

HIGHER ORDER AUTOMATIC DIFFERENTIATION OF HIGHER ORDER FUNCTIONS

MATHIEU HUOT^a, SAM STATON^a, AND MATTHIJS VÁKÁR^b

^a University of Oxford

^b Utrecht University

ABSTRACT. We present semantic correctness proofs of automatic differentiation (AD). We consider a forward-mode AD method on a higher order language with algebraic data types, and we characterise it as the unique structure preserving macro given a choice of derivatives for basic operations. We describe a rich semantics for differentiable programming, based on diffeological spaces. We show that it interprets our language, and we phrase what it means for the AD method to be correct with respect to this semantics. We show that our characterisation of AD gives rise to an elegant semantic proof of its correctness based on a gluing construction on diffeological spaces. We explain how this is, in essence, a logical relations argument. Throughout, we show how the analysis extends to AD methods for computing higher order derivatives using a Taylor approximation.

1. INTRODUCTION

Automatic differentiation (AD), loosely speaking, is the process of taking a program describing a function, and constructing the derivative of that function by applying the chain rule across the program code. As gradients play a central role in many aspects of machine learning, so too do automatic differentiation systems such as TensorFlow [AAB⁺16], PyTorch [PGC⁺17] or Stan [CHB⁺15].




Figure 1: Overview of semantics/correctness of AD.

Key words and phrases: automatic differentiation, software correctness, denotational semantics.
The authors contributed equally to this work.

Differentiation has a well-developed mathematical theory in terms of differential geometry. The aim of this paper is to formalize this connection between differential geometry and the syntactic operations of AD, particularly for AD methods that calculate higher order derivatives. In this way we achieve two things: (1) a compositional, denotational understanding of differentiable programming and AD; (2) an explanation of the correctness of AD.

This intuitive correspondence (summarized in Fig. 1) is in fact rather complicated. In this paper, we focus on resolving the following problem: higher order functions play a key role in programming, and yet they have no counterpart in traditional differential geometry. Moreover, we resolve this problem while retaining the compositionality of denotational semantics.

1.0.1. *Higher order functions and differentiation.* A major application of higher order functions is to support disciplined code reuse. Code reuse is particularly acute in machine learning. For example, a multi-layer neural network might be built of millions of near-identical neurons, as follows.

$$\text{neuron}_n : (\mathbf{real}^n * (\mathbf{real}^n * \mathbf{real})) \rightarrow \mathbf{real}$$


$$\text{neuron}_n \stackrel{\text{def}}{=} \lambda \langle x, \langle w, b \rangle \rangle. \zeta(w \cdot x + b)$$

$$\text{layer}_n : ((\tau_1 * P) \rightarrow \tau_2) \rightarrow (\tau_1 * P^n) \rightarrow \tau_2^n$$

$$\text{layer}_n \stackrel{\text{def}}{=} \lambda f. \lambda \langle x, \langle p_1, \dots, p_n \rangle \rangle. \langle f \langle x, p_1 \rangle, \dots, f \langle x, p_n \rangle \rangle$$

$$\text{comp} : (((\tau_1 * P) \rightarrow \tau_2) * ((\tau_2 * Q) \rightarrow \tau_3)) \rightarrow (\tau_1 * (P * Q)) \rightarrow \tau_3$$

$$\text{comp} \stackrel{\text{def}}{=} \lambda \langle f, g \rangle. \lambda \langle x, \langle p, q \rangle \rangle. g \langle f \langle x, p \rangle, q \rangle$$



(Here $\zeta(x) \stackrel{\text{def}}{=} \frac{1}{1+e^{-x}}$ is the sigmoid function, as illustrated.) We can use these functions to build a network as follows (see also Fig. 2):

$$\text{comp} \langle \text{layer}_{m_n}(\text{neuron}_k), \text{comp} \langle \text{layer}_n(\text{neuron}_m), \text{neuron}_n \rangle \rangle : (\mathbf{real}^k * P) \rightarrow \mathbf{real} \quad (1.1)$$

Here $P \cong \mathbf{real}^p$ with $p = (m(k+1) + n(m+1) + n + 1)$. This program (1.1) describes a smooth (infinitely differentiable) function. The goal of automatic differentiation is to find its derivative.

If we β -reduce all the λ 's, we end up with a very long function expression just built from the sigmoid function and linear algebra. We can then find a program for calculating its derivative by applying the chain rule. However, automatic differentiation can also be expressed without first β -reducing, in a compositional way, by explaining how higher order functions like (layer) and (comp) propagate derivatives. This paper is a semantic analysis of this compositional approach.

The general idea of denotational semantics is to interpret types as spaces and programs as functions between the spaces. In this paper, we propose to use diffeological spaces and smooth functions [Sou80, IZ13] to this end. These satisfy the following three desiderata:




Figure 2: The network in (1.1) with k inputs and two hidden layers.

- \mathbb{R} is a space, and the smooth functions $\mathbb{R} \rightarrow \mathbb{R}$ are exactly the functions that are infinitely differentiable;
- The set of smooth functions $X \rightarrow Y$ between spaces again forms a space, so we can interpret function types.
- The disjoint union of a sequence of spaces again forms a space, and this enables us to interpret variant types and inductive types, e.g. lists of reals form the space $\biguplus_{i=0}^{\infty} \mathbb{R}^i$.

We emphasise that the most standard formulation of differential geometry, using manifolds, does not support spaces of functions. Diffeological spaces seem to us the simplest notion of space that satisfies these conditions, but there are other candidates [BH11, Sta11]. A diffeological space is in particular a set X equipped with a chosen set of curves $C_X \subseteq X^{\mathbb{R}}$ and a smooth map $f : X \rightarrow Y$ must be such that if $\gamma \in C_X$ then $\gamma; f \in C_Y$. This is reminiscent of the method of logical relations.

1.0.2. *From smoothness to automatic derivatives at higher types.* Our denotational semantics in diffeological spaces guarantees that all definable functions are smooth. But we need more than just to know that a definable function happens to have a mathematical derivative: we need to be able to find that derivative.

In this paper we focus on forward mode automatic differentiation methods for computing higher derivatives, which are macro translations on syntax (called \vec{D} in Section 3). We are able to show that they are correct, using our denotational semantics.

Here there is one subtle point that is central to our development. Although differential geometry provides established derivatives for first order functions (such as neuron above), there is no canonical notion of derivative for higher order functions (such as layer and comp) in the theory of diffeological spaces (e.g. [CW14]). We propose a new way to resolve this, by interpreting types as triples (X, X', S) where, intuitively, X is a space of inhabitants of the type, X' is a space serving as a chosen bundle of tangents (or jets, in the case of higher order derivatives) over X , and $S \subseteq X^{\mathbb{R}} \times X'^{\mathbb{R}}$ is a binary relation between curves, informally relating curves in X with their tangent (resp. jet) curves in X' . This new model gives a denotational semantics for higher order automatic differentiation on a language with higher order functions.

In Section 4 we boil this new approach down to a straightforward and elementary logical relations argument for the correctness of higher order automatic differentiation. The approach is explained in detail in Section 6. We explore some subtleties of non-uniqueness of derivatives of higher order functions in Section 7.

1.0.3. *Related work and context.* AD has a long history and has many implementations. AD was perhaps first phrased in a functional setting in [PS08], and there are now a number of teams working on AD in the functional setting (e.g. [WWE⁺19, SFVPJ19, Ell18]), some providing efficient implementations. Although that work does not involve formal semantics, it is inspired by intuitions from differential geometry and category theory.

This paper adds to a very recent body of work on verified automatic differentiation. In the first order setting, there are recent accounts based on denotational semantics in manifolds [FST19, LYRY20] and based on synthetic differential geometry [CGM19], work making a categorical abstraction [CCG⁺20] and work connecting operational semantics with denotational semantics [AP20, Plo18], as well as work focussing on how to correctly differentiate programs that operate on tensors [BML⁺20] and programs that make use of

quantum computing [ZHCW20]. Recently there has also been significant progress at higher types. The work of Brunel et al. [BMP20] and Mazza and Pagani [MP21] give formal correctness proofs for reverse-mode derivatives on a linear λ -calculus with a particular operational semantics. The work of Barthe et al. [BCLG20] provides a general discussion of some new syntactic logical relations arguments including one very similar to our syntactic proof of Theorem 4.6. Sherman et al. [SMC20] discuss a differential programming technique that works at higher types, based on exact real arithmetic and relate it to a computable semantics. We understand that the authors of [CGM19] are working on higher types. Vákár [Vák21, VS21, LNV21] phrase and prove correct a reverse mode AD technique on a higher order language based on a similar gluing technique. Vákár [Vák20] extends a standard λ -calculus with type recursion, and proves correct a forward-mode AD on such a higher-order language, also using a gluing argument.

The differential λ -calculus [ER03] is related to AD, and explicit connections are made in [MO20, Man12]. One difference is that the differential λ -calculus allows the addition of terms at all types, and hence vector space models are suitable to interpret all types. This choice would appear peculiar with the variant and inductive types that we consider here, as the dimension of a disjoint union of spaces is only defined locally.

This paper builds on our previous work [HSV20a, Vák20] in which we gave denotational correctness proofs for forward mode AD algorithms for computing first derivatives. Here, we explain how these techniques extend to methods that calculate higher derivatives.

The Faà di Bruno construction has also been investigated [CS11] in the context of Cartesian differential categories.

The idea of directly calculating higher order derivatives by using automatic differentiation methods that work with Taylor approximations (also known as jets in differential geometry) is well-known [GUW00] and it has recently gained renewed interest [Bet18, BJD19]. So far, such “Taylor-mode AD” methods have only been applied to first order functional languages, however. This paper shows how to extend these higher order AD methods to languages with support for higher order functions and algebraic data types.

The two main methods for implementing AD are operator overloading and, the method used in this paper, source code transformation [VMBBL18]. Taylor-mode AD has been seen to be significantly faster than iterated AD in the context of operator overloading [BJD19] in Jax [FJL18]. There are other notable implementations of forward Taylor-mode [BS96, BS97, Kar01, PS07, WGP16]. Some of them are implemented in a functional language [Kar01, PS07]. Taylor-mode implementations use the rich algebraic structure of derivatives to avoid a lot of redundant computations occurring via iterated first order methods and share of a lot of redundant computations. Perhaps the simplest example to see this is with the sin function, whose iterated derivatives only involve sin, cos, and negation. Importantly, most AD tools have the right complexity up to a constant factor, but this constant is quite important in practice and Taylor-mode helps achieve better performance. Another stunning result of a version of Taylor-mode was achieved in [LMG18], where a gain of performance of up to two orders of magnitude was achieved for computing certain Hessian-vector products using Ricci calculus. In essence, the algorithm used is mixed-mode that is derived via jets in [Bet18]. This is further improved in [LMG20]. Taylor-mode can also be useful for ODE solvers and hence will be important for neural differential equations [CRBD18].

Finally, we emphasise that we have chosen the neural network (1.1) as our running example mainly for its simplicity. Indeed one would typically use reverse-mode AD to train neural networks in practice. There are many other examples of AD outside the

neural networks literature: AD is useful whenever derivatives need to be calculated on high dimensional spaces. This includes optimization problems more generally, where the derivative is passed to a gradient descent method (e.g. [RM51, KW⁺52, Qia99, KB14, DHS11, LN89]). Optimization problems involving higher order functions naturally show up in the calculus of variations and its applications in physics, where one typically looks for a function minimizing a certain integral [GSS00]. Other applications of AD are in advanced *integration* methods, since derivatives play a role in Hamiltonian Monte Carlo [Nea11, HG14] and variational inference [KTR⁺17]. Second order methods for gradient-descent have also been extensively studied. As the basic second order Newton method requires inverting a high dimensional hessian matrix, several alternatives and approximations have been studied. Some of them still require Taylor-like modes of differentiation and require a matrix-vector product where the matrix resembles the hessian or inverse hessian [KK04, Mar10, Ama12].

1.0.4. *Summary of contributions.* We have provided a semantic analysis of higher order automatic differentiation. Our syntactic starting point are higher order forward-mode AD macros on a typed higher order language that extend their well-known first order equivalent (e.g. [SFVPJ19, WWE⁺19, HSV20a]). We present these in Section 3 for function types, and in Section 5 we extend them to inductive types and variants. The main contributions of this paper are as follows.

- We give a denotational semantics for the language in diffeological spaces, showing that every definable expression is smooth (Section 4).
- We show correctness of the higher order AD macros by a logical relations argument (Th. 4.6).
- We give a categorical analysis of this correctness argument with two parts: a universal property satisfied by the macro in terms of syntactic categories, and a new notion of glued space that abstracts the logical relation (Section 6).
- We then use this analysis to state and prove a correctness argument at all first order types (Th. 6.7).

Relation to previous work. This paper extends and develops the paper [HSV20a] presented at the 23rd International Conference on Foundations of Software Science and Computation Structure (FoSSaCS 2020). This version includes numerous elaborations, notably the extension of the definition, semantics and correctness of automatic differentiation methods for computing higher order derivatives (introduced in Section 2.2-2.4) and a novel discussion about derivatives of higher-order functions (Section 7).

2. RUDIMENTS OF DIFFERENTIATION: HOW TO CALCULATE WITH DUAL NUMBERS AND TAYLOR APPROXIMATIONS

2.1. First order differentiation: the chain rule and dual numbers. We will now recall the definition of gradient of a differentiable function, the goal of AD and, and what it means for AD to be correct. Recall that the derivative of a function $f : \mathbb{R} \rightarrow \mathbb{R}$, if it exists, is a function $\nabla f : \mathbb{R} \rightarrow \mathbb{R}$ such that for all a , $\nabla f(a)$ is the gradient of f at a in the sense that the function $x \mapsto f(a) + \nabla f(a) \cdot (x - a)$ gives the best linear approximation of f at a . (The gradient $\nabla f(a)$ is often written $\frac{df(x)}{dx}(a)$.)

The chain rule for differentiation tells us that we can calculate $\nabla(f;g)(a) = \nabla f(a) \cdot \nabla g(f(a))$. In that sense, the chain rule tells us how linear approximations to a function transform under post-composition with another function.

To find ∇f in a compositional way, using the chain rule, two generalizations are reasonable:

- We need both f and ∇f when calculating $\nabla(f;g)$ of a composition $f;g$, using the chain rule, so we are really interested in the pair $(f, \nabla f) : \mathbb{R} \rightarrow \mathbb{R} \times \mathbb{R}$;
- In building f we will need to consider functions of multiple arguments, such as $+$: $\mathbb{R}^2 \rightarrow \mathbb{R}$, and these functions should propagate derivatives.

Thus we are more generally interested in transforming a function $g : \mathbb{R}^n \rightarrow \mathbb{R}$ into a function $h : (\mathbb{R} \times \mathbb{R})^n \rightarrow \mathbb{R} \times \mathbb{R}$ in such a way that for any $f_1 \dots f_n : \mathbb{R} \rightarrow \mathbb{R}$,

$$(f_1, \nabla f_1, \dots, f_n, \nabla f_n); h = ((f_1, \dots, f_n); g, \nabla((f_1, \dots, f_n); g)). \quad (2.1)$$

Computing automatically the program representing h , given a program representing g , is the goal of automatic differentiation. An intuition for h is often given in terms of dual numbers. The transformed function operates on pairs of numbers, (x, x') , and it is common to think of such a pair as $x + x'\epsilon$ for an ‘infinitesimal’ ϵ . But while this is a helpful intuition, the formalization of infinitesimals can be intricate, and the development in this paper is focussed on the elementary formulation in (2.1).

A function h satisfying (2.1) encodes all the partial derivatives of g . For example, if $g : \mathbb{R}^2 \rightarrow \mathbb{R}$, then with $f_1(x) \stackrel{\text{def}}{=} x$ and $f_2(x) \stackrel{\text{def}}{=} x_2$, by applying (2.1) to x_1 we obtain $h(x_1, 1, x_2, 0) = (g(x_1, x_2), \frac{\partial g(x, x_2)}{\partial x}(x_1))$ and similarly $h(x_1, 0, x_2, 1) = (g(x_1, x_2), \frac{\partial g(x_1, x)}{\partial x}(x_2))$. And conversely, if g is differentiable in each argument, then a unique h satisfying (2.1) can be found by taking linear combinations of partial derivatives, for example:

$$h(x_1, x'_1, x_2, x'_2) = (g(x_1, x_2), x'_1 \cdot \frac{\partial g(x, x_2)}{\partial x}(x_1) + x'_2 \cdot \frac{\partial g(x_1, x)}{\partial x}(x_2)).$$

(Here, recall that the partial derivative $\frac{\partial g(x, x_2)}{\partial x}(x_1)$ is a particular notation for the gradient $\nabla(g(-, x_2))(x_1)$, i.e. with x_2 fixed.)

In summary, the idea of differentiation with dual numbers is to transform a differentiable function $g : \mathbb{R}^n \rightarrow \mathbb{R}$ to a function $h : \mathbb{R}^{2n} \rightarrow \mathbb{R}^2$ which captures g and all its partial derivatives. We packaged this up in (2.1) as an invariant which is useful for building derivatives of compound functions $\mathbb{R} \rightarrow \mathbb{R}$ in a compositional way. The idea of (first order) forward mode automatic differentiation is to perform this transformation at the source code level.

We say that a macro for AD is correct if, given a semantic model $sem-$, the program P representing $g = \llbracket P \rrbracket$ is transformed by the macro to a program P' representing $h = \llbracket P' \rrbracket$. This means in particular that P' computes correct partial derivatives of (the function represented by) P .

Smooth functions. In what follows we will often speak of *smooth* functions $\mathbb{R}^k \rightarrow \mathbb{R}$, which are functions that are continuous and differentiable, such that their derivatives are also continuous and differentiable, and so on.

2.2. Higher order differentiation: the Faà di Bruno formula and Taylor approximations. We now generalize the above in two directions:

- We look for the best local approximations to f with polynomials of some order R , generalizing the above use of linear functions ($R = 1$).
- We can work directly with multivariate functions $\mathbb{R}^k \rightarrow \mathbb{R}$ instead of functions of one variable $\mathbb{R} \rightarrow \mathbb{R}$ ($k = 1$).

To make this precise, we recall that, given a smooth function $f : \mathbb{R}^k \rightarrow \mathbb{R}$ and a natural number $R \geq 0$, the R -th order Taylor approximation of f at $a \in \mathbb{R}^k$ is defined in terms of the partial derivatives of f :

$$\begin{aligned} \mathbb{R}^k &\rightarrow \mathbb{R} \\ x &\mapsto \sum_{\{(\alpha_1, \dots, \alpha_k) \in \mathbb{N}^k \mid \alpha_1 + \dots + \alpha_k \leq R\}} \frac{1}{\alpha_1! \cdots \alpha_k!} \frac{\partial^{\alpha_1 + \dots + \alpha_k} f(x)}{\partial x_1^{\alpha_1} \cdots \partial x_k^{\alpha_k}}(a) \cdot (x_1 - a_1)^{\alpha_1} \cdots (x_k - a_k)^{\alpha_k}. \end{aligned}$$

This is an R -th order polynomial. Similarly to the case of first order derivatives, we can recover the partial derivatives of f up to the R -th order from its Taylor approximation by evaluating the series at basis vectors. See Section 2.3 below for an example.

Recall that the ordering of partial derivatives does not matter for smooth functions (Schwarz/Clairaut’s theorem). So there will be $\binom{R+k-1}{k-1}$ R -th order partial derivatives, and altogether there are $\binom{R+k}{k}$ summands in the R -th order Taylor approximation. (This can be seen by a ‘stars-and-bars’ argument.)

Since there are $\binom{R+k}{k}$ partial derivatives of f_i of order $\leq R$, we can store them in the Euclidean space $\mathbb{R}^{\binom{R+k}{k}}$, which can also be regarded as the space of k -variate polynomials of degree $\leq R$.

We use a convention of coordinates $(y_{\alpha_1 \dots \alpha_k} \in \mathbb{R})_{(\alpha_1, \dots, \alpha_k) \in \{(\alpha_1, \dots, \alpha_k) \in \mathbb{N}^k \mid 0 \leq \alpha_1 + \dots + \alpha_k \leq R\}}$ where $y_{\alpha_1 \dots \alpha_k}$ is intended to represent a partial derivative $\frac{\partial^{\alpha_1 + \dots + \alpha_k} f}{\partial x_1^{\alpha_1} \cdots \partial x_k^{\alpha_k}}(a)$ for some function $f : \mathbb{R}^k \rightarrow \mathbb{R}$. We will choose these coordinates in lexicographic order of the multi-indices $(\alpha_1, \dots, \alpha_k)$, that is, the indexes in the Euclidean space $\mathbb{R}^{\binom{R+k}{k}}$ will typically range from $(0, \dots, 0)$ to $(R, 0, \dots, 0)$.

The (k, R) -Taylor representation of a function $g : \mathbb{R}^n \rightarrow \mathbb{R}$ is a function $h : \left(\mathbb{R}^{\binom{R+k}{k}}\right)^n \rightarrow \mathbb{R}^{\binom{R+k}{k}}$ that transforms the partial derivatives of $f : \mathbb{R}^k \rightarrow \mathbb{R}^n$ of order $\leq R$ under postcomposition with g :

$$\begin{aligned} &\left(\left(\frac{\partial^{\alpha_1 + \dots + \alpha_k} f_j(x)}{\partial x_1^{\alpha_1} \cdots \partial x_k^{\alpha_k}} \right)_{(\alpha_1, \dots, \alpha_k) = (0, \dots, 0)}^{(R, 0, \dots, 0)} \right)_{j=1}^n ; h = \\ &\left(\left(\frac{\partial^{\alpha_1 + \dots + \alpha_k} ((f_1, \dots, f_n); g)(x)}{\partial x_1^{\alpha_1} \cdots \partial x_k^{\alpha_k}} \right)_{(\alpha_1, \dots, \alpha_k) = (0, \dots, 0)}^{(R, 0, \dots, 0)} \right)_{j=1}^n. \end{aligned} \tag{2.2}$$

Thus the Taylor representation generalizes the dual numbers representation ($R = k = 1$).

To explicitly calculate the Taylor representation for a smooth function, we recall a generalization of the chain rule to higher derivatives. The chain rule tells us how the coefficients of linear approximations transform under composition of the functions. The *Faà di Bruno formula* [Sav06, EM03, CS96] tells us how coefficients of Taylor approximations —

that is, higher derivatives — transform under composition. We recall the multivariate form from [Sav06, Theorem 2.1]. Given functions $f = (f_1, \dots, f_l) : \mathbb{R}^k \rightarrow \mathbb{R}^l$ and $g : \mathbb{R}^l \rightarrow \mathbb{R}$, for $\alpha_1 + \dots + \alpha_k > 0$,

$$\begin{aligned} \frac{\partial^{\alpha_1 + \dots + \alpha_k} (f; g)(x)}{\partial x_1^{\alpha_1} \dots \partial x_k^{\alpha_k}}(a) &= \alpha_1! \cdot \dots \cdot \alpha_k! \cdot \sum_{\{(\beta_1, \dots, \beta_l) \in \mathbb{N}^l \mid 1 \leq \beta_1 + \dots + \beta_l \leq \alpha_1 + \dots + \alpha_k\}} \frac{\partial^{\beta_1 + \dots + \beta_l} g(y)}{\partial y_1^{\beta_1} \dots \partial y_l^{\beta_l}}(f(a)) \cdot \\ &\quad \sum_{\{(e_1^1, \dots, e_l^1), \dots, (e_1^q, \dots, e_l^q)\} \in (\mathbb{N}^l)^q \mid e_j^1 + \dots + e_j^q = \beta_j, (e_1^1 + \dots + e_l^1) \cdot \alpha_1^1 + \dots + (e_1^q + \dots + e_l^q) \cdot \alpha_l^q = \alpha_i\}} \\ &\quad \prod_{r=1}^q \prod_{j=1}^l \frac{1}{e_j^r!} \left(\frac{1}{\alpha_1^r! \cdot \dots \cdot \alpha_k^r!} \frac{\partial^{\alpha_1^r + \dots + \alpha_k^r} f_j(x)}{\partial x_1^{\alpha_1^r} \dots \partial x_k^{\alpha_k^r}}(a) \right)^{e_j^r}, \end{aligned}$$

where $(\alpha_1^1, \dots, \alpha_k^1), \dots, (\alpha_1^q, \dots, \alpha_k^q) \in \mathbb{N}^k$ are an enumeration of all the vectors $(\alpha_1^r, \dots, \alpha_k^r)$ of k natural numbers such that $\alpha_j^r \leq \alpha_j$ and $\alpha_1^r + \dots + \alpha_k^r > 0$ and we write q for the number of such vectors. The details of this formula reflect the complicated combinatorics that arise from repeated applications of the chain and product rules for differentiation that one uses to prove it. Conceptually, however, it is rather straightforward: it tells us that the coefficients of the R -th order Taylor approximation of $f; g$ can be expressed exclusively in terms of those of f and g .

Thus the Faà di Bruno formula uniquely determines the Taylor approximation $h : (\mathbb{R}^{\binom{R+k}{k}})^n \rightarrow \mathbb{R}^{\binom{R+k}{k}}$ in terms of the derivatives of $g : \mathbb{R}^n \rightarrow \mathbb{R}$ of order $\leq R$, and we can also recover all such derivatives from h .

2.3. Example: a two-dimensional second order Taylor series. As an example, we can specialize the Faà di Bruno formula above to the second order Taylor series of a function $f : \mathbb{R}^2 \rightarrow \mathbb{R}^l$ and its behaviour under postcomposition with a smooth function $g : \mathbb{R}^l \rightarrow \mathbb{R}$:

$$\frac{\partial^2 (f; g)(x)}{\partial x_i \partial x_{i'}}(a) = \sum_{j=1}^l \frac{\partial g(y)}{\partial y_j}(f(a)) \frac{\partial^2 f_j(x)}{\partial x_i \partial x_{i'}}(a) + \sum_{j, j'=1}^l \frac{\partial^2 g(y)}{\partial y_j \partial y_{j'}}(f(a)) \frac{\partial f_{j'}(x)}{\partial x_i}(a) \frac{\partial f_j(x)}{\partial x_{i'}}(a),$$

where $i, i' \in \{1, 2\}$ might either coincide or be distinct.

Rather than working with the full $(2, 2)$ -Taylor representation of g , we ignore the non-mixed second order derivatives $y_{02}^j = \frac{\partial^2 f_j(x)}{\partial x_2^2}$ and $y_{20}^j = \frac{\partial^2 f_j(x)}{\partial x_1^2}$ for the moment, and we represent the derivatives of order ≤ 2 of $f_j : \mathbb{R}^2 \rightarrow \mathbb{R}$ (at some point a) as the numbers

$$(y_{00}^j, y_{01}^j, y_{10}^j, y_{11}^j) = \left(f_j(a), \frac{\partial f_j(x)}{\partial x_2}(a), \frac{\partial f_j(x)}{\partial x_1}(a), \frac{\partial^2 f_j(x)}{\partial x_1 \partial x_2}(a) \right) \in \mathbb{R}^4$$

and we can choose a similar representation for the derivatives of $(f; g)$. Then, we observe that the Faà di Bruno formula induces the function $h : (\mathbb{R}^4)^l \rightarrow \mathbb{R}^4$

$$h((y_{00}^1, y_{01}^1, y_{10}^1, y_{11}^1), \dots, (y_{00}^l, y_{01}^l, y_{10}^l, y_{11}^l)) =$$

$$\left(\begin{array}{l} g(y_{00}^1, \dots, y_{00}^l) \\ \sum_{j=1}^l \frac{\partial g(y^1, \dots, y^l)}{\partial y_j}(y_{00}^1, \dots, y_{00}^l) \cdot y_{01}^j \\ \sum_{j=1}^l \frac{\partial g(y^1, \dots, y^l)}{\partial y_j}(y_{00}^1, \dots, y_{00}^l) \cdot y_{10}^j \\ \sum_{j=1}^l \frac{\partial g(y^1, \dots, y^l)}{\partial y_j}(y_{00}^1, \dots, y_{00}^l) \cdot y_{11}^j + \sum_{j,j'=1}^l \frac{\partial^2 g(y^1, \dots, y^l)}{\partial y_j \partial y_{j'}}(y_{00}^1, \dots, y_{00}^l) \cdot y_{10}^j \cdot y_{01}^{j'} \end{array} \right).$$

In particular, we can note that

$$h((y_{00}^1, y_{01}^1, y_{10}^1, 0), \dots, (y_{00}^l, y_{01}^l, y_{10}^l, 0)) = \left(\begin{array}{l} g(y_{00}^1, \dots, y_{00}^l) \\ \sum_{j=1}^l \frac{\partial g(y^1, \dots, y^l)}{\partial y_j}(y_{00}^1, \dots, y_{00}^l) \cdot y_{01}^j \\ \sum_{j=1}^l \frac{\partial g(y^1, \dots, y^l)}{\partial y_j}(y_{00}^1, \dots, y_{00}^l) \cdot y_{10}^j \\ \sum_{j,j'=1}^l \frac{\partial^2 g(y^1, \dots, y^l)}{\partial y_j \partial y_{j'}}(y_{00}^1, \dots, y_{00}^l) \cdot y_{10}^j \cdot y_{01}^{j'} \end{array} \right).$$

We see can use this method to calculate any directional first and second order derivative of g in one pass. For example, if $l = 3$, so $g : \mathbb{R}^3 \rightarrow \mathbb{R}$, then the last component of $h((x, x', x'', 0), (y, y', y'', 0), (z, z', z'', 0))$ is the result of taking the first derivative in direction (x', y', z') and the second derivative in direction (x'', y'', z'') , and evaluating at (x, y, z) .

In the proper Taylor representation we explicitly include the non-mixed second order derivatives as inputs and outputs, leading to a function $h' : (\mathbb{R}^6)^l \rightarrow \mathbb{R}^6$. Above we have followed a common trick to avoid some unnecessary storage and computation, since these extra inputs and outputs are not required for computing the second order derivatives of g . For instance, if $l = 2$ then the last component of $h((x, 1, 1, 0), (y, 0, 0, 0))$ computes $\frac{\partial^2 g(x, y)}{\partial x^2}(x, y)$.

2.4. Example: a one-dimensional second order Taylor series. As opposed to (2,2)-AD, (1,2)-AD computes the first and second order derivatives in the same direction. For example, if $g : \mathbb{R}^2 \rightarrow \mathbb{R}$ is a smooth function, then $h : (\mathbb{R}^3)^2 \rightarrow \mathbb{R}^3$. An intuition for h can be given in terms of triple numbers. The transformed function operates on triples of numbers, (x, x', x'') , and it is common to think of such a triple as $x + x'\epsilon + x''\epsilon^2$ for an ‘infinitesimal’ ϵ which has the property that $\epsilon^3 = 0$. For instance we have

$$\begin{aligned} h((x_1, 1, 0), (x_2, 0, 0)) &= (g(x_1, x_2), \frac{\partial g(x, x_2)}{\partial x}(x_1), \frac{\partial^2 g(x, x_2)}{\partial x^2}(x_1)) \\ h((x_1, 0, 0), (x_2, 1, 0)) &= (g(x_1, x_2), \frac{\partial g(x_1, x)}{\partial x}(x_2), \frac{\partial^2 g(x_1, x)}{\partial x^2}(x_2)) \\ h((x_1, 1, 0), (x_2, 1, 0)) &= \\ &(g(x_1, x_2), \frac{\partial g(x_1, x)}{\partial x}(x_2), \frac{\partial^2 g(x, x_2)}{\partial x^2}(x_1) + \frac{\partial^2 g(x_1, x)}{\partial x^2}(x_2) + 2\frac{\partial^2 g(x, y)}{\partial x \partial y}(x_1, x_2)) \end{aligned}$$

We see that we directly get non-mixed second-order partial derivatives but not the mixed-ones. We can recover $\frac{\partial^2 g(x, y)}{\partial x \partial y}(x_1, x_2)$ as $\frac{1}{2}(h((x_1, 1, 0), (x_2, 1, 0)) - h((x_1, 1, 0), (x_2, 0, 0)) - h((x_1, 0, 0), (x_2, 1, 0)))$.

More generally, if $g : \mathbb{R}^l \rightarrow \mathbb{R}$, then $h : (\mathbb{R}^3)^l \rightarrow \mathbb{R}^3$ satisfies:

$$h((x_1, x'_1, 0), \dots, (x_l, x'_l, 0)) = \left(\begin{array}{l} g(x_1, \dots, x_l) \\ \sum_{i=1}^l \frac{\partial g(x_1, \dots, x_l)}{\partial x_i}(x_1, \dots, x_l) \cdot x'_i \\ \sum_{i,j=1}^l \frac{\partial^2 g(x_1, \dots, x_l)}{\partial x_i \partial x_j}(x_1, \dots, x_l) \cdot x'_i \cdot x'_j \end{array} \right).$$

We can always recover the mixed second order partial derivatives from this but this requires several computations involving h . This is thus different from the (2,2) method which was more direct.

2.5. Remark. In the rest of this article, we study forward-mode (k, R) -automatic differentiation for a language with higher-order functions. The reader may like to fix $k = R = 1$ for a standard automatic differentiation with first-order derivatives, based on dual numbers. This is the approach taken in the conference version of this paper [HSV20b]. But the generalization to higher-order derivatives with arbitrary k and R flows straightforwardly through the whole narrative.

3. A HIGHER ORDER FORWARD-MODE AD TRANSLATION

3.1. A simple language of smooth functions. We consider a standard higher order typed language with a first order type **real** of real numbers. The types (τ, σ) and terms (t, s) are as follows.

$\tau, \sigma, \rho ::=$	types
real	real numbers
$(\tau_1 * \dots * \tau_n)$	finite product
$\tau \rightarrow \sigma$	function
$t, s, r ::=$	terms
x	variable
$\text{op}(t_1, \dots, t_n)$	operations (including constants)
$\langle t_1, \dots, t_n \rangle \mid \text{case } t \text{ of } \langle x_1, \dots, x_n \rangle \rightarrow s$	tuples/pattern matching
$\lambda x. t \mid t s$	function abstraction/application

The typing rules are in Figure 3. We have included some abstract basic n -ary operations $\text{op} \in \text{Op}_n$ for every $n \in \mathbb{N}$. These are intended to include the usual (smooth) mathematical operations that are used in programs to which automatic differentiation is applied. For example,

- for any real constant $c \in \mathbb{R}$, we typically include a constant $\underline{c} \in \text{Op}_0$; we slightly abuse notation and will simply write \underline{c} for $\underline{c}()$ in our examples;
- we include some unary operations such as $\zeta \in \text{Op}_1$ which we intend to stand for the usual sigmoid function, $\zeta(x) \stackrel{\text{def}}{=} \frac{1}{1+e^{-x}}$;
- we include some binary operations such as addition and multiplication $(+), (*) \in \text{Op}_2$;

We add some simple syntactic sugar $t - u \stackrel{\text{def}}{=} t + \underline{(-1)} * u$ and, for some natural number n ,

$$n \cdot t \stackrel{\text{def}}{=} \overbrace{t + \dots + t}^{n \text{ times}} \quad \text{and} \quad t^n \stackrel{\text{def}}{=} \overbrace{t * \dots * t}^{n \text{ times}}$$

Similarly, we will frequently denote repeated sums and products using \sum - and \prod -signs, respectively: for example, we write $t_1 + \dots + t_n$ as $\sum_{i \in \{1, \dots, n\}} t_i$ and $t_1 * \dots * t_n$ as $\prod_{i \in \{1, \dots, n\}} t_i$. This in addition to programming sugar such as **let** $x = t$ **in** s for $(\lambda x. s) t$ and $\lambda \langle x_1, \dots, x_n \rangle. t$ for $\lambda x. \text{case } x \text{ of } \langle x_1, \dots, x_n \rangle \rightarrow t$.

3.2. Syntactic automatic differentiation: a functorial macro. The aim of higher order forward mode AD is to find the (k, R) -Taylor representation of a function by syntactic manipulations, for some choice of (k, R) that we fix. For our simple language, we implement this as the following inductively defined macro $\vec{\mathcal{D}}_{(k,R)}$ on both types and terms (see also [WWE⁺19, SFVJP19]). For the sake of legibility, we simply write $\vec{\mathcal{D}}_{(k,R)}$ as $\vec{\mathcal{D}}$ here and leave the dimension k and order R of the Taylor representation implicit. The following definition is for general k and R , but we treat specific cases afterwards in Example 3.1.

$$\begin{aligned} \vec{\mathcal{D}}(\tau \rightarrow \sigma) &\stackrel{\text{def}}{=} \vec{\mathcal{D}}(\tau) \rightarrow \vec{\mathcal{D}}(\sigma) & \vec{\mathcal{D}}(\tau_1 * \dots * \tau_n) &\stackrel{\text{def}}{=} \vec{\mathcal{D}}(\tau_1) * \dots * \vec{\mathcal{D}}(\tau_n) \\ \vec{\mathcal{D}}(\mathbf{real}) &\stackrel{\text{def}}{=} \mathbf{real}^{\binom{R+k}{k}} \quad (\text{i.e., the type of tuples of reals of length } \binom{R+k}{k}) \end{aligned}$$

$$\begin{aligned} \vec{\mathcal{D}}(x) &\stackrel{\text{def}}{=} x & \vec{\mathcal{D}}(\underline{c}) &\stackrel{\text{def}}{=} \langle \underline{c}, 0 \rangle \\ \vec{\mathcal{D}}(\lambda x.t) &\stackrel{\text{def}}{=} \lambda x.\vec{\mathcal{D}}(t) & \vec{\mathcal{D}}(ts) &\stackrel{\text{def}}{=} \vec{\mathcal{D}}(t) \vec{\mathcal{D}}(s) & \vec{\mathcal{D}}(\langle t_1, \dots, t_n \rangle) &\stackrel{\text{def}}{=} \langle \vec{\mathcal{D}}(t_1), \dots, \vec{\mathcal{D}}(t_n) \rangle \\ \vec{\mathcal{D}}(\mathbf{case } t \mathbf{ of } \langle x_1, \dots, x_n \rangle \rightarrow s) &\stackrel{\text{def}}{=} \mathbf{case } \vec{\mathcal{D}}(t) \mathbf{ of } \langle x_1, \dots, x_n \rangle \rightarrow \vec{\mathcal{D}}(s) \\ \vec{\mathcal{D}}(\mathbf{op}(t_1, \dots, t_n)) &\stackrel{\text{def}}{=} \mathbf{case } \vec{\mathcal{D}}(t_1) \mathbf{ of } \langle x_{0\dots 0}^1, \dots, x_{R,0\dots 0}^1 \rangle \rightarrow \\ &\vdots \\ &\mathbf{case } \vec{\mathcal{D}}(t_n) \mathbf{ of } \langle x_{0\dots 0}^n, \dots, x_{R,0\dots 0}^n \rangle \rightarrow \\ &\langle D^{0\dots 0} \mathbf{op}(x_{0\dots 0}^1, \dots, x_{R,0\dots 0}^1, \dots, x_{0\dots 0}^n, \dots, x_{R,0\dots 0}^n), \\ &\quad \dots \\ &D^{R\dots 0} \mathbf{op}(x_{0\dots 0}^1, \dots, x_{R,0\dots 0}^1, \dots, x_{0\dots 0}^n, \dots, x_{R,0\dots 0}^n) \rangle \end{aligned}$$

where $D^{0\dots 0} \mathbf{op}(x_{0\dots 0}^1, \dots, x_{R,0\dots 0}^1, \dots, x_{0\dots 0}^n, \dots, x_{R,0\dots 0}^n) \stackrel{\text{def}}{=} \mathbf{op}(x_{0\dots 0}^1, \dots, x_{0\dots 0}^n)$

$$\begin{aligned} D^{\alpha_1 \dots \alpha_k} \mathbf{op}(x_{0\dots 0}^1, \dots, x_{R,0\dots 0}^1, \dots, x_{0\dots 0}^n, \dots, x_{R,0\dots 0}^n) &\stackrel{\text{def}}{=} \quad (\text{for } \alpha_1 + \dots + \alpha_k > 0) \\ \alpha_1! \cdot \dots \cdot \alpha_k! \cdot \sum_{\{(\beta_1, \dots, \beta_n) \in \mathbb{N}^n \mid 1 \leq \beta_1 + \dots + \beta_n \leq \alpha_1 + \dots + \alpha_k\}} &\partial_{\beta_1 \dots \beta_n} \mathbf{op}(x_{0\dots 0}^1, \dots, x_{0\dots 0}^n) * \\ \sum_{\{(e_1^1, \dots, e_i^1), \dots, (e_1^q, \dots, e_i^q)\} \in (\mathbb{N}^l)^q \mid e_j^1 + \dots + e_j^q = \beta_j, (e_1^1 + \dots + e_i^1) \cdot \alpha_i^1 + \dots + (e_1^q + \dots + e_i^q) \cdot \alpha_i^q = \alpha_i\}} & \\ \prod_{r=1}^q \prod_{j=1}^l \frac{1}{e_j^r!} \cdot \left(\frac{1}{\alpha_1^r! \dots \alpha_k^r!} \cdot x_{\alpha_1 \dots \alpha_k}^j \right)^{e_j^r} & . \end{aligned}$$

$\frac{\Gamma \vdash t_1 : \mathbf{real} \quad \dots \quad \Gamma \vdash t_n : \mathbf{real}}{\Gamma \vdash \mathbf{op}(t_1, \dots, t_n) : \mathbf{real}} \quad (\mathbf{op} \in \mathbf{Op}_n)$	
$\frac{\Gamma \vdash t_1 : \tau_1 \quad \dots \quad \Gamma \vdash t_n : \tau_n}{\Gamma \vdash \langle t_1, \dots, t_n \rangle : (\tau_1 * \dots * \tau_n)}$	$\frac{\Gamma \vdash t : (\sigma_1 * \dots * \sigma_n) \quad \Gamma, x_1 : \sigma_1, \dots, x_n : \sigma_n \vdash s : \tau}{\Gamma \vdash \mathbf{case } t \mathbf{ of } \langle x_1, \dots, x_n \rangle \rightarrow s : \tau}$
$\frac{}{\Gamma \vdash x : \tau} \quad ((x : \tau) \in \Gamma)$	$\frac{\Gamma, x : \tau \vdash t : \sigma}{\Gamma \vdash \lambda x : \tau. t : \tau \rightarrow \sigma} \quad \frac{\Gamma \vdash t : \sigma \rightarrow \tau \quad \Gamma \vdash s : \sigma}{\Gamma \vdash ts : \tau}$

Figure 3: Typing rules for the simple language.

Here, $(\partial_{\beta_1 \dots \beta_n} \mathbf{op})(x_1, \dots, x_n)$ are some chosen terms of type **real** in the language with free variables from x_1, \dots, x_n . We think of these terms as implementing the partial derivative $\frac{\partial^{\beta_1 + \dots + \beta_n} \llbracket \mathbf{op} \rrbracket (x_1, \dots, x_n)}{\partial x_1^{\beta_1} \dots \partial x_n^{\beta_n}}$ of the smooth function $\llbracket \mathbf{op} \rrbracket : \mathbb{R}^n \rightarrow \mathbb{R}$ that \mathbf{op} implements. For example, we could choose the following representations of derivatives of order ≤ 2 of our example operations

$$\begin{aligned} \partial_{01}(+)(x_1, x_2) &= \underline{1} & \partial_{02}(+)(x_1, x_2) &= \underline{0} \\ \partial_{10}(+)(x_1, x_2) &= \underline{1} & \partial_{11}(+)(x_1, x_2) &= \underline{0} \\ \partial_{20}(+)(x_1, x_2) &= \underline{0} & & \\ \\ \partial_{01}(*)(x_1, x_2) &= x_1 & \partial_{02}(*)(x_1, x_2) &= \underline{0} \\ \partial_{10}(*)(x_1, x_2) &= x_2 & \partial_{11}(*)(x_1, x_2) &= \underline{1} \\ \partial_{20}(*)(x_1, x_2) &= \underline{0} & & \\ \\ \partial_1(\varsigma)(x) &= \mathbf{let } y = \varsigma(x) \mathbf{ in } y * (\underline{1} - y) & \partial_2(\varsigma)(x) &= \mathbf{let } y = \varsigma(x) \mathbf{ in} \\ & & & \mathbf{let } z = y * (\underline{1} - y) \mathbf{ in } z * (\underline{1} - \underline{2} * y) \end{aligned}$$

Note that our rules, in particular, imply that $\vec{\mathcal{D}}(\underline{c}) = \langle \underline{c}, \underline{0}, \dots, \underline{0} \rangle$.

Example 3.1 ((1,1)- and (2,2)-AD). Our choices of partial derivatives of the example operations are sufficient to implement (k, R) -Taylor forward AD with $R \leq 2$. To be explicit, the distinctive formulas for (1,1)- and (2,2)-AD methods (specializing our abstract definition of $\vec{\mathcal{D}}_{(k,R)}$ above) are

$$\begin{aligned} \vec{\mathcal{D}}_{(1,1)}(\mathbf{real}) &= (\mathbf{real} * \mathbf{real}) \\ \vec{\mathcal{D}}_{(1,1)}(\mathbf{op}(t_1, \dots, t_n)) &= \\ &\mathbf{case } \vec{\mathcal{D}}_{(1,1)}(t_1) \mathbf{ of } \langle x_0^1, x_1^1 \rangle \rightarrow \dots \rightarrow \mathbf{case } \vec{\mathcal{D}}_{(1,1)}(t_n) \mathbf{ of } \langle x_0^n, x_1^n \rangle \rightarrow \\ &\langle \mathbf{op}(x_0^1, \dots, x_0^n), \sum_{i=1}^n x_1^i * \partial_i \mathbf{op}(x_0^1, \dots, x_0^n) \rangle \\ \\ \vec{\mathcal{D}}_{(2,2)}(\mathbf{real}) &= \mathbf{real}^6 \\ \vec{\mathcal{D}}_{(2,2)}(\mathbf{op}(t_1, \dots, t_n)) &= \\ &\mathbf{case } \vec{\mathcal{D}}_{(2,2)}(t_1) \mathbf{ of } \langle x_{00}^1, x_{01}^1, x_{02}^1, x_{10}^1, x_{11}^1, x_{20}^1 \rangle \rightarrow \\ &\vdots \\ &\mathbf{case } \vec{\mathcal{D}}_{(2,2)}(t_n) \mathbf{ of } \langle x_{00}^n, x_{01}^n, x_{02}^n, x_{10}^n, x_{11}^n, x_{20}^n \rangle \rightarrow \\ &\langle \mathbf{op}(x_{00}^1, \dots, x_{00}^n), \\ &\sum_{i=1}^n x_{01}^i * \partial_i \mathbf{op}(x_{00}^1, \dots, x_{00}^n), \\ &\sum_{i=1}^n x_{02}^i * \partial_i \mathbf{op}(x_{00}^1, \dots, x_{00}^n) + \sum_{i,j=1}^n x_{01}^i * x_{01}^j * \partial_{\widehat{i,j}} \mathbf{op}(x_{00}^1, \dots, x_{00}^n), \\ &\sum_{i=1}^n x_{10}^i * \partial_i \mathbf{op}(x_{00}^1, \dots, x_{00}^n), \\ &\sum_{i=1}^n x_{11}^i * \partial_i \mathbf{op}(x_{00}^1, \dots, x_{00}^n) + \sum_{i,j=1}^n x_{10}^i * x_{01}^j * \partial_{\widehat{i,j}} \mathbf{op}(x_{00}^1, \dots, x_{00}^n), \\ &\sum_{i=1}^n x_{20}^i * \partial_i \mathbf{op}(x_{00}^1, \dots, x_{00}^n) + \sum_{i,j=1}^n x_{10}^i * x_{10}^j * \partial_{\widehat{i,j}} \mathbf{op}(x_{00}^1, \dots, x_{00}^n) \rangle \end{aligned}$$

where we informally write \hat{i} for the one-hot encoding of i (the sequence of length n consisting exclusively of zeros except in position i where it has a 1) and $\widehat{i,j}$ for the two-hot encoding of i and j (the sequence of length n consisting exclusively of zeros except in positions i and j where it has a 1 if $i \neq j$ and a 2 if $i = j$)

As noted in Section 2, it is often unnecessary to include all components of the (2,2)-algorithm, for example when computing a second order directional derivative. In that case,

we may define a restricted (2,2)-AD algorithm that drops the non-mixed second order derivatives from the definitions above and defines $\vec{\mathcal{D}}_{(2,2)'}(\mathbf{real}) = \mathbf{real}^4$ and

$$\begin{aligned} \vec{\mathcal{D}}_{(2,2)'}(\mathbf{op}(t_1, \dots, t_n)) &= \mathbf{case} \vec{\mathcal{D}}_{(2,2)'}(t_1) \mathbf{of} \langle x_{00}^1, x_{01}^1, x_{10}^1, x_{11}^1 \rangle \rightarrow \\ &\quad \vdots \\ &\quad \mathbf{case} \vec{\mathcal{D}}_{(2,2)'}(t_n) \mathbf{of} \langle x_{00}^n, x_{01}^n, x_{10}^n, x_{11}^n \rangle \rightarrow \\ &\quad \langle \mathbf{op}(x_{00}^1, \dots, x_{00}^n), \\ &\quad \sum_{i=1}^n x_{01}^i * \partial_i \mathbf{op}(x_{00}^1, \dots, x_{00}^n), \\ &\quad \sum_{i=1}^n x_{10}^i * \partial_i \mathbf{op}(x_{00}^1, \dots, x_{00}^n), \\ &\quad \sum_{i=1}^n x_{11}^i * \partial_i \mathbf{op}(x_{00}^1, \dots, x_{00}^n) \\ &\quad + \sum_{i,i'=1}^n x_{10}^i * x_{01}^{i'} * \partial_{i,i'} \mathbf{op}(x_{00}^1, \dots, x_{00}^n) \rangle. \end{aligned}$$

We extend $\vec{\mathcal{D}}$ to contexts: $\vec{\mathcal{D}}(\{x_1:\tau_1, \dots, x_n:\tau_n\}) \stackrel{\text{def}}{=} \{x_1:\vec{\mathcal{D}}(\tau_1), \dots, x_n:\vec{\mathcal{D}}(\tau_n)\}$. This turns $\vec{\mathcal{D}}$ into a well-typed, functorial macro in the following sense.

Lemma 3.2 (Functorial macro). *If $\Gamma \vdash t : \tau$ then $\vec{\mathcal{D}}(\Gamma) \vdash \vec{\mathcal{D}}(t) : \vec{\mathcal{D}}(\tau)$.*

If $\Gamma, x : \sigma \vdash t : \tau$ and $\Gamma \vdash s : \sigma$ then $\vec{\mathcal{D}}(\Gamma) \vdash \vec{\mathcal{D}}(t[s/x]) = \vec{\mathcal{D}}(t)[\vec{\mathcal{D}}(s)/x]$.

Proof. By induction on the structure of typing derivations. □

Example 3.3 (Inner products). Let us write τ^n for the n -fold product $(\tau * \dots * \tau)$. Then, given $\Gamma \vdash t, s : \mathbf{real}^n$ we can define their inner product

$$\Gamma \vdash t \cdot_n s \stackrel{\text{def}}{=} \mathbf{case} t \mathbf{of} \langle z_1, \dots, z_n \rangle \rightarrow \mathbf{case} s \mathbf{of} \langle y_1, \dots, y_n \rangle \rightarrow z_1 * y_1 + \dots + z_n * y_n : \mathbf{real}$$

To illustrate the calculation of $\vec{\mathcal{D}}_{(1,1)}$, let us expand (and β -reduce) $\vec{\mathcal{D}}_{(1,1)}(t \cdot_2 s)$:

$$\begin{aligned} &\mathbf{case} \vec{\mathcal{D}}_{(1,1)}(t) \mathbf{of} \langle z_1, z_2 \rangle \rightarrow \mathbf{case} \vec{\mathcal{D}}_{(1,1)}(s) \mathbf{of} \langle y_1, y_2 \rangle \rightarrow \\ &\mathbf{case} z_1 \mathbf{of} \langle z_{1,1}, z_{1,2} \rangle \rightarrow \mathbf{case} y_1 \mathbf{of} \langle y_{1,1}, y_{1,2} \rangle \rightarrow \\ &\mathbf{case} z_2 \mathbf{of} \langle z_{2,1}, z_{2,2} \rangle \rightarrow \mathbf{case} y_2 \mathbf{of} \langle y_{2,1}, y_{2,2} \rangle \rightarrow \\ &\langle z_{1,1} * y_{1,1} + z_{2,1} * y_{2,1}, z_{1,1} * y_{1,2} + z_{1,2} * y_{1,1} + z_{2,1} * y_{2,2} + z_{2,2} * y_{2,1} \rangle \end{aligned}$$

Let us also expand the calculation of $\vec{\mathcal{D}}_{(2,2)'}(t \cdot_2 s)$:

$$\begin{aligned} &\mathbf{case} \vec{\mathcal{D}}_{(2,2)'}(t) \mathbf{of} \langle z_1, z_2 \rangle \rightarrow \mathbf{case} \vec{\mathcal{D}}_{(2,2)'}(s) \mathbf{of} \langle y_1, y_2 \rangle \rightarrow \\ &\mathbf{case} z_1 \mathbf{of} \langle z_1, z'_{1,1}, z'_{1,2}, z''_1 \rangle \rightarrow \mathbf{case} y_1 \mathbf{of} \langle y_1, y'_{1,1}, y'_{1,2}, y''_1 \rangle \rightarrow \\ &\mathbf{case} z_2 \mathbf{of} \langle z_2, z'_{2,1}, z'_{2,2}, z''_2 \rangle \rightarrow \mathbf{case} y_2 \mathbf{of} \langle y_2, y'_{2,1}, y'_{2,2}, y''_2 \rangle \rightarrow \\ &\langle z_1 * y_1 + z_2 * y_2, \\ &\quad z_1 * y'_{1,1} + z'_{1,1} * y_1 + z_2 * y'_{2,1} + z'_{2,1} * y_2, \\ &\quad z_1 * y'_{1,2} + z'_{1,2} * y_1 + z_2 * y'_{2,2} + z'_{2,2} * y_2, \\ &\quad z''_1 * y_1 + z''_2 * y_2 + y''_1 * z_1 + y''_2 * z_2 + \\ &\quad y'_{1,1} * y'_{1,2} + y'_{2,1} * y'_{2,2} + z'_{1,1} * z'_{1,2} + z'_{2,1} * z'_{2,2} \rangle \end{aligned}$$

Example 3.4 (Neural networks). In our introduction, we provided a program (1.1) in our language to build a neural network out of expressions `neuron`, `layer`, `comp`; this program makes use of the inner product of Ex. 3.3. We can similarly calculate the derivatives of deep neural nets by mechanically applying the macro $\vec{\mathcal{D}}$.

4. SEMANTICS OF DIFFERENTIATION

Consider for a moment the first order fragment of the language in Section 3, with only one type, **real**, and no λ 's or pairs. This has a simple semantics in the category of cartesian spaces and smooth maps. Indeed, a term $x_1 \dots x_n : \mathbf{real} \vdash t : \mathbf{real}$ has a natural reading as a function $\llbracket t \rrbracket : \mathbb{R}^n \rightarrow \mathbb{R}$ by interpreting our operation symbols by the well-known operations on $\mathbb{R}^n \rightarrow \mathbb{R}$ with the corresponding name. In fact, the functions that are definable in this first order fragment are smooth. Let us write **CartSp** for this category of cartesian spaces (\mathbb{R}^n for some n) and smooth functions.

The category **CartSp** has cartesian products, and so we can also interpret product types, tupling and pattern matching, giving us a useful syntax for constructing functions into and out of products of \mathbb{R} . For example, the interpretation of (neuron_n) in (1.1) becomes

$$\mathbb{R}^n \times \mathbb{R}^n \times \mathbb{R} \xrightarrow{\llbracket \cdot \rrbracket_n \times \text{id}_{\mathbb{R}}} \mathbb{R} \times \mathbb{R} \xrightarrow{\llbracket + \rrbracket} \mathbb{R} \xrightarrow{\llbracket \varsigma \rrbracket} \mathbb{R}$$

where $\llbracket \cdot \rrbracket_n$, $\llbracket + \rrbracket$ and $\llbracket \varsigma \rrbracket$ are the usual inner product, addition and the sigmoid function on \mathbb{R} , respectively.

Inside this category, we can straightforwardly study the first order language without λ 's, and automatic differentiation. In fact, we can prove the following by plain induction on the syntax:

The interpretation of the (syntactic) forward AD $\vec{D}(t)$ of a first order term t equals the usual (semantic) derivative of the interpretation of t as a smooth function.

However, as is well-known, the category **CartSp** does not support function spaces. To see this, notice that we have polynomial terms

$$x_1, \dots, x_d : \mathbf{real} \vdash \lambda y. \sum_{n=1}^d x_n y^n : \mathbf{real} \rightarrow \mathbf{real}$$

for each d , and so if we could interpret $(\mathbf{real} \rightarrow \mathbf{real})$ as a Euclidean space \mathbb{R}^p then, by interpreting these polynomial expressions, we would be able to find continuous injections $\mathbb{R}^d \rightarrow \mathbb{R}^p$ for every d , which is topologically impossible for any p , for example as a consequence of the Borsuk-Ulam theorem (see Appx. A).

This lack of function spaces means that we cannot interpret the functions (layer) and (comp) from (1.1) in **CartSp**, as they are higher order functions, even though they are very useful and innocent building blocks for differential programming! Clearly, we could define neural nets such as (1.1) directly as smooth functions without any higher order subcomponents, though that would quickly become cumbersome for deep networks. A problematic consequence of the lack of a semantics for higher order differential programs is that we have no obvious way of establishing compositional semantic correctness of \vec{D} for the given implementation of (1.1).

We now show that every definable function is smooth, and then in Section 4.2 we show that the \vec{D} macro witnesses its derivatives.

4.1. Smoothness at higher types and diffeologies. The aim of this section is to introduce diffeological spaces as a semantic model for the simple language in Section 3. By way of motivation, we begin with a standard set theoretic semantics, where types are interpreted as follows

$$\llbracket \mathbf{real} \rrbracket \stackrel{\text{def}}{=} \mathbb{R} \quad \llbracket (\tau_1 * \dots * \tau_n) \rrbracket \stackrel{\text{def}}{=} \prod_{i=1}^n \llbracket \tau_i \rrbracket \quad \llbracket \tau \rightarrow \sigma \rrbracket \stackrel{\text{def}}{=} (\llbracket \tau \rrbracket \rightarrow \llbracket \sigma \rrbracket)$$

and a term $x_1 : \tau_1, \dots, x_n : \tau_n \vdash t : \sigma$ is interpreted as a function $\prod_{i=1}^n \llbracket \tau_i \rrbracket \rightarrow \llbracket \sigma \rrbracket$, mapping a valuation of the context to a result.

We can show that the interpretation of a term $x_1 : \mathbf{real}, \dots, x_n : \mathbf{real} \vdash t : \mathbf{real}$ is always a smooth function $\mathbb{R}^n \rightarrow \mathbb{R}$, even if it has higher order subterms. We begin with a fairly standard logical relations proof of this, and then move from this to the semantic model of diffeological spaces.

Proposition 4.1. *If $x_1 : \mathbf{real}, \dots, x_n : \mathbf{real} \vdash t : \mathbf{real}$ then the function $\llbracket t \rrbracket : \mathbb{R}^n \rightarrow \mathbb{R}$ is smooth.*

Proof. For each type τ define a set $Q_\tau \subseteq [\mathbb{R}^k \rightarrow \llbracket \tau \rrbracket]$ by induction on the structure of types:

$$\begin{aligned} Q_{\mathbf{real}} &= \{f : \mathbb{R}^k \rightarrow \mathbb{R} \mid f \text{ is smooth}\} \\ Q_{(\tau_1 * \dots * \tau_n)} &= \{f : \mathbb{R}^k \rightarrow \prod_{i=1}^n \llbracket \tau_i \rrbracket \mid \forall \vec{r} \in \mathbb{R}^k. \forall i. f_i(\vec{r}) \in Q_{\tau_i}\} \\ Q_{\tau \rightarrow \sigma} &= \{f : \mathbb{R}^k \rightarrow \llbracket \tau \rrbracket \rightarrow \llbracket \sigma \rrbracket \mid \forall g \in Q_\tau. \lambda(\vec{r}). f(\vec{r})(g(\vec{r})) \in Q_\sigma\} \end{aligned}$$

Now we show the fundamental lemma: if $x_1 : \tau_1, \dots, x_n : \tau_n \vdash u : \sigma$ and $g_1 \in Q_{\tau_1} \dots g_n \in Q_{\tau_n}$ then $((g_1 \dots g_n); \llbracket u \rrbracket) \in Q_\sigma$. This is shown by induction on the structure of typing derivations. The only interesting step here is that the basic operations (+, *, ζ etc.) are smooth. We deduce the statement of the theorem by putting $u = t$, $k = n$, and letting $g_i : \mathbb{R}^n \rightarrow \mathbb{R}$ be the projections. \square

At higher types, the logical relations Q show that we can only define functions that send smooth functions to smooth functions, meaning that we can never use them to build first order functions that are not smooth. For example, (comp) in (1.1) has this property.

This logical relations proof suggests to build a semantic model by interpreting types as sets with structure: for each type we have a set X together with a set $Q_X^{\mathbb{R}^k} \subseteq [\mathbb{R}^k \rightarrow X]$ of plots.

Definition 4.2. A *diffeological space* (X, \mathcal{P}_X) consists of a set X together with, for each n and each open subset U of \mathbb{R}^n , a set $\mathcal{P}_X^U \subseteq [U \rightarrow X]$ of functions, called *plots*, such that

- all constant functions are plots;
- if $f : V \rightarrow U$ is a smooth function and $p \in \mathcal{P}_X^U$, then $f; p \in \mathcal{P}_X^V$;
- if $(p_i \in \mathcal{P}_X^{U_i})_{i \in I}$ is a compatible family of plots ($x \in U_i \cap U_j \Rightarrow p_i(x) = p_j(x)$) and $(U_i)_{i \in I}$ covers U , then the gluing $p : U \rightarrow X : x \in U_i \mapsto p_i(x)$ is a plot.

We call a function $f : X \rightarrow Y$ between diffeological spaces *smooth* if, for all plots $p \in \mathcal{P}_X^U$, we have that $p; f \in \mathcal{P}_Y^U$. We write $\mathbf{Diff}(X, Y)$ for the set of smooth maps from X to Y . Smooth functions compose, and so we have a category \mathbf{Diff} of diffeological spaces and smooth functions.

A diffeological space is thus a set equipped with structure. Many constructions of sets carry over straightforwardly to diffeological spaces.

Example 4.3 (Cartesian diffeologies). Each open subset U of \mathbb{R}^n can be given the structure of a diffeological space by taking all the smooth functions $V \rightarrow U$ as \mathcal{P}_U^V . Smooth functions from $V \rightarrow U$ in the traditional sense coincide with smooth functions in the sense of diffeological spaces [IZ13]. Thus diffeological spaces have a profound relationship with ordinary calculus.

In categorical terms, this gives a full embedding of \mathbf{CartSp} in \mathbf{Diff} .

Example 4.4 (Product diffeologies). Given a family $(X_i)_{i \in I}$ of diffeological spaces, we can equip the product $\prod_{i \in I} X_i$ of sets with the *product diffeology* in which U -plots are precisely the functions of the form $(p_i)_{i \in I}$ for $p_i \in \mathcal{P}_{X_i}^U$.

This gives us the categorical product in **Diff**.

Example 4.5 (Functional diffeology). We can equip the set $\mathbf{Diff}(X, Y)$ of smooth functions between diffeological spaces with the *functional diffeology* in which U -plots consist of functions $f : U \rightarrow \mathbf{Diff}(X, Y)$ such that $(u, x) \mapsto f(u)(x)$ is an element of $\mathbf{Diff}(U \times X, Y)$.

This specifies the categorical function object in **Diff**.

We can now give a denotational semantics to our language from Section 3 in the category of diffeological spaces. We interpret each type τ as a set $\llbracket \tau \rrbracket$ equipped with the relevant diffeology, by induction on the structure of types:

$$\begin{aligned} \llbracket \mathbf{real} \rrbracket &\stackrel{\text{def}}{=} \mathbb{R} && \text{with the standard diffeology} \\ \llbracket (\tau_1 * \dots * \tau_n) \rrbracket &\stackrel{\text{def}}{=} \prod_{i=1}^n \llbracket \tau_i \rrbracket && \text{with the product diffeology} \\ \llbracket \tau \rightarrow \sigma \rrbracket &\stackrel{\text{def}}{=} \mathbf{Diff}(\llbracket \tau \rrbracket, \llbracket \sigma \rrbracket) && \text{with the functional diffeology} \end{aligned}$$

A context $\Gamma = (x_1 : \tau_1 \dots x_n : \tau_n)$ is interpreted as a diffeological space $\llbracket \Gamma \rrbracket \stackrel{\text{def}}{=} \prod_{i=1}^n \llbracket \tau_i \rrbracket$. Now well typed terms $\Gamma \vdash t : \tau$ are interpreted as smooth functions $\llbracket t \rrbracket : \llbracket \Gamma \rrbracket \rightarrow \llbracket \tau \rrbracket$, giving a meaning for t for every valuation of the context. This is routinely defined by induction on the structure of typing derivations once we choose a smooth function $\llbracket \text{op} \rrbracket : \mathbb{R}^n \rightarrow \mathbb{R}$ to interpret each n -ary operation $\text{op} \in \mathbf{Op}_n$. For example, constants $\underline{c} : \mathbf{real}$ are interpreted as constant functions; and the first order operations $(+, *, \varsigma)$ are interpreted by composing with the corresponding functions, which are smooth: e.g., $\llbracket \varsigma(t) \rrbracket(\rho) \stackrel{\text{def}}{=} \varsigma(\llbracket t \rrbracket(\rho))$, where $\rho \in \llbracket \Gamma \rrbracket$. Variables are interpreted as $\llbracket x_i \rrbracket(\rho) \stackrel{\text{def}}{=} \rho_i$. The remaining constructs are interpreted as follows, and it is straightforward to show that smoothness is preserved.

$$\begin{aligned} \llbracket \langle t_1, \dots, t_n \rangle \rrbracket(\rho) &\stackrel{\text{def}}{=} (\llbracket t_1 \rrbracket(\rho), \dots, \llbracket t_n \rrbracket(\rho)) && \llbracket \lambda x : \tau. t \rrbracket(\rho)(a) \stackrel{\text{def}}{=} \llbracket t \rrbracket(\rho, a) \quad (a \in \llbracket \tau \rrbracket) \\ \llbracket \mathbf{case } t \text{ of } \langle \dots \rangle \rightarrow s \rrbracket(\rho) &\stackrel{\text{def}}{=} \llbracket s \rrbracket(\rho, \llbracket t \rrbracket(\rho)) && \llbracket t s \rrbracket(\rho) \stackrel{\text{def}}{=} \llbracket t \rrbracket(\rho)(\llbracket s \rrbracket(\rho)) \end{aligned}$$

The logical relations proof of Proposition 4.1 is reminiscent of diffeological spaces. We now briefly remark on the suitability of the axioms of diffeological spaces (Def 4.2) for a semantic model of smooth programs. The first axiom says that we only consider reflexive logical relations. From the perspective of the interpretation, it recognizes in particular that the semantics of an expression of type $(\mathbf{real} \rightarrow \mathbf{real}) \rightarrow \mathbf{real}$ is defined by its value on smooth functions rather than arbitrary arguments. That is to say, the set-theoretic semantics at the beginning of this section, $\llbracket (\mathbf{real} \rightarrow \mathbf{real}) \rightarrow \mathbf{real} \rrbracket$, is different to the diffeological semantics, $\llbracket (\mathbf{real} \rightarrow \mathbf{real}) \rightarrow \mathbf{real} \rrbracket$. The second axiom for diffeological spaces ensures that the smooth maps in $\mathbf{Diff}(U, X)$ are exactly the plots in \mathcal{P}_X^U . The third axiom ensures that categories of manifolds fully embed into **Diff**; it will not play a visible role in this paper — in fact, [BCLG20] prove similar results for a simple language like ours by using plain logical relations (over **Set**) and without demanding the diffeology axioms. However, we expect the third axiom to be crucial for programming with other smooth structures or partiality.

4.2. Correctness of AD. We have shown that a term $x_1 : \mathbf{real}, \dots, x_n : \mathbf{real} \vdash t : \mathbf{real}$ is interpreted as a smooth function $\llbracket t \rrbracket : \mathbb{R}^n \rightarrow \mathbb{R}$, even if t involves higher order functions (like (1.1)). Moreover, the macro differentiation $\vec{\mathcal{D}}_{(k,R)}(t)$ is a function $\llbracket \vec{\mathcal{D}}_{(k,R)}(t) \rrbracket :$

$(\mathbb{R}^{\binom{R+k}{k}})^n \rightarrow \mathbb{R}^{\binom{R+k}{k}}$ (Proposition 3.2). This enables us to state a limited version of our main correctness theorem:

Theorem 4.6 (Semantic correctness of $\vec{\mathcal{D}}$ (limited)). *For any term $x_1 : \mathbf{real}, \dots, x_n : \mathbf{real} \vdash t : \mathbf{real}$, the function $\llbracket \vec{\mathcal{D}}_{(k,R)}(t) \rrbracket$ is the (k, R) -Taylor representation (2.2) of $\llbracket t \rrbracket$. In detail: for any smooth functions $f_1 \dots f_n : \mathbb{R}^k \rightarrow \mathbb{R}$,*

$$\left(\left(\frac{\partial^{\alpha_1 + \dots + \alpha_k} f_j(x)}{\partial x_1^{\alpha_1} \dots \partial x_k^{\alpha_k}} \right)_{(\alpha_1, \dots, \alpha_k) = (0, \dots, 0)}^{(R, 0, \dots, 0)} \right)_{j=1}^n ; \llbracket \vec{\mathcal{D}}_{(k,R)}(t) \rrbracket = \left(\left(\frac{\partial^{\alpha_1 + \dots + \alpha_k} ((f_1, \dots, f_n); \llbracket t \rrbracket)(x)}{\partial x_1^{\alpha_1} \dots \partial x_k^{\alpha_k}} \right)_{(\alpha_1, \dots, \alpha_k) = (0, \dots, 0)}^{(R, 0, \dots, 0)} \right)_{j=1}^n.$$

For instance, if $n = 2$, then $\llbracket \vec{\mathcal{D}}_{(1,1)}(t) \rrbracket(x_1, 1, x_2, 0) = (\llbracket t \rrbracket(x_1, x_2), \frac{\partial \llbracket t \rrbracket(x_1, x_2)}{\partial x}(x_1))$.

Proof. We prove this by logical relations. A categorical version of this proof is in Section 6.2.

For each type τ , we define a binary relation S_τ between (open) k -dimensional plots in $\llbracket \tau \rrbracket$ and (open) k -dimensional plots in $\llbracket \vec{\mathcal{D}}_{(k,R)}(\tau) \rrbracket$, i.e. $S_\tau \subseteq \mathcal{P}_{\llbracket \tau \rrbracket}^{\mathbb{R}^k} \times \mathcal{P}_{\llbracket \vec{\mathcal{D}}_{(k,R)}(\tau) \rrbracket}^{\mathbb{R}^k}$, by induction on τ :

$$\begin{aligned} S_{\mathbf{real}} &\stackrel{\text{def}}{=} \left\{ \left(f, \left(\frac{\partial^{\alpha_1 + \dots + \alpha_k} f(x)}{\partial x_1^{\alpha_1} \dots \partial x_k^{\alpha_k}} \right)_{(\alpha_1, \dots, \alpha_k) = (0, \dots, 0)}^{(R, 0, \dots, 0)} \right) \mid f : \mathbb{R}^k \rightarrow \mathbb{R} \text{ smooth} \right\} \\ S_{(\tau_1 * \dots * \tau_n)} &\stackrel{\text{def}}{=} \{ ((f_1, \dots, f_n), (g_1, \dots, g_n)) \mid (f_1, g_1) \in S_{\tau_1}, \dots, (f_n, g_n) \in S_{\tau_n} \} \\ S_{\tau \rightarrow \sigma} &\stackrel{\text{def}}{=} \{ (f_1, f_2) \mid \forall (g_1, g_2) \in S_\tau. (x \mapsto f_1(x)(g_1(x)), x \mapsto f_2(x)(g_2(x))) \in S_\sigma \} \end{aligned}$$

Then, we establish the following ‘fundamental lemma’:

If $x_1 : \tau_1, \dots, x_n : \tau_n \vdash t : \sigma$ and, for all $1 \leq i \leq n$, $f_i : \mathbb{R}^k \rightarrow \llbracket \tau_i \rrbracket$ and $g_i : \mathbb{R}^k \rightarrow \llbracket \vec{\mathcal{D}}_{(k,R)}(\tau_i) \rrbracket$ are such that (f_i, g_i) is in S_{τ_i} , then we have that

$$\left((f_1, \dots, f_n); \llbracket t \rrbracket, (g_1, \dots, g_n); \llbracket \vec{\mathcal{D}}_{(k,R)}(t) \rrbracket \right)$$

is in S_σ .

This is proved routinely by induction on the typing derivation of t . The case for $\mathbf{op}(t_1, \dots, t_n)$ relies on the precise definition of $\vec{\mathcal{D}}_{(k,R)}(\mathbf{op}(t_1, \dots, t_n))$.

We conclude the theorem from the fundamental lemma by considering the case where $\tau_i = \sigma = \mathbf{real}$, $m = n$ and $s_i = y_i$. \square

5. EXTENDING THE LANGUAGE: VARIANT AND INDUCTIVE TYPES

In this section, we show that the definition of forward AD and the semantics generalize if we extend the language of Section 3 with variants and inductive types. As an example of inductive types, we consider lists. This specific choice is only for expository purposes and the whole development works at the level of generality of arbitrary algebraic data types generated as initial algebras of (polynomial) type constructors formed by finite products and variants. These types are easily interpreted in the category of diffeological spaces in

much the same way. The categorically minded reader may regard this as a consequence of **Diff** being a concrete Grothendieck quasitopos, e.g. [BH11], and hence is complete and cocomplete.

5.1. **Language.** We additionally consider the following types and terms:

$\tau, \sigma, \rho ::= \dots$	types
$\{\ell_1 \tau_1 \mid \dots \mid \ell_n \tau_n\}$	variant
list (τ)	list
$t, s, r ::= \dots$	terms
$\tau.l t$	variant constructor
$[] \mid t :: s$	empty list and cons
case t of $\{\ell_1 x_1 \rightarrow s_1 \mid \dots \mid \ell_n x_n \rightarrow s_n\}$	pattern matching: variants
fold $(x_1, x_2).t$ over s from r	list fold

We extend the type system according to the rules of Fig. 4. We can then extend $\vec{\mathcal{D}}_{(k,R)}$ (again, writing it as $\vec{\mathcal{D}}$, for legibility) to our new types and terms by

$$\begin{aligned}
\vec{\mathcal{D}}(\{\ell_1 \tau_1 \mid \dots \mid \ell_n \tau_n\}) &\stackrel{\text{def}}{=} \{\ell_1 \vec{\mathcal{D}}(\tau_1) \mid \dots \mid \ell_n \vec{\mathcal{D}}(\tau_n)\} \\
\vec{\mathcal{D}}(\mathbf{list}(\tau)) &\stackrel{\text{def}}{=} \mathbf{list}(\vec{\mathcal{D}}(\tau)) \\
\vec{\mathcal{D}}(\tau.l t) &\stackrel{\text{def}}{=} \vec{\mathcal{D}}(\tau).l \vec{\mathcal{D}}(t) \\
\vec{\mathcal{D}}([]) &\stackrel{\text{def}}{=} [] \\
\vec{\mathcal{D}}(t :: s) &\stackrel{\text{def}}{=} \vec{\mathcal{D}}(t) :: \vec{\mathcal{D}}(s) \\
\vec{\mathcal{D}}(\mathbf{case } t \mathbf{ of } \{\ell_1 x_1 \rightarrow s_1 \mid \dots \mid \ell_n x_n \rightarrow s_n\}) &\stackrel{\text{def}}{=} \\
&\quad \mathbf{case } \vec{\mathcal{D}}(t) \mathbf{ of } \{\ell_1 x_1 \rightarrow \vec{\mathcal{D}}(s_1) \mid \dots \mid \ell_n x_n \rightarrow \vec{\mathcal{D}}(s_n)\} \\
\vec{\mathcal{D}}(\mathbf{fold } (x_1, x_2).t \mathbf{ over } s \mathbf{ from } r) &\stackrel{\text{def}}{=} \mathbf{fold } (x_1, x_2).\vec{\mathcal{D}}(t) \mathbf{ over } \vec{\mathcal{D}}(s) \mathbf{ from } \vec{\mathcal{D}}(r)
\end{aligned}$$

To demonstrate the practical use of expressive type systems for differential programming, we consider the following two examples.

$\frac{\Gamma \vdash t : \tau_i \quad ((\ell_i \tau_i) \in \tau)}{\Gamma \vdash \tau.l_i t : \tau}$	$\frac{}{\Gamma \vdash [] : \mathbf{list}(\tau)}$	$\frac{\Gamma \vdash t : \tau \quad \Gamma \vdash s : \mathbf{list}(\tau)}{\Gamma \vdash t :: s : \mathbf{list}(\tau)}$
$\frac{\Gamma \vdash t : \{\ell_1 \tau_1 \mid \dots \mid \ell_n \tau_n\} \quad \text{for each } 1 \leq i \leq n: \Gamma, x_i : \tau_i \vdash s_i : \tau}{\Gamma \vdash \mathbf{case } t \mathbf{ of } \{\ell_1 x_1 \rightarrow s_1 \mid \dots \mid \ell_n x_n \rightarrow s_n\} : \tau}$		
$\frac{\Gamma \vdash s : \mathbf{list}(\tau) \quad \Gamma \vdash r : \sigma \quad \Gamma, x_1 : \tau, x_2 : \sigma \vdash t : \sigma}{\Gamma \vdash \mathbf{fold } (x_1, x_2).t \mathbf{ over } s \mathbf{ from } r : \sigma}$		

Figure 4: Additional typing rules for the extended language.

Example 5.1 (Lists of inputs for neural nets). Usually, we run a neural network on a large data set, the size of which might be determined at runtime. To evaluate a neural network on multiple inputs, in practice, one often sums the outcomes. This can be coded in our extended language as follows. Suppose that we have a network $f : (\mathbf{real}^n * P) \rightarrow \mathbf{real}$ that operates on single input vectors. We can construct one that operates on lists of inputs as follows:

$$g \stackrel{\text{def}}{=} \lambda \langle l, w \rangle. \mathbf{fold} (x_1, x_2). f \langle x_1, w \rangle + x_2 \mathbf{over} l \mathbf{from} 0 : (\mathbf{list}(\mathbf{real}^n) * P) \rightarrow \mathbf{real}$$

Example 5.2 (Missing data). In practically every application of statistics and machine learning, we face the problem of *missing data*: for some observations, only partial information is available.

In an expressive typed programming language like we consider, we can model missing data conveniently using the data type $\mathbf{maybe}(\tau) = \{\mathbf{Nothing} () \mid \mathbf{Just} \tau\}$. In the context of a neural network, one might use it as follows. First, define some helper functions

$$\mathbf{fromMaybe} \stackrel{\text{def}}{=} \lambda x. \lambda m. \mathbf{case} m \mathbf{of} \{\mathbf{Nothing} _ \rightarrow x \mid \mathbf{Just} x' \rightarrow x'\}$$

$$\begin{aligned} \mathbf{fromMaybe}^n \stackrel{\text{def}}{=} & \lambda \langle x_1, \dots, x_n \rangle. \lambda \langle m_1, \dots, m_n \rangle. \langle \mathbf{fromMaybe} x_1 m_1, \dots, \mathbf{fromMaybe} x_n m_n \rangle \\ & : (\mathbf{maybe}(\mathbf{real}))^n \rightarrow \mathbf{real}^n \rightarrow \mathbf{real}^n \end{aligned}$$

$$\mathbf{map} \stackrel{\text{def}}{=} \lambda f. \lambda l. \mathbf{fold} (x_1, x_2). f x_1 :: x_2 \mathbf{over} l \mathbf{from} [] : (\tau \rightarrow \sigma) \rightarrow \mathbf{list}(\tau) \rightarrow \mathbf{list}(\sigma)$$

Given a neural network $f : (\mathbf{list}(\mathbf{real}^k) * P) \rightarrow \mathbf{real}$, we can build a new one that operates on a data set for which some covariates (features) are missing, by passing in default values to replace the missing covariates:

$$\lambda \langle l, \langle m, w \rangle \rangle. f \langle \mathbf{map} (\mathbf{fromMaybe}^k m) l, w \rangle : (\mathbf{list}((\mathbf{maybe}(\mathbf{real}))^k) * (\mathbf{real}^k * P)) \rightarrow \mathbf{real}$$

Then, given a data set l with missing covariates, we can perform automatic differentiation on this network to optimize, simultaneously, the ordinary network parameters w and the default values for missing covariates m .

5.2. Semantics. In Section 4 we gave a denotational semantics for the simple language in diffeological spaces. This extends to the language in this section, as follows. As before, each type τ is interpreted as a diffeological space, which is a set equipped with a family of plots:

- A variant type $\{\ell_1 \tau_1 \mid \dots \mid \ell_n \tau_n\}$ is inductively interpreted as the disjoint union of the semantic spaces, $\llbracket \{\ell_1 \tau_1 \mid \dots \mid \ell_n \tau_n\} \rrbracket \stackrel{\text{def}}{=} \uplus_{i=1}^n \llbracket \tau_i \rrbracket$, with U -plots

$$\mathcal{P}_{\llbracket \{\ell_1 \tau_1 \mid \dots \mid \ell_n \tau_n\} \rrbracket}^U \stackrel{\text{def}}{=} \left\{ \left[U_j \xrightarrow{f_j} \llbracket \tau_j \rrbracket \rightarrow \uplus_{i=1}^n \llbracket \tau_i \rrbracket \right]_{j=1}^n \mid U = \uplus_{j=1}^n U_j, f_j \in \mathcal{P}_{\llbracket \tau_j \rrbracket}^{U_j} \right\}.$$

- A list type $\mathbf{list}(\tau)$ is interpreted as the union of the sets of length i tuples for all natural numbers i , $\llbracket \mathbf{list}(\tau) \rrbracket \stackrel{\text{def}}{=} \uplus_{i=0}^{\infty} \llbracket \tau \rrbracket^i$ with U -plots

$$\mathcal{P}_{\llbracket \mathbf{list}(\tau) \rrbracket}^U \stackrel{\text{def}}{=} \left\{ \left[U_j \xrightarrow{f_j} \llbracket \tau \rrbracket^j \rightarrow \uplus_{i=0}^{\infty} \llbracket \tau \rrbracket^i \right]_{j=0}^{\infty} \mid U = \uplus_{j=0}^{\infty} U_j, f_j \in \mathcal{P}_{\llbracket \tau \rrbracket^j}^{U_j} \right\}$$

The constructors and destructors for variants and lists are interpreted as in the usual set theoretic semantics.

It is routine to show inductively that these interpretations are smooth. Thus every term $\Gamma \vdash t : \tau$ in the extended language is interpreted as a smooth function $\llbracket t \rrbracket : \llbracket \Gamma \rrbracket \rightarrow \llbracket \tau \rrbracket$ between diffeological spaces. List objects as initial algebras are computed as usual in a cocomplete category (e.g. [JR11]). More generally, the interpretation for algebraic data types follows exactly the usual categorical semantics of variant types and inductive types (e.g. [Pit95]).

6. CATEGORICAL ANALYSIS OF (HIGHER ORDER) FORWARD AD AND ITS CORRECTNESS

This section has three parts. First, we give a categorical account of the functoriality of AD (Ex. 6.3). Then we introduce our gluing construction, and relate it to the correctness of AD (dgm. (6.1)). Finally, we state and prove a correctness theorem for all first order types by considering a category of manifolds (Th. 6.7).

6.1. Syntactic categories. The key contribution of this subsection is that the AD macro translation (Section 3.2) has a canonical status as a unique functor between categories with structure. To this end, we build a syntactic category **Syn** from our language, which has the property of a *free* category with certain structure. This means that for any category \mathcal{C} with this structure, there is a unique structure-preserving functor $\mathbf{Syn} \rightarrow \mathcal{C}$, which is an interpretation of our language in that category. Generally speaking, this is the categorical view of denotational semantics (e.g. [Pit95]). But in this particular setting, the category **Syn** itself admits alternative forms of this structure, given by the dual numbers interpretation, the triple numbers interpretation etc. of Section 2. This gives canonical functors $\mathbf{Syn} \rightarrow \mathbf{Syn}$ translating the language into itself, which are the AD macro translations (Section 3.2). A key point is that **Syn** is almost entirely determined by universal properties (for example, cartesian closure for the function space); the only freedom is in the choice of interpretation of

- (1) the real numbers **real**, which can be taken as the plain type **real**, or as the dual numbers interpretation $\mathbf{real} * \mathbf{real}$ etc.;
- (2) the primitive operations **op**, which can be taken as the operation **op** itself, or as the derivative of the operation etc..

In more detail, our language induces a syntactic category as follows.

Definition 6.1. Let **Syn** be the category whose objects are types, and where a morphism $\tau \rightarrow \sigma$ is a term in context $x : \tau \vdash t : \sigma$ modulo the $\beta\eta$ -laws (Fig. 5). Composition is by substitution.

For simplicity, we do not impose identities involving the primitive operations, such as the arithmetic identity $x + y = y + x$ in **Syn**. As is standard, this category has the following universal property.

Lemma 6.2 (e.g. [Pit95]). *For every bicartesian closed category \mathcal{C} with list objects, and every choice of an object $F(\mathbf{real}) \in \mathcal{C}$ and morphisms $F(\mathbf{op}) \in \mathcal{C}(F(\mathbf{real})^n, F(\mathbf{real}))$ for all $\mathbf{op} \in \mathbf{Op}_n$ and $n \in \mathbb{N}$, in \mathcal{C} , there is a unique functor $F : \mathbf{Syn} \rightarrow \mathcal{C}$ respecting the interpretation and preserving the bicartesian closed structure as well as list objects.*

Proof notes. The functor $F : \mathbf{Syn} \rightarrow \mathcal{C}$ is a canonical denotational semantics for the language, interpreting types as objects of \mathcal{C} and terms as morphisms. For instance, $F(\tau \rightarrow \sigma) \stackrel{\text{def}}{=} (F\tau \rightarrow F\sigma)$, the function space in the category \mathcal{C} , and $F(ts)$ is the composite (Ft, Fs) ; *eval*. \square

When $\mathcal{C} = \mathbf{Diff}$, the denotational semantics of the language in diffeological spaces (Section 4.5.2) can be understood as the unique structure preserving functor $\llbracket - \rrbracket : \mathbf{Syn} \rightarrow \mathbf{Diff}$ satisfying $\llbracket \mathbf{real} \rrbracket = \mathbb{R}$, $\llbracket \varsigma \rrbracket = \varsigma$ and so on.

Example 6.3 (Canonical definition of forward AD). The forward AD macro $\vec{\mathcal{D}}_{(k,R)}$ (Section 3.5.1) arises as a canonical bicartesian closed functor on \mathbf{Syn} that preserves list objects. Consider the unique bicartesian closed functor $F : \mathbf{Syn} \rightarrow \mathbf{Syn}$ that preserves list objects such that $F(\mathbf{real}) = \mathbf{real}^{\binom{R+k}{k}}$ and

$$F(\mathbf{op}) = z : (F(\mathbf{real}) * \dots * F(\mathbf{real})) \vdash \mathbf{case } z \mathbf{ of } \langle x_1, \dots, x_n \rangle \rightarrow \vec{\mathcal{D}}_{(k,R)}(\mathbf{op}(x_1, \dots, x_n)) : F(\mathbf{real}).$$

Then for any type τ , $F(\tau) = \vec{\mathcal{D}}_{(k,R)}(\tau)$, and for any term $x : \tau \vdash t : \sigma$, $F(t) = \vec{\mathcal{D}}_{(k,R)}(t)$ as morphisms $F(\tau) \rightarrow F(\sigma)$ in the syntactic category.

This observation is a categorical counterpart to Lemma 3.2.

6.2. Categorical gluing and logical relations. Gluing is a method for building new categorical models which has been used for many purposes, including logical relations and realizability [MS92]. Our logical relations argument in the proof of Theorem 4.6 can be understood in this setting. (In fact we originally found the proof of Theorem 4.6 in this way.) In this subsection, for the categorically minded, we explain this, and in doing so we quickly recover a correctness result for the more general language in Section 5 and for arbitrary first order types.

The general, established idea of categorical logical relations starts from the observation that that logical relations are defined by induction on the structure of types. Types have universal properties in a categorical semantics (e.g. cartesian closure for the function space),

$$\begin{aligned}
& \mathbf{case } \langle t_1, \dots, t_n \rangle \mathbf{ of } \langle x_1, \dots, x_n \rangle \rightarrow s = s^{[t_1/x_1, \dots, t_n/x_n]} \\
& s^{[t/y]} \#_{x_1, \dots, x_n} \mathbf{case } t \mathbf{ of } \langle x_1, \dots, x_n \rangle \rightarrow s^{[\langle x_1, \dots, x_n \rangle / y]} \\
& \mathbf{case } \ell_i t \mathbf{ of } \{ \ell_1 x_1 \rightarrow s_1 \mid \dots \mid \ell_n x_n \rightarrow s_n \} = s_i^{[t/x_i]} \\
& s^{[t/y]} \#_{x_1, \dots, x_n} \mathbf{case } t \mathbf{ of } \{ \ell_1 x_1 \rightarrow s^{[\ell_1 x_1 / y]} \mid \dots \mid \ell_n x_n \rightarrow s^{[\ell_n x_n / y]} \} \\
& \mathbf{fold } (x_1, x_2).t \mathbf{ over } [] \mathbf{ from } r = r \\
& \mathbf{fold } (x_1, x_2).t \mathbf{ over } s_1 :: s_2 \mathbf{ from } r = t^{[s_1/x_1, \mathbf{fold } (x_1, x_2).t \mathbf{ over } s_2 \mathbf{ from } r / x_2]} \\
& u = s^{[[]/y]}, r^{[s/x_2]} = s^{[x_1 :: y / y]} \Rightarrow s^{[t/y]} \#_{x_1, x_2} \mathbf{fold } (x_1, x_2).r \mathbf{ over } t \mathbf{ from } u \\
& (\lambda x.t) s = t^{[s/x]} \\
& t \#_x \lambda x.t x \\
& \text{We write } \#_{x_1, \dots, x_n} \text{ to indicate that the variables are free in the left hand side}
\end{aligned}$$

Figure 5: Standard $\beta\eta$ -laws (e.g. [Pit95]) for products, functions, variants and lists.

and so we can organize the logical relations argument by defining some category \mathcal{C} of relations and observing that it has the requisite categorical structure. The interpretation of types as relations can then be understood as coming from a unique structure preserving map $\mathbf{Syn} \rightarrow \mathcal{C}$. In this paper, our logical relations are not quite as simple as a binary relation on sets; rather they are relations between plots. Nonetheless, this still forms a category with the appropriate structure, which follows because it can still be regarded as arising from a gluing construction, as we now explain.

We define a category \mathbf{Gl}_k whose objects are triples (X, X', S) where X and X' are diffeological spaces and $S \subseteq \mathcal{P}_X^{\mathbb{R}^k} \times \mathcal{P}_{X'}^{\mathbb{R}^k}$ is a relation between their k -dimensional plots. A morphism $(X, X', S) \rightarrow (Y, Y', T)$ is a pair of smooth functions $f: X \rightarrow Y$, $f': X' \rightarrow Y'$, such that if $(g, g') \in S$ then $(g; f, g'; f') \in T$. The idea is that this is a semantic domain where we can simultaneously interpret the language and its automatic derivatives.

Proposition 6.4. *The category \mathbf{Gl}_k is bicartesian closed, has list objects, and the projection functor $\text{proj}: \mathbf{Gl}_k \rightarrow \mathbf{Diff} \times \mathbf{Diff}$ preserves this structure.*

Proof notes. The category \mathbf{Gl}_k is a full subcategory of the comma category $\text{id}_{\mathbf{Set}} \downarrow \mathbf{Diff}(\mathbb{R}^k, -) \times \mathbf{Diff}(\mathbb{R}^k, -)$. The result thus follows by the general theory of categorical gluing (e.g. [JLS07, Lemma 15]). □

We give a semantics $\llbracket - \rrbracket = (\llbracket - \rrbracket_0, \llbracket - \rrbracket_1, S_-)$ for the language in \mathbf{Gl}_k , interpreting types τ as objects $(\llbracket \tau \rrbracket_0, \llbracket \tau \rrbracket_1, S_\tau)$, and terms as morphisms. We let $(\mathbf{real})_0 \stackrel{\text{def}}{=} \mathbb{R}$ and $(\mathbf{real})_1 \stackrel{\text{def}}{=} \mathbb{R}^{\binom{R+k}{k}}$, with the relation

$$S_{\mathbf{real}} \stackrel{\text{def}}{=} \left\{ \left(f, \left(\frac{\partial^{\alpha_1 + \dots + \alpha_k} f(x)}{\partial x_1^{\alpha_1} \dots \partial x_k^{\alpha_k}} \right)_{(\alpha_1, \dots, \alpha_k) = (0, \dots, 0)}^{(R, 0, \dots, 0)} \right) \mid f: \mathbb{R}^k \rightarrow \mathbb{R} \text{ smooth} \right\}.$$

We interpret the operations op according to $\llbracket \text{op} \rrbracket$ in $\llbracket - \rrbracket_0$, but according to the (k, R) -Taylor representation of $\llbracket \text{op} \rrbracket$ in $\llbracket - \rrbracket_1$. For instance, when $k = 2$ and $r = 2$, $(*)_1: \mathbb{R}^2 \times \mathbb{R}^2 \rightarrow \mathbb{R}^2$ is

$$\begin{aligned} & (*)_1((x_{00}, x_{01}, x_{02}, x_{10}, x_{11}, x_{20}), (y_{00}, y_{01}, y_{02}, y_{10}, y_{11}, y_{20})) \stackrel{\text{def}}{=} \\ & (x_{00}y_{00}, \\ & x_{00}y_{01} + x_{01}y_{00}, \\ & x_{02}y_{00} + 2x_{01}y_{01} + x_{00}y_{02}, \\ & x_{00}y_{10} + x_{10}y_{00}, \\ & x_{11}y_{00} + x_{01}y_{10} + x_{10}y_{01} + x_{00}y_{11}, \\ & x_{20}y_{00} + 2x_{10}y_{10} + x_{00}y_{20}). \end{aligned}$$

At this point one checks that these interpretations are indeed morphisms in \mathbf{Gl}_k . This is equivalent to the statement that $\llbracket \text{op} \rrbracket_1$ is the (k, R) -Taylor representation of $\llbracket \text{op} \rrbracket$ (2.2). The remaining constructions of the language are interpreted using the categorical structure of \mathbf{Gl}_k , following Lemma 6.2.

Notice that the diagram below commutes. One can check this by hand or note that it follows from the initiality of **Syn** (Lemma 6.2): all the functors preserve all the structure.

$$\begin{array}{ccc}
 \mathbf{Syn} & \xrightarrow{(\text{id}, \vec{\mathcal{D}}_{(k,R)}(-))} & \mathbf{Syn} \times \mathbf{Syn} \\
 \downarrow \llbracket - \rrbracket & & \downarrow \llbracket - \rrbracket \times \llbracket - \rrbracket \\
 \mathbf{Gl}_k & \xrightarrow{\text{proj}} & \mathbf{Diff} \times \mathbf{Diff}
 \end{array} \tag{6.1}$$

We thus arrive at a restatement of the correctness theorem (Th. 4.6), which holds even for the extended language with variants and lists, because for any $x_1 : \mathbf{real}, \dots, x_n : \mathbf{real} \vdash t : \mathbf{real}$, the interpretations $(\llbracket t \rrbracket, \llbracket \vec{\mathcal{D}}_{(k,R)}(t) \rrbracket)$ are in the image of the projection $\mathbf{Gl}_k \rightarrow \mathbf{Diff} \times \mathbf{Diff}$, and hence $\llbracket \vec{\mathcal{D}}_{(k,R)}(t) \rrbracket$ is a (k, R) -Taylor representation of $\llbracket t \rrbracket$.

6.3. Correctness at all first order types, via manifolds. We now generalize Theorem 4.6 to hold at all first order types, not just the reals.

So far, we have shown that our macro translation (Section 3.2) gives correct derivatives to functions of the real numbers, even if other types are involved in the definitions of the functions (Theorem 4.6 and Section 6.2). We can state this formally because functions of the real numbers have well understood derivatives (Section 2). There are no established mathematical notions of derivatives at higher types, and so we cannot even begin to argue that our syntactic derivatives of functions $(\mathbf{real} \rightarrow \mathbf{real}) \rightarrow (\mathbf{real} \rightarrow \mathbf{real})$ match with some existing mathematical notion (see also Section 7).

However, for functions of first order type, like $\mathbf{list}(\mathbf{real}) \rightarrow \mathbf{list}(\mathbf{real})$, there *are* established mathematical notions of derivative, because we can understand $\mathbf{list}(\mathbf{real})$ as the *manifold* of all tuples of reals, and then appeal to the well-known theory of manifolds and jet bundles. We do this now, to achieve a correctness theorem for all first order types (Theorem 6.7). The key high level points are that

- manifolds support a notion of differentiation, and an interpretation of all first order types, but not an interpretation of higher types;
- diffeological spaces support all types, including higher types, but not an established notion of differentiation in general;
- manifolds and smooth maps embed full and faithfully in diffeological spaces, preserving the interpretation of first order types, so we can use the two notions together.

We now explain this development in more detail.

For our purposes, a smooth manifold M is a second-countable Hausdorff topological space together with a smooth atlas. In more detail, a topological space X is second-countable when there exists a collection $U := \{U_i\}_{i \in \mathbb{N}}$ of open subsets of X such that any open subset of X can be written as a union of elements of U . A topological space X is Hausdorff if for every distinct points x and y , there exists disjoint open subsets U, V of X such that $x \in U, y \in V$. A smooth atlas of a topological space X is an open cover \mathcal{U} together with homeomorphisms $(\phi_U : U \rightarrow \mathbb{R}^{n(U)})_{U \in \mathcal{U}}$ (called charts, or local coordinates) such that $\phi_U^{-1}; \phi_V$ is smooth on its domain of definition for all $U, V \in \mathcal{U}$. A function $f : M \rightarrow N$ between manifolds is smooth if $\phi_U^{-1}; f; \psi_V$ is smooth for all charts ϕ_U and ψ_V of M and N , respectively. Let us write **Man** for this category. This definition of manifolds is a slight generalisation of the more usual one from differential geometry because different charts in an atlas may have different

finite dimensions $n(U)$. Thus we consider manifolds with dimensions that are potentially unbounded, albeit locally finite.

Each open subset of \mathbb{R}^n can be regarded as a manifold. This lets us regard the category of manifolds **Man** as a full subcategory of the category of diffeological spaces. We consider a manifold $(X, \{\phi_U\}_U)$ as a diffeological space with the same carrier set X and where the plots \mathcal{P}_X^U , called the *manifold diffeology*, are the smooth functions in $\mathbf{Man}(U, X)$. A function $X \rightarrow Y$ is smooth in the sense of manifolds if and only if it is smooth in the sense of diffeological spaces [IZ13]. For the categorically minded reader, this means that we have a full embedding of **Man** into **Diff**. Moreover, the natural interpretation of the first order fragment of our language in **Man** coincides with that in **Diff**. That is, the embedding of **Man** into **Diff** preserves finite products and countable coproducts (hence initial algebras of polynomial endofunctors).

Proposition 6.5. *Suppose that a type τ is first order, i.e. it is just built from reals, products, variants, and lists (or, again, arbitrary inductive types), and not function types. Then the diffeological space $\llbracket \tau \rrbracket$ is a manifold.*

Proof notes. This is proved by induction on the structure of types. In fact, one may show that every such $\llbracket \tau \rrbracket$ is isomorphic to a manifold of the form $\biguplus_{i=1}^n \mathbb{R}^{d_i}$ where the bound n is either finite or ∞ , but this isomorphism is typically not an identity function. \square

We recall how the Taylor representation of any morphism $f : M \rightarrow N$ of manifolds is given by its action on jets [KSM99, Chapter IV]. For each point x in a manifold M , define the (k, R) -jet space $\mathcal{J}_x^{(k,R)} M$ to be the set $\{\gamma \in \mathbf{Man}(\mathbb{R}^k, M) \mid \gamma(0) = x\} / \sim$ of equivalence classes $[\gamma]$ of k -dimensional plots γ in M based at x , where we identify $\gamma_1 \sim \gamma_2$ iff all partial derivatives of order $\leq R$ coincide in the sense that

$$\frac{\partial^{\alpha_1 + \dots + \alpha_k}(\gamma_1; f)(x)}{\partial x_1^{\alpha_1} \dots \partial x_k^{\alpha_k}}(0) = \frac{\partial^{\alpha_1 + \dots + \alpha_k}(\gamma_2; f)(x)}{\partial x_1^{\alpha_1} \dots \partial x_k^{\alpha_k}}(0)$$

for all smooth $f : M \rightarrow \mathbb{R}$ and all multi-indices $(\alpha_1, \dots, \alpha_k) = (0, \dots, 0), \dots, (R, 0, \dots, 0)$. In the case of $(k, R) = (1, 1)$, a (k, R) -jet space is better known as a *tangent space*. The (k, R) -jet bundle (a.k.a. *tangent bundle*, in case $(k, R) = (1, 1)$) of M is the set $\mathcal{J}^{(k,R)}(M) \stackrel{\text{def}}{=} \biguplus_{x \in M} \mathcal{J}_x^{(k,R)}(M)$. The charts of M equip $\mathcal{J}^{(k,R)}(M)$ with a canonical manifold structure. The (manifold) diffeology of these jet bundles can be concisely summarized by the plots $\mathcal{P}_{\mathcal{J}^{(k,R)}(M)}^U = \left\{ f : U \rightarrow |\mathcal{J}^{(k,R)}(M)| \mid \exists g \in \mathcal{P}_M^{U \times \mathbb{R}^k} . \forall u \in U. (g(u, 0), [v \mapsto g(u, v)]) = f(u) \right\}$.

Then $\mathcal{J}^{(k,R)}$ acts on smooth maps $f : M \rightarrow N$ to give $\mathcal{J}^{(k,R)}(f) : \mathcal{J}^{(k,R)}(M) \rightarrow \mathcal{J}^{(k,R)}(N)$ is defined as $\mathcal{J}^{(k,R)}(f)(x, [\gamma]) \stackrel{\text{def}}{=} (f(x), [\gamma; f])$. In local coordinates, this action $\mathcal{J}^{(k,R)}(f)$ is seen to coincide precisely with the (k, R) -Taylor representation of f given by the Faà di Bruno formula [Mer04]. All told, the (k, R) -jet bundle is a functor $\mathcal{J}^{(k,R)} : \mathbf{Man} \rightarrow \mathbf{Man}$ [KSM99].

We can understand the jet bundle of a composite space in terms of that of its parts.

Lemma 6.6. *There are canonical isomorphisms $\mathcal{J}^{(k,R)}(\biguplus_{i=1}^{\infty} M_i) \cong \biguplus_{i=1}^{\infty} \mathcal{J}^{(k,R)}(M_i)$ and $\mathcal{J}^{(k,R)}(M_1 \times \dots \times M_n) \cong \mathcal{J}^{(k,R)}(M_1) \times \dots \times \mathcal{J}^{(k,R)}(M_n)$.*

Proof notes. For disjoint unions, notice that that smooth morphisms from \mathbb{R}^k into a disjoint union of manifolds always factor over a single inclusion, because \mathbb{R}^k is connected. For products, it is well-known that partial derivatives of a morphism (f_1, \dots, f_n) are calculated component-wise [Lee13, ex. 3-2]. \square

We define a canonical isomorphism $\phi_{\tau}^{\vec{\mathcal{D}}\mathcal{J}} : \llbracket \vec{\mathcal{D}}_{(k,R)}(\tau) \rrbracket \rightarrow \mathcal{J}^{(k,R)}(\llbracket \tau \rrbracket)$ for every type τ , by induction on the structure of types. We let $\phi_{\mathbf{real}}^{\vec{\mathcal{D}}\mathcal{J}} : \llbracket \vec{\mathcal{D}}_{(k,R)}(\mathbf{real}) \rrbracket \rightarrow \mathcal{J}^{(k,R)}(\llbracket \mathbf{real} \rrbracket)$ be given by $\phi_{\mathbf{real}}^{\vec{\mathcal{D}}\mathcal{J}}(x, x') \stackrel{\text{def}}{=} (x, [t \mapsto x + x't])$. For the other types, we use Lemma 6.6. We can now phrase correctness at all first order types.

Theorem 6.7 (Semantic correctness of $\vec{\mathcal{D}}_{(k,R)}$ (full)). *For any ground τ , any first order context Γ and any term $\Gamma \vdash t : \tau$, the syntactic translation $\vec{\mathcal{D}}_{(k,R)}$ coincides with the (k, R) -jet bundle functor, modulo these canonical isomorphisms:*

$$\begin{array}{ccc} \llbracket \vec{\mathcal{D}}_{(k,R)}(\Gamma) \rrbracket & \xrightarrow{\llbracket \vec{\mathcal{D}}_{(k,R)}(t) \rrbracket} & \llbracket \vec{\mathcal{D}}_{(k,R)}(\tau) \rrbracket \\ \phi_{\Gamma}^{\vec{\mathcal{D}}\mathcal{J}} \downarrow \cong & & \cong \downarrow \phi_{\tau}^{\vec{\mathcal{D}}\mathcal{J}} \\ \mathcal{J}^{(k,R)}(\llbracket \Gamma \rrbracket) & \xrightarrow{\mathcal{J}^{(k,R)}(\llbracket t \rrbracket)} & \mathcal{J}^{(k,R)}(\llbracket \tau \rrbracket) \end{array}$$

Proof notes. For any k -dimensional plot $\gamma \in \mathbf{Man}(\mathbb{R}^k, M)$, let $\bar{\gamma} \in \mathbf{Man}(\mathbb{R}^k, \mathcal{J}^{(k,R)}(M))$ be the (k, R) -jet curve, given by $\bar{\gamma}(x) = (\gamma(x), [t \mapsto \gamma(x + t)])$. First, we note that a smooth map $h : \mathcal{J}^{(k,R)}(M) \rightarrow \mathcal{J}^{(k,R)}(N)$ is of the form $\mathcal{J}^{(k,R)}(g)$ for some $g : M \rightarrow N$ if for all smooth $\gamma : \mathbb{R}^k \rightarrow M$ we have $\bar{\gamma}; h = \overline{(\gamma; g)} : \mathbb{R}^k \rightarrow \mathcal{J}^{(k,R)}(N)$. This generalizes (2.2). Second, for any first order type τ , $S_{\llbracket \tau \rrbracket} = \{(f, \tilde{f}) \mid \tilde{f}; \phi_{\tau}^{\vec{\mathcal{D}}\mathcal{J}} = \bar{f}\}$. This is shown by induction on the structure of types. We conclude the theorem from diagram (6.1), by putting these two observations together. □

7. DISCUSSION: WHAT ARE DERIVATIVES OF HIGHER ORDER FUNCTIONS?

In our gluing categories \mathbf{Gl}_k of Section 6.2, we have avoided the question of what semantic derivatives should be associated with higher order functions. Our syntactic macro $\vec{\mathcal{D}}$ provides a specific derivative for every definable function, but in the model \mathbf{Gl}_k there is only a *relation* between plots and their corresponding Taylor representations, and this relation is not necessarily single-valued. Our approach has been rather indifferent about what “the” correct derivative of a higher order function should be. Instead, all we have cared about is that we are using “a” derivative that is correct in the sense that it can never be used to produce incorrect derivatives for first order functions, where we do have an unambiguous notion of correct derivative.

7.1. Automatic derivatives of higher order functions may not be unique! For a concrete example to show that derivatives of higher order functions might not be unique in our framework, let us consider the case $(k, R) = (1, 1)$ and focus on first derivatives of the evaluation function

$$\begin{aligned} \text{ev} : \mathbb{R} &\rightarrow \llbracket (\mathbf{real} \rightarrow \mathbf{real}) \rightarrow \mathbf{real} \rrbracket = \mathbb{R} \rightarrow (\mathbb{R} \Rightarrow \mathbb{R}) \Rightarrow \mathbb{R}; \\ r &\mapsto (f \mapsto f(r)). \end{aligned}$$

Our macro $\vec{\mathcal{D}}$ will return $\lambda a : \mathbb{R}. \lambda f : \mathbb{R} \times \mathbb{R} \Rightarrow \mathbb{R} \times \mathbb{R}. f(a, 1)$. In this section we show that the lambda term $\lambda a : \mathbb{R}. \lambda f : \mathbb{R} \times \mathbb{R} \Rightarrow \mathbb{R} \times \mathbb{R}. \text{sort} f(a, 1)$ is also a valid derivative of the evaluation map, where $\text{sort} : (\mathbb{R} \times \mathbb{R} \Rightarrow \mathbb{R} \times \mathbb{R}) \Rightarrow (\mathbb{R} \times \mathbb{R} \Rightarrow \mathbb{R} \times \mathbb{R})$ is defined by

$$\text{sort} := \lambda f. \lambda ((r, -). (\pi_1(f(r, 0)), \pi_2 \vec{\mathcal{D}}((- , 0); f; \pi_1)(r))).$$

This map is idempotent and it converts any map $\mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R} \times \mathbb{R}$ into the dual-numbers representation of its first component. For example, $(\text{sort}(\text{swap}))$ is the constantly $(0, 0)$ function, where we write

$$\begin{aligned} \text{swap} : \mathbb{R} \times \mathbb{R} &\rightarrow \mathbb{R} \times \mathbb{R} \\ (r, r') &\mapsto (r', r). \end{aligned}$$

According to our gluing semantics, a function $g : (\tau)_1 \rightarrow (\sigma)_1$ defines a correct (k, R) -Taylor representation of a function $f : (\tau)_0 \rightarrow (\sigma)_0$ iff (f, g) defines a morphism $(\tau) \rightarrow (\sigma)$ in \mathbf{Gl}_k . In particular, there is no guarantee that every f has a *unique* correct (k, R) -Taylor representation g . (Although such Taylor representations are, in fact, unique when τ, σ are first order types.) The gluing relation $(\langle \mathbf{real} \rightarrow \mathbf{real} \rangle \rightarrow \mathbf{real})$ in \mathbf{Gl}_1 relates curves in $\gamma : \mathbb{R} \rightarrow (\mathbb{R} \Rightarrow \mathbb{R}) \Rightarrow \mathbb{R}$ to “tangent curves” $\gamma' : \mathbb{R} \rightarrow (\mathbb{R} \times \mathbb{R} \Rightarrow \mathbb{R} \times \mathbb{R}) \Rightarrow \mathbb{R} \times \mathbb{R}$. In this relation, the function ev is related to at least two different tangent curves.

Lemma 7.1. *We have a smooth map*

$$\begin{aligned} \text{sort} : (\mathbb{R} \times \mathbb{R} \Rightarrow \mathbb{R} \times \mathbb{R}) &\rightarrow (\mathbb{R} \times \mathbb{R} \Rightarrow \mathbb{R} \times \mathbb{R}) \\ f &\mapsto ((r, -) \mapsto (\pi_1(f(r, 0)), \nabla((- , 0); f; \pi_1)(r))). \end{aligned}$$

Proof. Let $f \in \mathcal{P}_{\mathbb{R} \times \mathbb{R} \Rightarrow \mathbb{R} \times \mathbb{R}}^U$ and let $\gamma_1, \gamma_2 \in \mathcal{P}_{\mathbb{R}}^U$. Then, also $u \mapsto (\gamma_1(u), 0); f(u); \pi_1 \in \mathcal{P}_{\mathbb{R}}^U = \mathbf{Man}(U, \mathbb{R})$ by definition of the exponential in \mathbf{Diff} . Therefore, we also have that $u \mapsto \nabla(u' \mapsto (\gamma_1(u'), 0); f(u'); \pi_1) \in \mathbf{Man}(U, \mathbb{R}) = \mathcal{P}_{\mathbb{R}}^U$, as we are working with infinitely differentiable smooth maps. Consequently,

$u \mapsto (f; \text{sort})(u)(\gamma_1(u), \gamma_2(u)) = (\pi_1(f(u))(\gamma_1(u), 0), \nabla(u' \mapsto (\gamma_1(u'), 0); f(u'); \pi_1)) \in \mathcal{P}_{\mathbb{R} \times \mathbb{R}}^U$, by definition of the product in \mathbf{Diff} . It follows that $(f; \text{sort}) \in \mathcal{P}_{\mathbb{R} \times \mathbb{R} \Rightarrow \mathbb{R} \times \mathbb{R}}^U$. \square

Proposition 7.2. *We have that both $(\text{ev}, \text{ev}'_1) \in (\langle \mathbf{real} \rightarrow \mathbf{real} \rangle \rightarrow \mathbf{real})$ and $(\text{ev}, \text{ev}'_2) \in (\langle \mathbf{real} \rightarrow \mathbf{real} \rangle \rightarrow \mathbf{real})$ for*

$$\begin{aligned} \text{ev}'_1 : \mathbb{R} &\rightarrow (\mathbb{R} \times \mathbb{R} \Rightarrow \mathbb{R} \times \mathbb{R}) \Rightarrow \mathbb{R} \times \mathbb{R} \\ a &\mapsto (f \mapsto f(a, 1)) \\ \text{ev}'_2 : \mathbb{R} &\rightarrow (\mathbb{R} \times \mathbb{R} \Rightarrow \mathbb{R} \times \mathbb{R}) \Rightarrow \mathbb{R} \times \mathbb{R} \\ a &\mapsto (f \mapsto (\text{sort} f)(a, 1)). \end{aligned}$$

Proof. By definition of $(\langle - \rangle)$, we need to show that for any $(\gamma, \gamma') \in (\langle \mathbf{real} \rightarrow \mathbf{real} \rangle)$, we have that $(x \mapsto \text{ev}(x)(\gamma(x)), x \mapsto \text{ev}'_i(x)(\gamma'(x))) \in (\langle \mathbf{real} \rangle)$. This means that we need to show that for $i = 1, 2$

$$x \mapsto \text{ev}'_i(x)(\gamma'(x)) = (x \mapsto \text{ev}(x)(\gamma(x)), \nabla(x \mapsto \text{ev}(x)(\gamma(x))))$$

Unrolling further, this means we need to show that for any $\gamma : \mathbb{R} \rightarrow \mathbb{R} \Rightarrow \mathbb{R}$ and $\gamma' : \mathbb{R} \rightarrow \mathbb{R} \times \mathbb{R} \Rightarrow \mathbb{R} \times \mathbb{R}$ such that for any $(\delta, \delta') \in (\langle \mathbf{real} \rangle)$ (which means that $\delta : \mathbb{R} \rightarrow \mathbb{R}$ and $\delta' = (\delta, \nabla \delta)$), we have that

$$(r \mapsto \gamma(r)(\delta(r)), r \mapsto \gamma'(r)(\delta'(r))) \in (\langle \mathbf{real} \rangle)$$

The latter part finally means that we need to show that

$$r \mapsto \gamma'(r)(\delta(r), \nabla\delta(r)) = (r \mapsto \gamma(r)(\delta(r)), \nabla(r \mapsto \gamma(r)(\delta(r))))$$

Now, focussing on ev'_1 : we need to show that

$$x \mapsto \text{ev}'_1(x)(\gamma'(x)) = (x \mapsto \text{ev}(x)(\gamma(x)), \nabla(x \mapsto \text{ev}(x)(\gamma(x))))$$

Inlining the definition of ev'_1 : we need to show that

$$x \mapsto \gamma'(x)(x, 1) = (x \mapsto \gamma(x)(x), \nabla(x \mapsto \gamma(x)(x)))$$

This follows by assumption by choosing $\delta(r) = r$, and hence $\delta'(r) = (r, 1)$.

Focussing on ev'_2 : we need to show that

$$x \mapsto \text{ev}'_2(x)(\gamma'(x)) = (x \mapsto \text{ev}(x)(\gamma(x)), \nabla(x \mapsto \text{ev}(x)(\gamma(x))))$$

Inlining ev'_2 's definition: we need to show that

$$\begin{aligned} & \left(x \mapsto ((\pi_1(\gamma'(x)(x, 0)), x \mapsto \nabla((- , 0); \gamma'(x); \pi_1(x)))) \right) = \\ & \left(x \mapsto ((r, _) \mapsto (\pi_1(\gamma'(x)(r, 0)), \nabla((- , 0); \gamma'(x); \pi_1(r)))(x, 1) \right) \end{aligned}$$

is equal to

$$\begin{aligned} & x \mapsto (\text{sort } \gamma'(x))(x, 1) = \\ & \left(x \mapsto \gamma(x)(x), \nabla(x \mapsto \gamma(x)(x)) \right) \end{aligned}$$

That is, we need to show that $\pi_1(\gamma'(x)(x, 0)) = \gamma(x)(x)$ for all $x \in \mathbb{R}$, which holds by the assumption that $(\gamma, \gamma') \in (\mathbf{real} \rightarrow \mathbf{real})$ by choosing $\delta(x') = x$ (and hence $\delta'(x') = (x, 0)$) and then specializing to $x = x'$. \square

Yet, $\text{ev}'_1 \neq \text{ev}'_2$ as $\text{ev}'_1(a)(\text{swap}) = (1, a)$ and $\text{ev}'_2(a)(\text{swap}) = (0, 0)$. This shows that $\text{ev}'_1 \neq \text{ev}'_2$ are both “valid” semantic derivatives of the evaluation function (ev) in our framework. In particular, it shows that semantic derivatives of higher order functions might not be unique. Our macro \vec{D} will return ev'_1 , but everything would still work just as well if it instead returned ev'_2 .

7.2. Canonical derivatives of higher order functions? Differential geometers and analysts have long pursued notions of a canonical derivative of various higher order functions arising, for example, in the calculus of variations and in the study of infinite dimensional Lie groups [KM97]. Such an uncontroversial notion of derivative exists on various (infinite dimensional) spaces of functions that form suitable (so-called convenient) vector spaces, or, manifolds locally modelled on such vector spaces. At the level of generality of diffeological spaces, however, various natural notions of derivative that coincide in convenient vector spaces start to diverge and it is no longer clear what the best definition of a derivative is [CW14]. Another, fundamentally different setting that defines canonical derivatives of many higher order functions is given by synthetic differential geometry [Koc06].

While derivatives of higher order functions are of deep interest and have rightly been studied in their own right in differential geometry, we believe the situation is subtly different in computer science:

- (1) In programming applications, we use higher order programs only to construct the first order functions that we ultimately end up running and calculating derivatives of. Automatic differentiation methods can exploit this freedom: derivatives of higher order functions only matter in so far as they can be used to construct the correct derivatives of first order functions, so we can choose a simple and cheap notion of derivative among the valid options. As such, the fact that our semantics does not commit to a single notion of derivative of higher order functions can be seen as a *feature rather than bug* that models the pragmatics of programming practice.
- (2) While function spaces in differential geometry are typically infinite dimensional objects that are unsuitable for representation in the finite memory of a computer, higher order functions as used in programming are much more restricted: all they can do is call a function on finitely many arguments and analyse the function outputs. As such, function types in programming can be thought of as (locally) finite dimensional. In case a canonical notion of automatic derivative of higher order function is really desired, it may be worth pursuing a more intentional notion of semantics such as one based on game semantics. Such intentional techniques could capture the computational notion of higher order function better than our current (and other) extensional semantics using existing techniques from differential geometry. We hope that an exploration of such techniques might lead to an appropriate notion of computable derivative, even for higher order functions.

8. DISCUSSION AND FUTURE WORK

8.1. Summary. We have shown that diffeological spaces provide a denotational semantics for a higher order language with variants and inductive types (Section 4,5). We have used this to show correctness of simple forward-mode AD translations for calculating higher derivatives (Theorem 4.6, Theorem 6.7).

The structure of our elementary correctness argument for Theorem 4.6 is a typical logical relations proof over a denotational semantics. As explained in Section 6, this can equivalently be understood as a denotational semantics in a new kind of space obtained by categorical gluing.

Overall, then, there are two logical relations at play. One is in diffeological spaces, which ensures that all definable functions are smooth. The other is in the correctness proof (equivalently in the categorical gluing), which explicitly tracks the derivative of each function, and tracks the syntactic AD even at higher types.

8.2. Connection to the state of the art in AD implementation. As is common in denotational semantics research, we have here focused on an idealized language and simple translations to illustrate the main aspects of the method. There are a number of points where our approach is simplistic compared to the advanced current practice, as we now explain.

8.2.1. *Representation of vectors.* In our examples we have treated n -vectors as tuples of length n . This style of programming does not scale to large n . A better solution would be to use array types, following [SFVPJ19]. As demonstrated by [CJS20], our categorical semantics and correctness proofs straightforwardly extend to cover them, in a similar way to our treatment of lists. In fact, [CJS20] formalizes our correctness arguments in Coq and extends them to apply to the system of [SFVPJ19].

8.2.2. *Efficient forward-mode AD.* For AD to be useful, it must be fast. The $(1, 1)$ -AD macro $\vec{D}_{(1,1)}$ that we use is the basis of an efficient AD library [SFVPJ19]. Numerous optimizations are needed, ranging from algebraic manipulations, to partial evaluations, to the use of an optimizing C compiler, but the resulting implementation is performant in experiments [SFVPJ19]. The Coq formalization [CJS20] validates some of these manipulations using a similar semantics to ours. We believe the implementation in [SFVPJ19] can be extended to apply to the more general (k, R) -AD methods we described in this paper through minor changes.

8.2.3. *Reverse-mode and mixed-mode AD.* While forward-mode AD methods are useful, many applications require reverse-mode AD, or even mixed-mode AD for efficiency. In [HSV20a], we described how our correctness proof applies to a continuation-based AD technique that closely resembles reverse-mode AD, but only has the correct complexity under a non-standard operational semantics [BMP20] (in particular, the linear factoring rule is crucial). It remains to be seen whether this technique and its correctness proof can be adapted to yield genuine reverse AD under a standard operational semantics.

Alternatively, by relying on a variation of our techniques, [Vák21] gives a correctness proof of a rather different $(1, 1)$ -reverse AD algorithm that stores the (primal, adjoint)-vector pair as a struct-of-arrays rather than as an array-of-structs. Future work could explore extended its analysis to (k, R) -reverse AD and mixed-mode AD for efficiently computing higher order derivatives.

8.2.4. *Other language features.* The idealized languages that we considered so far do not touch on several useful language constructs. For example: the use of functions that are partial (such as division) or partly-smooth (such as ReLU); phenomena such as iteration, recursion; and probabilities. Recent work by MV [Vák20] shows how our analysis of $(1, 1)$ -AD extends to apply to partiality, iteration, and recursion. This development is orthogonal to the one in this paper: its methods combine directly with those in the present paper to analyze (k, R) -forward mode AD of recursive programs. We leave the analysis of AD of probabilistic programs for future work.

ACKNOWLEDGMENT

We have benefited from discussing this work with many people, including M. Betancourt, B. Carpenter, O. Kammar, C. Mak, L. Ong, B. Pearlmutter, G. Plotkin, A. Shaikhha, J. Sigal, and others. In the course of this work, MV has also been employed at Oxford (EPSRC Project EP/M023974/1) and at Columbia in the Stan development team. This project has also received funding from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 895827; a Royal Society

University Research Fellowship; the ERC BLAST grant; the Air Force Office of Scientific Research under award number FA9550-21-1-0038; and a Facebook Research Award.

REFERENCES

- [AAB⁺16] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. Tensorflow: A system for large-scale machine learning. In *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*, pages 265–283, 2016.
- [Ama12] Shun-ichi Amari. *Differential-geometrical methods in statistics*, volume 28. Springer Science & Business Media, 2012.
- [AP20] Martín Abadi and Gordon D Plotkin. A simple differentiable programming language. In *Proc. POPL 2020*. ACM, 2020.
- [BCLG20] Gilles Barthe, Raphaëlle Crubillé, Ugo Dal Lago, and Francesco Gavazzo. On the versatility of open logical relations: Continuity, automatic differentiation, and a containment theorem. In *Proc. ESOP 2020*. Springer, 2020. To appear.
- [Bet18] Michael Betancourt. A geometric theory of higher-order automatic differentiation. *arXiv preprint arXiv:1812.11592*, 2018.
- [BH11] John Baez and Alexander Hoffnung. Convenient categories of smooth spaces. *Transactions of the American Mathematical Society*, 363(11):5789–5825, 2011.
- [BJD19] Jesse Bettencourt, Matthew J Johnson, and David Duvenaud. Taylor-mode automatic differentiation for higher-order derivatives in JAX. 2019.
- [BML⁺20] Gilbert Bernstein, Michael Mara, Tzu-Mao Li, Dougal Maclaurin, and Jonathan Ragan-Kelley. Differentiating a tensor language. *arXiv preprint arXiv:2008.11256*, 2020.
- [BMP20] Alois Brunel, Damiano Mazza, and Michele Pagani. Backpropagation in the simply typed lambda-calculus with linear negation. In *Proc. POPL 2020*, 2020.
- [BS96] Claus Bendtsen and Ole Stauning. Fadbad, a flexible C++ package for automatic differentiation. Technical report, Technical Report IMM-REP-1996-17, Department of Mathematical Modelling, Technical University of Denmark, Lyngby, 1996.
- [BS97] Claus Bendtsen and Ole Stauning. Tadiiff, a flexible c++ package for automatic differentiation. *TU of Denmark, Department of Mathematical Modelling, Lyngby. Technical report IMM-REP-1997-07*, 1997.
- [CCG⁺20] J. Robin B. Cockett, Geoff S. H. Cruttwell, Jonathan Gallagher, Jean-Simon Pacaud Lemay, Benjamin MacAdam, Gordon D. Plotkin, and Dorette Pronk. Reverse derivative categories. In *Proc. CSL 2020*, 2020.
- [CGM19] Geoff Cruttwell, Jonathan Gallagher, and Ben MacAdam. Towards formalizing and extending differential programming using tangent categories. In *Proc. ACT 2019*, 2019.
- [CHB⁺15] Bob Carpenter, Matthew D Hoffman, Marcus Brubaker, Daniel Lee, Peter Li, and Michael Betancourt. The Stan math library: Reverse-mode automatic differentiation in C++. *arXiv preprint arXiv:1509.07164*, 2015.
- [CJS20] Curtis Chin Jen Sem. Formalized correctness proofs of automatic differentiation in Coq. *Master’s Thesis, Utrecht University*, 2020. Thesis: <https://dspace.library.uu.nl/handle/1874/400790>. Coq code: <https://github.com/crtschin/thesis>.
- [CRBD18] Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. In *Advances in Neural Information Processing Systems*, pages 6571–6583, 2018.
- [CS96] G Constantine and T Savits. A multivariate Faa di Bruno formula with applications. *Transactions of the American Mathematical Society*, 348(2):503–520, 1996.
- [CS11] J Robin B Cockett and Robert AG Seely. The Faa di Bruno construction. *Theory and Applications of Categories*, 25(15):394–425, 2011.

- [CW14] J Daniel Christensen and Enxin Wu. Tangent spaces and tangent bundles for diffeological spaces. *arXiv preprint arXiv:1411.5425*, 2014.
- [DHS11] John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12(Jul):2121–2159, 2011.
- [Ell18] Conal Elliott. The simple essence of automatic differentiation. *Proceedings of the ACM on Programming Languages*, 2(ICFP):70, 2018.
- [EM03] L Hernández Encinas and J Muñoz Masque. A short proof of the generalized Faà di Bruno’s formula. *Applied Mathematics Letters*, 16(6):975–979, 2003.
- [ER03] Thomas Ehrhard and Laurent Regnier. The differential lambda-calculus. *Theoretical Computer Science*, 309(1-3):1–41, 2003.
- [FJL18] Roy Frostig, Matthew James Johnson, and Chris Leary. Compiling machine learning programs via high-level tracing. *Systems for Machine Learning*, 2018.
- [FST19] Brendan Fong, David Spivak, and Rémy Tuyéras. Backprop as functor: A compositional perspective on supervised learning. In *2019 34th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 1–13. IEEE, 2019.
- [GSS00] Izrail Moiseevitch Gelfand, Richard A Silverman, and Richard A Silverman. *Calculus of variations*. Courier Corporation, 2000.
- [GUW00] Andreas Griewank, Jean Utke, and Andrea Walther. Evaluating higher derivative tensors by forward propagation of univariate Taylor series. *Mathematics of Computation*, 69(231):1117–1130, 2000.
- [HG14] Matthew D Hoffman and Andrew Gelman. The No-U-Turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research*, 15(1):1593–1623, 2014.
- [HSV20a] Mathieu Huot, Sam Staton, and Matthijs Vákár. Correctness of automatic differentiation via diffeologies and categorical gluing. In *FoSSaCS*, pages 319–338, 2020.
- [HSV20b] Mathieu Huot, Sam Staton, and Matthijs Vákár. Correctness of automatic differentiation via diffeologies and categorical gluing. Full version, 2020. arxiv:2001.02209.
- [IZ13] Patrick Iglesias-Zemmour. *Diffeology*. American Mathematical Soc., 2013.
- [JLS07] Peter T Johnstone, Stephen Lack, and P Sobocinski. Quasitoposes, quasiadhesive categories and Artin glueing. In *Proc. CALCO 2007*, 2007.
- [JR11] Bart Jacobs and JMMM Rutten. An introduction to (co)algebras and (co)induction. In *Advanced Topics in Bisimulation and Coinduction*, pages 38–99. CUP, 2011.
- [Kar01] Jerzy Karczmarczuk. Functional differentiation of computer programs. *Higher-Order and Symbolic Computation*, 14(1):35–57, 2001.
- [KB14] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [KK04] Dana A Knoll and David E Keyes. Jacobian-free Newton–Krylov methods: a survey of approaches and applications. *Journal of Computational Physics*, 193(2):357–397, 2004.
- [KM97] Andreas Kriegl and Peter W Michor. *The convenient setting of global analysis*, volume 53. American Mathematical Soc., 1997.
- [Koc06] Anders Kock. *Synthetic differential geometry*, volume 333. Cambridge University Press, 2006.
- [KSM99] Ivan Kolár, Jan Slovák, and Peter W Michor. Natural operations in differential geometry. 1999.
- [KTR⁺17] Alp Kucukelbir, Dustin Tran, Rajesh Ranganath, Andrew Gelman, and David M Blei. Automatic differentiation variational inference. *The Journal of Machine Learning Research*, 18(1):430–474, 2017.
- [KW⁺52] Jack Kiefer, Jacob Wolfowitz, et al. Stochastic estimation of the maximum of a regression function. *The Annals of Mathematical Statistics*, 23(3):462–466, 1952.
- [Lee13] John M Lee. Smooth manifolds. In *Introduction to Smooth Manifolds*, pages 1–31. Springer, 2013.
- [LMG18] Sören Laue, Matthias Mitterreiter, and Joachim Giesen. Computing higher order derivatives of matrix and tensor expressions. *Advances in Neural Information Processing Systems*, 31:2750–2759, 2018.
- [LMG20] Sören Laue, Matthias Mitterreiter, and Joachim Giesen. A simple and efficient tensor calculus. In *AAAI*, pages 4527–4534, 2020.

- [LN89] Dong C Liu and Jorge Nocedal. On the limited memory BFGS method for large scale optimization. *Mathematical programming*, 45(1-3):503–528, 1989.
- [LNV21] Fernando Lucatelli Nunes and Matthijs Vákár. CHAD for expressive total languages. *arXiv e-prints*, pages arXiv–2110, 2021.
- [LYRY20] Wonyeol Lee, Hangyeol Yu, Xavier Rival, and Hongseok Yang. On correctness of automatic differentiation for non-differentiable functions. In *Advances in Neural Information Processing Systems*, 2020.
- [Man12] Oleksandr Manzyuk. A simply typed λ -calculus of forward automatic differentiation. In *Proc. MFPS 2012*, 2012.
- [Mar10] James Martens. Deep learning via Hessian-free optimization. In *ICML*, volume 27, pages 735–742, 2010.
- [Mer04] Joel Merker. Four explicit formulas for the prolongations of an infinitesimal lie symmetry and multivariate Faa di Bruno formulas. *arXiv preprint math/0411650*, 2004.
- [MO20] Carol Mak and Luke Ong. A differential-form pullback programming language for higher-order reverse-mode automatic differentiation. arxiv:2002.08241, 2020.
- [MP21] Damiano Mazza and Michele Pagani. Automatic differentiation in PCF. *Proc. ACM Program. Lang.*, 5(POPL):1–27, 2021. doi:10.1145/3434309.
- [MS92] John C Mitchell and Andre Scedrov. Notes on scoring and relators. In *International Workshop on Computer Science Logic*, pages 352–378. Springer, 1992.
- [Nea11] Radford M Neal. MCMC using Hamiltonian dynamics. In *Handbook of Markov Chain Monte Carlo*, chapter 5. Chapman & Hall / CRC Press, 2011.
- [PGC⁺17] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017.
- [Pit95] Andrew M Pitts. Categorical logic. Technical report, University of Cambridge, Computer Laboratory, 1995.
- [Plo18] Gordon D Plotkin. Some principles of differential programming languages. Invited talk, POPL 2018, 2018.
- [PS07] Barak A Pearlmutter and Jeffrey Mark Siskind. Lazy multivariate higher-order forward-mode ad. *ACM SIGPLAN Notices*, 42(1):155–160, 2007.
- [PS08] Barak A Pearlmutter and Jeffrey Mark Siskind. Reverse-mode AD in a functional framework: Lambda the ultimate backpropagator. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 30(2):7, 2008.
- [Qia99] Ning Qian. On the momentum term in gradient descent learning algorithms. *Neural networks*, 12(1):145–151, 1999.
- [RM51] Herbert Robbins and Sutton Monro. A stochastic approximation method. *The Annals of Mathematical Statistics*, pages 400–407, 1951.
- [Sav06] Thomas H Savits. Some statistical applications of Faa di Bruno. *Journal of Multivariate Analysis*, 97(10):2131–2140, 2006.
- [SFVPJ19] Amir Shaikhha, Andrew Fitzgibbon, Dimitrios Vytiniotis, and Simon Peyton Jones. Efficient differentiable programming in a functional array-processing language. *Proceedings of the ACM on Programming Languages*, 3(ICFP):97, 2019.
- [SMC20] Benjamin Sherman, Jesse Michel, and Michael Carbin. λ_S : Computable semantics for differentiable programming with higher-order functions and datatypes. *arXiv preprint arXiv:2007.08017*, 2020.
- [Sou80] Jean-Marie Souriau. Groupes différentiels. In *Differential geometrical methods in mathematical physics*, pages 91–128. Springer, 1980.
- [Sta11] Andrew Stacey. Comparative smoothology. *Theory Appl. Categ.*, 25(4):64–117, 2011.
- [Vák20] Matthijs Vákár. Denotational correctness of forward-mode automatic differentiation for iteration and recursion. *arXiv preprint arXiv:2007.05282*, 2020.
- [Vák21] Matthijs Vákár. Reverse AD at higher types: Pure, principled and denotationally correct. In *ESOP*, pages 607–634, 2021.
- [VMBBL18] Bart Van Merriënboer, Olivier Breuleux, Arnaud Bergeron, and Pascal Lamblin. Automatic differentiation in ML: Where we are and where we should be going. In *Advances in Neural Information Processing Systems*, pages 8757–8767, 2018.

- [VS21] Matthijs Vákár and Tom Smeding. CHAD: Combinatory homomorphic automatic differentiation. *arXiv preprint arXiv:2103.15776*, 2021.
- [WGP16] Mu Wang, Assefaw Gebremedhin, and Alex Pothen. Capitalizing on live variables: new algorithms for efficient hessian computation via automatic differentiation. *Mathematical Programming Computation*, 8(4):393–433, 2016.
- [WWE⁺19] Fei Wang, Xilun Wu, Gregory Essertel, James Decker, and Tiark Rompf. Demystifying differentiable programming: Shift/reset the penultimate backpropagator. *Proceedings of the ACM on Programming Languages*, 3(ICFP), 2019.
- [ZHCW20] Shaopeng Zhu, Shih-Han Hung, Shouvanik Chakrabarti, and Xiaodi Wu. On the principles of differentiable quantum programming languages. In *Proceedings of the 41st ACM SIGPLAN International Conference on Programming Language Design and Implementation, PLDI 2020, London, UK, June 15-20, 2020*, pages 272–285. ACM, 2020. doi:10.1145/3385412.3386011.

APPENDIX A. **CartSp** AND **Man** ARE NOT CARTESIAN CLOSED CATEGORIES

Lemma A.1. *There is no continuous injection $\mathbb{R}^{d+1} \rightarrow \mathbb{R}^d$.*

Proof. If there were, it would restrict to a continuous injection $S^d \rightarrow \mathbb{R}^d$. The Borsuk-Ulam theorem, however, tells us that every continuous $f : S^d \rightarrow \mathbb{R}^d$ has some $x \in S^d$ such that $f(x) = f(-x)$, which is a contradiction. \square

Let us define the terms:

$$x_0 : \mathbf{real}, \dots, x_n : \mathbf{real} \vdash t_n = \lambda y. x_0 + x_1 * y + \dots + x_n * y^n : \mathbf{real} \rightarrow \mathbf{real}$$

Assuming that **CartSp/Man** is cartesian closed, observe that these get interpreted as injective continuous (because smooth) functions $\mathbb{R}^n \rightarrow \llbracket \mathbf{real} \rightarrow \mathbf{real} \rrbracket$ in **CartSp** and **Man**.

Theorem A.2. ***CartSp** is not cartesian closed.*

Proof. In case **CartSp** were cartesian closed, we would have $\llbracket \mathbf{real} \rightarrow \mathbf{real} \rrbracket = \mathbf{real}^n$ for some n . Then, we would get, in particular a continuous injection $\llbracket t_{n+1} \rrbracket : \mathbb{R}^{n+1} \rightarrow \mathbb{R}^n$, which contradicts Lemma A.1. \square

Theorem A.3. ***Man** is not cartesian closed.*

Proof. Observe that we have $\iota_n : \mathbb{R}^n \rightarrow \mathbb{R}^{n+1}; \langle a_0, \dots, a_n \rangle \mapsto \langle a_0, \dots, a_n, 0 \rangle$ and that $\iota_n; \llbracket t_{n+1} \rrbracket = \llbracket t_n \rrbracket$. Let us write A_n for the image of $\llbracket t_n \rrbracket$ and $A = \cup_{n \in \mathbb{N}} A_n$. Then, A_n is connected because it is the continuous image of a connected set. Similarly, A is connected because it is the non-disjoint union of connected sets. This means that A lies in a single connected component of $\llbracket \mathbf{real} \rightarrow \mathbf{real} \rrbracket$, which is a manifold with some finite dimension, say d .

Take some $x \in \mathbb{R}^{d+1}$ (say, 0), take some open d -ball U around $\llbracket t_{d+1} \rrbracket(x)$, and take some open $d+1$ -ball V around x in $\llbracket t_{d+1} \rrbracket^{-1}(U)$. Then, $\llbracket t_{d+1} \rrbracket$ restricts to a continuous injection from V to U , or equivalently, \mathbb{R}^{d+1} to \mathbb{R}^d , which contradicts Lemma A.1. \square