected open subsets of \mathbb{R} , such that every g_j is not everywhere-zero and $N_{f,k} \subseteq \bigcup_j Z_{g_j}$. Since any connected open subset of \mathbb{R} is an open interval, g_j satisfies the assumption of (i), and thus each Z_{g_j} is countable by the above result. Hence, $N_{f,k}$ is countable, since a countable union of countable sets is countable.

Suppose $n > 1$. Consider $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ defined by $f(x) = |x_1 - x_2|$, and $k = 1$. Then, f is PAP, since the following is a PAP representation of f :

$$\left\{ \begin{array}{l} \langle \{x \in \mathbb{R}^2 \mid x_1 > x_2\}, x \in \mathbb{R}^2 \mapsto x_1 - x_2 \rangle, \\ \langle \{x \in \mathbb{R}^2 \mid x_1 = x_2\}, x \in \mathbb{R}^2 \mapsto 0 \rangle, \\ \langle \{x \in \mathbb{R}^2 \mid x_1 < x_2\}, x \in \mathbb{R}^2 \mapsto x_2 - x_1 \rangle \end{array} \right\}.$$

However, $N_{f,k} = \{\langle x, y \rangle \in \mathbb{R}^2 \mid x = y\}$ is uncountable. □