

CHAD: Combinatory Homomorphic Automatic Differentiation

MATTHIJS VÁKÁR and TOM SMEDING, Utrecht University, Netherlands

We introduce Combinatory Homomorphic Automatic Differentiation (CHAD), a principled, pure, provably correct define-then-run method for performing forward and reverse mode automatic differentiation (AD) on programming languages with expressive features. It implements AD as a compositional, type-respecting source-code transformation that generates purely functional code. This code transformation is principled in the sense that it is the unique homomorphic (structure preserving) extension to expressive languages of Elliott's well-known and unambiguous definitions of AD for a first-order functional language. Correctness of the method follows by a (compositional) logical relations argument that shows that the semantics of the syntactic derivative is the usual calculus derivative of the semantics of the original program.

In their most elegant formulation, the transformations generate code with linear types. However, the code transformations can be implemented in a standard functional language lacking linear types: While the correctness proof requires tracking of linearity, the actual transformations do not. In fact, even in a standard functional language, we can get all of the type-safety that linear types give us: We can implement all linear types used to type the transformations as abstract types by using a basic module system.

In this article, we detail the method when applied to a simple higher-order language for manipulating statically sized arrays. However, we explain how the methodology applies, more generally, to functional languages with other expressive features. Finally, we discuss how the scope of CHAD extends beyond applications in AD to other dynamic program analyses that accumulate data in a commutative monoid.

This article provides an extended version of Reference [42], augmenting it with

- examples and interpretation of AD at higher-order types (notably in Section 2);
- an explicit definition of the AD algorithm, directly phrased on a λ -calculus rather than categorical combinators (Section 7);
- a simplified semantics and correctness proof of the algorithm, based on sets rather than diffeological spaces, to make the article more accessible (Sections 5 and 8);
- proofs (Section 8);
- an extended discussion of how to implement the proposed algorithm (Section 9);
- a reference implementation of CHAD, in Haskell (Section 9), available at <https://github.com/VMatthijs/CHAD> (under continuous improvement by Tom Smeding, Matthijs Vákár, and others);
- the concept of CHAD as a more broadly applicable technique for AD (and even other program analyses) on expressive functional languages (Section 11);
- various major rewrites throughout.

This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 895827. We thank Michael Betancourt, Philip de Bruin, Bob Carpenter, Mathieu Huot, Danny de Jong, Ohad Kammar, Gabriele Keller, Pieter Knops, Fernando Lucatelli Nunes, Curtis Chin Jen Sem, Amir Shaikhha, and Sam Staton for helpful discussions about automatic differentiation.

Authors' address: M. Vákár and T. Smeding, Utrecht University, Buys Ballotgebouw Princetonplein 5 3584 CC Utrecht Netherlands; emails: {m.i.l.vakar, t.j.smeding}@uu.nl.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.

0164-0925/2022/08-ART20 \$15.00

<https://doi.org/10.1145/3527634>

CCS Concepts: • **Theory of computation** → **Categorical semantics**; • **Mathematics of computing** → **Automatic differentiation**; • **Software and its engineering** → **Functional languages**;

Additional Key Words and Phrases: Automatic differentiation, software correctness, denotational semantics, functional programming

ACM Reference format:

Matthijs Vákár and Tom Smeding. 2022. CHAD: Combinatory Homomorphic Automatic Differentiation. *ACM Trans. Program. Lang. Syst.* 44, 3, Article 20 (August 2022), 49 pages.
<https://doi.org/10.1145/3527634>

1 INTRODUCTION

Automatic differentiation (AD) is a technique for transforming code that implements a function f into code that computes f 's derivative, essentially by using the chain rule for derivatives. Due to its efficiency and numerical stability, AD is the technique of choice whenever we need to compute derivatives of functions that are implemented as programs, particularly in high-dimensional settings. Optimization and Monte Carlo integration algorithms, such as gradient descent and Hamiltonian Monte Carlo methods, rely crucially on the calculation of derivatives. These algorithms are used in virtually every machine learning and computational statistics application, and the calculation of derivatives is usually the computational bottleneck. These applications explain the recent surge of interest in AD, which has resulted in the proliferation of popular AD systems such as TensorFlow [1], PyTorch [36], and Stan Math [9].

AD, roughly speaking, comes in two modes: forward mode and reverse mode. When differentiating a function $\mathbb{R}^n \rightarrow \mathbb{R}^m$, forward mode tends to be more efficient if $m \gg n$, while reverse mode generally is more performant if $n \gg m$. As most applications reduce to optimization or Monte Carlo integration of an objective function $\mathbb{R}^n \rightarrow \mathbb{R}$ with n very large (at the time of this article, in the order of 10^4 – 10^7), reverse mode AD is in many ways the more interesting algorithm [5].

However, reverse AD is also more complicated to understand and implement than forward AD. Forward AD can be implemented as a structure-preserving program transformation, even on languages with complex features [38]. As such, it admits an elegant proof of correctness [21]. By contrast, reverse AD is only well understood as a compile-time source-code transformation that does not rely on using a runtime interpreter (also called *define-then-run* style AD) on limited programming languages, such as first-order functional languages. Typically, its implementations on more expressive languages that have features such as higher-order functions make use of interpreted *define-by-run* approaches. These approaches first build a computation graph during runtime, effectively evaluating the program until a straight-line first-order program is left, and then they perform automatic differentiation on this new program [9, 36]. Such approaches first have the severe downside that they can suffer from interpretation overhead. Second, the differentiated code cannot benefit as well from existing optimizing compiler architectures. As such, these AD libraries need to be implemented using carefully, manually optimized code, that, for example, does not contain any common subexpressions. This implementation process is precarious and labour intensive. Furthermore, some whole-program optimizations that a compiler would detect go entirely unused in such systems.

Similarly, correctness proofs of reverse AD have taken a define-by-run approach or have relied on non-standard operational semantics, using forms of symbolic execution [2, 8, 31]. Most work that treats reverse AD as a source-code transformation does so by making use of complex transformations that introduce mutable state and/or non-local control flow [37, 44]. As a result, we are not sure whether and why such techniques are correct. Furthermore, AD applications (e.g., in machine

learning) tend to be run on parallel hardware, which can be easier to target with purely functional code. Another approach has been to compile high-level languages to a low-level imperative representation first and then to perform AD at that level [22] using mutation and jumps. This approach has the downside that we might lose important opportunities for compiler optimizations, such as map-fusion and embarrassingly parallel maps, which we can exploit if we perform define-then-run AD on a high-level functional representation.

A notable exception to these define-by-run and non-functional approaches to AD is Elliott’s work [16], which presents an elegant, purely functional, define-then-run version of reverse AD. Unfortunately, their techniques are limited to first-order programs over tuples of real numbers. The workshop paper [43] by Vytiniotis, Belov, Wei, Plotkin, and Abadi proposes two possible extensions of Elliott’s functional AD to accommodate higher-order functions. However, it does not address whether or why these extensions would be correct or establish a more general methodology for applying AD to languages with expressive features.

This article introduces **Combinatory Homomorphic Automatic Differentiation (CHAD)** and its proof of correctness. CHAD is based on the observation that Elliott’s work [16] has a unique structure preserving extension that lets us perform AD on various expressive programming language features. We see purely functional higher-order (parallel) array processing languages such as Accelerate [10] and Futhark [19] as particularly relevant platforms for the machine learning applications that AD tends to be used for. With that in mind, we detail CHAD when applied to higher-order functional programs over (primitive) arrays of reals. This article includes the following:

- We introduce CHAD, a categorical perspective on AD, that lets us see AD as a uniquely determined homomorphic (structure-preserving) functor from the syntax of its source programming language (Section 3) to the syntax of its target language (Section 4).
- We explain, from this categorical setting, precisely in what sense reverse AD is the “mirror image” of forward AD (Section 6).
- We detail how this technique lets us define purely functional define-then-run reverse mode AD on a higher-order language (Section 7).
- We present an elegant proof of semantic correctness of the resulting AD transformations, based on a semantic logical relations argument, demonstrating that the transformations calculate the derivatives of the program in the usual mathematical sense (Sections 5 and 8).
- We show that the AD definitions and correctness proof are extensible to higher-order primitives such as a `map`-operation over our primitive arrays (Section 10).
- We show how our techniques are readily implementable in standard functional languages to give purely functional, principled, semantically correct, compositional, define-then-run reverse mode AD (Section 9).
- Finally, we place CHAD in a broader context and explain how it applies, more generally, to dynamic program analyses that accumulate information in a commutative monoid (Section 11).

We start by giving a high-level overview of the main insights and theorems in this article in Section 2.

For a review of the basics of AD, we refer the reader to References [5, 32]. We discuss recent related work studying AD from a programming languages perspective in Section 12.

2 KEY IDEAS

We start by providing a high-level overview of the article, highlighting the main insights and theorems underlying our contributions.

2.1 Aims of Automatic Differentiation

The basic challenge that automatic differentiation aims to solve is the following. We are given a program $x : \mathbf{real}^n \vdash t : \mathbf{real}^m$ that takes an n -dimensional array of (floating point) real numbers as input and produces an m -dimensional array of reals as output. That is, t computes some mathematical function $\llbracket t \rrbracket : \mathbb{R}^n \rightarrow \mathbb{R}^m$. We want to transform the code of t to¹

- a program $\vec{\mathcal{D}}(t)_2$ that computes the derivative $D\llbracket t \rrbracket : \mathbb{R}^n \rightarrow \underline{\mathbb{R}}^n \multimap \underline{\mathbb{R}}^m$, in the case of forward AD;
- a program $\overleftarrow{\mathcal{D}}(t)_2$ that computes the transposed derivative $D\llbracket t \rrbracket^t : \mathbb{R}^n \rightarrow \underline{\mathbb{R}}^m \multimap \underline{\mathbb{R}}^n$, in the case of reverse AD.

Here, we write $\underline{\mathbb{R}}^n$ for the space of (co)tangent vectors to \mathbb{R}^n ; we regard $\underline{\mathbb{R}}^n$ as a commutative monoid under elementwise addition. We write \multimap for a linear function type to emphasize that derivatives are linear in the sense of being monoid homomorphisms.

Furthermore, we have some more desiderata for these code transformations:

- (1) we want these code transformations to be defined compositionally, so we can easily extend the source programming language we apply the transformations to with new primitives;
- (2) we want these transformations to apply to a wide range of programming techniques, so we are not limited in our programming style even if we want our code to be differentiated;
- (3) we want the transformations to generate purely functional code so we can easily prove its correctness and deploy it on parallel hardware;
- (4) we want the code size of $\vec{\mathcal{D}}(t)_2$ and $\overleftarrow{\mathcal{D}}(t)_2$ to grow linearly in the size of t , so we can apply the technique to large code-bases;
- (5) we want the time complexity of $\vec{\mathcal{D}}(t)_2$ and $\overleftarrow{\mathcal{D}}(t)_2$ to be proportional to that of t and, generally, as low as possible; this means that we can use forward AD to efficiently compute a column of the Jacobian matrix of partial derivatives, while reverse AD efficiently computes a row of the Jacobian.

In this article, we demonstrate how the CHAD technique of automatic differentiation satisfies desiderata (1)–(4); we leave (5) to future work. It achieves this by taking seriously the mathematical structure of programming languages as freely generated categories and by observing that differentiation is compositional according to the chain rule.

2.2 The Chain Rule: Pairing and Sharing of Primals and Derivatives

To achieve desideratum (1) of compositionality, it is tempting to examine the chain rule, the key property of compositionality of derivatives. Given $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$, we write

$$\begin{aligned} \mathcal{T}f : \mathbb{R}^n &\rightarrow \mathbb{R}^m \times (\underline{\mathbb{R}}^n \multimap \underline{\mathbb{R}}^m) \\ x &\mapsto (f(x), \quad v \mapsto Df(x)(v)) \end{aligned}$$

for the function that pairs up the primal function value $f(x)$ with the derivative $Df(x)$ of f at x that acts on tangent vectors v . The chain rule then gives the following formula for the derivative of the composition $f;g$ of f and g :

$$\mathcal{T}(f;g)(x) = (\mathcal{T}_1g(\mathcal{T}_1f(x)), \mathcal{T}_2f(x); \mathcal{T}_2g(\mathcal{T}_1f(x))),$$

¹The program transformations are really called $\vec{\mathcal{D}}(-)_2$ and $\overleftarrow{\mathcal{D}}(-)_2$ here. In Section 2.2, we discuss that it is better to define our actual program transformations to have a slightly different type. The second half of those transformations (defined in Section 2.3) corresponds to these $\vec{\mathcal{D}}(-)_2$ and $\overleftarrow{\mathcal{D}}(-)_2$.

where we write $\mathcal{T}_1 f \stackrel{\text{def}}{=} \mathcal{T}f; \pi_1$ and $\mathcal{T}_2 f \stackrel{\text{def}}{=} \mathcal{T}f; \pi_2$ for the first and second components of $\mathcal{T}f$, respectively. We make two observations:

- (1) the derivative of the composition $f;g$ does not only depend on the derivatives of g and f but also on the primal value of f ;
- (2) the primal value of f is used twice: once in the primal value of $f;g$ and once in its derivative; we want to share these repeated subcomputations, to address desiderata (4) and (5).

INSIGHT 1. *It is wise to pair up computations of primal function values and derivatives and to share computation between both if we want to calculate derivatives of functions compositionally and efficiently.*

Similarly, we can pair up f 's transposed (adjoint) derivative Df^t , which propagates not tangent vectors but cotangent vectors:

$$\begin{aligned} \mathcal{T}^* f : \mathbb{R}^n &\rightarrow \mathbb{R}^m \times (\mathbb{R}^m \multimap \mathbb{R}^n) \\ x &\mapsto (f(x), \quad v \mapsto Df^t(x)(v)). \end{aligned}$$

It then satisfies the following chain rule, which follows from the usual chain rule above together with the fact that $(A;B)^t = B^t;A^t$ for linear maps A and B (transposition is contravariant—note the resulting reversed order of $\mathcal{T}_2^* f$ and $\mathcal{T}_2^* g$ for reverse AD):

$$\mathcal{T}^*(f;g)(x) = (\mathcal{T}_1^* g(\mathcal{T}_1^* f(x)), \mathcal{T}_2^* g(\mathcal{T}_1^* f(x)); \mathcal{T}_2^* f(x)).$$

Again, pairing and sharing the primal and (transposed) derivative computations is beneficial.

CHAD directly implements the operations \mathcal{T} and \mathcal{T}^* as source code transformations $\overrightarrow{\mathcal{T}}$ and $\overleftarrow{\mathcal{T}}$ on a functional language to implement forward² and reverse mode AD, respectively. These code transformations are defined compositionally through structural induction on the syntax by exploiting the chain rules above combined with the categorical structure of programming languages.

2.3 CHAD on a First-order Functional Language

Here, we outline how CHAD looks when applied to programs written in a first-order functional language. We consider this material known, as it is essentially the algorithm of Reference [16]. However, we present it in terms of a λ -calculus rather than categorical combinators, by applying the well-known mechanical translations between both formalisms [13]. We hope that this presentation makes the algorithm easier to apply in practice.

We consider a source programming language (see Section 3) where we write τ, σ, ρ for types that are either statically sized arrays of n real numbers \mathbf{real}^n or tuples $\tau * \sigma$ of types τ, σ . These types will be called *first-order types* in this section.³ We consider programs t of type σ in a typing context $\Gamma = x_1 : \tau_1, \dots, x_n : \tau_n$, where x_i are identifiers. We write such typing of programs in a context as $\Gamma \vdash t : \sigma$. As long as our language has certain primitive operations (which we represent

²For forward AD, we can also choose to implement instead

$$\begin{aligned} \mathcal{T}' f : (\mathbb{R}^n \times \mathbb{R}^n) &\rightarrow (\mathbb{R}^m \times \mathbb{R}^m) \\ (x, v) &\mapsto (f(x), Df(x)(v)), \end{aligned}$$

together with its chain rule as code transformations. This leads to a different style of forward AD based on a *dual numbers representation*. Reference [21] gives an analysis of this style of forward AD, similar to the treatment of reverse AD and (non-dual number) forward AD in this article. Although dual numbers forward AD is more efficient in its memory use and preferable in practical implementations, it does not have an obvious reverse mode variant. See Section 11.1 for more discussion.

³In the rest of the article, we also consider the unit type 1 a first-order type. These types are also called *ground types*.

schematically),

$$\frac{\Gamma \vdash t_1 : \mathbf{real}^{n_1} \quad \dots \quad \Gamma \vdash t_k : \mathbf{real}^{n_k}}{\Gamma \vdash \text{op}(t_1, \dots, t_k) : \mathbf{real}^m},$$

such as constants (as nullary operations), (elementwise) addition and multiplication of arrays, inner products and certain non-linear functions like sigmoid functions, we can write complex programs by sequencing together such operations. Figure 1(a) and (b) give some examples of some programs we can write, where we write **real** for \mathbf{real}^1 and indicate shared subcomputations with let-bindings.

CHAD transforms the types and programs of this source language into types and programs of a suitably chosen target language (see Section 4) that is a superset of the source language. CHAD associates to each source language type τ types of

- forward mode primal values $\vec{\mathcal{D}}(\tau)_1$;
we define $\vec{\mathcal{D}}(\mathbf{real}^n) \stackrel{\text{def}}{=} \mathbf{real}^n$ and $\vec{\mathcal{D}}(\tau * \sigma)_1 \stackrel{\text{def}}{=} \vec{\mathcal{D}}(\tau)_1 * \vec{\mathcal{D}}(\sigma)_1$; that is, for now $\vec{\mathcal{D}}(\tau)_1 = \tau$;
- reverse mode primal values $\overleftarrow{\mathcal{D}}(\tau)_1$;
we define $\overleftarrow{\mathcal{D}}(\mathbf{real}^n) \stackrel{\text{def}}{=} \mathbf{real}^n$ and $\overleftarrow{\mathcal{D}}(\tau * \sigma)_1 \stackrel{\text{def}}{=} \overleftarrow{\mathcal{D}}(\tau)_1 * \overleftarrow{\mathcal{D}}(\sigma)_1$; that is, for now $\overleftarrow{\mathcal{D}}(\tau)_1 = \tau$;
- forward mode tangent values $\vec{\mathcal{D}}(\tau)_2$;
we define $\vec{\mathcal{D}}(\mathbf{real}^n)_2 \stackrel{\text{def}}{=} \mathbf{real}^n$ and $\vec{\mathcal{D}}(\tau * \sigma)_2 \stackrel{\text{def}}{=} \vec{\mathcal{D}}(\tau)_2 * \vec{\mathcal{D}}(\sigma)_2$;
- reverse mode cotangent values $\overleftarrow{\mathcal{D}}(\tau)_2$;
we define $\overleftarrow{\mathcal{D}}(\mathbf{real}^n)_2 \stackrel{\text{def}}{=} \mathbf{real}^n$ and $\overleftarrow{\mathcal{D}}(\tau * \sigma)_2 \stackrel{\text{def}}{=} \overleftarrow{\mathcal{D}}(\tau)_2 * \overleftarrow{\mathcal{D}}(\sigma)_2$.

The types $\vec{\mathcal{D}}(\tau)_1$ and $\overleftarrow{\mathcal{D}}(\tau)_1$ of primals are Cartesian types, which we can think of as denoting sets, while the types $\vec{\mathcal{D}}(\tau)_2$ and $\overleftarrow{\mathcal{D}}(\tau)_2$ are linear types that denote commutative monoids. That is, such linear types in our language need to have a commutative monoid structure $(\underline{0}, +)$. For example, \mathbf{real}^n is the commutative monoid over \mathbf{real}^n with $\underline{0}$ is the zero vector and $(+)$ as elementwise addition of vectors. Derivatives and transposed derivatives are then linear functions, that is, homomorphisms of this $(\underline{0}, +)$ -monoid structure. As we will see, we use the monoid structure to initialize and accumulate (co)tangents in the definition of CHAD.

We extend these operations $\vec{\mathcal{D}}$ and $\overleftarrow{\mathcal{D}}$ to act not only on types but also on typing contexts $\Gamma = x_1 : \tau_1, \dots, x_n : \tau_n$ to produce primal contexts and (co)tangent types:

$$\begin{aligned} \vec{\mathcal{D}}(x_1 : \tau_1, \dots, x_n : \tau_n)_1 &= x_1 : \vec{\mathcal{D}}(\tau_1)_1, \dots, x_n : \vec{\mathcal{D}}(\tau_n)_1 & \vec{\mathcal{D}}(x_1 : \tau_1, \dots, x_n : \tau_n)_2 &= \vec{\mathcal{D}}(\tau_1)_2 * \dots * \vec{\mathcal{D}}(\tau_n)_2 \\ \overleftarrow{\mathcal{D}}(x_1 : \tau_1, \dots, x_n : \tau_n)_1 &= x_1 : \overleftarrow{\mathcal{D}}(\tau_1)_1, \dots, x_n : \overleftarrow{\mathcal{D}}(\tau_n)_1 & \overleftarrow{\mathcal{D}}(x_1 : \tau_1, \dots, x_n : \tau_n)_2 &= \overleftarrow{\mathcal{D}}(\tau_1)_2 * \dots * \overleftarrow{\mathcal{D}}(\tau_n)_2. \end{aligned}$$

To each program $\Gamma \vdash t : \sigma$, CHAD then associates programs calculating the forward mode and reverse mode derivatives $\vec{\mathcal{D}}_{\vec{\Gamma}}(t)$ and $\overleftarrow{\mathcal{D}}_{\vec{\Gamma}}(t)$, whose definitions make use of the list $\vec{\Gamma}$ of identifiers that occur in Γ :

$$\begin{aligned} \vec{\mathcal{D}}(\Gamma)_1 \vdash \vec{\mathcal{D}}_{\vec{\Gamma}}(t) : \vec{\mathcal{D}}(\sigma)_1 * (\vec{\mathcal{D}}(\Gamma)_2 \multimap \vec{\mathcal{D}}(\sigma)_2) \\ \overleftarrow{\mathcal{D}}(\Gamma)_1 \vdash \overleftarrow{\mathcal{D}}_{\vec{\Gamma}}(t) : \overleftarrow{\mathcal{D}}(\sigma)_1 * (\overleftarrow{\mathcal{D}}(\sigma)_2 \multimap \overleftarrow{\mathcal{D}}(\Gamma)_2). \end{aligned}$$

Observing that each program t computes a differentiable (infinitely differentiable) function $\llbracket t \rrbracket$ between Euclidean spaces, as long as all primitive operations op are differentiable, the key property that we prove for these code transformations is that they actually calculate derivatives:

THEOREM A (CORRECTNESS OF CHAD, THEOREM 8.3). *For any well-typed program (where τ_i and σ are first-order types, i.e., \mathbf{real}^n and tuples of such types)*

$$x_1 : \tau_1, \dots, x_n : \tau_n \vdash t : \sigma$$

we have that $\llbracket \vec{\mathcal{D}}(t) \rrbracket = \mathcal{T} \llbracket t \rrbracket$ and $\llbracket \overleftarrow{\mathcal{D}}(t) \rrbracket = \mathcal{T}^ \llbracket t \rrbracket$.*

Once we fix a semantics for the source and target languages, we can show that this theorem holds if we define $\overrightarrow{\mathcal{D}}$ and $\overleftarrow{\mathcal{D}}$ on programs using the chain rule. The proof works by plain induction on the syntax.

For example, we can correctly define reverse mode CHAD on a first-order language as follows (see Section 7):

$$\begin{aligned}
\overleftarrow{\mathcal{D}}_{\Gamma}(\text{op}(t_1, \dots, t_k)) &\stackrel{\text{def}}{=} \text{let } \langle x_1, x'_1 \rangle = \overleftarrow{\mathcal{D}}_{\Gamma}(t_1) \text{ in } \dots \text{let } \langle x_k, x'_k \rangle = \overleftarrow{\mathcal{D}}_{\Gamma}(t_k) \text{ in} \\
&\quad \langle \text{op}(x_1, \dots, x_k), \underline{\lambda}v. \text{let } v = \text{Dop}^t(x_1, \dots, x_k; v) \text{ in } x'_1 \bullet (\text{proj}_1 v) + \dots + x'_k \bullet (\text{proj}_k v) \rangle \\
\overleftarrow{\mathcal{D}}_{\Gamma}(x) &\stackrel{\text{def}}{=} \langle x, \underline{\lambda}v. \text{coproj}_{\text{idx}(x; \overline{\Gamma})}(v) \rangle \\
\overleftarrow{\mathcal{D}}_{\Gamma}(\text{let } x = t \text{ in } s) &\stackrel{\text{def}}{=} \text{let } \langle x, x' \rangle = \overleftarrow{\mathcal{D}}_{\Gamma}(t) \text{ in let } \langle y, y' \rangle = \overleftarrow{\mathcal{D}}_{\Gamma, x}(s) \text{ in } \langle y, \underline{\lambda}v. \text{let } v = y' \bullet v \text{ in fst } v + x' \bullet (\text{snd } v) \rangle \\
\overleftarrow{\mathcal{D}}_{\Gamma}(\langle t, s \rangle) &\stackrel{\text{def}}{=} \text{let } \langle x, x' \rangle = \overleftarrow{\mathcal{D}}_{\Gamma}(t) \text{ in let } \langle y, y' \rangle = \overleftarrow{\mathcal{D}}_{\Gamma}(s) \text{ in } \langle \langle x, y \rangle, \underline{\lambda}v. x' \bullet (\text{fst } v) + y' \bullet (\text{snd } v) \rangle \\
\overleftarrow{\mathcal{D}}_{\Gamma}(\text{fst } t) &\stackrel{\text{def}}{=} \text{let } \langle x, x' \rangle = \overleftarrow{\mathcal{D}}_{\Gamma}(t) \text{ in } \langle \text{fst } x, \underline{\lambda}v. x' \bullet \langle v, \underline{0} \rangle \rangle \\
\overleftarrow{\mathcal{D}}_{\Gamma}(\text{snd } t) &\stackrel{\text{def}}{=} \text{let } \langle x, x' \rangle = \overleftarrow{\mathcal{D}}_{\Gamma}(t) \text{ in } \langle \text{snd } x, \underline{\lambda}v. x' \bullet \langle \underline{0}, v \rangle \rangle.
\end{aligned}$$

Here, we write $\underline{\lambda}v. t$ for a linear function abstraction (merely a notational convention – it can simply be thought of as a plain function abstraction) and $t \bullet s$ for a linear function application of $t : \tau \multimap \sigma$ to the argument $s : \tau$ (which again can be thought of as a plain function application). Furthermore, given a program t of tuple type $\underline{\sigma}_1 * \dots * \underline{\sigma}_n$, we write $\text{proj}_i t$ for its i th projection of type $\underline{\sigma}_i$. Similarly, given a program t of linear type $\underline{\sigma}_i$, we write $\text{coproj}_i(t)$ for the i th coprojection $\langle \underline{0}, \dots, \underline{0}, t, \underline{0}, \dots, \underline{0} \rangle$ of type $\underline{\sigma}_1 * \dots * \underline{\sigma}_n$ and we write $\text{idx}(x_i; x_1, \dots, x_n) = i$ for the index of an identifier in a list of identifiers. Finally, Dop^t here is a linear operation that implements the transposed derivative of the primitive operation op . We note that we crucially need the commutative monoid structure on linear types to correctly define the reverse mode derivatives of programs that involve tuples (or n -ary operations for $n \neq 1$). Intuitively, matrix transposition (of derivatives) flips the copying-deleting comonoid structure provided by tuples into the addition-zero monoid structure.

INSIGHT 2. *In functional define-then-run reverse AD, we need to have a commutative monoid structure on types of cotangents to mirror the comonoid structure coming from tuples: copying fan-out in the original program gets translated into fan-in in the transposed derivative, for accumulating incoming cotangents. This leads to linear types of cotangents.*

Furthermore, observe that CHAD pairs up primal and (co)tangent values and shares common subcomputations, as desired. We see that what CHAD achieves is a compositional efficient reverse mode AD algorithm that computes the (transposed) derivatives of a composite program in terms of the (transposed) derivatives Dop^t of the basic building blocks op . Finally, it does so in a way that satisfies desiderata (1)–(5).

For example, Figure 1(c) and (d) display the code that forward and reverse mode CHAD, respectively, generate for the source programs of Figure 1(a) and (b). This is the code that is actually generated by the CHAD code transformations in our Haskell implementation followed by some very basic simplifications that do not affect time complexity and whose only purpose here is to aid legibility. For more information about how exactly this code relates to the output one gets when applying the forward and reverse AD macros in this article to the source programs, see Appendix B.

2.4 Intermezzo: The Categorical Structure of CHAD

While this definition of CHAD on a first-order language straightforwardly follows from the mathematics of derivatives, it is not immediately clear how to should be extended to source

```

let y = 2 * x
    z = x * y
    w = cos z
    v = ⟨y, z, w⟩ in
    v

```

(a) Original program $x : \mathbf{real} \vdash t : \mathbf{real} * \mathbf{real} * \mathbf{real}$ computing a function $\llbracket t \rrbracket : \mathbb{R} \rightarrow \mathbb{R}^3; x \mapsto v$.

```

let y = 2 * x
    z = x * y
    w = cos z
    v = ⟨y, z, w⟩ in
    ⟨v, λx'.
      let y' = 2 * x'
          z' = x' * y + x * y'
          w' = -sin z * z'
          v' = ⟨y', z', w'⟩ in
          v'⟩

```

} primals

} tangents

(c) Forward AD transformed program $x : \mathbf{real} \vdash \overline{\mathcal{D}}(t) : (\mathbf{real} * \mathbf{real}) * (\mathbf{real} \multimap \mathbf{real} * \mathbf{real} * \mathbf{real})$ computing the derivative $\llbracket \overline{\mathcal{D}}(t) \rrbracket = \mathcal{T}\llbracket t \rrbracket : \mathbb{R} \rightarrow \mathbb{R}^3 \times (\mathbb{R} \multimap \mathbb{R}^3); x \mapsto (v, x' \mapsto v')$.

```

let y = x1 * x4 + 2 * x2
    z = y * x3
    w = z + x4
    v = sin w in
    v

```

(b) Original program $x1 : \mathbf{real}, x2 : \mathbf{real}, x3 : \mathbf{real}, x4 : \mathbf{real} \vdash s : \mathbf{real}$ computing a function $\llbracket s \rrbracket : \mathbb{R}^4 \rightarrow \mathbb{R}; (x1, x2, x3, x4) \mapsto v$.

```

let y = x1 * x4
    + 2 * x2
    z = y * x3
    w = z + x4
    v = sin w in
    ⟨v, λv'.
      let w' = cos w * v'
          z' = w'
          y' = x3 * z'
          x1' = y' * x4
          x2' = 2 * y'
          x3' = y * z'
          x4' = x1 * y'
          + w' in
          ⟨x1', x2', x3', x4'⟩

```

} primals

} cotangents

(d) Reverse AD transformed program $x1 : \mathbf{real}, x2 : \mathbf{real}, x3 : \mathbf{real}, x4 : \mathbf{real} \vdash \overline{\mathcal{D}}(s) : \mathbf{real} * (\mathbf{real} \multimap \mathbf{real} * \mathbf{real} * \mathbf{real} * \mathbf{real})$ computing the transposed derivative $\llbracket \overline{\mathcal{D}}(s) \rrbracket = \mathcal{T}^* \llbracket s \rrbracket : \mathbb{R}^4 \rightarrow \mathbb{R} \times (\mathbb{R} \multimap \mathbb{R}^4); (x1, x2, x3, x4) \mapsto (v, v' \mapsto (x1', x2', x3', x4'))$. We see that the copying of $x4$ is translated into the addition of cotangents in $x4'$.

Fig. 1. Forward and reverse AD illustrated on simple first-order functional programs.

languages with more expressive features such as higher-order functions. Indeed, we do not typically consider derivatives of higher-order functions in calculus. In fact, it is not even clear what a tangent or cotangent to a function type should be, or, for that matter, what a primal associated with a value of function type is. To solve this mystery, we employ some category theory.

Observe that the first-order source language we consider can be viewed as a category \mathbf{Syn} with products (see Section 3): Its objects are types τ, σ, ρ and morphisms $t \in \mathbf{Syn}(\tau, \sigma)$ are programs $x : \tau \vdash t : \sigma$ modulo standard $\beta\eta$ -program equivalence (identities are given by variables and composition is done through **let**-bindings). This category is freely generated by the objects \mathbf{real}^n and morphisms op in the sense that any consistent assignment of objects $F(\mathbf{real}^n)$ and morphisms $F(\text{op})$ in a category with products \mathcal{C} extends to a unique product preserving functor $F : \mathbf{Syn} \rightarrow \mathcal{C}$.

Suppose that we are given a categorical model $\mathcal{L} : \mathcal{C}^{op} \rightarrow \mathbf{Cat}$ of linear logic (a so-called locally indexed category—see, for example, Section 9.3.4 in Reference [29]), where we think of the objects and morphisms of \mathcal{C} as the semantics of Cartesian types and their programs and of the objects and morphisms of \mathcal{L} as the semantics of linear types and their programs. We observe that we can define categories $\Sigma_{\mathcal{C}}\mathcal{L}$ and $\Sigma_{\mathcal{C}}\mathcal{L}^{op}$ (their so-called Grothendieck constructions, or Σ -types, see Section 6) with objects that are pairs (A_1, A_2) with A_1 an object of \mathcal{C} and A_2 an object of \mathcal{L} and homsets

$$\Sigma_{\mathcal{C}}\mathcal{L}((A_1, A_2), (B_1, B_2)) \stackrel{\text{def}}{=} \mathcal{C}(A_1, B_1) \times \mathcal{L}(A_1)(A_2, B_2) \cong \mathcal{C}(A_1, B_1 \times (A_2 \multimap B_2))$$

$$\Sigma_C \mathcal{L}^{op}((A_1, A_2), (B_1, B_2)) \stackrel{\text{def}}{=} C(A_1, B_1) \times \mathcal{L}(A_1)(B_2, A_2) \cong C(A_1, B_1 \times (B_2 \multimap A_2)).$$

We prove that these categories have finite products, provided that some conditions are satisfied: namely that C has finite products and \mathcal{L} has indexed finite biproducts (or, equivalently, has indexed finite products and is enriched over commutative monoids). Indeed, then $\prod_{i \in I} (A_{1i}, A_{2i}) = (\prod_{i \in I} A_{1i}, \prod_{i \in I} A_{2i})$. In other words, it is sufficient if our model of linear logic is biadditive. In particular, the categorical model of linear logic that we can build from the syntax of our target language for CHAD, $\mathbf{LSyn} : \mathbf{CSyn}^{op} \rightarrow \mathbf{Cat}$, satisfies our conditions (in fact, it is the initial model that does so), so $\Sigma_{\mathbf{CSyn}} \mathbf{LSyn}$ and $\Sigma_{\mathbf{CSyn}} \mathbf{LSyn}^{op}$ have finite products. By the universal property of the source language \mathbf{Syn} , we obtain a canonical definition of CHAD.

THEOREM B (CHAD FROM A UNIVERSAL PROPERTY, COROLLARY 7.1). *Forward and reverse mode CHAD are the unique structure preserving functors*

$$\vec{\mathcal{D}}(-) : \mathbf{Syn} \rightarrow \Sigma_{\mathbf{CSyn}} \mathbf{LSyn} \qquad \overleftarrow{\mathcal{D}}(-) : \mathbf{Syn} \rightarrow \Sigma_{\mathbf{CSyn}} \mathbf{LSyn}^{op}$$

from the syntactic category \mathbf{Syn} of the source language to (opposite) Grothendieck construction of the target language $\mathbf{LSyn} : \mathbf{CSyn}^{op} \rightarrow \mathbf{Cat}$ that send primitive operations op to their derivative Dop and transposed derivative Dop^t , respectively.

The definitions following from this universal property reproduce the definitions of CHAD that we have given so far. Intuitively, the linear types represent commutative monoids, implementing the idea that (transposed) derivatives are linear functions in the sense that $Df(x)(0) = 0$ and $Df(x)(v + v') = Df(x)(v) + Df(x)(v')$. We have seen that this commutative monoid structure is important when writing down the definitions of AD as a source-code transformation.

Seeing that a higher-order language can be viewed as a freely generated *Cartesian closed* category \mathbf{Syn} , it is tempting to find a suitable target language such that $\Sigma_{\mathbf{CSyn}} \mathbf{LSyn}$ and $\Sigma_{\mathbf{CSyn}} \mathbf{LSyn}^{op}$ are Cartesian closed. Then, we can *define* CHAD on this higher-order language via Theorem B.

INSIGHT 3. *To understand how to perform CHAD on a source language with language feature X (e.g., higher-order functions), we need to understand the categorical semantics of language feature X (e.g., categorical exponentials) in categories of the form $\Sigma_C \mathcal{L}$ and $\Sigma_C \mathcal{L}^{op}$. Giving sufficient conditions on a model of linear logic $\mathcal{L} : C^{op} \rightarrow \mathbf{Cat}$ for such a semantics to exist yields a suitable target language for CHAD as the initial such model $\mathbf{LSyn} : \mathbf{CSyn}^{op} \rightarrow \mathbf{Cat}$, with the definition of the algorithm falling from the universal property of the source language.*

2.5 Cartesian Closure of $\Sigma_C \mathcal{L}$ and $\Sigma_C \mathcal{L}^{op}$ and CHAD of Higher-order Functions

Having had this insight, we identify conditions on a locally indexed category $\mathcal{L} : C^{op} \rightarrow \mathbf{Cat}$ that are enough to guarantee that $\Sigma_C \mathcal{L}$ and $\Sigma_C \mathcal{L}^{op}$ are Cartesian closed (see Section 6).

THEOREM C (CARTESIAN CLOSURE OF $\Sigma_C \mathcal{L}$ AND $\Sigma_C \mathcal{L}^{op}$, THEOREMS 6.1–6.2). *Suppose that a locally indexed category $\mathcal{L} : C^{op} \rightarrow \mathbf{Cat}$ supports (we are intentionally a bit vague here for the sake of legibility)*

- *linear $!(-) \otimes (-)$ -types (copowers);*
- *linear $(-) \Rightarrow (-)$ -types (powers);*
- *Cartesian $(-) \multimap (-)$ -types (types of linear functions);*
- *linear biproduct types (or equivalently, linear (additive) product types and enrichment of \mathcal{L} over commutative monoids);*
- *Cartesian tuple and function types.*

Then, $\Sigma_C \mathcal{L}$ and $\Sigma_C \mathcal{L}^{op}$ are Cartesian closed with, respective, exponentials:

$$(A_1, A_2) \Rightarrow \Sigma_C \mathcal{L}(B_1, B_2) = (A_1 \Rightarrow B_1 \times (A_2 \multimap B_2), A_1 \Rightarrow B_2)$$

$$(A_1, A_2) \Rightarrow \Sigma_C \mathcal{L}^{op}(B_1, B_2) = (A_1 \Rightarrow B_1 \times (B_2 \multimap A_2), !A_1 \otimes B_2).$$

In particular, if we extend our target language with (linear) powers, (linear) copowers and (Cartesian) function types, then $\mathbf{LSyn} : \mathbf{CSyn}^{op} \rightarrow \mathbf{Cat}$ satisfies the conditions of Theorem C, so we can extend Theorem B to our higher-order source language. In particular, we find the following definitions of CHAD for primals and (co)tangents to function types:

$$\begin{aligned} \vec{\mathcal{D}}(\tau \rightarrow \sigma)_1 &\stackrel{\text{def}}{=} \vec{\mathcal{D}}(\tau)_1 \rightarrow \vec{\mathcal{D}}(\sigma) * (\vec{\mathcal{D}}(\tau)_2 \multimap \vec{\mathcal{D}}(\sigma)_2) & \vec{\mathcal{D}}(\tau \rightarrow \sigma)_2 &\stackrel{\text{def}}{=} \vec{\mathcal{D}}(\tau)_1 \rightarrow \vec{\mathcal{D}}(\sigma)_2 \\ \overleftarrow{\mathcal{D}}(\tau \rightarrow \sigma)_1 &\stackrel{\text{def}}{=} \overleftarrow{\mathcal{D}}(\tau)_1 \rightarrow \overleftarrow{\mathcal{D}}(\sigma) * (\overleftarrow{\mathcal{D}}(\sigma)_2 \multimap \overleftarrow{\mathcal{D}}(\tau)_2) & \overleftarrow{\mathcal{D}}(\tau \rightarrow \sigma)_2 &\stackrel{\text{def}}{=} !\overleftarrow{\mathcal{D}}(\tau)_1 \otimes \overleftarrow{\mathcal{D}}(\sigma)_2. \end{aligned}$$

Interestingly, we see that for higher-order programs, the primal transformations are no longer the identity. Indeed, the primals $\vec{\mathcal{D}}(\tau \rightarrow \sigma)_1$ and $\overleftarrow{\mathcal{D}}(\tau \rightarrow \sigma)_1$ of the function type $\tau \rightarrow \sigma$ store not only the primal function itself, but also its derivative with respect to its argument. The other half of a function's derivative, namely the derivative with respect to the context variables over which it closes, is stored in the tangent space $\vec{\mathcal{D}}(\tau \rightarrow \sigma)_2$ and cotangent space $\overleftarrow{\mathcal{D}}(\tau \rightarrow \sigma)_2$ of the function type $\tau \rightarrow \sigma$.

INSIGHT 4. *A forward (respectively, reverse) mode primal to a function type $\tau \rightarrow \sigma$ keeps track of both the function and its derivative with respect to its argument (respectively, transposed derivative). For reverse AD, a cotangent at function type $\tau \rightarrow \sigma$ (to be propagated back to the enclosing context of the function), keeps track of the incoming cotangents v of type $\overleftarrow{\mathcal{D}}(\sigma)_2$ for each a primal x of type $\overleftarrow{\mathcal{D}}(\tau)_1$ on which we call the function. We store these pairs (x, v) in the type $!\overleftarrow{\mathcal{D}}(\tau)_1 \otimes \overleftarrow{\mathcal{D}}(\sigma)_2$ (which we will see is essentially a quotient of a list of pairs of type $\overleftarrow{\mathcal{D}}(\tau)_1 * \overleftarrow{\mathcal{D}}(\sigma)_2$). Less surprisingly, for forward AD, a tangent at function type $\tau \rightarrow \sigma$ (propagated forward from the enclosing context of the function) consists of a function sending each argument primal of type $\vec{\mathcal{D}}(\tau)_1$ to the outgoing tangent of type $\vec{\mathcal{D}}(\sigma)_2$.*

On programs, we obtain the following extensions of our definitions for reverse AD:

$$\begin{aligned} \vec{\mathcal{D}}_{\Gamma}(\lambda x. t) &\stackrel{\text{def}}{=} \text{let } y = \lambda x. \vec{\mathcal{D}}_{\Gamma, x}(t) \text{ in } \langle \lambda x. \text{let } \langle z, z' \rangle = y \ x \text{ in } \langle z, \underline{\lambda} v. \text{snd}(z' \bullet v) \rangle, \underline{\lambda} v. \text{case } v \text{ of } !x \otimes v \rightarrow \text{fst}((\text{snd}(y \ x)) \bullet v) \rangle \\ \overleftarrow{\mathcal{D}}_{\Gamma}(t \ s) &\stackrel{\text{def}}{=} \langle \text{let } \langle x, x'_{\text{ctx}} \rangle = \overleftarrow{\mathcal{D}}_{\Gamma}(t) \text{ in let } \langle y, y' \rangle = \overleftarrow{\mathcal{D}}_{\Gamma}(s) \text{ in let } \langle z, x'_{\text{arg}} \rangle = x \ y \text{ in } \langle z, \underline{\lambda} v. x'_{\text{ctx}} \bullet (!y \otimes v) + y' \bullet (x'_{\text{arg}} \bullet v) \rangle. \end{aligned}$$

With regards to $\vec{\mathcal{D}}_{\Gamma}(\lambda x. t)$: suppose that $(\lambda x. t) : \tau \rightarrow \sigma$. Note then that we have $\Gamma, x : \tau \vdash t : \sigma$ and hence we get for t 's derivative that $\vec{\mathcal{D}}(\Gamma)_1, x : \vec{\mathcal{D}}(\tau)_1 \vdash \vec{\mathcal{D}}_{\Gamma, x}(t) : \vec{\mathcal{D}}(\sigma)_1 * (\vec{\mathcal{D}}(\tau)_2 \multimap \vec{\mathcal{D}}(\sigma)_2)$. Calling the transposed derivative function for t (z' in the primal, $\text{snd}(y \ x)$ in the dual) hence gives us *both* halves of the transposed derivative (the derivative with respect to the function argument and the context variables, that is) of the function; we then select the appropriate ones using projections. Similarly, in $\overleftarrow{\mathcal{D}}_{\Gamma}(t \ s)$ we extract the transposed derivative x'_{ctx} of t with respect to the context variables from the cotangent of t and obtain the transposed derivative x'_{arg} of t with respect to its function arguments from t 's primal. We combine these two halves of the transposed derivative with s 's transposed derivative (which we get from its cotangent) to get the correct transposed derivative for the function application $t \ s$.

2.6 Proving CHAD Correct

With these definitions in place, we turn to the correctness of the source-code transformations. To phrase correctness, we first need to construct a suitable semantics with an uncontroversial notion of semantic differentiation (see Section 5). We choose to work with a semantics in terms of the category **Set** of sets and functions,⁴ noting that any function $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$ has a unique derivative

⁴In Reference [42], we worked with a semantics in terms of diffeological spaces and differentiable functions, instead, to ensure that any first-order function is differentiable. This choice amounted to a separation between the proof that every first-order denotation is differentiable and that AD computes the correct derivative. To make the presentation of this article

as long as f is differentiable. We will only be interested in this semantic notion of derivative of first-order functions for the sake of correctness of AD, and we will not concern ourselves with semantic derivatives of higher-order functions. We interpret the required linear types in the category **CMon** of commutative monoids and homomorphisms.

By the universal properties of the syntax, we obtain canonical, structure-preserving (homomorphic) functors $\llbracket - \rrbracket : \mathbf{Syn} \rightarrow \mathbf{Set}$, $\llbracket - \rrbracket : \mathbf{CSyn} \rightarrow \mathbf{Set}$ and $\llbracket - \rrbracket : \mathbf{LSyn} \rightarrow \mathbf{CMon}$ once we fix interpretations \mathbb{R}^n of \mathbf{real}^n and well-typed (differentiable) interpretations $\llbracket \text{op} \rrbracket$ for each operation op . These functors define a concrete denotational semantics for our source and target languages.

Having constructed the semantics, we can turn to the correctness proof (of Section 8). Because calculus does not provide an unambiguous notion of derivative at function spaces, we cannot prove that the AD transformations correctly implement mathematical derivatives by plain induction on the syntax. Instead, we use a logical relations argument over the semantics.

INSIGHT 5. *Once we show that the (transposed) derivatives of primitive operations op are correctly implemented, correctness of (transposed) derivatives of all other programs follows from a standard logical relations construction over the semantics that relates a curve to its primal and (co)tangent curve. By the chain rule for (transposed) derivatives, all CHAD transformed programs respect the logical relations. By basic calculus results, CHAD therefore has to compute the (transposed) derivative.*

In Section 8, we present an elegant high-level formulation of this correctness argument, using categorical logical relations techniques (subconing). To make this argument accessible to a wider audience of readers, we present here a low-level description of the logical relations argument. The reader may note that these arguments look significantly different from the usual definitions of logical relations. That difference is caused by the non-standard Cartesian closed structure of $\Sigma_C \mathcal{L}$ and $\Sigma_C \mathcal{L}^{op}$ and the proof is entirely standard when viewed from the higher level of abstraction that subconing gives us.

We first sketch the correctness argument for forward mode CHAD. By induction on the structure of types, writing (f, f') for the product pairing of f and f' , we construct a logical relation P_τ on types τ as

$$\begin{aligned} P_\tau &\subseteq (\mathbb{R}^d \Rightarrow \llbracket \tau \rrbracket) \times (\mathbb{R}^d \Rightarrow (\llbracket \vec{\mathcal{D}}(\tau)_1 \rrbracket \times (\mathbb{R}^d \multimap \llbracket \vec{\mathcal{D}}(\tau)_2 \rrbracket))) \\ P_{\mathbf{real}^n} &\stackrel{\text{def}}{=} \{(f, g) \mid f \text{ is differentiable and } g = \mathcal{T}f\} \\ P_1 &\stackrel{\text{def}}{=} \{(x \mapsto (), x \mapsto ((, r \mapsto ())))\} \\ P_{\tau * \sigma} &\stackrel{\text{def}}{=} \{((f, f'), ((g, g'), x \mapsto r \mapsto (h(x)(r), h'(x)(r)))) \mid (f, (g, h)) \in P_\tau, (f', (g', h')) \in P_\sigma\} \\ P_{\tau \rightarrow \sigma} &\stackrel{\text{def}}{=} \{(f, (g, h)) \mid \forall (f', (g', h')) \in P_\tau. (x \mapsto f(x)(f'(x)), (x \mapsto \pi_1(g(x)(g'(x))), \\ &\quad x \mapsto r \mapsto (\pi_2(g(x)(g'(x)))(h'(x)(r)) + h(x)(r)(g'(x)))) \in P_\sigma\}. \end{aligned}$$

We extend the logical relation to typing contexts $\Gamma = x_1 : \tau_1, \dots, x_n : \tau_n$ as $P_\Gamma \stackrel{\text{def}}{=} P_{\tau_1 * \dots * \tau_n}$. Then, we establish the following fundamental lemma, which says that all well-typed source language programs t respect the logical relation.

LEMMA 2.1. *For any source language program $\Gamma \vdash t : \sigma$ and any $f : \mathbb{R}^d \rightarrow \llbracket \Gamma \rrbracket$, $g : \mathbb{R}^d \rightarrow \llbracket \vec{\mathcal{D}}(\Gamma)_1 \rrbracket$, $h : \mathbb{R}^d \rightarrow \mathbb{R}^d \multimap \llbracket \vec{\mathcal{D}}(\Gamma)_2 \rrbracket$ such that $(f, (g, h)) \in P_\Gamma$, we have that $(f; \llbracket t \rrbracket, (g; \llbracket \vec{\mathcal{D}}_\Gamma(t) \rrbracket; \pi_1, x \mapsto r \mapsto \pi_2(\llbracket \vec{\mathcal{D}}_\Gamma(t) \rrbracket)(g(x))(h(x)(r)))) \in P_\sigma$.*

more accessible, we have chosen to simply work with sets and functions and to prove differentiability of every first-order denotation simultaneously in the proof that AD computes the correct derivative.

The proof goes via induction on the typing derivation of t . The main remaining step in the argument is to note that any tangent vector at $\llbracket \tau_1 * \dots * \tau_n \rrbracket \cong \mathbb{R}^N$, for first-order τ_i , can be represented by a curve $\mathbb{R} \rightarrow \llbracket \tau_1 * \dots * \tau_n \rrbracket$.

Similarly, for reverse mode CHAD, we define, by induction on the structure of types, a logical relation P_τ on types τ (and, as before, we also define $P_\Gamma = P_{\tau_1 * \dots * \tau_n}$ for typing contexts $\Gamma = x_1 : \tau_1, \dots, x_n : \tau_n$):

$$\begin{aligned} P_\tau &\subseteq (\mathbb{R}^d \rightarrow \llbracket \tau \rrbracket) \times (\mathbb{R}^d \rightarrow (\llbracket \widetilde{\mathcal{D}}(\tau)_1 \rrbracket \times (\llbracket \widetilde{\mathcal{D}}(\tau)_2 \rrbracket \multimap \mathbb{R}^d))) \\ P_{\text{real}^n} &\stackrel{\text{def}}{=} \{(f, g) \mid f \text{ is differentiable and } g = \mathcal{T}^* f\} \\ P_1 &\stackrel{\text{def}}{=} \{(x \mapsto (), x \mapsto ((), v \mapsto 0))\} \\ P_{\tau * \sigma} &\stackrel{\text{def}}{=} \{((f, f'), ((g, g'), x \mapsto v \mapsto h(x)(\pi_1 v) + h'(x)(\pi_2 v))) \mid (f, (g, h)) \in P_\tau, (f', (g', h')) \in P_\sigma\} \\ P_{\tau \multimap \sigma} &\stackrel{\text{def}}{=} \{(f, (g, h)) \mid \forall (f', (g', h')) \in P_\tau. (x \mapsto f(x)(f'(x)), (x \mapsto \pi_1(g(x)(g'(x))), \\ &\quad x \mapsto v \mapsto h(x)(!g'(x) \otimes v) + h'(x)((\pi_2(g(x)(g'(x))))v)) \in P_\sigma\}. \end{aligned}$$

Then, we establish the following fundamental lemma.

LEMMA 2.2. *For any source language program $\Gamma \vdash t : \sigma$ and any $f : \mathbb{R}^d \rightarrow \llbracket \Gamma \rrbracket$, $g : \mathbb{R}^d \rightarrow \llbracket \widetilde{\mathcal{D}}(\Gamma)_1 \rrbracket$, $h : \mathbb{R}^d \rightarrow \llbracket \widetilde{\mathcal{D}}(\Gamma)_2 \rrbracket \multimap \mathbb{R}^d$ such that $(f, (g, h)) \in P_\Gamma$, we have that $(f; \llbracket t \rrbracket, (g; \llbracket \widetilde{\mathcal{D}}_\Gamma(t) \rrbracket; \pi_1, x \mapsto v \mapsto h(x)(\pi_2(\llbracket \widetilde{\mathcal{D}}_\Gamma(t) \rrbracket(g(x)))(v)))) \in P_\sigma$.*

The proof goes via induction on the typing derivation of t . Correctness follows from the fundamental lemma by observing that any cotangent vector at $\llbracket \tau_1 * \dots * \tau_n \rrbracket \cong \mathbb{R}^N$, for first-order τ_i , can be represented by a curve $\mathbb{R} \rightarrow \llbracket \tau_1 * \dots * \tau_n \rrbracket$.

We obtain our main theorem, Theorem A, but now for our CHAD algorithms applied to a higher-order source language.

2.7 A Practical Implementation in Haskell

Next, we address the practicality of our method (in Section 9). The code transformations we employ are not too daunting to implement and they are well behaved in the sense that the derivative code they generate grows linearly in the size of the original source code. However, the implementation of the required linear types presents a challenge. Indeed, types like $!(-) \otimes (-)$ and $(-) \multimap (-)$ are absent from languages such as Haskell and OCaml. Luckily, in this instance, we can implement them using abstract data types by using a (basic) module system:

INSIGHT 6. *Under the hood, $! \tau \otimes \sigma$ can consist of a list of values of type $\tau * \sigma$. Its API ensures that the list order and the difference between $xs \mathrel{++} [(t, s), (t, s')]$ and $xs \mathrel{++} [(t, s + s')]$ (or xs and $xs \mathrel{++} [(t, 0)]$) cannot be observed: as such, it is a quotient type. Meanwhile, $\tau \multimap \sigma$ can be implemented as a standard function type $\tau \rightarrow \sigma$ with a limited API that enforces that we can only ever construct linear functions: as such, it is a subtype.*

This idea leads to our reference implementation of CHAD in Haskell (available at <https://github.com/VMatthijs/CHAD>), which generates perfectly standard simply typed functional code that is given a bit of extra type-safety with the linear types, implemented as abstract types. To illustrate what our method does in practice, we consider two programs of our higher-order source language, that we may want to differentiate in Figure 2(a) and (b). The forward and reverse mode derivatives that our CHAD implementation generates for these programs are listed in Figure 2(c)

```

let f  =  $\lambda z. x * z + 1$ 
      zs = replicate x
      ys = map f zs in
      ys

```

(a) Original program $x : \text{real} \vdash t : \text{real}^n$ computing a function $\llbracket t \rrbracket : \mathbb{R} \rightarrow \mathbb{R}^n; x \mapsto ys$. replicate x copies x n times.

```

let f  =  $\lambda x2i. x1 * x2i$ 
      ys = map f x2
      w  = sum ys in
      w

```

(b) Original program $x1 : \text{real}, x2 : \text{real}^n \vdash s : \text{real}$ computing a function $\llbracket s \rrbracket : \mathbb{R} \times \mathbb{R}^n \rightarrow \mathbb{R}; (x1, x2) \mapsto w$.

```

let f  =  $\lambda z. \langle x * z + 1, \lambda z'. x * z' \rangle$ 
      zs = replicate x
      ys = map ( $\lambda z. \text{fst}(f z)$ ) zs in
       $\langle ys, \lambda x'.$ 
        let f' =  $\lambda z. x' * z$ 
            zs' = replicate x'
            ys' = map f' zs + zipWith ( $\lambda z. \text{snd}(f z)$ ) zs zs' in
            ys'
       $\rangle$ 

```

} primals

} tangents

(c) Forward AD transformed program $x : \text{real} \vdash \vec{D}(t) : \text{real}^n * (\underline{\text{real}} \multimap \underline{\text{real}}^n)$ computing the derivative $\llbracket \vec{D}(t) \rrbracket = \mathcal{T}\llbracket t \rrbracket : \mathbb{R} \rightarrow \mathbb{R}^n \times (\mathbb{R} \multimap \mathbb{R}^n); x \mapsto (ys, x' \mapsto ys')$.

```

let f  =  $\lambda x2i. \langle x1 * x2i, \lambda y'. x1 * y' \rangle$ 
      ys = map ( $\lambda x2i. \text{fst}(f x2i)$ ) x2
      w  = sum ys in
       $\langle w, \lambda w'.$ 
        let ys' = replicate w'
            f' = zip x2 ys'
            x1' = sum (map ( $\lambda \langle x2i, y' \rangle. y' * x2i$ ) f')
            x2' = zipWith ( $\lambda x2i. \text{snd}(f x2i)$ ) x2 ys' in
             $\langle x1', x2' \rangle$ 
       $\rangle$ 

```

} primals

} cotangents

(d) Reverse AD transformed program $x1 : \text{real}, x2 : \text{real}^n \vdash \overleftarrow{D}(s) : \text{real} * (\underline{\text{real}} \multimap \underline{\text{real}} * \underline{\text{real}}^n)$ computing the transposed derivative $\llbracket \overleftarrow{D}(s) \rrbracket = \mathcal{T}^* \llbracket s \rrbracket : \mathbb{R} \times \mathbb{R}^n \rightarrow \mathbb{R} \times (\mathbb{R} \multimap \mathbb{R} \times \mathbb{R}^n); (x1, x2) \mapsto (w, w' \mapsto (x1', x2'))$.

Fig. 2. Forward and reverse AD illustrated on higher-order functional array processing programs. The parts of the programs that involve AD on higher-order functions are marked in blue. Observe that in (c) (respectively, (d)), the primal value associated with the function f from program (a) (respectively, (b)) computes both the original function f as well as its derivative (respectively, transposed derivative) with respect to its argument z (respectively, $x2i$). In (c), the tangent f' to f is *produced* by propagating *forwards* the tangent x' to the context variable x that f captures, by using f 's derivative with respect to x . This lets us correctly propagate forwards the contributions to ys' from both f 's dependence on its argument z and on its context variable x . Dually, in (d), the cotangent to f , which we construct from the cotangent ys' , is *consumed* by propagating it *backwards* to the cotangent $x1'$ to the context variable $x1$ that f captures, by using f 's transposed derivative with respect to $x1$. Meanwhile, the adjoint $x2'$ is constructed using the part of the primal of f that captures f 's transposed derivative with respect to $x2i$.

and (d), again modulo minor simplifications that aid legibility but have no significant runtime implications.⁵

In Section 9, we also phrase the correctness proof of the AD transformations in elementary terms, such that it holds in the applied setting where we use abstract types to implement linear types. We show that our correctness results are meaningful, as they make use of a denotational semantics that is adequate with respect to the standard operational semantics. Furthermore, to stress

⁵For information on the exact simplifications performed, see Appendix B.

the applicability of our method, we show in Section 10 that it extends to higher-order (primitive) operations, such as `map`.

Finally (in Section 11), we zoom out and reflect on how this method generalizes. The crux of CHAD is in the following steps:

- view the source language as a freely generated category \mathbf{Syn} with some appropriate structure \mathcal{S} (such as Cartesian closure, coproducts, (co)inductive types, iteration), generated from objects \mathbf{real}^n and morphisms `op`;
- find a suitable target language \mathbf{LSyn} (with linear types arising from the effect of commutative monoids) for the translation such that $\Sigma_{\mathbf{CSyn}}\mathbf{LSyn}$ and $\Sigma_{\mathbf{CSyn}}\mathbf{LSyn}^{op}$ are categories with the structure \mathcal{S} ; in our experience, this is possible for most common choices of \mathcal{S} corresponding to programming language constructs;
- then, by the universal property of \mathbf{Syn} , we obtain unique structure preserving (homomorphic) functors $\vec{\mathcal{D}} : \mathbf{Syn} \rightarrow \Sigma_{\mathbf{CSyn}}\mathbf{LSyn}$ and $\overleftarrow{\mathcal{D}} : \mathbf{Syn} \rightarrow \Sigma_{\mathbf{CSyn}}\mathbf{LSyn}^{op}$ defining forward and reverse mode AD transformations, as soon as we fix their action on `op` (and \mathbf{real}^n) to implement the derivative of the operations;
- the correctness of these AD methods follows by a standard categorical logical relations argument as the subscones $\overrightarrow{\mathbf{SScone}}$ and $\overleftarrow{\mathbf{SScone}}$ tend to also be categories with the structure \mathcal{S} for most choices of \mathcal{S} .

INSIGHT 7. *The definition and correctness proof of forward and reverse AD on expressive programming languages follow automatically, by viewing the algorithms as structure preserving functors $\vec{\mathcal{D}} : \mathbf{Syn} \rightarrow \Sigma_{\mathbf{CSyn}}\mathbf{LSyn}$ and $\overleftarrow{\mathcal{D}} : \mathbf{Syn} \rightarrow \Sigma_{\mathbf{CSyn}}\mathbf{LSyn}^{op}$.*

We conclude by observing that, in this sense, CHAD is not specific to automatic differentiation at all. We can choose different generators than \mathbf{real}^n and `op` for \mathbf{Syn} and different mappings of these generators under $\vec{\mathcal{D}}$ and $\overleftarrow{\mathcal{D}}$. Doing so lets CHAD derive various other dynamic program analyses that accumulate data in a commutative monoid, together with their correctness proofs by logical relations (see Section 11.3).

3 λ -CALCULUS AS A SOURCE LANGUAGE FOR AUTOMATIC DIFFERENTIATION

As a source language for our AD translations, we can begin with a standard, simply typed λ -calculus that has ground types \mathbf{real}^n of statically sized⁶ arrays of n real numbers, for all $n \in \mathbb{N}$, and sets $\mathbf{Op}_{n_1, \dots, n_k}^m$ of primitive operations `op` for all $k, m, n_1, \dots, n_k \in \mathbb{N}$. These operations will be interpreted as differentiable⁷ functions $(\mathbb{R}^{n_1} \times \dots \times \mathbb{R}^{n_k}) \rightarrow \mathbb{R}^m$. Examples to keep in mind for `op` include

- constants $\underline{c} \in \mathbf{Op}^n$ for each $c \in \mathbb{R}^n$, for which we slightly abuse notation and write $\underline{c}()$ as \underline{c} ;
- elementwise addition and product $(+), (*) \in \mathbf{Op}_{n,n}^n$ and matrix-vector product $(\star) \in \mathbf{Op}_{n \cdot m, m}^n$;
- operations for summing all the elements in an array: $\mathbf{sum} \in \mathbf{Op}_n^1$;
- some non-linear functions like the sigmoid function $\varsigma \in \mathbf{Op}_1^1$.

We intentionally present operations in a schematic way, as primitive operations tend to form a collection that is added to in a by-need fashion, as an AD library develops. The precise operations

⁶Here, we work with statically sized arrays to simplify the theoretical development. However, in our implementation, we show that CHAD applies equally well to types of varying dimension such as dynamically sized arrays.

⁷Observe that this restriction does not exclude, for example, functions that are differentiable almost everywhere like ReLU in a meaningful way, as such functions can be approximated with differentiable functions. Given how coarse an approximation real numbers already are to the reality of floating point arithmetic, the distinction between everywhere differentiable and almost-everywhere differentiable is not meaningful, in practice.

$\frac{((x : \tau) \in \Gamma)}{\Gamma \vdash x : \tau}$	$\frac{\Gamma \vdash t : \tau \quad \Gamma, x : \tau \vdash s : \sigma}{\Gamma \vdash \text{let } x = t \text{ in } s : \sigma}$	$\frac{\{\Gamma \vdash t_i : \mathbf{real}^{n_i} \mid i = 1, \dots, k\} \quad (\text{op} \in \text{Op}_{n_1, \dots, n_k}^m)}{\Gamma \vdash \text{op}(t_1, \dots, t_n) : \mathbf{real}^m}$	$\frac{}{\Gamma \vdash \langle \rangle : 1}$
$\frac{\Gamma \vdash t : \tau \quad \Gamma \vdash s : \sigma}{\Gamma \vdash \langle t, s \rangle : \tau * \sigma}$	$\frac{\Gamma \vdash t : \tau * \sigma}{\Gamma \vdash \text{fst } t : \tau}$	$\frac{\Gamma \vdash t : \tau * \sigma}{\Gamma \vdash \text{snd } t : \sigma}$	$\frac{\Gamma, x : \tau \vdash t : \sigma \quad \Gamma \vdash t : \sigma \rightarrow \tau \quad \Gamma \vdash s : \sigma}{\Gamma \vdash t s : \tau}$

Fig. 3. Typing rules for the AD source language.

needed will depend on the applications. In statistics and machine learning applications, Op tends to include a mix of multi-dimensional linear algebra operations and mostly one-dimensional non-linear functions. A typical library for use in machine learning would work with multi-dimensional arrays (sometimes called “tensors”). We focus here on one-dimensional arrays as the issues of how precisely to represent the arrays are orthogonal to the concerns of our development.

The types τ, σ, ρ and terms t, s, r of our AD source language are as follows:

$\tau, \sigma, \rho ::=$	types	$\tau_1 * \tau_2$	binary product
\mathbf{real}^n	real arrays	$\tau \rightarrow \sigma$	function
1	nullary product		
$t, s, r ::=$	terms	$\langle \rangle \mid \langle t, s \rangle$	product tuples
x	identifier	$\text{fst } t \mid \text{snd } t$	product projections
$\text{let } x = t \text{ in } s$	let-bindings	$\lambda x. t$	function abstraction
$\text{op}(t_1, \dots, t_k)$	operations	$t s$	function application

The typing rules are in Figure 3. We use the usual conventions for free and bound variables and write the capture-avoiding substitution of x with s in t as $t[s/x]$. We employ the usual syntactic sugar $\lambda\langle x, y \rangle. t \stackrel{\text{def}}{=} \lambda z. t[\text{fst } z/x, \text{snd } z/y]$, and we write \mathbf{real} for \mathbf{real}^1 . As Figure 4 displays, we consider the standard $\beta\eta$ -equational theory for our language, where equations hold on pairs of terms of the same type in the same context. We could consider further equations for our operations, but we do not as we will not need them.

This standard λ -calculus is widely known to be equivalent to the free Cartesian closed category **Syn** generated by the objects \mathbf{real}^n and the morphisms op (see Reference [27]).

- **Syn** has types τ, σ, ρ as objects;
- **Syn** has morphisms $t \in \mathbf{Syn}(\tau, \sigma)$ that are in 1-1 correspondence with terms $x : \tau \vdash t : \sigma$ up to $\beta\eta$ -equivalence (which includes α -equivalence);
- identities are represented by $x : \tau \vdash x : \tau$;
- composition of $x : \tau \vdash t : \sigma$ and $y : \sigma \vdash s : \rho$ is represented by $x : \tau \vdash \text{let } y = t \text{ in } s : \rho$;
- 1 and $\tau * \sigma$ represent nullary and binary product, while $\tau \rightarrow \sigma$ is the categorical exponential.

Syn has the following well-known universal property.

PROPOSITION 3.1 (UNIVERSAL PROPERTY OF **Syn).** *For any Cartesian closed category $(C, 1, \times, \Rightarrow)$, we obtain a unique Cartesian closed functor $F : \mathbf{Syn} \rightarrow C$, once we choose objects $F(\mathbf{real}^n)$ of C as well as, for each $\text{op} \in \text{Op}_{n_1, \dots, n_k}^m$, make well-typed choices of C -morphisms $F(\text{op}) : (F(\mathbf{real}^{n_1}) \times \dots \times F(\mathbf{real}^{n_k})) \rightarrow F(\mathbf{real}^m)$.*

4 LINEAR λ -CALCULUS AS AN IDEALISED AD TARGET LANGUAGE

As a target language for our AD source code transformations, we consider a language that extends the language of Section 3 with limited linear types. We could opt to work with a full linear logic as in Reference [3] or Reference [6]. Instead, however, we will only include the bare minimum of linear

$$\text{let } x = t \text{ in } s = s[t/x] \quad t = \langle \rangle \quad \text{fst } \langle t, s \rangle = t \quad \text{snd } \langle t, s \rangle = s \quad t = \langle \text{fst } t, \text{snd } t \rangle \quad (\lambda x. t) s = t[s/x] \quad t \stackrel{\#x}{=} \lambda x. t x$$

Fig. 4. Standard $\beta\eta$ -laws for products and functions. We write $\stackrel{\#x_1, \dots, x_n}{=}$ to indicate that the variables x_1, \dots, x_n need to be fresh in the left hand side. As usual, we only distinguish terms up to α -renaming of bound variables.

type formers that we actually need to phrase the AD transformations. The resulting language is closely related to, but more minimal than, the Enriched Effect Calculus of Reference [14]. We limit our language in this way because we want to stress that the resulting code transformations can easily be implemented in existing functional languages such as Haskell or OCaml. As we discuss in Section 9, the idea will be to make use of a module system to implement the required linear types as abstract data types.

In our idealised target language, we consider *linear types* (aka computation types) $\underline{\tau}, \underline{\sigma}, \underline{\rho}$, in addition to the *Cartesian types* (value types) τ, σ, ρ that we have considered so far. We think of Cartesian types as denoting sets and linear types as denoting sets equipped with an algebraic structure. The Cartesian types will be used to represent sets of primals. The relevant algebraic structure on linear types, in this instance, turns out to be that of a commutative monoid, as this algebraic structure is needed to phrase automatic differentiation algorithms. Indeed, we will use the linear types to denote sets of (co)tangent vectors. These (co)tangents form a commutative monoid under addition.

Concretely, we extend the types and terms of our language as follows:

$\underline{\tau}, \underline{\sigma}, \underline{\rho} ::=$	linear types	$\underline{\tau} * \underline{\sigma}$	binary product
$\underline{\text{real}}^n$	real array	$\underline{\tau} \rightarrow \underline{\sigma}$	power
$\underline{1}$	unit type	$! \underline{\tau} \otimes \underline{\sigma}$	copower
$\tau, \sigma, \rho ::=$	Cartesian types	$\underline{\tau} \multimap \underline{\sigma}$	linear function
\dots	as in Section 3		
$t, s, r ::=$	terms	$\text{lop}(t_1, \dots, t_k; s)$	linear operation
\dots	as in Section 3	$! t \otimes s \mid \text{case } t \text{ of } ! y \otimes v \rightarrow s$	copower intro/elim
v	linear identifier	$\underline{\lambda} v. t \mid t \bullet s$	abstraction/application
$\text{let } v = t \text{ in } s$	linear let-binding	$\underline{0} \mid t + s$	monoid structure.

We work with linear operations $\text{lop} \in \text{LOp}_{n_1, \dots, n_k; n'_1, \dots, n'_l}^{m_1, \dots, m_r}$, which are intended to represent differentiable functions

$$(\mathbb{R}^{n_1} \times \dots \times \mathbb{R}^{n_k} \times \mathbb{R}^{n'_1} \times \dots \times \mathbb{R}^{n'_l}) \rightarrow \mathbb{R}^{m_1} \times \dots \times \mathbb{R}^{m_r}$$

that are linear (in the sense of respecting $\underline{0}$ and $+$) in the last l arguments but not in the first k . We write

$$\text{LDom}(\text{lop}) \stackrel{\text{def}}{=} \text{real}^{n'_1} * \dots * \text{real}^{n'_l} \quad \text{and} \quad \text{CDom}(\text{lop}) \stackrel{\text{def}}{=} \text{real}^{m_1} * \dots * \text{real}^{m_r}$$

for $\text{lop} \in \text{LOp}_{n_1, \dots, n_k; n'_1, \dots, n'_l}^{m_1, \dots, m_r}$. These operations can include dense and sparse matrix-vector multiplications, for example. Their purpose is to serve as primitives to implement derivatives $\text{Dop}(x_1, \dots, x_k; v)$ and transposed derivatives $\text{Dop}^t(x_1, \dots, x_k; v)$ of the operations op from the source language as terms with free variables x_1, \dots, x_k, v that are linear in v . In fact, one can also opt to directly include primitive linear operations

$$\text{Dop} \in \text{LOp}_{n_1, \dots, n_k; n_1, \dots, n_k}^m \quad \text{Dop}^t \in \text{LOp}_{n_1, \dots, n_k; m}^{n_1, \dots, n_k}$$

in LOp as the derivatives of each (Cartesian) operation $\text{op} \in \text{Op}$.

$\frac{}{\Gamma; v : \underline{\tau} \vdash v : \underline{\tau}}$	$\frac{\Gamma \vdash t : \tau \quad \Gamma, x : \tau; v : \underline{\sigma} \vdash s : \underline{\rho}}{\Gamma; v : \underline{\sigma} \vdash \text{let } x = t \text{ in } s : \underline{\rho}}$	$\frac{\Gamma; v : \underline{\tau} \vdash t : \underline{\sigma} \quad \Gamma; v : \underline{\sigma} \vdash s : \underline{\rho}}{\Gamma; v : \underline{\tau} \vdash \text{let } v = t \text{ in } s : \underline{\rho}}$	
$\frac{\{\Gamma \vdash t_i : \mathbf{real}^{n_i} \mid i = 1, \dots, k\} \quad \Gamma; v : \underline{\tau} \vdash s : \mathbf{LDom}(\text{lop}) \quad (\text{lop} \in \text{LOp}_{n_1, \dots, n_k; n'_1, \dots, n'_l}^{m_1, \dots, m_r})}{\Gamma; v : \underline{\tau} \vdash \text{lop}(t_1, \dots, t_k; s) : \mathbf{CDom}(\text{lop})}$			
$\frac{}{\Gamma; v : \underline{\tau} \vdash \langle \rangle : \underline{1}}$	$\frac{\Gamma; v : \underline{\tau} \vdash t : \underline{\sigma} \quad \Gamma; v : \underline{\tau} \vdash s : \underline{\rho}}{\Gamma; v : \underline{\tau} \vdash \langle t, s \rangle : \underline{\sigma * \rho}}$	$\frac{\Gamma; v : \underline{\tau} \vdash t : \underline{\sigma * \rho}}{\Gamma; v : \underline{\tau} \vdash \text{fst } t : \underline{\sigma}}$	$\frac{\Gamma; v : \underline{\tau} \vdash t : \underline{\sigma * \rho}}{\Gamma; v : \underline{\tau} \vdash \text{snd } t : \underline{\rho}}$
$\frac{\Gamma, y : \sigma; v : \underline{\tau} \vdash t : \underline{\rho}}{\Gamma; v : \underline{\tau} \vdash \lambda y. t : \sigma \rightarrow \underline{\rho}}$	$\frac{\Gamma; v : \underline{\tau} \vdash t : \sigma \rightarrow \underline{\rho} \quad \Gamma \vdash s : \sigma}{\Gamma; v : \underline{\tau} \vdash t s : \underline{\rho}}$	$\frac{\Gamma \vdash t : \sigma \quad \Gamma; v : \underline{\tau} \vdash s : \underline{\rho}}{\Gamma; v : \underline{\tau} \vdash !t \otimes s : !\sigma \otimes \underline{\rho}}$	
$\frac{\Gamma; v : \underline{\tau} \vdash t : !\sigma \otimes \underline{\rho} \quad \Gamma, y : \sigma; z : \underline{\rho} \vdash s : \underline{\rho}'}{\Gamma; v : \underline{\tau} \vdash \text{case } t \text{ of } !y \otimes v \rightarrow s : \underline{\rho}'}$		$\frac{\Gamma; v : \underline{\tau} \vdash t : \underline{\sigma}}{\Gamma \vdash \lambda v. t : \underline{\tau} \multimap \underline{\sigma}}$	
$\frac{\Gamma \vdash t : \underline{\rho} \multimap \underline{\sigma} \quad \Gamma; v : \underline{\tau} \vdash s : \underline{\rho}}{\Gamma; v : \underline{\tau} \vdash t \bullet s : \underline{\sigma}}$	$\frac{}{\Gamma; v : \underline{\tau} \vdash \underline{0} : \underline{\sigma}}$	$\frac{\Gamma; v : \underline{\tau} \vdash t : \underline{\sigma} \quad \Gamma; v : \underline{\tau} \vdash s : \underline{\sigma}}{\Gamma; v : \underline{\tau} \vdash t + s : \underline{\sigma}}$	

Fig. 5. Typing rules for the idealised AD target language with linear types, which we consider on top of the rules of Figure 3.

$\text{let } v = t \text{ in } s = s[t/v]$	
$\text{case } !t \otimes s \text{ of } !x \otimes v \mapsto r = r[t/x, s/v]$	$t[s/x] \stackrel{\#y, z}{=} \text{case } s \text{ of } !y \otimes v \rightarrow t[!y \otimes v/x]$
$(\lambda v. t) \bullet s = t[s/v]$	$t = \lambda v. t \bullet v$
$t + \underline{0} = t \quad \underline{0} + t = t$	$(t + s) + r = t + (s + r) \quad t + s = s + t$
$(\Gamma; v : \underline{\tau} \vdash t : \underline{\sigma}) \Rightarrow t[\underline{0}/v] = \underline{0}$	$(\Gamma; v : \underline{\tau} \vdash t : \underline{\sigma}) \Rightarrow t[s+r/v] = t[s/v] + t[r/v]$

Fig. 6. Equational rules for the idealised, linear AD language, which we use on top of the rules of Figure 4. In addition to standard $\beta\eta$ -rules for $!(-) \otimes (-)$ and \multimap -types, we add rules making $(\underline{0}, +)$ into a commutative monoid on the terms of each linear type as well as rules that say that terms of linear types are homomorphisms in their linear variable.

In addition to the judgement $\Gamma \vdash t : \tau$, which we encountered in Section 3, we now consider an additional judgement $\Gamma; v : \underline{\tau} \vdash t : \underline{\sigma}$. While we think of the former as denoting a function between sets, we think of the latter as a function from the set that Γ denotes to the set of monoid homomorphisms from the denotation of $\underline{\tau}$ to that of $\underline{\sigma}$.

Figures 3 and 5 display the typing rules of our language.

We consider the $\beta\eta$ -equational theory of Figures 4 and 6 for our language, where equations hold on pairs of terms of the same type in the same context. It includes $\beta\eta$ -rules as well as commutative monoid and homomorphism laws.

5 SEMANTICS OF THE SOURCE AND TARGET LANGUAGES

5.1 Preliminaries

5.1.1 Category Theory. We assume familiarity with categories, functors, natural transformations, and their theory of (co)limits and adjunctions. We write:

- unary, binary, and I -ary products as $\mathbb{1}$, $X_1 \times X_2$, and $\prod_{i \in I} X_i$, writing π_i for the projections and $()$, (x_1, x_2) , and $(x_i)_{i \in I}$ for the tupling maps;
- unary, binary, and I -ary coproducts as 0 , $X_1 + X_2$, and $\sum_{i \in I} X_i$, writing ι_i for the injections and $[]$, $[x_1, x_2]$, and $[x_i]_{i \in I}$ for the cotupling maps;
- exponentials as $Y \Rightarrow X$, writing Λ and ev for currying and evaluation.

5.1.2 Commutative Monoids. A *monoid* $(|X|, 0_X, +_X)$ consists of a set $|X|$ with an element $0_X \in |X|$ and a function $(+_X) : |X| \times |X| \rightarrow |X|$ such that $0_X +_X x = x = x +_X 0_X$ for any $x \in |X|$ and $x +_X (x' +_X x'') = (x +_X x') +_X x''$ for any $x, x', x'' \in |X|$. A *monoid* $(|X|, 0_X, +_X)$ is called *commutative* if $x +_X x' = x' +_X x$ for all $x, x' \in |X|$. Given monoids X and Y , a function $f : |X| \rightarrow |Y|$ is called a *homomorphism of monoids* if $f(0_X) = 0_Y$ and $f(x +_X x') = f(x) +_Y f(x')$. We write **CMon** for the category of commutative monoids and their homomorphisms. We will frequently simply write 0 for 0_X and $+$ for $+_X$, if X is clear from context. We will sometimes write $\sum_{i=1}^n x_i$ for $(x_1 + x_2) + \dots + x_n$.

Example 5.1. The real numbers \mathbb{R} form a commutative monoid with 0 and $+$ equal to the number 0 and addition of numbers.

Example 5.2. Given commutative monoids $(X_i)_{i \in I}$, we can form the *product monoid* $\prod_{i \in I} X_i$ with underlying set $\prod_{i \in I} |X_i|$, $0 = (0_{X_i})_{i \in I}$ and $(x_i)_{i \in I} + (y_i)_{i \in I} \stackrel{\text{def}}{=} (x_i + y_i)_{i \in I}$. Given a set I and a commutative monoid X , we can form the *power monoid* $I \Rightarrow X \stackrel{\text{def}}{=} \prod_{i \in I} X$ as the I -fold self-product monoid.

Example 5.2 gives the categorical product in **CMon**. We can, for example, construct a commutative monoid structure on any Euclidean space $\mathbb{R}^k \stackrel{\text{def}}{=} \{0, \dots, k-1\} \Rightarrow \mathbb{R}$ by combining the one on \mathbb{R} with the power monoid structure.

Example 5.3. Given commutative monoids $(X_i)_{i \in I}$, we can form the *coproduct monoid* $\sum_{i \in I} X_i$ with underlying set $\{(x_i)_{i \in I} \in \prod_{i \in I} X_i \mid \{j \in I \mid x_j \neq 0_{X_j}\} \text{ is finite}\}$, $0 = (0_{X_i})_{i \in I}$ and $(x_i)_{i \in I} + (y_i)_{i \in I} \stackrel{\text{def}}{=} (x_i + y_i)_{i \in I}$. Given a set I and a commutative monoid X , we can form the *copower monoid* $!I \otimes X \stackrel{\text{def}}{=} \sum_{i \in I} X$ as the I -fold self-coproduct monoid. We will often write $!i \otimes x \stackrel{\text{def}}{=}$ (if $j = i$ then x else 0_X) $_{j \in I} \in !I \otimes X$.

Example 5.3 gives the categorical coproduct in **CMon**.

Example 5.4. Given commutative monoids X and Y , we can form the commutative monoid $X \multimap Y$ of *homomorphisms from X to Y* . We define $|X \multimap Y| \stackrel{\text{def}}{=} \mathbf{CMon}(X, Y)$, $0_{X \multimap Y} \stackrel{\text{def}}{=} (x \mapsto 0_Y)$, and $f +_{X \multimap Y} g \stackrel{\text{def}}{=} (x \mapsto f(x) +_Y g(x))$.

Example 5.4 gives the categorical internal hom in **CMon**. Commutative monoid homomorphisms $!I \otimes X \rightarrow Y$ are in 1-1-correspondence with functions $I \rightarrow |X \multimap Y|$.

Finally, a category \mathcal{C} is called **CMon**-enriched if we have a commutative monoid structure on each homset $\mathcal{C}(C, C')$ and function composition gives monoid homomorphisms $\mathcal{C}(C, C') \rightarrow \mathcal{C}(C', C'') \multimap \mathcal{C}(C, C'')$. In a category \mathcal{C} with finite products, these products are well-known to be biproducts (i.e., simultaneously products and coproducts) if and only if \mathcal{C} is **CMon**-enriched (for more details, see, for example, Reference [17]): define $[] \stackrel{\text{def}}{=} 0$ and $[f, g] \stackrel{\text{def}}{=} \pi_1; f + \pi_2; g$ and, conversely, $0 \stackrel{\text{def}}{=} []$ and $f + g \stackrel{\text{def}}{=} (\text{id}, \text{id}); [f, g]$.

5.2 Abstract Denotational Semantics

By the universal property of **Syn** (Proposition 3.1), the language of Section 3 has a canonical interpretation in any Cartesian closed category $(C, \mathbb{1}, \times, \Rightarrow)$, once we fix C -objects $\llbracket \mathbf{real}^n \rrbracket$ to interpret \mathbf{real}^n and C -morphisms $\llbracket \mathbf{op} \rrbracket \in C(\llbracket \mathbf{Dom}(\mathbf{op}) \rrbracket, \llbracket \mathbf{real}^m \rrbracket)$ to interpret $\mathbf{op} \in \mathbf{Op}_{n_1, \dots, n_k}^m$. That is, any Cartesian closed category with such a choice of objects and morphisms is a *categorical model of the source language* of Section 3. We interpret types τ and contexts Γ as C -objects $\llbracket \tau \rrbracket$ and $\llbracket \Gamma \rrbracket$:

$$\llbracket x_1 : \tau_1, \dots, x_n : \tau_n \rrbracket \stackrel{\text{def}}{=} \llbracket \tau_1 \rrbracket \times \dots \times \llbracket \tau_n \rrbracket \quad \llbracket \mathbb{1} \rrbracket \stackrel{\text{def}}{=} \mathbb{1} \quad \llbracket \tau * \sigma \rrbracket \stackrel{\text{def}}{=} \llbracket \tau \rrbracket \times \llbracket \sigma \rrbracket \quad \llbracket \tau \rightarrow \sigma \rrbracket \stackrel{\text{def}}{=} \llbracket \tau \rrbracket \Rightarrow \llbracket \sigma \rrbracket.$$

We interpret terms $\Gamma \vdash t : \tau$ as morphisms $\llbracket t \rrbracket$ in $C(\llbracket \Gamma \rrbracket, \llbracket \tau \rrbracket)$:

$$\begin{aligned} \llbracket \mathbf{op}(t_1, \dots, t_k) \rrbracket &\stackrel{\text{def}}{=} (\llbracket t_1 \rrbracket, \dots, \llbracket t_k \rrbracket); \llbracket \mathbf{op} \rrbracket \\ \llbracket x_1 : \tau_1, \dots, x_n : \tau_n \vdash x_k : \tau_k \rrbracket &\stackrel{\text{def}}{=} \pi_k & \llbracket \mathbf{let } x = t \mathbf{ in } s \rrbracket &\stackrel{\text{def}}{=} (\text{id}, \llbracket t \rrbracket); \llbracket s \rrbracket \\ \llbracket \langle \rangle \rrbracket &\stackrel{\text{def}}{=} () & \llbracket \langle t, s \rangle \rrbracket &\stackrel{\text{def}}{=} (\llbracket t \rrbracket, \llbracket s \rrbracket) \\ \llbracket \mathbf{fst } t \rrbracket &\stackrel{\text{def}}{=} \llbracket t \rrbracket; \pi_1 & \llbracket \mathbf{snd } t \rrbracket &\stackrel{\text{def}}{=} \llbracket t \rrbracket; \pi_2 & \llbracket \lambda x. t \rrbracket &\stackrel{\text{def}}{=} \Lambda(\llbracket t \rrbracket) & \llbracket t s \rrbracket &\stackrel{\text{def}}{=} (\llbracket t \rrbracket, \llbracket s \rrbracket); \mathbf{ev}. \end{aligned}$$

We discuss how to extend $\llbracket - \rrbracket$ to apply to the full target language of Section 4 by defining an appropriate notion of *categorical model for the target language* of Section 4.

Definition 5.5 (Categorical Model of the Target Language). By a categorical model of the target language, we mean the following data:

- A categorical model C of the source language.
- A locally indexed category (see, for example, Section 9.3.4 in Reference [29]) $\mathcal{L} : C^{op} \rightarrow \mathbf{Cat}$, i.e.,⁸ a (strict) contravariant functor from C to the category \mathbf{Cat} of categories, such that $\text{ob } \mathcal{L}(C) = \text{ob } \mathcal{L}(C')$ and $\mathcal{L}(f)(L) = L$ for any object L of $\text{ob } \mathcal{L}(C)$ and any $f : C' \rightarrow C$ in C .
- \mathcal{L} is *biadditive*: each category $\mathcal{L}(C)$ has (chosen) finite biproducts $(\mathbb{1}, \times)$ and $\mathcal{L}(f)$ preserves them, for any $f : C' \rightarrow C$ in C , in the sense that $\mathcal{L}(f)(\mathbb{1}) = \mathbb{1}$ and $\mathcal{L}(f)(L \times L') = \mathcal{L}(f)(L) \times \mathcal{L}(f)(L')$.
- \mathcal{L} *supports* $!(-) \otimes (-)$ -types and \Rightarrow -types: $\mathcal{L}(\pi_1)$ has a left adjoint $!C' \otimes_C -$ and a right adjoint functor $C' \Rightarrow_C -$, for each product projection $\pi_1 : C \times C' \rightarrow C$ in C , satisfying a Beck-Chevalley condition⁹: $!C' \otimes_C L = !C' \otimes_{C''} L$ and $C' \Rightarrow_C L = C' \Rightarrow_{C''} L$ for any $C, C'' \in \text{ob } C$. We simply write $!C' \otimes L$ and $C' \Rightarrow L$. We write Φ and Ψ for the natural isomorphisms $\mathcal{L}(C)(!C' \otimes L, L') \xrightarrow{\cong} \mathcal{L}(C \times C')(L, L')$ and $\mathcal{L}(C \times C')(L, L') \xrightarrow{\cong} \mathcal{L}(C)(L, C' \Rightarrow L')$.
- \mathcal{L} *supports Cartesian* \multimap -types: the functor $C^{op} \rightarrow \mathbf{Set}; C \mapsto \mathcal{L}(C)(L, L')$ is representable for any objects L, L' of \mathcal{L} . That is, we have objects $L \multimap L'$ of C with isomorphisms $\underline{\Lambda} : \mathcal{L}(C)(L, L') \xrightarrow{\cong} C(C, L \multimap L')$, natural in C .
- \mathcal{L} interprets primitive types and operations: we have a choice $\llbracket \mathbf{real}^n \rrbracket \in \text{ob } \mathcal{L}$ to interpret \mathbf{real}^n and, for each $\mathbf{lop} \in \mathbf{LOp}_{n_1, \dots, n_k; n'_1, \dots, n'_l}^{m_1, \dots, m_r}$, compatible \mathcal{L} -morphisms $\llbracket \mathbf{lop} \rrbracket$ in $\mathcal{L}(\llbracket \mathbf{real}^{n_1} \rrbracket \times \dots \times \llbracket \mathbf{real}^{n_k} \rrbracket)(\llbracket \mathbf{LDom}(\mathbf{lop}) \rrbracket, \llbracket \mathbf{CDom}(\mathbf{lop}) \rrbracket)$.

In particular, any biadditive model of intuitionistic linear/non-linear logic [6, 17, 35] is such a categorical model, as long as we choose interpretations for primitive types and operations.

⁸A locally C -indexed category \mathcal{L} can be equivalently defined as a category \mathcal{L} enriched over the presheaf category $[C^{op}, \mathbf{Set}]$. We prefer to consider locally indexed categories as special cases of indexed categories, instead, as CHAD's natural generalization to data types of varying dimension, such as unsized arrays or sum types, requires us to work with more general (non-locally) indexed categories [30].

⁹This condition says that the types $!C' \otimes_C L$ and $C' \Rightarrow_C L$ do not depend on C . We need to add this condition to match the syntax of the target language, in which copowers and powers only depend on two argument types.

Next, we turn to the interpretation of our target language in such models, which gives an operational intuition of the different components of a categorical model. We can interpret linear types $\underline{\tau}$ as objects $\llbracket \underline{\tau} \rrbracket$ of \mathcal{L} :

$$\llbracket \mathbf{1} \rrbracket \stackrel{\text{def}}{=} \mathbf{1} \quad \llbracket \underline{\tau} * \underline{\sigma} \rrbracket \stackrel{\text{def}}{=} \llbracket \underline{\tau} \rrbracket \times \llbracket \underline{\sigma} \rrbracket \quad \llbracket \underline{\tau} \rightarrow \underline{\sigma} \rrbracket \stackrel{\text{def}}{=} \llbracket \underline{\tau} \rrbracket \Rightarrow \llbracket \underline{\sigma} \rrbracket \quad \llbracket !\underline{\tau} \otimes \underline{\sigma} \rrbracket \stackrel{\text{def}}{=} !\llbracket \underline{\tau} \rrbracket \otimes \llbracket \underline{\sigma} \rrbracket.$$

We can interpret $\underline{\tau} \multimap \underline{\sigma}$ as the C -object $\llbracket \underline{\tau} \multimap \underline{\sigma} \rrbracket \stackrel{\text{def}}{=} \llbracket \underline{\tau} \rrbracket \multimap \llbracket \underline{\sigma} \rrbracket$. Finally, we can interpret terms $\Gamma \vdash t : \tau$ as morphisms $\llbracket t \rrbracket$ in $C(\llbracket \Gamma \rrbracket, \llbracket \tau \rrbracket)$ and terms $\Gamma; \underline{v} : \underline{\tau} \vdash t : \underline{\sigma}$ as $\llbracket t \rrbracket$ in $\mathcal{L}(\llbracket \Gamma \rrbracket)(\llbracket \underline{\tau} \rrbracket, \llbracket \underline{\sigma} \rrbracket)$:

$$\begin{aligned} \llbracket \text{lop}(t_1, \dots, t_k; s) \rrbracket &\stackrel{\text{def}}{=} \llbracket s \rrbracket; \mathcal{L}(\llbracket t_1 \rrbracket, \dots, \llbracket t_k \rrbracket)(\llbracket \text{lop} \rrbracket) \\ \llbracket \Gamma; \underline{v} : \underline{\tau} \vdash \underline{v} : \underline{\tau} \rrbracket &\stackrel{\text{def}}{=} \text{id}_{\llbracket \underline{\tau} \rrbracket} \quad \llbracket \text{let } x = t \text{ in } s \rrbracket \stackrel{\text{def}}{=} \mathcal{L}(\text{id}, \llbracket t \rrbracket)(\llbracket s \rrbracket) \quad \llbracket \text{let } \underline{v} = t \text{ in } s \rrbracket \stackrel{\text{def}}{=} \llbracket t \rrbracket; \llbracket s \rrbracket \\ \llbracket \langle \rangle \rrbracket &\stackrel{\text{def}}{=} () \quad \llbracket \langle t, s \rangle \rrbracket \stackrel{\text{def}}{=} (\llbracket t \rrbracket, \llbracket s \rrbracket) \quad \llbracket \text{fst } t \rrbracket \stackrel{\text{def}}{=} \llbracket t \rrbracket; \pi_1 \quad \llbracket \text{snd } t \rrbracket \stackrel{\text{def}}{=} \llbracket t \rrbracket; \pi_2 \\ \llbracket \lambda x. t \rrbracket &\stackrel{\text{def}}{=} \Psi(\llbracket t \rrbracket) \quad \llbracket t \ s \rrbracket \stackrel{\text{def}}{=} \mathcal{L}(\text{id}, \llbracket s \rrbracket)(\Psi^{-1}(\llbracket t \rrbracket)) \\ \llbracket !t \otimes s \rrbracket &\stackrel{\text{def}}{=} \mathcal{L}(\text{id}, \llbracket t \rrbracket)(\Phi(\text{id})); (!\llbracket \sigma \rrbracket \otimes \llbracket s \rrbracket) \quad \llbracket \text{case } t \text{ of } !y \otimes v \rightarrow s \rrbracket \stackrel{\text{def}}{=} \llbracket t \rrbracket; \Phi^{-1}(\llbracket s \rrbracket) \\ \llbracket \underline{\lambda} v. t \rrbracket &\stackrel{\text{def}}{=} \underline{\Delta}(\llbracket t \rrbracket) \quad \llbracket t \bullet s \rrbracket \stackrel{\text{def}}{=} \underline{\Delta}^{-1}(\llbracket t \rrbracket); \llbracket s \rrbracket \quad \llbracket \mathbf{0} \rrbracket \stackrel{\text{def}}{=} [] \quad \llbracket t + s \rrbracket \stackrel{\text{def}}{=} (\text{id}, \text{id}); \llbracket t \rrbracket, \llbracket s \rrbracket. \end{aligned}$$

Observe that we interpret $\mathbf{0}$ and $+$ using the biproduct structure of \mathcal{L} .

PROPOSITION 5.6. *The interpretation $\llbracket - \rrbracket$ of the language of Section 4 in categorical models is both sound and complete with respect to the $\beta\eta$ -equational theory: $t \stackrel{\beta\eta}{=} s$ iff $\llbracket t \rrbracket = \llbracket s \rrbracket$ in each such model.*

The proof is a minor variation of syntax-semantics correspondences developed in detail in chapters 3 and 5 of Reference [41], where we use the well-known result that finite products in a category are biproducts iff the category is enriched over commutative monoids [17]. Soundness follows by case analysis on the $\beta\eta$ -rules. Completeness follows by the construction of the syntactic model $\text{LSyn} : \text{CSyn}^{\text{op}} \rightarrow \text{Cat}$:

- CSyn extends its full subcategory Syn with Cartesian \multimap -types;
- Objects of $\text{LSyn}(\tau)$ are linear types $\underline{\sigma}$ of our target language.
- Morphisms in $\text{LSyn}(\tau)(\underline{\sigma}, \underline{\rho})$ are terms $x : \tau; \underline{v} : \underline{\sigma} \vdash t : \underline{\rho}$ modulo $(\alpha)\beta\eta$ -equivalence.
- Identities in $\text{LSyn}(\tau)$ are represented by the terms $x : \tau; \underline{v} : \underline{\sigma} \vdash \underline{v} : \underline{\sigma}$.
- Composition of $x : \tau; \underline{v} : \underline{\sigma}_1 \vdash t : \underline{\sigma}_2$ and $x : \tau; \underline{v} : \underline{\sigma}_2 \vdash s : \underline{\sigma}_3$ in $\text{LSyn}(\tau)$ is represented by $x : \tau; \underline{v} : \underline{\sigma}_1 \vdash \text{let } \underline{v} = t \text{ in } s : \underline{\sigma}_3$.
- Change of base $\text{LSyn}(t) : \text{LSyn}(\tau) \rightarrow \text{LSyn}(\tau')$ along $(x' : \tau' \vdash t : \tau) \in \text{CSyn}(\tau', \tau)$ is defined $\text{LSyn}(t)(x : \tau; \underline{v} : \underline{\sigma} \vdash s : \underline{\rho}) \stackrel{\text{def}}{=} x' : \tau'; \underline{v} : \underline{\sigma} \vdash \text{let } x = t \text{ in } s : \underline{\rho}$.
- All type formers are interpreted as one expects based on their notation, using introduction and elimination rules for the required structural isomorphisms.

5.3 Concrete Denotational Semantics

5.3.1 Sets and Commutative Monoids. Throughout this article, we have a particularly simple instance of the abstract semantics of our languages in mind, as we intend to interpret real^n as the usual Euclidean space \mathbb{R}^n (considered as a set) and to interpret each program $x_1 : \text{real}^{n_1}, \dots, x_k : \text{real}^{n_k} \vdash t : \text{real}^m$ as a function $\mathbb{R}^{n_1} \times \dots \times \mathbb{R}^{n_k} \rightarrow \mathbb{R}^m$. Similarly, we intend to interpret real^n as the commutative monoid \mathbb{R}^n and each program $x_1 : \text{real}^{n_1}, \dots, x_k : \text{real}^{n_k}; y : \text{real}^m \vdash t : \text{real}^r$ as a function $\mathbb{R}^{n_1} \times \dots \times \mathbb{R}^{n_k} \rightarrow \mathbb{R}^m \multimap \mathbb{R}^r$. That is, we will work with a concrete denotational semantics in terms of sets and commutative monoids.

Some readers will immediately recognize that the free-forgetful adjunction $\mathbf{Set} \rightleftarrows \mathbf{CMon}$ gives a model of full intuitionistic linear logic [35]. In fact, seeing that \mathbf{CMon} is \mathbf{CMon} -enriched, the model is biadditive [17].

However, we do not need such a rich type system. For us, the following suffices. Define $\mathbf{CMon}(X)$, for $X \in \mathbf{obSet}$, to have the objects of \mathbf{CMon} and homsets $\mathbf{CMon}(X)(Y, Z) \stackrel{\text{def}}{=} \mathbf{Set}(X, Y \multimap Z)$. Identities are defined as $x \mapsto (y \mapsto y)$ and composition $f;_{\mathbf{CMon}(X)} g$ is given by $x \mapsto (f(x);_{\mathbf{CMon}} g(x))$. Given $f \in \mathbf{Set}(X, X')$, we define change-of-base $\mathbf{CMon}(X') \rightarrow \mathbf{CMon}(X)$ as $\mathbf{CMon}(f)(g) \stackrel{\text{def}}{=} f;_{\mathbf{Set}} g$. $\mathbf{CMon}(-)$ defines a locally indexed category. By taking $C = \mathbf{Set}$ and $\mathcal{L}(-) = \mathbf{CMon}(-)$, we obtain a concrete instance of our abstract semantics. Indeed, we have natural isomorphisms

$$\begin{aligned} \mathbf{CMon}(X)(!X' \otimes Y, Z) &\xrightarrow{\Phi} \mathbf{CMon}(X \times X')(Y, Z) & \mathbf{CMon}(X \times X')(Y, Z) &\xrightarrow{\Psi} \mathbf{CMon}(X)(Y, X' \Rightarrow Z) \\ \Phi(f)(x, x')(y) &\stackrel{\text{def}}{=} f(x)(!x' \otimes y) & \Psi(f)(x)(y)(x') &\stackrel{\text{def}}{=} f(x, x')(y) \\ \Phi^{-1}(f)(x) \left(\sum_{i=1}^n (!x'_i \otimes y_i) \right) &\stackrel{\text{def}}{=} \sum_{i=1}^n f(x, x'_i)(y_i) & \Psi^{-1}(f)(x, x')(y) &\stackrel{\text{def}}{=} f(x)(y)(x'). \end{aligned}$$

The prime motivating examples of morphisms in this category are derivatives. Recall that the *derivative at x* , $Df(x)$, and *transposed derivative at x* , $Df^t(x)$, of a differentiable function $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$ are defined as the unique functions $Df(x) : \mathbb{R}^n \rightarrow \mathbb{R}^m$ and $Df^t(x) : \mathbb{R}^m \rightarrow \mathbb{R}^n$ satisfying

$$Df(x)(v) = \lim_{\delta \rightarrow 0} \frac{f(x + \delta \cdot v) - f(x)}{\delta} \quad Df^t(x)(w) \odot v = w \odot Df(x)(v),$$

where we write $v \odot v'$ for the inner product $\sum_{i=1}^n (\pi_i v) \cdot (\pi_i v')$ of vectors $v, v' \in \mathbb{R}^n$. Now, for differentiable $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$, Df and Df^t give maps in $\mathbf{CMon}(\mathbb{R}^n)(\mathbb{R}^n, \mathbb{R}^m)$ and $\mathbf{CMon}(\mathbb{R}^n)(\mathbb{R}^m, \mathbb{R}^n)$, respectively. Indeed, derivatives $Df(x)$ of f at x are linear functions, as are transposed derivatives $Df^t(x)$. Both depend differentiably on x in case f is differentiable. Note that the derivatives are not merely linear in the sense of preserving 0 and $+$. They are also multiplicative in the sense that $(Df)(x)(c \cdot v) = c \cdot (Df)(x)(v)$. We could have captured this property by working with vector spaces rather than commutative monoids. However, we will not need this property to phrase or establish correctness of AD. Therefore, we restrict our attention to the more straightforward structure of commutative monoids.

Defining $\llbracket \mathbf{real}^n \rrbracket \stackrel{\text{def}}{=} \mathbb{R}^n$ and interpreting each $\text{lop} \in \mathbf{LOp}$ as the (differentiable) function $\llbracket \text{lop} \rrbracket : (\mathbb{R}^{n_1} \times \dots \times \mathbb{R}^{n_k}) \rightarrow (\mathbb{R}^{n'_1} \times \dots \times \mathbb{R}^{n'_l}) \multimap (\mathbb{R}^{m_1} \times \dots \times \mathbb{R}^{m_r})$ it is intended to represent, we obtain a canonical interpretation of our target language in \mathbf{CMon} .

5.4 Operational Semantics

In this section, we describe an operational semantics for our source and target languages. We consider call-by-value evaluation, but similar results can be obtained for call-by-name evaluation.¹⁰ We present this semantics in big-step style. Finally, we show that our denotational semantics are adequate with respect to this operational semantics, showing that the denotational semantics are sound tools for reasoning about our programs.

¹⁰In fact, we conjecture our target language to be pure in the sense that reductions are confluent.

$(\Gamma; v : \underline{\text{real}}^{n_1} * \dots * \underline{\text{real}}^{n_k} \not\vdash v : \underline{\text{real}}^{n_1} * \dots * \underline{\text{real}}^{n_k})$	$(\Gamma; v : \underline{\text{real}}^{n_1} * \dots * \underline{\text{real}}^{n_k} \vdash v : \underline{\text{real}}^{n_1} * \dots * \underline{\text{real}}^{n_k})$	
$v \Downarrow v$	$v \Downarrow \text{id}(v)$	$\underline{lc}(v) \Downarrow \underline{lc}(v)$
$\frac{t \Downarrow v \quad s[v/x] \Downarrow w}{\text{let } x = t \text{ in } s \Downarrow w}$	$\frac{t_1 \Downarrow c_1 \quad \dots \quad t_n \Downarrow c_n}{\text{op}(t_1, \dots, t_n) \Downarrow \llbracket \text{op} \rrbracket(c_1, \dots, c_n)}$	$\frac{t \Downarrow v \quad s \Downarrow w}{\langle t, s \rangle \Downarrow \langle v, w \rangle} \quad \frac{t \Downarrow \langle v, w \rangle}{\text{fst } t \Downarrow v} \quad \frac{t \Downarrow \langle v, w \rangle}{\text{snd } t \Downarrow w}$
$\frac{t \Downarrow \lambda x. r \quad s \Downarrow v \quad r[v/x] \Downarrow w}{t \Downarrow s \Downarrow w}$	$\frac{t \Downarrow v \quad s[v/v] \Downarrow w}{\text{let } v = t \text{ in } s \Downarrow w}$	$\frac{t_1 \Downarrow c_1 \quad \dots \quad t_n \Downarrow c_n \quad s \Downarrow \underline{lc}(v)}{\text{lop}(t_1, \dots, t_n; s) \Downarrow \llbracket \text{lop} \rrbracket(c_1, \dots, c_n; \underline{lc}(v))} \quad \frac{t \Downarrow v \quad s \Downarrow w}{!t \otimes s \Downarrow !v \otimes w}$
$\frac{t \Downarrow !v_1 \otimes w_1 + (!v_2 \otimes w_2 + (\dots !v_n \otimes w_n) \dots) \quad \{s[v_i/y, w_i/v] \Downarrow u_i\}_{i=1}^n \quad u_1 + (u_2 + (\dots u_n) \dots) \Downarrow u}{\text{case } t \text{ of } !y \otimes v \rightarrow s \Downarrow u}$		
$\frac{t \Downarrow \lambda v. v \quad s \Downarrow w \quad v[w/v] \Downarrow u}{t \bullet s \Downarrow u}$	$(\Gamma; v : \underline{\text{real}}^{n_1} * \dots * \underline{\text{real}}^{n_k} \vdash 0 : \underline{\text{real}}^{m_1} * \dots * \underline{\text{real}}^{m_l})$	
	$0 \Downarrow 0(v)$	
$\frac{t_1 \Downarrow \underline{lc}_1(v) \quad t_2 \Downarrow \underline{lc}_2(v)}{t_1 + t_2 \Downarrow \underline{lc}(v)}$	$\frac{t_1 \Downarrow \langle \rangle \quad t_2 \Downarrow \langle \rangle}{t_1 + t_2 \Downarrow \langle \rangle}$	$\frac{t_1 \Downarrow \langle v_1, w_1 \rangle \quad t_2 \Downarrow \langle v_2, w_2 \rangle \quad v_1 + v_2 \Downarrow v \quad w_1 + w_2 \Downarrow w}{t_1 + t_2 \Downarrow \langle v, w \rangle}$
$\frac{t_1 \Downarrow \lambda x. s_1 \quad t_2 \Downarrow \lambda x. s_2}{t_1 + t_2 \Downarrow \lambda x. s_1 + s_2}$		
$\frac{t_1 \Downarrow !v_1 \otimes w_1 + (!v_2 \otimes w_2 + (\dots !v_n \otimes w_n) \dots) \quad t_2 \Downarrow !v_{n+1} \otimes w_{n+1} + (!v_{n+2} \otimes w_{n+2} + (\dots !v_{n+m} \otimes w_{n+m}) \dots)}{t_1 + t_2 \Downarrow !v_1 \otimes w_1 + (!v_2 \otimes w_2 + (\dots !v_{n+m} \otimes w_{n+m}) \dots)}$		

Fig. 7. The big-step call-by-value operational semantics $t \Downarrow v$ for the source and target languages. In the first rule, we intend to indicate that $v \Downarrow v$ unless v is a linear function between tuples of real arrays in the sense that $\Gamma; v : \underline{\text{real}}^{n_1} * \dots * \underline{\text{real}}^{n_k} \vdash v : \underline{\text{real}}^{n_1} * \dots * \underline{\text{real}}^{n_k}$.

We consider the following programs *values*, where we write \underline{c} for $\text{op}()$ and $\underline{lc}(v)$ for $\text{lop}(v)$:

$v, w, u ::=$	values	v
x		$\lambda x. t$
\underline{c}		$\lambda v. t$
$\langle \rangle$		$!v_1 \otimes w_1 + (!v_2 \otimes w_2 + (\dots !v_n \otimes w_n) \dots)$
$\langle v, w \rangle$		$\underline{lc}(v).$

We then define the big-step reduction relation $t \Downarrow v$, which says that a program t evaluates to the value v , in Figure 7. To be able to define this semantics, we assume that our languages contain, at least, nullary operations $\text{op} = \underline{c}$ for all constants $c \in \mathbb{R}^n$ and nullary linear operations \underline{lc} for all linear maps (matrices) $lc \in \mathbb{R}^n \multimap \mathbb{R}^m$. For all operations op and linear operations lop , we assume that an intended semantics $\llbracket \text{op} \rrbracket$ and $\llbracket \text{lop} \rrbracket$ is specified as (functions on) vectors of reals. As a side-note, we observe that this operational semantics has the following basic properties:

LEMMA 5.7 (SUBJECT REDUCTION, TERMINATION, DETERMINISM). *If $\Gamma \vdash t : \tau$, then there is a unique value v such that $t \Downarrow v$. Then, $\Gamma \vdash v : \tau$. Similarly, if $\Gamma; v : \underline{\tau} \vdash t : \underline{\sigma}$, then there is a unique value v , such that $t \Downarrow v$. Then, $\Gamma, v : \underline{\tau} \vdash v : \underline{\sigma}$.*

Subject reduction and termination are proved through a standard logical relations argument similar to the ones in Reference [39]. Determinism is observed by simply noting that all rules in the definition of \Downarrow have conclusions $t \Downarrow v$ with disjoint t .

In fact, seeing that every well-typed program t has a unique value v such that $t \Downarrow v$, we write $\Downarrow t$ for this v .

We assume that only first-order types are observable (i.e., have decidable equality on their values):

$$\begin{array}{ll}
 \phi, \psi ::= & \text{first-order Cartesian types} \\
 | & \mathbf{real}^n \\
 | & \mathbf{1} \\
 | & \phi * \psi \\
 \phi, \psi ::= & \text{first-order linear types} \\
 | & \mathbf{real}^n \\
 | & \mathbf{1} \\
 | & \phi * \psi.
 \end{array}$$

We define *program contexts* $C[_]$ to be programs $C[_]$ that use the variable $_$ exactly once. We call such program contexts *of first-order type* if they satisfy the typing judgement $_ : \tau \vdash C[_] : \phi$ for first-order Cartesian type ϕ or $_ : \tau; \nu : \underline{\psi} \vdash C[_] : \underline{\phi}$ for first-order linear types $\underline{\psi}$ and $\underline{\phi}$. We write $C[t]$ for the capturing substitution of t for $_$ in $C[_]$. This operational semantics and notion of observable types lead us to define *observational equivalence* (aka contextual equivalence) $t \approx s$ of programs $\cdot \vdash t, s : \tau$, where we say that $t \approx s$ holds if $\Downarrow C[t] = \Downarrow C[s]$ for all program contexts of first-order type. Similarly, we call two programs $\cdot; \nu : \underline{\tau} \vdash t, s : \underline{\sigma}$ of linear type observationally equivalent (write also $t \approx s$) if $\lambda \underline{v}. t \approx \lambda \underline{v}. s$.

Note that we consider values $\cdot; \nu : \underline{\psi} \vdash v : \underline{\phi}$ for first-order linear types $\underline{\psi}$ and $\underline{\phi}$ to be observable, seeing that linear functions between finite-dimensional spaces are finite-dimensional objects that can be fully observed by evaluating them on a (finite) basis for their domain type $\underline{\psi}$. Indeed, such values v are always of the form $\underline{lc}(v)$ for some $lc : \mathbb{R}^n \rightarrow \mathbb{R}^m$ hence are effectively matrices.

We first show two standard lemmas.

LEMMA 5.8 (COMPOSITIONALITY OF $\llbracket - \rrbracket$). *For any two terms $\Gamma \vdash t, s : \tau$ and any type-compatible program context $C[_]$ we have that $\llbracket t \rrbracket = \llbracket s \rrbracket$ implies $\llbracket C[t] \rrbracket = \llbracket C[s] \rrbracket$.*

This is proved by induction on the structure of terms.

LEMMA 5.9 (SOUNDNESS OF \Downarrow). *In case t , we have that $\llbracket t \rrbracket = \llbracket \Downarrow t \rrbracket$.*

This is proved by induction on the definition of \Downarrow : note that every operational rule is also an equation in the semantics. Then, adequacy follows.

THEOREM 5.10 (ADEQUACY). *In case $\llbracket t \rrbracket = \llbracket s \rrbracket$, it follows that $t \approx s$.*

PROOF. Suppose that $\llbracket t \rrbracket = \llbracket s \rrbracket$ and let $C[_]$ be a type-compatible program context of first-order type. Then, $\llbracket \Downarrow C[t] \rrbracket = \llbracket C[t] \rrbracket = \llbracket C[s] \rrbracket = \llbracket \Downarrow C[s] \rrbracket$ by the previous two lemmas. Finally, as values of observable types are easily seen to be faithfully (injectively) interpreted in our denotational semantics, it follows that $\Downarrow C[t] = \Downarrow C[s]$. Therefore, $t \approx s$. \square

That is, the denotational semantics is a sound means for proving observational equivalences of the operational semantics.

6 PAIRING PRIMALS WITH (CO)TANGENTS, CATEGORICALLY

In this section, we show that any categorical model $\mathcal{L} : C^{op} \rightarrow \mathbf{Cat}$ of our target language gives rise to two Cartesian closed categories $\Sigma_C \mathcal{L}$ and $\Sigma_C \mathcal{L}^{op}$. We believe that these observations of Cartesian closure are novel. Surprisingly, they are highly relevant for obtaining a principled understanding of AD on a higher-order language: the former for forward AD and the latter for reverse AD. Applying these constructions to the syntactic category $\mathbf{LSyn} : \mathbf{CSyn}^{op} \rightarrow \mathbf{Cat}$ of our target language, we produce a canonical definition of the AD macros, as the canonical interpretation of the λ -calculus in the Cartesian closed categories $\Sigma_{\mathbf{CSyn}} \mathbf{LSyn}$ and $\Sigma_{\mathbf{CSyn}} \mathbf{LSyn}^{op}$. In addition, when we apply this construction to the denotational semantics $\mathbf{CMon} : \mathbf{Set}^{op} \rightarrow \mathbf{Cat}$ and invoke a categorical logical relations technique, known as *subconing*, we find an elegant correctness proof of the source code transformations. The abstract construction delineated in this section is in many ways the theoretical crux of this article.

6.1 Grothendieck Constructions on Strictly Indexed Categories

Recall that for any strictly indexed category, i.e., a (strict) functor $\mathcal{L} : C^{op} \rightarrow \mathbf{Cat}$, we can consider its total category (or Grothendieck construction) $\Sigma_C \mathcal{L}$, which is a fibered category over C (see Sections A1.1.7 and B1.3.1 in Reference [23]). We can view it as a Σ -type of categories, which generalizes the Cartesian product. Concretely, its objects are pairs (A_1, A_2) of objects A_1 of C and A_2 of $\mathcal{L}(A_1)$. Its morphisms $(A_1, A_2) \rightarrow (B_1, B_2)$ are pairs (f_1, f_2) of a morphism $f_1 : A_1 \rightarrow B_1$ in C and a morphism $f_2 : A_2 \rightarrow \mathcal{L}(f_1)(B_2)$ in $\mathcal{L}(A_1)$. Identities are $\text{id}_{(A_1, A_2)} \stackrel{\text{def}}{=} (\text{id}_{A_1}, \text{id}_{A_2})$ and composition is $(f_1, f_2); (g_1, g_2) \stackrel{\text{def}}{=} (f_1; g_1, f_2; \mathcal{L}(f_1)(g_2))$. Furthermore, given a strictly indexed category $\mathcal{L} : C^{op} \rightarrow \mathbf{Cat}$, we can consider its fibrewise dual category $\mathcal{L}^{op} : C^{op} \rightarrow \mathbf{Cat}$, which is defined as the composition $C^{op} \xrightarrow{\mathcal{L}} \mathbf{Cat} \xrightarrow{op} \mathbf{Cat}$. Thus, we can apply the same construction to \mathcal{L}^{op} to obtain a category $\Sigma_C \mathcal{L}^{op}$.

6.2 Structure of $\Sigma_C \mathcal{L}$ and $\Sigma_C \mathcal{L}^{op}$ for Locally Indexed Categories

Section 6.1 applies, in particular, to the locally indexed categories of Section 5. In this case, we will analyze the categorical structure of $\Sigma_C \mathcal{L}$ and $\Sigma_C \mathcal{L}^{op}$. For reference, we first give a concrete description.

$\Sigma_C \mathcal{L}$ is the following category:

- objects are pairs (A_1, A_2) of objects A_1 of C and A_2 of \mathcal{L} ;
- morphisms $(A_1, A_2) \rightarrow (B_1, B_2)$ are pairs (f_1, f_2) with $f_1 : A_1 \rightarrow B_1 \in C$ and $f_2 : A_2 \rightarrow B_2 \in \mathcal{L}(A_1)$;
- composition of $(A_1, A_2) \xrightarrow{(f_1, f_2)} (B_1, B_2)$ and $(B_1, B_2) \xrightarrow{(g_1, g_2)} (C_1, C_2)$ is given by $(f_1; g_1, f_2; \mathcal{L}(f_1)(g_2))$ and identities $\text{id}_{(A_1, A_2)}$ are $(\text{id}_{A_1}, \text{id}_{A_2})$.

$\Sigma_C \mathcal{L}^{op}$ is the following category:

- objects are pairs (A_1, A_2) of objects A_1 of C and A_2 of \mathcal{L} ;
- morphisms $(A_1, A_2) \rightarrow (B_1, B_2)$ are pairs (f_1, f_2) with $f_1 : A_1 \rightarrow B_1 \in C$ and $f_2 : B_2 \rightarrow A_2 \in \mathcal{L}(A_1)$;
- composition of $(A_1, A_2) \xrightarrow{(f_1, f_2)} (B_1, B_2)$ and $(B_1, B_2) \xrightarrow{(g_1, g_2)} (C_1, C_2)$ is given by $(f_1; g_1, \mathcal{L}(f_1)(g_2); f_2)$ and identities $\text{id}_{(A_1, A_2)}$ are $(\text{id}_{A_1}, \text{id}_{A_2})$.

These categories are relevant to automatic differentiation for the following reason. Let us write \mathbf{CartSp} for the category of Cartesian spaces \mathbb{R}^n and differentiable functions between them. Observe that for any categorical model $\mathcal{L} : C^{op} \rightarrow \mathbf{Cat}$ of the target language,

$$\begin{aligned} \Sigma_C \mathcal{L}((A_1, A_2), (B_1, B_2)) &= C(A_1, B_1) \times \mathcal{L}(A_1)(A_2, B_2) \cong C(A_1, B_1 \times (A_2 \multimap B_2)) \\ \Sigma_C \mathcal{L}^{op}((A_1, A_2), (B_1, B_2)) &= C(A_1, B_1) \times \mathcal{L}(A_1)(B_2, A_2) \cong C(A_1, B_1 \times (B_2 \multimap A_2)). \end{aligned}$$

Then, observing that the composition in these Σ -types of categories is precisely the chain rule, we see that the paired-up derivative \mathcal{T} and transposed derivative \mathcal{T}^* of Section 2.2 define functors

$$\mathcal{T} : \mathbf{CartSp} \rightarrow \Sigma_{\mathbf{Set}} \mathbf{CMon} \qquad \mathcal{T}^* : \mathbf{CartSp} \rightarrow \Sigma_{\mathbf{Set}} \mathbf{CMon}^{op}.$$

As we will see in Section 7, we can implement (higher-order extensions of) these functors as code transformations

$$\vec{\mathcal{D}} : \mathbf{Syn} \rightarrow \Sigma_{\mathbf{CSyn}} \mathbf{LSyn} \qquad \overleftarrow{\mathcal{D}} : \mathbf{Syn} \rightarrow \Sigma_{\mathbf{CSyn}} \mathbf{LSyn}^{op}.$$

As we will see, we can derive these code transformations by examining the categorical structure present in $\Sigma_C \mathcal{L}$ and $\Sigma_C \mathcal{L}^{op}$ for categorical models $\mathcal{L} : C^{op} \rightarrow \mathbf{Cat}$ of the target language in the

sense of Section 5. We believe the existence of this categorical structure is a novel observation. We will make heavy use of it to define our AD algorithms and to prove them correct.

THEOREM 6.1. *For a categorical model $\mathcal{L} : \mathcal{C}^{op} \rightarrow \mathbf{Cat}$ of the target language, $\Sigma_C \mathcal{L}$ has:*

- *terminal object $\mathbb{1} = (\mathbb{1}, \mathbb{1})$ and binary products $(A_1, A_2) \times (B_1, B_2) = (A_1 \times B_1, A_2 \times B_2)$;*
- *exponentials $(A_1, A_2) \Rightarrow (B_1, B_2) = (A_1 \Rightarrow (B_1 \times (A_2 \multimap B_2)), A_1 \Rightarrow B_2)$.*

PROOF. We have (natural) bijections

$$\Sigma_C \mathcal{L}((A_1, A_2), (\mathbb{1}, \mathbb{1})) = C(A_1, \mathbb{1}) \times \mathcal{L}(A_1)(A_2, \mathbb{1}) \cong \mathbb{1} \times \mathbb{1} \cong \mathbb{1} \quad \{ \mathbb{1} \text{ terminal in } C \text{ and } \mathcal{L}(A_1) \}$$

$$\begin{aligned} \Sigma_C \mathcal{L}((A_1, A_2), (B_1 \times C_1, B_2 \times C_2)) &= C(A_1, B_1 \times C_1) \times \mathcal{L}(A_1)(A_2, B_2 \times C_2) \\ &\cong C(A_1, B_1) \times C(A_1, C_1) \times \mathcal{L}(A_1)(A_2, B_2) \times \mathcal{L}(A_1)(A_2, C_2) \quad \{ \times \text{ product in } C \text{ and } \mathcal{L}(A_1) \} \\ &\cong \Sigma_C \mathcal{L}((A_1, A_2), (B_1, B_2)) \times \Sigma_C \mathcal{L}((A_1, A_2), (C_1, C_2)) \end{aligned}$$

$$\begin{aligned} \Sigma_C \mathcal{L}((A_1, A_2) \times (B_1, B_2), (C_1, C_2)) &= \Sigma_C \mathcal{L}((A_1 \times B_1, A_2 \times B_2), (C_1, C_2)) \\ &= C(A_1 \times B_1, C_1) \times \mathcal{L}(A_1 \times B_1)(A_2 \times B_2, C_2) \\ &\cong C(A_1 \times B_1, C_1) \times \mathcal{L}(A_1 \times B_1)(A_2, C_2) \times \mathcal{L}(A_1 \times B_1)(B_2, C_2) \quad \{ \times \text{ coproducts in } \mathcal{L}(A_1 \times B_1) \} \\ &\cong C(A_1 \times B_1, C_1) \times \mathcal{L}(A_1)(A_2, B_1 \Rightarrow C_2) \times \mathcal{L}(A_1 \times B_1)(B_2, C_2) \quad \{ \Rightarrow\text{-types in } \mathcal{L} \} \\ &\cong C(A_1 \times B_1, C_1) \times \mathcal{L}(A_1)(A_2, B_1 \Rightarrow C_2) \times C(A_1 \times B_1, B_2 \multimap C_2) \quad \{ \text{Cartesian } \multimap\text{-types} \} \\ &\cong C(A_1 \times B_1, C_1 \times (B_2 \multimap C_2)) \times \mathcal{L}(A_1)(A_2, B_1 \Rightarrow C_2) \quad \{ \times \text{ is product in } C \} \\ &\cong C(A_1, B_1 \Rightarrow (C_1 \times (B_2 \multimap C_2))) \times \mathcal{L}(A_1)(A_2, B_1 \Rightarrow C_2) \quad \{ \Rightarrow \text{ is exponential in } C \} \\ &= \Sigma_C \mathcal{L}((A_1, A_2), (B_1 \Rightarrow (C_1 \times (B_2 \multimap C_2)), B_1 \Rightarrow C_2)) \\ &= \Sigma_C \mathcal{L}((A_1, A_2), (B_1, B_2) \Rightarrow (C_1, C_2)). \end{aligned} \quad \square$$

We observe that we need \mathcal{L} to have biproducts (equivalently, to be **CMon** enriched) to show Cartesian closure. Furthermore, we need linear \Rightarrow -types and Cartesian \multimap -types to construct exponentials. Codually, we also obtain the Cartesian closure of $\Sigma_C \mathcal{L}^{op}$. However, for concreteness, we write out the proof by hand.

THEOREM 6.2. *For a categorical model $\mathcal{L} : \mathcal{C}^{op} \rightarrow \mathbf{Cat}$ of the target language, $\Sigma_C \mathcal{L}^{op}$ has*

- *terminal object $\mathbb{1} = (\mathbb{1}, \mathbb{1})$ and binary products $(A_1, A_2) \times (B_1, B_2) = (A_1 \times B_1, A_2 \times B_2)$;*
- *exponentials $(A_1, A_2) \Rightarrow (B_1, B_2) = (A_1 \Rightarrow (B_1 \times (B_2 \multimap A_2)), !A_1 \otimes B_2)$.*

PROOF. We have (natural) bijections

$$\Sigma_C \mathcal{L}^{op}((A_1, A_2), (\mathbb{1}, \mathbb{1})) = C(A_1, \mathbb{1}) \times \mathcal{L}(A_1)(\mathbb{1}, A_2) \cong \mathbb{1} \times \mathbb{1} \cong \mathbb{1} \quad \{ \mathbb{1} \text{ terminal in } C, \text{ initial in } \mathcal{L}(A_1) \}$$

$$\begin{aligned} \Sigma_C \mathcal{L}^{op}((A_1, A_2), (B_1 \times C_1, B_2 \times C_2)) &= C(A_1, B_1 \times C_1) \times \mathcal{L}(A_1)(B_2 \times C_2, A_2) \\ &\cong C(A_1, B_1) \times C(A_1, C_1) \times \mathcal{L}(A_1)(B_2, A_2) \times \mathcal{L}(A_1)(C_2, A_2) \quad \{ \times \text{ product in } C, \text{ coproduct in } \mathcal{L}(A_1) \} \\ &= \Sigma_C \mathcal{L}^{op}((A_1, A_2), (B_1, B_2)) \times \Sigma_C \mathcal{L}^{op}((A_1, A_2), (C_1, C_2)) \end{aligned}$$

$$\begin{aligned} \Sigma_C \mathcal{L}^{op}((A_1, A_2) \times (B_1, B_2), (C_1, C_2)) &= \Sigma_C \mathcal{L}^{op}((A_1 \times B_1, A_2 \times B_2), (C_1, C_2)) \\ &= C(A_1 \times B_1, C_1) \times \mathcal{L}(A_1 \times B_1)(C_2, A_2 \times B_2) \\ &\cong C(A_1 \times B_1, C_1) \times \mathcal{L}(A_1 \times B_1)(C_2, A_2) \times \mathcal{L}(A_1 \times B_1)(C_2, B_2) \quad \{ \times \text{ is product in } \mathcal{L}(A_1 \times B_1) \} \\ &\cong C(A_1 \times B_1, C_1) \times C(A_1 \times B_1, C_2 \multimap B_2) \times \mathcal{L}(A_1 \times B_1)(C_2, A_2) \quad \{ \text{Cartesian } \multimap\text{-types} \} \\ &\cong C(A_1 \times B_1, C_1 \times (C_2 \multimap B_2)) \times \mathcal{L}(A_1 \times B_1)(C_2, A_2) \quad \{ \times \text{ is product in } C \} \end{aligned}$$

$$\begin{aligned}
&\cong C(A_1, B_1 \Rightarrow (C_1 \times (C_2 \multimap B_2))) \times \mathcal{L}(A_1 \times B_1)(C_2, A_2) & \{ \quad \Rightarrow \text{is exponential in } C \quad \} \\
&\cong C(A_1, B_1 \Rightarrow (C_1 \times (C_2 \multimap B_2))) \times \mathcal{L}(A_1)(!B_1 \otimes C_2, A_2) & \{ \quad !(-) \otimes (-)\text{-types} \quad \} \\
&= \Sigma_C \mathcal{L}^{op}((A_1, A_2), (B_1 \Rightarrow (C_1 \times (C_2 \multimap B_2)), !B_1 \otimes C_2)) \\
&= \Sigma_C \mathcal{L}^{op}((A_1, A_2), (B_1, B_2) \Rightarrow (C_1, C_2)). \quad \square
\end{aligned}$$

Observe that we need the biproduct structure of \mathcal{L} to construct finite products in $\Sigma_C \mathcal{L}^{op}$. Furthermore, we need Cartesian \multimap -types and $!(-) \otimes (-)$ -types, but not biproducts, to construct exponentials.

Interestingly, we observe that the exponentials in $\Sigma_C \mathcal{L}$ and $\Sigma_C \mathcal{L}^{op}$ are not fibered over C (unlike their products, for example). Indeed $(A_1, A_2) \Rightarrow (B_1, B_2)$ has first component not equal to $A_1 \Rightarrow B_1$. In the context of automatic differentiation, this has the consequence that primals associated with values f of function type are not equal to f itself. Instead, as we will see, they include both a copy of f and a copy of its (transposed) derivative. These primals at higher-order types can be contrasted with the situation at first-order types, where values are equal to their associated primal, as a result of the finite products being fibered.

7 NOVEL AD ALGORITHMS AS SOURCE-CODE TRANSFORMATIONS

As $\Sigma_{\text{CSyn}} \text{LSyn}$ and $\Sigma_{\text{CSyn}} \text{LSyn}^{op}$ are both Cartesian closed categories by Theorems 6.1 and 6.2, the universal property of the source language (Proposition 3.1) gives us the following definition of forward and reverse mode CHAD as a canonical homomorphic functors.

COROLLARY 7.1 (CANONICAL DEFINITION OF CHAD). *Once we fix compatible definitions $\vec{\mathcal{D}}(\text{real}^n)$ and $\vec{\mathcal{D}}(\text{op})$ (resp. $\overleftarrow{\mathcal{D}}(\text{real}^n)$ and $\overleftarrow{\mathcal{D}}(\text{op})$), we obtain a unique structure-preserving functor*

$$\vec{\mathcal{D}}(-) : \text{Syn} \rightarrow \Sigma_{\text{CSyn}} \text{LSyn} \quad (\text{resp. } \overleftarrow{\mathcal{D}}(-) : \text{Syn} \rightarrow \Sigma_{\text{CSyn}} \text{LSyn}^{op}).$$

In this section, we discuss

- the interpretation of the above functors as a type-respecting code transformation;
- how we can define the basic definitions $\vec{\mathcal{D}}(\text{real}^n)$, $\overleftarrow{\mathcal{D}}(\text{real}^n)$, $\vec{\mathcal{D}}(\text{op})$ and $\overleftarrow{\mathcal{D}}(\text{op})$;
- what the induced AD definitions $\vec{\mathcal{D}}(t)$ and $\overleftarrow{\mathcal{D}}(t)$ are for arbitrary source language programs t ;
- some consequences of the sharing of subexpressions that we have employed when defining the code transformations.

7.1 Some Notation

In the rest of this section, we use the following syntactic sugar:

- a notation for (linear) n -ary tuple types: $(\underline{\tau}_1 * \dots * \underline{\tau}_n) \stackrel{\text{def}}{=} (((\underline{\tau}_1 * \underline{\tau}_2) \dots * \underline{\tau}_{n-1}) * \underline{\tau}_n)$;
- a notation for n -ary tuples: $\langle t_1, \dots, t_n \rangle \stackrel{\text{def}}{=} \langle \langle \langle t_1, t_2 \rangle \dots, t_{n-1} \rangle, t_n \rangle$;
- given $\Gamma; \nu : \underline{\tau} \vdash t : (\underline{\sigma}_1 * \dots * \underline{\sigma}_n)$, we write $\Gamma; \nu : \underline{\tau} \vdash \text{proj}_i(t) : \underline{\sigma}_i$ for the obvious i th projection of t , which is constructed by repeatedly applying **fst** and **snd** to t ;
- given $\Gamma; \nu : \underline{\tau} \vdash t : \underline{\sigma}_i$, we write the i th coprojection $\Gamma; \nu : \underline{\tau} \vdash \text{coproj}_i(t) \stackrel{\text{def}}{=} \langle \underline{0}, \dots, \underline{0}, t, \underline{0}, \dots, \underline{0} \rangle : (\underline{\sigma}_1 * \dots * \underline{\sigma}_n)$;
- for a list x_1, \dots, x_n of distinct identifiers, we write $\text{idx}(x_i; x_1, \dots, x_n) \stackrel{\text{def}}{=} i$ for the index of the identifier x_i in this list;
- a **let**-binding for tuples: $\text{let } \langle x, y \rangle = t \text{ in } s \stackrel{\text{def}}{=} \text{let } z = t \text{ in let } x = \text{fst } z \text{ in let } y = \text{snd } z \text{ in } s$, where z is a fresh variable.

Furthermore, all variables used in the source code transformations below are assumed to be freshly chosen.

7.2 $\vec{\mathcal{D}}(-)$ and $\overleftarrow{\mathcal{D}}(-)$ as Type-Respecting Code Transformations

Writing out the definitions of the categories \mathbf{Syn} , $\Sigma_{\mathbf{CSyn}}\mathbf{LSyn}$, $\Sigma_{\mathbf{CSyn}}\mathbf{LSyn}^{op}$, $\vec{\mathcal{D}}(-)$ and $\overleftarrow{\mathcal{D}}(-)$ provide for each type τ of the source language (Section 3) the following types in the target language (Section 4):

- a Cartesian type $\vec{\mathcal{D}}(\tau)_1$ of forward mode primals;
- a linear type $\vec{\mathcal{D}}(\tau)_2$ of forward mode tangents;
- a Cartesian type $\overleftarrow{\mathcal{D}}(\tau)_1$ of reverse mode primals;
- a linear type $\overleftarrow{\mathcal{D}}(\tau)_2$ of reverse mode cotangents.

We can extend the actions of $\vec{\mathcal{D}}(-)$ and $\overleftarrow{\mathcal{D}}(-)$ to typing contexts $\Gamma = x_1 : \tau_1, \dots, x_n : \tau_n$ as

$$\begin{aligned} \vec{\mathcal{D}}(\Gamma)_1 &\stackrel{\text{def}}{=} x_1 : \vec{\mathcal{D}}(\tau_1)_1, \dots, x_n : \vec{\mathcal{D}}(\tau_n)_1 && \text{(a Cartesian typing context)} \\ \vec{\mathcal{D}}(\Gamma)_2 &\stackrel{\text{def}}{=} (\vec{\mathcal{D}}(\tau_1)_2 * \dots * \vec{\mathcal{D}}(\tau_n)_2) && \text{(a linear type)} \\ \overleftarrow{\mathcal{D}}(\Gamma)_1 &\stackrel{\text{def}}{=} x_1 : \overleftarrow{\mathcal{D}}(\tau_1)_1, \dots, x_n : \overleftarrow{\mathcal{D}}(\tau_n)_1 && \text{(a Cartesian typing context)} \\ \overleftarrow{\mathcal{D}}(\Gamma)_2 &\stackrel{\text{def}}{=} (\overleftarrow{\mathcal{D}}(\tau_1)_2 * \dots * \overleftarrow{\mathcal{D}}(\tau_n)_2) && \text{(a linear type).} \end{aligned}$$

Similarly, $\vec{\mathcal{D}}(-)$ and $\overleftarrow{\mathcal{D}}(-)$ associate to each source language program $\Gamma \vdash t : \tau$ the following programs in the target language (Section 4):

- a forward mode primal computation $\vec{\mathcal{D}}(\Gamma)_1 \vdash \vec{\mathcal{D}}_{\Gamma}(t)_1 : \vec{\mathcal{D}}(\tau)_1$;
- a forward mode tangent computation $\vec{\mathcal{D}}(\Gamma)_1; v : \vec{\mathcal{D}}(\Gamma)_2 \vdash \vec{\mathcal{D}}_{\Gamma}(t)_2 : \vec{\mathcal{D}}(\tau)_2$;
- a reverse mode primal computation $\overleftarrow{\mathcal{D}}(\Gamma)_1 \vdash \overleftarrow{\mathcal{D}}_{\Gamma}(t)_1 : \overleftarrow{\mathcal{D}}(\tau)_1$;
- a reverse mode cotangent computation $\overleftarrow{\mathcal{D}}(\Gamma)_1; v : \overleftarrow{\mathcal{D}}(\Gamma)_2 \vdash \overleftarrow{\mathcal{D}}_{\Gamma}(t)_2 : \overleftarrow{\mathcal{D}}(\tau)_2$.

Here, we write $\bar{\Gamma}$ for the list of identifiers x_1, \dots, x_n that occur in the typing context $\Gamma = x_1 : \tau_1, \dots, x_n : \tau_n$. As we will see, we need to know these context identifiers to define the code transformation. Equivalently, we can pair up the primal and (co)tangent computations as

- a combined forward mode primal and tangent computation $\vec{\mathcal{D}}(\Gamma)_1 \vdash \vec{\mathcal{D}}_{\Gamma}(t) : \vec{\mathcal{D}}(\tau)_1 * (\vec{\mathcal{D}}(\Gamma)_2 \multimap \vec{\mathcal{D}}(\tau)_2)$, where $\vec{\mathcal{D}}_{\Gamma}(t) \stackrel{\beta\eta^+}{=} \langle \vec{\mathcal{D}}_{\Gamma}(t)_1, \underline{\lambda}v. \vec{\mathcal{D}}_{\Gamma}(t)_2 \rangle$;
- a combined reverse mode primal and cotangent computation $\overleftarrow{\mathcal{D}}(\Gamma)_1 \vdash \overleftarrow{\mathcal{D}}_{\Gamma}(t) : \overleftarrow{\mathcal{D}}(\tau)_1 * (\overleftarrow{\mathcal{D}}(\Gamma)_2 \multimap \overleftarrow{\mathcal{D}}(\tau)_2)$, where $\overleftarrow{\mathcal{D}}_{\Gamma}(t) \stackrel{\beta\eta^+}{=} \langle \overleftarrow{\mathcal{D}}_{\Gamma}(t)_1, \underline{\lambda}v. \overleftarrow{\mathcal{D}}_{\Gamma}(t)_2 \rangle$.

We prefer to work with these combined primal and (co)tangent code transformations as this allows us to share common subexpressions between the primal and (co)tangent computations using **let**-bindings. Indeed, note that the universal property of \mathbf{Syn} only defines the code transformations $\vec{\mathcal{D}}(-)$ and $\overleftarrow{\mathcal{D}}(-)$ up to $\stackrel{\beta\eta^+}{=}$. In writing down the definitions of CHAD on programs, we make sure to choose sensible representatives of these $\beta\eta^+$ -equivalence classes: ones that share common subexpressions through **let**-bindings. While these **let**-bindings naturally do not affect correctness of the transformation, they let us avoid code explosion at compile-time and unnecessary recomputation at runtime.

Finally, we note that, due to their definition from a universal property, our code transformations automatically respect equational reasoning in the sense that $\Gamma \vdash t \stackrel{\beta\eta}{=} s : \tau$ implies that $\vec{\mathcal{D}}_{\Gamma}(t) \stackrel{\beta\eta^+}{=} \vec{\mathcal{D}}_{\Gamma}(s)$ and $\overleftarrow{\mathcal{D}}_{\Gamma}(t) \stackrel{\beta\eta^+}{=} \overleftarrow{\mathcal{D}}_{\Gamma}(s)$.

7.3 The Basic Definitions: $\vec{\mathcal{D}}(\text{real}^n)$, $\overleftarrow{\mathcal{D}}(\text{real}^n)$, $\vec{\mathcal{D}}(\text{op})$, and $\overleftarrow{\mathcal{D}}(\text{op})$

In Section 4, we have assumed that there are suitable terms (for example, linear operations)

$$\begin{aligned} x_1 : \text{real}^{n_1}, \dots, x_k : \text{real}^{n_k} \ ; \ v : \underline{\text{real}}^{n_1} * \dots * \underline{\text{real}}^{n_k} &\vdash \text{Dop}(x_1, \dots, x_k; v) : \underline{\text{real}}^m \\ x_1 : \text{real}^{n_1}, \dots, x_k : \text{real}^{n_k} \ ; \ v : \underline{\text{real}}^m &\vdash \text{Dop}^t(x_1, \dots, x_k; v) : \underline{\text{real}}^{n_1} * \dots * \underline{\text{real}}^{n_k} \end{aligned}$$

to represent the forward and reverse mode derivatives of the primitive operations $\text{op} \in \text{Op}_{n_1, \dots, n_k}^m$. Using these, we define

$$\begin{aligned} \vec{\mathcal{D}}(\text{real}^n)_1 &\stackrel{\text{def}}{=} \text{real}^n \\ \vec{\mathcal{D}}(\text{real}^n)_2 &\stackrel{\text{def}}{=} \underline{\text{real}}^n \\ \vec{\mathcal{D}}_{\Gamma}(\text{op}(t_1, \dots, t_k)) &\stackrel{\text{def}}{=} \text{let } \langle x_1, x'_1 \rangle = \vec{\mathcal{D}}_{\Gamma}(t_1) \text{ in } \dots \text{let } \langle x_k, x'_k \rangle = \vec{\mathcal{D}}_{\Gamma}(t_k) \text{ in} \\ &\quad \langle \text{op}(x_1, \dots, x_k), \underline{\lambda} v. \text{Dop}(x_1, \dots, x_k; \langle x'_1 \bullet v, \dots, x'_k \bullet v \rangle) \rangle \\ \overleftarrow{\mathcal{D}}(\text{real}^n)_1 &\stackrel{\text{def}}{=} \text{real}^n \\ \overleftarrow{\mathcal{D}}(\text{real}^n)_2 &\stackrel{\text{def}}{=} \underline{\text{real}}^n \\ \overleftarrow{\mathcal{D}}_{\Gamma}(\text{op}(t_1, \dots, t_k)) &\stackrel{\text{def}}{=} \text{let } \langle x_1, x'_1 \rangle = \overleftarrow{\mathcal{D}}_{\Gamma}(t_1) \text{ in } \dots \text{let } \langle x_k, x'_k \rangle = \overleftarrow{\mathcal{D}}_{\Gamma}(t_k) \text{ in} \\ &\quad \langle \text{op}(x_1, \dots, x_k), \underline{\lambda} v. \text{let } v = \text{Dop}^t(x_1, \dots, x_k; v) \text{ in } x'_1 \bullet (\text{proj}_1 v) + \dots + x'_k \bullet (\text{proj}_k v) \rangle \end{aligned}$$

These basic definitions of CHAD for primitive operations implement the well-known multivariate chain rules for (transposed) derivatives of Section 2.2.

For the AD transformations to be correct, it is important that these derivatives of language primitives are implemented correctly in the sense that

$$\llbracket x_1, \dots, x_k; y \vdash \text{Dop}(x_1, \dots, x_k; v) \rrbracket = D[\llbracket \text{op} \rrbracket] \quad \llbracket x_1, \dots, x_k; v \vdash \text{Dop}^t(x_1, \dots, x_k; v) \rrbracket = D[\llbracket \text{op} \rrbracket]^t.$$

For example, for elementwise multiplication $(*) \in \text{Op}_{n,n}^n$, which we interpret as the usual elementwise product $\llbracket (*) \rrbracket \stackrel{\text{def}}{=} (*) : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}^n$, we need, by the product rule for differentiation, that

$$\begin{aligned} \llbracket D(*) (x_1, x_2; v) \rrbracket ((a_1, a_2), (b_1, b_2)) &= a_1 * b_2 + a_2 * b_1 \\ \llbracket D(*)^t (x_1, x_2; v) \rrbracket ((a_1, a_2), b) &= (a_2 * b, a_1 * b). \end{aligned}$$

By Proposition 3.1, the extension of the AD transformations $\vec{\mathcal{D}}$ and $\overleftarrow{\mathcal{D}}$ to the full source language are now canonically determined, as the unique Cartesian closed functors that extend these basic definitions.

7.4 The Implied Forward Mode CHAD Definitions

We define the types of (forward mode) primals $\vec{\mathcal{D}}(\tau)_1$ and tangents $\vec{\mathcal{D}}(\tau)_2$ associated with a type τ as follows:

$$\begin{aligned} \vec{\mathcal{D}}(1)_1 &\stackrel{\text{def}}{=} 1 & \vec{\mathcal{D}}(1)_2 &\stackrel{\text{def}}{=} \underline{1} \\ \vec{\mathcal{D}}(\tau * \sigma)_1 &\stackrel{\text{def}}{=} \vec{\mathcal{D}}(\tau)_1 * \vec{\mathcal{D}}(\sigma)_1 & \vec{\mathcal{D}}(\tau * \sigma)_2 &\stackrel{\text{def}}{=} \vec{\mathcal{D}}(\tau)_2 * \vec{\mathcal{D}}(\sigma)_2 \\ \vec{\mathcal{D}}(\tau \rightarrow \sigma)_1 &\stackrel{\text{def}}{=} \vec{\mathcal{D}}(\tau)_1 \rightarrow (\vec{\mathcal{D}}(\sigma)_1 * (\vec{\mathcal{D}}(\tau)_2 \multimap \vec{\mathcal{D}}(\sigma)_2)) & \vec{\mathcal{D}}(\tau \rightarrow \sigma)_2 &\stackrel{\text{def}}{=} \vec{\mathcal{D}}(\tau)_1 \rightarrow \vec{\mathcal{D}}(\sigma)_2. \end{aligned}$$

Observe that the type of primals associated with a function type is not equal to the original type. This is a consequence of the non-fibered nature of the exponentials in $\Sigma\text{CSynLSyn}$ of categories (Section 6).

For programs t , we define their efficient CHAD transformation $\vec{\mathcal{D}}_{\overline{\Gamma}}(t)$ as follows:

$$\begin{aligned}
\vec{\mathcal{D}}_{\overline{\Gamma}}(x) &\stackrel{\text{def}}{=} \langle x, \underline{\lambda}v. \text{proj}_{\text{idx}(x;\overline{\Gamma})}(v) \rangle \\
\vec{\mathcal{D}}_{\overline{\Gamma}}(\text{let } x = t \text{ in } s) &\stackrel{\text{def}}{=} \text{let } \langle x, x' \rangle = \vec{\mathcal{D}}_{\overline{\Gamma}}(t) \text{ in let } \langle y, y' \rangle = \vec{\mathcal{D}}_{\overline{\Gamma},x}(s) \text{ in } \langle y, \underline{\lambda}v. y' \bullet \langle x', x' \bullet v \rangle \rangle \\
\vec{\mathcal{D}}_{\overline{\Gamma}}(\langle \rangle) &\stackrel{\text{def}}{=} \langle \langle \rangle, \underline{\lambda}v. \langle \rangle \rangle \\
\vec{\mathcal{D}}_{\overline{\Gamma}}(\langle t, s \rangle) &\stackrel{\text{def}}{=} \text{let } \langle x, x' \rangle = \vec{\mathcal{D}}_{\overline{\Gamma}}(t) \text{ in let } \langle y, y' \rangle = \vec{\mathcal{D}}_{\overline{\Gamma}}(s) \text{ in } \langle \langle x, y \rangle, \underline{\lambda}v. \langle x' \bullet v, y' \bullet v \rangle \rangle \\
\vec{\mathcal{D}}_{\overline{\Gamma}}(\text{fst } t) &\stackrel{\text{def}}{=} \text{let } \langle x, x' \rangle = \vec{\mathcal{D}}_{\overline{\Gamma}}(t) \text{ in } \langle \text{fst } x, \underline{\lambda}v. \text{fst}(x' \bullet v) \rangle \\
\vec{\mathcal{D}}_{\overline{\Gamma}}(\text{snd } t) &\stackrel{\text{def}}{=} \text{let } \langle x, x' \rangle = \vec{\mathcal{D}}_{\overline{\Gamma}}(t) \text{ in } \langle \text{snd } x, \underline{\lambda}v. \text{snd}(x' \bullet v) \rangle \\
\vec{\mathcal{D}}_{\overline{\Gamma}}(\lambda x. t) &\stackrel{\text{def}}{=} \text{let } y = \lambda x. \vec{\mathcal{D}}_{\overline{\Gamma},x}(t) \text{ in } \langle \lambda x. \text{let } \langle z, z' \rangle = y \ x \text{ in } \langle z, \underline{\lambda}v. z' \bullet \langle \underline{0}, v \rangle \rangle, \underline{\lambda}v. \lambda x. (\text{snd}(y \ x)) \bullet \langle v, \underline{0} \rangle \rangle \\
\vec{\mathcal{D}}_{\overline{\Gamma}}(t \ s) &\stackrel{\text{def}}{=} \text{let } \langle x, x_{\text{ctx}} \rangle = \vec{\mathcal{D}}_{\overline{\Gamma}}(t) \text{ in let } \langle y, y' \rangle = \vec{\mathcal{D}}_{\overline{\Gamma}}(s) \text{ in let } \langle z, x'_{\text{arg}} \rangle = x \ y \text{ in} \\
&\quad \langle z, \underline{\lambda}v. (x'_{\text{ctx}} \bullet v) \ y + x'_{\text{arg}} \bullet (y' \bullet v) \rangle.
\end{aligned}$$

We explain and justify these transformations in the next subsection after discussing the transformations for reverse CHAD.

7.5 The Implied Reverse Mode CHAD Definitions

We define the types of (reverse mode) primals $\overleftarrow{\mathcal{D}}(\tau)_1$ and cotangents $\overleftarrow{\mathcal{D}}(\tau)_2$ associated with a type τ as follows:

$$\begin{aligned}
\overleftarrow{\mathcal{D}}(1)_1 &\stackrel{\text{def}}{=} 1 & \overleftarrow{\mathcal{D}}(1)_2 &\stackrel{\text{def}}{=} \underline{1} \\
\overleftarrow{\mathcal{D}}(\tau * \sigma)_1 &\stackrel{\text{def}}{=} \overleftarrow{\mathcal{D}}(\tau)_1 * \overleftarrow{\mathcal{D}}(\sigma)_1 & \overleftarrow{\mathcal{D}}(\tau * \sigma)_2 &\stackrel{\text{def}}{=} \overleftarrow{\mathcal{D}}(\tau)_2 * \overleftarrow{\mathcal{D}}(\sigma)_2 \\
\overleftarrow{\mathcal{D}}(\tau \rightarrow \sigma)_1 &\stackrel{\text{def}}{=} \overleftarrow{\mathcal{D}}(\tau)_1 \rightarrow (\overleftarrow{\mathcal{D}}(\sigma)_1 * (\overleftarrow{\mathcal{D}}(\sigma)_2 \multimap \overleftarrow{\mathcal{D}}(\tau)_2)) & \overleftarrow{\mathcal{D}}(\tau \rightarrow \sigma)_2 &\stackrel{\text{def}}{=} !\overleftarrow{\mathcal{D}}(\tau)_1 \otimes \overleftarrow{\mathcal{D}}(\sigma)_2.
\end{aligned}$$

Again, we associate a non-trivial type of primals to function types as exponentials are not fibered in $\Sigma\text{CSynLSyn}^{op}$ (Section 6).

For programs t , we define their efficient CHAD transformation $\overleftarrow{\mathcal{D}}_{\overline{\Gamma}}(t)$ as follows:

$$\begin{aligned}
\overleftarrow{\mathcal{D}}_{\overline{\Gamma}}(x) &\stackrel{\text{def}}{=} \langle x, \underline{\lambda}v. \text{coproj}_{\text{idx}(x;\overline{\Gamma})}(v) \rangle \\
\overleftarrow{\mathcal{D}}_{\overline{\Gamma}}(\text{let } x = t \text{ in } s) &\stackrel{\text{def}}{=} \text{let } \langle x, x' \rangle = \overleftarrow{\mathcal{D}}_{\overline{\Gamma}}(t) \text{ in let } \langle y, y' \rangle = \overleftarrow{\mathcal{D}}_{\overline{\Gamma},x}(s) \text{ in } \langle y, \underline{\lambda}v. \text{let } v = y' \bullet v \text{ in fst } v + x' \bullet (\text{snd } v) \rangle \\
\overleftarrow{\mathcal{D}}_{\overline{\Gamma}}(\langle \rangle) &\stackrel{\text{def}}{=} \langle \langle \rangle, \underline{\lambda}v. \underline{0} \rangle \\
\overleftarrow{\mathcal{D}}_{\overline{\Gamma}}(\langle t, s \rangle) &\stackrel{\text{def}}{=} \text{let } \langle x, x' \rangle = \overleftarrow{\mathcal{D}}_{\overline{\Gamma}}(t) \text{ in let } \langle y, y' \rangle = \overleftarrow{\mathcal{D}}_{\overline{\Gamma}}(s) \text{ in } \langle \langle x, y \rangle, \underline{\lambda}v. x' \bullet \langle \text{fst } v \rangle + y' \bullet (\text{snd } v) \rangle \\
\overleftarrow{\mathcal{D}}_{\overline{\Gamma}}(\text{fst } t) &\stackrel{\text{def}}{=} \text{let } \langle x, x' \rangle = \overleftarrow{\mathcal{D}}_{\overline{\Gamma}}(t) \text{ in } \langle \text{fst } x, \underline{\lambda}v. x' \bullet \langle v, \underline{0} \rangle \rangle \\
\overleftarrow{\mathcal{D}}_{\overline{\Gamma}}(\text{snd } t) &\stackrel{\text{def}}{=} \text{let } \langle x, x' \rangle = \overleftarrow{\mathcal{D}}_{\overline{\Gamma}}(t) \text{ in } \langle \text{snd } x, \underline{\lambda}v. x' \bullet \langle \underline{0}, v \rangle \rangle \\
\overleftarrow{\mathcal{D}}_{\overline{\Gamma}}(\lambda x. t) &\stackrel{\text{def}}{=} \text{let } y = \lambda x. \overleftarrow{\mathcal{D}}_{\overline{\Gamma},x}(t) \text{ in} \\
&\quad \langle \lambda x. \text{let } \langle z, z' \rangle = y \ x \text{ in } \langle z, \underline{\lambda}v. \text{snd}(z' \bullet v) \rangle, \underline{\lambda}v. \text{case } v \text{ of } !x \otimes v \rightarrow \text{fst}((\text{snd}(y \ x)) \bullet v) \rangle \\
\overleftarrow{\mathcal{D}}_{\overline{\Gamma}}(t \ s) &\stackrel{\text{def}}{=} \text{let } \langle x, x'_{\text{ctx}} \rangle = \overleftarrow{\mathcal{D}}_{\overline{\Gamma}}(t) \text{ in let } \langle y, y' \rangle = \overleftarrow{\mathcal{D}}_{\overline{\Gamma}}(s) \text{ in let } \langle z, x'_{\text{arg}} \rangle = x \ y \text{ in} \\
&\quad \langle z, \underline{\lambda}v. x'_{\text{ctx}} \bullet (!y \otimes v) + y' \bullet (x'_{\text{arg}} \bullet v) \rangle.
\end{aligned}$$

We now explain and justify the forward and reverse CHAD transformations. The transformations for variables, tuples and projections implement the well-known multivariate calculus facts about (transposed) derivatives of differentiable functions into and out of products of spaces. The transformations for **let**-bindings add to that the chain rules for \mathcal{T} and \mathcal{T}^* of Section 2.2. The transformations for λ -abstractions split the derivative of a closure $\lambda x. t$ into the derivative z' with respect to the function argument x and the derivative $\text{snd}(y x)$ with respect to the captured context variables; they store z' together with the primal computation z of $\lambda x. t$ in the primal associated with the closure and they store $\text{snd}(y x)$ in the (co)tangent associated with the closure. Conversely, the transformations for evaluations extracts those two components of the (transposed) derivative x'_{ctx} (w.r.t. context variables) and x'_{arg} (w.r.t. function argument) from the (co)tangent and primal, respectively, and recombine them to correctly propagate (co)tangent contributions from both sources.

7.6 Sharing of Common Subexpressions

Through careful use of **let**-bindings, we have taken care to ensure that the CHAD code transformations we specified have the following good property: for every program former $C[t_1, \dots, t_n]$ that takes n subprograms t_1, \dots, t_n (for example, function application $t s$ takes two subprograms t and s), we have that $\vec{\mathcal{D}}_{\overline{F}}(C[t_1, \dots, t_n])$ uses $\vec{\mathcal{D}}_{\overline{F}_i}(t_i)$ exactly once in its definition for each subprogram t_i , for some list of identifiers \overline{F}_i . Similarly, $\overleftarrow{\mathcal{D}}_{\overline{F}}(C[t_1, \dots, t_n])$ uses $\overleftarrow{\mathcal{D}}_{\overline{F}_i}(t_i)$ exactly once in its definition for each subprogram t_i , which demonstrates the following.

COROLLARY 7.2 (NO CODE EXPLOSION). *The code sizes of the forward and reverse CHAD transformed programs $\vec{\mathcal{D}}_{\overline{F}}(t)$ and $\overleftarrow{\mathcal{D}}_{\overline{F}}(t)$ both grow linearly in the size of the original source program t .*

This compile-time complexity property is crucial if we are to keep compilation times and executable sizes manageable when performing AD on large code-bases.

Of course, our use of **let**-bindings has the additional benefit at runtime that repeated subcomputations are performed only once and their stored results are shared, rather than recomputing every time their results are needed. We have taken care to avoid any unnecessary computation in this way, which we hope will benefit the performance of CHAD in practice. However, we leave a proper complexity and practical performance analysis to future work.

8 PROVING REVERSE AND FORWARD AD SEMANTICALLY CORRECT

In this section, we show that the CHAD code transformations described in Section 7 correctly compute mathematical derivatives (Theorem 8.3). The proof mainly consists of an (open) logical relations argument over the semantics in the Cartesian closed categories $\mathbf{Set} \times \Sigma_{\mathbf{Set}} \mathbf{CMon}$ and $\mathbf{Set} \times \Sigma_{\mathbf{Set}} \mathbf{CMon}^{op}$. The intuition behind the proof is as follows:

- the logical relations relate a differentiable functions $\mathbb{R}^d \rightarrow \llbracket \tau \rrbracket$ to associated primal and (co)tangent functions;
- the semantics $\llbracket t \rrbracket \times \llbracket \vec{\mathcal{D}}_{\overline{F}}(t) \rrbracket$ and $\llbracket t \rrbracket \times \llbracket \overleftarrow{\mathcal{D}}_{\overline{F}}(t) \rrbracket$ of forward and reverse mode CHAD respect the logical relations;
- therefore, by basic results in calculus, they must equal the derivative and transposed derivative of $\llbracket t \rrbracket$.

This logical relations proof can be phrased in elementary terms, but the resulting argument is technical and would be hard to discover. Instead, we prefer to phrase it in terms of a categorical subconing construction, a more abstract and elegant perspective on logical relations. We discovered the proof by taking this categorical perspective, and, while we have verified the elementary argument (see Section 2.6), we would not otherwise have come up with it.

8.1 Preliminaries

8.1.1 Subconing. Logical relations arguments provide a powerful proof technique for demonstrating properties of typed programs. The arguments proceed by induction on the structure of types. Here, we briefly review the basics of categorical logical relations arguments, or *subconing constructions*. We restrict to the level of generality that we need here, but we would like to point out that the theory applies much more generally.

Consider a Cartesian closed category $(C, \mathbb{1}, \times, \Rightarrow)$. Suppose that we are given a functor $F : C \rightarrow \mathbf{Set}$ to the category \mathbf{Set} of sets and functions that preserves finite products in the sense that $F(\mathbb{1}) \cong \mathbb{1}$ and $F(C \times C') \cong F(C) \times F(C')$. Then, we can form the *subcone* of F , or category of logical relations over F , which is Cartesian closed, with a faithful Cartesian closed functor π_1 to C that forgets about the predicates [24]:

- objects are pairs (C, P) of an object C of C and a predicate $P \subseteq FC$;
- morphisms $(C, P) \rightarrow (C', P')$ are C morphisms $f : C \rightarrow C'$ that respect the predicates in the sense that $F(f)(P) \subseteq P'$;
- identities and composition are as in C ;
- $(\mathbb{1}, F\mathbb{1})$ is the terminal object, and binary products and exponentials are given by

$$(C, P) \times (C', P') = (C \times C', \{\alpha \in F(C \times C') \mid F(\pi_1)(\alpha) \in P, F(\pi_2)(\alpha) \in P'\})$$

$$(C, P) \Rightarrow (C', P') = (C \Rightarrow C', \{F(\pi_1)(\gamma) \mid \gamma \in F((C \Rightarrow C') \times C) \text{ s.t. } F(\pi_2)(\gamma) \in P \text{ implies } F(\text{ev})(\gamma) \in P'\}).$$

In typical applications, C can be the syntactic category of a language (like **Syn**), the codomain of a denotational semantics $\llbracket - \rrbracket$ (like **Set**), or a product of the above, if we want to consider n -ary logical relations. Typically, F tends to be a hom-functor (which always preserves products), like $C(\mathbb{1}, -)$ or $C(C_0, -)$, for some important object C_0 . When applied to the syntactic category **Syn** and $F = \mathbf{Syn}(\mathbb{1}, -)$, the formulae for products and exponentials in the subcone clearly reproduce the usual recipes in traditional, syntactic logical relations arguments. As such, subconing generalises standard logical relations methods.

8.2 Subconing for Correctness of AD

We apply the subconing construction above to

$$\begin{aligned} C &= \mathbf{Set} \times \Sigma_{\mathbf{Set}} \mathbf{CMon} & F &= \mathbf{Set} \times \Sigma_{\mathbf{Set}} \mathbf{CMon}((\mathbb{R}^d, (\mathbb{R}^d, \underline{\mathbb{R}}^d)), -) && \text{(forward AD)} \\ C &= \mathbf{Set} \times \Sigma_{\mathbf{Set}} \mathbf{CMon}^{op} & F &= \mathbf{Set} \times \Sigma_{\mathbf{Set}} \mathbf{CMon}^{op}((\mathbb{R}^d, (\mathbb{R}^d, \underline{\mathbb{R}}^d)), -) && \text{(reverse AD),} \end{aligned}$$

where we note that \mathbf{Set} , $\Sigma_{\mathbf{Set}} \mathbf{CMon}$, and $\Sigma_{\mathbf{Set}} \mathbf{CMon}^{op}$ are Cartesian closed (given the arguments of Sections 5 and 6) and that the product of Cartesian closed categories is again Cartesian closed. Let us write $\overrightarrow{\mathbf{SScone}}$ and $\overleftarrow{\mathbf{SScone}}$, respectively, for the resulting categories of logical relations.

Seeing that $\overrightarrow{\mathbf{SScone}}$ and $\overleftarrow{\mathbf{SScone}}$ are Cartesian closed, we obtain unique Cartesian closed functors $(-)^f : \mathbf{Syn} \rightarrow \overrightarrow{\mathbf{SScone}}$ and $(-)^r : \mathbf{Syn} \rightarrow \overleftarrow{\mathbf{SScone}}$, by the universal property of **Syn** (Section 3), once we fix an interpretation of **real**^{*n*} and all operations *op*. We write P_τ^f and P_τ^r , respectively, for the relations $\pi_2(\tau)^f$ and $\pi_2(\tau)^r$. Let us interpret

$$(\mathbf{real}^n)^f \stackrel{\text{def}}{=} (((\mathbb{R}^n, (\mathbb{R}^n, \underline{\mathbb{R}}^n)), \{(f, (g, h)) \mid f \text{ is differentiable, } f = g \text{ and } h = Df\}))$$

$$(\mathbf{real}^n)^r \stackrel{\text{def}}{=} (((\mathbb{R}^n, (\mathbb{R}^n, \underline{\mathbb{R}}^n)), \{(f, (g, h)) \mid f \text{ is differentiable, } f = g \text{ and } h = Df^t\}))$$

$$(\mathbf{op})^f \stackrel{\text{def}}{=} (\llbracket \mathbf{op} \rrbracket, (\llbracket \mathbf{op} \rrbracket, \llbracket \mathbf{Dop} \rrbracket)) \quad (\mathbf{op})^r \stackrel{\text{def}}{=} (\llbracket \mathbf{op} \rrbracket, (\llbracket \mathbf{op} \rrbracket, \llbracket \mathbf{Dop}^t \rrbracket)),$$

where we write Df for the semantic derivative of f and $(-)^t$ for the matrix transpose (see Section 5).

LEMMA 8.1. *These definitions extend uniquely to define Cartesian closed functors*

$$(\llbracket - \rrbracket)^f : \mathbf{Syn} \rightarrow \overline{\mathbf{SScone}} \quad \text{and} \quad (\llbracket - \rrbracket)^r : \mathbf{Syn} \rightarrow \overleftarrow{\mathbf{SScone}}.$$

PROOF. This follows from the universal property of \mathbf{Syn} (Proposition 3.1) once we verify that $(\llbracket \text{op} \rrbracket, (\llbracket \text{op} \rrbracket, \llbracket \text{Dop} \rrbracket))$ and $(\llbracket \text{op} \rrbracket, (\llbracket \text{op} \rrbracket, \llbracket \text{Dop}^t \rrbracket))$ respect the logical relations P^f and P^r , respectively. This respecting of relations follows immediately from the chain rule for multivariate differentiation, as long as we have implemented our derivatives correctly for the basic operations op , in the sense that

$$\llbracket x; y \vdash \text{Dop}(x; y) \rrbracket = D\llbracket \text{op} \rrbracket \quad \text{and} \quad \llbracket x; y \vdash \text{Dop}^t(x; y) \rrbracket = D\llbracket \text{op} \rrbracket^t.$$

Writing $\mathbf{real}^{n_1, \dots, n_k} \stackrel{\text{def}}{=} \mathbf{real}^{n_1} * \dots * \mathbf{real}^{n_k}$ and $\mathbb{R}^{n_1, \dots, n_k} \stackrel{\text{def}}{=} \mathbb{R}^{n_1} \times \dots \times \mathbb{R}^{n_k}$, we compute

$$\begin{aligned} (\llbracket \mathbf{real}^{n_1, \dots, n_k} \rrbracket)^f &= ((\mathbb{R}^{n_1, \dots, n_k}, (\mathbb{R}^{n_1, \dots, n_k}, \underline{\mathbb{R}}^{n_1, \dots, n_k})), \{(f, (g, h)) \mid f \text{ is differentiable, } f = g, h = Df\}) \\ (\llbracket \mathbf{real}^{n_1, \dots, n_k} \rrbracket)^r &= ((\mathbb{R}^{n_1, \dots, n_k}, (\mathbb{R}^{n_1, \dots, n_k}, \underline{\mathbb{R}}^{n_1, \dots, n_k})), \{(f, (g, h)) \mid f \text{ is differentiable, } f = g, h = Df^t\}), \end{aligned}$$

since derivatives of tuple-valued functions are computed component-wise. (In fact, the corresponding facts hold more generally for any first-order type, as an iterated product of \mathbf{real}^n .) Suppose that $(f, (g, h)) \in P_{\mathbf{real}^{n_1, \dots, n_k}}^f$, i.e., f is differentiable, $g = f$ and $h = Df$. Then, using the chain rule in the last step, we have

$$\begin{aligned} &(f, (g, h)); (\llbracket \text{op} \rrbracket, (\llbracket \text{op} \rrbracket, \llbracket \text{Dop} \rrbracket)) \\ &= (f, (f, Df)); (\llbracket \text{op} \rrbracket, (\llbracket \text{op} \rrbracket, \llbracket x; y \vdash \text{Dop}(x; y) \rrbracket)) \\ &= (f, (f, Df)); (\llbracket \text{op} \rrbracket, (\llbracket \text{op} \rrbracket, D\llbracket \text{op} \rrbracket)) \\ &= (f; \llbracket \text{op} \rrbracket, (f; \llbracket \text{op} \rrbracket, x \mapsto r \mapsto D\llbracket \text{op} \rrbracket(f(x))(Df(x)(r)))) \\ &= (f; \llbracket \text{op} \rrbracket, (f; \llbracket \text{op} \rrbracket, D(f; \llbracket \text{op} \rrbracket))) \in P_{\mathbf{real}^m}^f. \end{aligned}$$

Similarly, if $(f, (g, h)) \in P_{\mathbf{real}^{n_1, \dots, n_k}}^r$, then by the chain rule and linear algebra

$$\begin{aligned} &(f, (g, h)); (\llbracket \text{op} \rrbracket, (\llbracket \text{op} \rrbracket, \llbracket \text{Dop}^t \rrbracket)) \\ &= (f, (f, Df^t)); (\llbracket \text{op} \rrbracket, (\llbracket \text{op} \rrbracket, \llbracket x; y \vdash \text{Dop}^t(x; y) \rrbracket)) \\ &= (f, (f, Df^t)); (\llbracket \text{op} \rrbracket, (\llbracket \text{op} \rrbracket, D\llbracket \text{op} \rrbracket^t)) \\ &= (f; \llbracket \text{op} \rrbracket, (f; \llbracket \text{op} \rrbracket, x \mapsto v \mapsto Df^t(x)(D\llbracket \text{op} \rrbracket^t(f(x))(v)))) \\ &= (f; \llbracket \text{op} \rrbracket, (f; \llbracket \text{op} \rrbracket, x \mapsto v \mapsto Df(x); D\llbracket \text{op} \rrbracket(f(x))^t(v))) \\ &= (f; \llbracket \text{op} \rrbracket, (f; \llbracket \text{op} \rrbracket, D(f; \llbracket \text{op} \rrbracket)^t)) \in P_{\mathbf{real}^m}^r. \end{aligned}$$

Consequently, we obtain our unique Cartesian closed functors $(\llbracket - \rrbracket)^f$ and $(\llbracket - \rrbracket)^r$. \square

Furthermore, observe that $\Sigma_{[-]} \llbracket - \rrbracket (t_1, t_2) \stackrel{\text{def}}{=} (\llbracket t_1 \rrbracket, \llbracket t_2 \rrbracket)$ defines a Cartesian closed functor $\Sigma_{[-]} \llbracket - \rrbracket : \Sigma_{\mathbf{CSyn}} \mathbf{LSyn} \rightarrow \Sigma_{\mathbf{Set}} \mathbf{CMon}$. Similarly, we get a Cartesian closed functor $\Sigma_{[-]} \llbracket - \rrbracket^{op} : \Sigma_{\mathbf{CSyn}} \mathbf{LSyn}^{op} \rightarrow \Sigma_{\mathbf{Set}} \mathbf{CMon}^{op}$. As a consequence, both squares below commute.

$$\begin{array}{ccc} \mathbf{Syn} & \xrightarrow{(\text{id}, \overline{\mathcal{D}})} & \mathbf{Syn} \times \Sigma_{\mathbf{CSyn}} \mathbf{LSyn} \\ \downarrow (\llbracket - \rrbracket)^f & & \downarrow \llbracket - \rrbracket \times \Sigma_{[-]} \llbracket - \rrbracket \\ \overline{\mathbf{SScone}} & \xrightarrow{\pi_1} & \mathbf{Set} \times \Sigma_{\mathbf{Set}} \mathbf{CMon} \end{array} \qquad \begin{array}{ccc} \mathbf{Syn} & \xrightarrow{(\text{id}, \overline{\mathcal{D}})} & \mathbf{Syn} \times \Sigma_{\mathbf{CSyn}} \mathbf{LSyn}^{op} \\ \downarrow (\llbracket - \rrbracket)^r & & \downarrow \llbracket - \rrbracket \times \Sigma_{[-]} \llbracket - \rrbracket^{op} \\ \overleftarrow{\mathbf{SScone}} & \xrightarrow{\pi_1} & \mathbf{Set} \times \Sigma_{\mathbf{Set}} \mathbf{CMon}^{op}. \end{array}$$

Indeed, going around the squares in both directions define Cartesian closed functors that agree on their action on the generators \mathbf{real}^n and \mathbf{op} of the Cartesian closed category \mathbf{Syn} .

COROLLARY 8.2. *For any source language (Section 3) program $\Gamma \vdash t : \tau$, $(\llbracket t \rrbracket, (\llbracket \vec{\mathcal{D}}_{\Gamma}(t) \rrbracket_1, \llbracket \vec{\mathcal{D}}_{\Gamma}(t) \rrbracket_2))$ is a morphism in $\overrightarrow{\mathbf{SScone}}$ and therefore respects the logical relations P^f . Similarly, $(\llbracket t \rrbracket, (\llbracket \overleftarrow{\mathcal{D}}_{\Gamma}(t) \rrbracket_1, \llbracket \overleftarrow{\mathcal{D}}_{\Gamma}(t) \rrbracket_2))$ is a morphism in $\overleftarrow{\mathbf{SScone}}$ and therefore respects the logical relations P^r .*

Most of the work is now in place to show correctness of AD. We finish the proof below. To ease notation, we work with terms in a context with a single type. Doing so is not a restriction as our language has products, and the theorem holds for arbitrary terms between first-order types.

THEOREM 8.3 (CORRECTNESS OF AD). *For programs $\Gamma \vdash t : \sigma$ where σ and all types τ_i in $\Gamma = x_1 : \tau_1, \dots, x_n : \tau_n$ are first-order types, $\llbracket t \rrbracket$ is differentiable and*

$$\llbracket \vec{\mathcal{D}}_{\Gamma}(t) \rrbracket_1 = \llbracket t \rrbracket \quad \llbracket \vec{\mathcal{D}}_{\Gamma}(t) \rrbracket_2 = D\llbracket t \rrbracket \quad \llbracket \overleftarrow{\mathcal{D}}_{\Gamma}(t) \rrbracket_1 = \llbracket t \rrbracket \quad \llbracket \overleftarrow{\mathcal{D}}_{\Gamma}(t) \rrbracket_2 = D\llbracket t \rrbracket^t,$$

where we write D and $(-)^t$ for the usual calculus derivative and matrix transpose. Hence,

$$\llbracket \vec{\mathcal{D}}_{\Gamma}(t) \rrbracket = (\llbracket t \rrbracket, D\llbracket t \rrbracket) \quad \text{and} \quad \llbracket \overleftarrow{\mathcal{D}}_{\Gamma}(t) \rrbracket = (\llbracket t \rrbracket, D\llbracket t \rrbracket^t).$$

PROOF. Seeing that our language has tuples, we may assume without loss of generality that $\Gamma = x : \tau$. We use the logical relations for general d to show differentiability of all programs. Next, the case of $d = 1$ suffices to show that CHAD computes correct derivatives.

First, we observe that $\llbracket t \rrbracket$ sends differentiable functions $\mathbb{R}^d \rightarrow \llbracket \tau \rrbracket$ to differentiable functions $\mathbb{R}^d \rightarrow \llbracket \sigma \rrbracket$, as t respects the logical relations. Observing that $\llbracket \tau \rrbracket \cong \mathbb{R}^N$ for some N , as τ is a first-order type, we can choose $d = N$. Then, P_{τ}^f contains $(f, (g, h))$ for a differentiable isomorphism f . It therefore follows that $\llbracket t \rrbracket$ is differentiable.

Second, we focus on the correctness of forward AD, $\vec{\mathcal{D}}$.

Let $x \in \llbracket \vec{\mathcal{D}}(\tau) \rrbracket_1 = \llbracket \tau \rrbracket \cong \mathbb{R}^N$ and $v \in \llbracket \vec{\mathcal{D}}(\tau) \rrbracket_2 \cong \mathbb{R}^N$ (for some N). Then, there is a differentiable curve $\gamma : \mathbb{R} \rightarrow \llbracket \tau \rrbracket$, such that $\gamma(0) = x$ and $D\gamma(0)(1) = v$. Clearly, $(\gamma, (\gamma, D\gamma)) \in P_{\tau}^f$ (for $d = 1$).

As $(\llbracket t \rrbracket, (\llbracket \vec{\mathcal{D}}_{\Gamma}(t) \rrbracket_1, \llbracket \vec{\mathcal{D}}_{\Gamma}(t) \rrbracket_2))$ respects the logical relation P^f by Corollary 8.2, we have

$$(\gamma; \llbracket t \rrbracket, (\gamma; \llbracket \vec{\mathcal{D}}_{\Gamma}(t) \rrbracket_1, x \mapsto r \mapsto \llbracket \vec{\mathcal{D}}_{\Gamma}(t) \rrbracket_2(\gamma(x))(D\gamma(x)(r)))) = (\gamma, (\gamma, D\gamma)); (\llbracket t \rrbracket, (\llbracket \vec{\mathcal{D}}_{\Gamma}(t) \rrbracket_1, \llbracket \vec{\mathcal{D}}_{\Gamma}(t) \rrbracket_2)) \in P_{\sigma}^f,$$

where we use the definition of composition in $\mathbf{Set} \times \Sigma_{\mathbf{Set}} \mathbf{CMon}$. Therefore,

$$\gamma; \llbracket t \rrbracket = \gamma; \llbracket \vec{\mathcal{D}}_{\Gamma}(t) \rrbracket_1$$

and, by the chain rule,

$$x \mapsto r \mapsto D\llbracket t \rrbracket(\gamma(x))(D\gamma(x)(r)) = D(\gamma; \llbracket t \rrbracket) = x \mapsto r \mapsto \llbracket \vec{\mathcal{D}}_{\Gamma}(t) \rrbracket_2(\gamma(x))(D\gamma(x)(r)).$$

Evaluating the former at 0 gives $\llbracket t \rrbracket(x) = \llbracket \vec{\mathcal{D}}_{\Gamma}(t) \rrbracket_1(x)$. Similarly, evaluating the latter at 0 and 1 gives $D\llbracket t \rrbracket(x)(v) = \llbracket \vec{\mathcal{D}}_{\Gamma}(t) \rrbracket_2(x)(v)$.

Third, we turn to the correctness of reverse AD, $\overleftarrow{\mathcal{D}}$.

Let $x \in \llbracket \overleftarrow{\mathcal{D}}(\tau) \rrbracket_1 = \llbracket \tau \rrbracket \cong \mathbb{R}^N$ and $v \in \llbracket \overleftarrow{\mathcal{D}}(\tau) \rrbracket_2 \cong \mathbb{R}^N$ (for some N). Let $\gamma_i : \mathbb{R} \rightarrow \llbracket \tau \rrbracket$ be a differentiable curve such that $\gamma_i(0) = x$ and $D\gamma_i(0)(1) = e_i$, where we write e_i for the i th standard basis vector of $\llbracket \overleftarrow{\mathcal{D}}(\tau) \rrbracket_2 \cong \mathbb{R}^N$. Clearly, $(\gamma_i, (\gamma_i, D\gamma_i^t)) \in P_{\tau}^r$ (for $d = 1$).

As $(\llbracket t \rrbracket, (\llbracket \overleftarrow{\mathcal{D}}_{\Gamma}(t) \rrbracket_1, \llbracket \overleftarrow{\mathcal{D}}_{\Gamma}(t) \rrbracket_2))$ respects the logical relation P^r by Corollary 8.2, we have

$$(\gamma_i; \llbracket t \rrbracket, (\gamma_i; \llbracket \overleftarrow{\mathcal{D}}_{\Gamma}(t) \rrbracket_1, x \mapsto w \mapsto D\gamma_i(x)^t(\llbracket \overleftarrow{\mathcal{D}}_{\Gamma}(t) \rrbracket_2(\gamma_i(x))(w)))) = (\gamma_i, (\gamma_i, D\gamma_i^t)); (\llbracket t \rrbracket, (\llbracket \overleftarrow{\mathcal{D}}_{\Gamma}(t) \rrbracket_1, \llbracket \overleftarrow{\mathcal{D}}_{\Gamma}(t) \rrbracket_2)) \in P_{\sigma}^r,$$

by using the definition of composition in $\mathbf{Set} \times \Sigma_{\mathbf{Set}} \mathbf{CMon}^{op}$. Consequently,

$$\gamma_i; \llbracket t \rrbracket = \gamma_i; \llbracket \overleftarrow{\mathcal{D}}_{\overline{\Gamma}}(t)_1 \rrbracket$$

and, by the chain rule,

$$x \mapsto w \mapsto D\gamma_i(x)^t (D\llbracket t \rrbracket(\gamma_i(x))^t(w)) = D(\gamma_i; \llbracket t \rrbracket)^t = x \mapsto w \mapsto D\gamma_i(x)^t (\llbracket \overleftarrow{\mathcal{D}}_{\overline{\Gamma}}(t)_2 \rrbracket(\gamma_i(x))(w)).$$

Evaluating the former at 0 gives $\llbracket t \rrbracket(x) = \llbracket \overleftarrow{\mathcal{D}}_{\overline{\Gamma}}(t)_1 \rrbracket(x)$. Similarly, evaluating the latter at 0 and v gives us $e_i \odot D\llbracket t \rrbracket(x)^t(v) = e_i \odot \llbracket \overleftarrow{\mathcal{D}}_{\overline{\Gamma}}(t)_2 \rrbracket(x)(v)$. As this equation holds for all basis vectors e_i of $\llbracket \overleftarrow{\mathcal{D}}(\tau) \rrbracket$, we find that

$$D\llbracket t \rrbracket(x)^t(v) = \sum_{i=1}^N (e_i \odot D\llbracket t \rrbracket(x)^t(v)) \cdot e_i = \sum_{i=1}^N (e_i \odot \llbracket \overleftarrow{\mathcal{D}}_{\overline{\Gamma}}(t)_2 \rrbracket(x)(v)) \cdot e_i = \llbracket \overleftarrow{\mathcal{D}}_{\overline{\Gamma}}(t)_2 \rrbracket(x)(v). \quad \square$$

9 PRACTICAL RELEVANCE AND IMPLEMENTATION IN FUNCTIONAL LANGUAGES

Most popular functional languages, such as Haskell and OCaml, do not natively support linear types. As such, the transformations described in this article may seem hard to implement. However, as we will argue in this section, we can easily implement the limited linear types used in phrasing the transformations as abstract data types by using merely a basic module system, such as that of Haskell. The key idea is that linear function types $\underline{\tau} \multimap \underline{\sigma}$ can be represented as plain functions $\tau \rightarrow \sigma$ and copowers $! \tau \otimes \underline{\sigma}$ can be represented as lists or arrays of pairs of type $\tau * \sigma$.

To substantiate that claim, we provide a reference implementation of CHAD operating on strongly typed deeply embedded DSLs Haskell at <https://github.com/VMatthijs/CHAD>. This section explains how that implementation relates to the theoretical development in the rest of this article. It is rather short, because our implementation almost exactly follows the theoretical development in Sections 3, 4, 5, and 7.

9.1 Implementing Linear Functions and Copowers as Abstract Types in Functional Languages

Based on the denotational semantics, $\underline{\tau} \multimap \underline{\sigma}$ -types should hold (representations of) functions f from $\underline{\tau}$ to $\underline{\sigma}$ that are homomorphisms of the monoid structures on $\underline{\tau}$ and $\underline{\sigma}$. We will see that these types can be implemented using an abstract data type that holds certain basic linear functions (extensible as the library evolves) and is closed under the identity, composition, argument swapping, and currying. Again, based on the semantics, $! \underline{\tau} \otimes \underline{\sigma}$ should contain (representations of) finite maps (associative arrays) $\sum_{i=1}^n !t_i \otimes s_i$ of pairs (t_i, s_i) , where t_i is of type $\underline{\tau}$, and s_i is of type $\underline{\sigma}$, and where we identify $xs + !t \otimes s + !t \otimes s'$ and $xs + !t \otimes (s + s')$.

To implement this idea, we consider abstract types $\mathbf{LFun}(\tau, \sigma)$ of linear functions and $\mathbf{Coprover}(\tau, \sigma)$ of copowers. Their programs are generated by the following grammar

$\tau, \sigma, \rho ::=$	types	$\mathbf{Coprover}(\tau, \sigma)$	copower types
...	as in Section 3	$\mathbf{LFun}(\tau, \sigma)$	linear function
$t, s, r ::=$	terms	$\mathbf{lswap} \ t$	swapping args
...	as in Section 3	\mathbf{leval}_t	linear evaluation
$\mathbf{lop}(t_1, \dots, t_n)$	linear operations	$\{(t, -)\}$	singletons
$\underline{0}_\tau$	zero	$\mathbf{lcopowfold} \ t$	Coprover -elimination
$t + s$	plus	\mathbf{lfst}	linear projection
\mathbf{lid}	linear identity	\mathbf{lsnd}	linear projection
$t;_\ell s$	linear composition	$\mathbf{lpair}(t, s)$	linear pairing,

and their API can be typed according to the rules of Figure 8.

$\frac{\{\Gamma \vdash t_i : \mathbf{real}^{n_i} \mid i = 1, \dots, k\} \quad (\text{lop}(t_1, \dots, t_k) \in \text{LOp}_{n_1, \dots, n_k; n'_1, \dots, n'_l}^{m_1, \dots, m_r})}{\Gamma \vdash \text{lop}(t) : \mathbf{LFun}(\mathbf{LDom}(\text{lop}), \mathbf{CDom}(\text{lop}))}$		$\frac{}{\Gamma \vdash \underline{0}_\tau : \tau}$	$\frac{\Gamma \vdash t : \tau \quad \Gamma \vdash s : \tau}{\Gamma \vdash t +_\tau s : \tau}$
$\frac{}{\Gamma \vdash \text{lid} : \mathbf{LFun}(\tau, \tau)}$		$\frac{\Gamma \vdash t : \mathbf{LFun}(\tau, \sigma) \quad \Gamma \vdash s : \mathbf{LFun}(\sigma, \rho)}{\Gamma \vdash t;_\ell s : \mathbf{LFun}(\tau, \rho)}$	
$\frac{\Gamma \vdash t : \tau \rightarrow \mathbf{LFun}(\sigma, \rho)}{\Gamma \vdash \text{lswap } t : \mathbf{LFun}(\sigma, \tau \rightarrow \rho)}$		$\frac{\Gamma \vdash t : \tau}{\Gamma \vdash \text{leval}_t : \mathbf{LFun}(\tau \rightarrow \sigma, \sigma)}$	
$\frac{\Gamma \vdash t : \tau}{\Gamma \vdash \{(t, -)\} : \mathbf{LFun}(\sigma, \mathbf{Copower}(\tau, \sigma))}$		$\frac{\Gamma \vdash t : \tau \rightarrow \mathbf{LFun}(\sigma, \rho)}{\Gamma \vdash \text{lcopowfold } t : \mathbf{LFun}(\mathbf{Copower}(\tau, \sigma), \rho)}$	
$\frac{}{\Gamma \vdash \text{lfst} : \mathbf{LFun}(\tau * \sigma, \tau)}$	$\frac{}{\Gamma \vdash \text{lsnd} : \mathbf{LFun}(\tau * \sigma, \sigma)}$	$\frac{\Gamma \vdash t : \mathbf{LFun}(\tau, \sigma) \quad \Gamma \vdash s : \mathbf{LFun}(\tau, \rho)}{\Gamma \vdash \text{lpair}(t, s) : \mathbf{LFun}(\tau, \sigma * \rho)}$	

Fig. 8. Typing rules for the applied target language, to extend the source language.

We note that these abstract types give us precisely the functionality and type safety of the linear function and copower types of our target language of Section 4. Indeed, we can define a semantics and type preserving translation $(-)^T$ from that target language to our source language extended with these $\mathbf{LFun}(\tau, \sigma)$ and $\mathbf{Copower}(\tau, \sigma)$ types, for which $(! \tau \otimes \underline{\sigma})^T \stackrel{\text{def}}{=} \mathbf{Copower}(\tau^T, \underline{\sigma}^T)$, $(\tau \multimap \underline{\sigma})^T \stackrel{\text{def}}{=} \mathbf{LFun}(\tau^T, \underline{\sigma}^T)$, $(\mathbf{real}^n)^T \stackrel{\text{def}}{=} \mathbf{real}^n$ and we extend $(-)^T$ structurally recursively, letting it preserve all other type formers. We then translate $(x_1 : \tau, \dots, x_n : \tau; y : \underline{\sigma} \vdash t : \rho)^T \stackrel{\text{def}}{=} x_1 : \tau^T, \dots, x_n : \tau^T \vdash t^T : (\underline{\sigma} \multimap \underline{\rho})^T$ and $(x_1 : \tau, \dots, x_n : \tau \vdash t : \sigma)^T \stackrel{\text{def}}{=} x_1 : \tau^T, \dots, x_n : \tau^T \vdash t^T : \sigma^T$. We believe an interested reader can fill in the details.

9.2 Implementing the API of $\mathbf{LFun}(\tau, \sigma)$ and $\mathbf{Copower}(\tau, \sigma)$ Types

We observe that we can implement this API of $\mathbf{LFun}(\tau, \sigma)$ and $\mathbf{Copower}(\tau, \sigma)$ types, as follows, in a language that extends the source language with types $\mathbf{List}(\tau)$ of lists (or arrays) of elements of type τ . Indeed, we implement $\mathbf{LFun}(\tau, \sigma)$ under the hood, for example, as $\tau \rightarrow \sigma$ and $\mathbf{Copower}(\tau, \sigma)$ as $\mathbf{List}(\tau * \sigma)$. The idea is that $\mathbf{LFun}(\tau, \sigma)$, which arose as a right adjoint in our linear language, is essentially a *subtype* of $\tau \rightarrow \sigma$. However, $\mathbf{Copower}(\tau, \sigma)$, which arose as a left adjoint, is a *quotient type* of $\mathbf{List}(\tau * \sigma)$. We achieve the desired subtyping and quotient typing by exposing only the API of Figure 8 and hiding the implementation. We can then implement this interface as follows.¹¹

$$\begin{aligned}
\text{lop} &\stackrel{\text{def}}{=} \overline{\text{lop}} & \underline{0}_1 &\stackrel{\text{def}}{=} \langle \rangle & t +_1 s &\stackrel{\text{def}}{=} \langle \rangle & \underline{0}_{\tau * \sigma} &\stackrel{\text{def}}{=} \langle \underline{0}_\tau, \underline{0}_\sigma \rangle & t +_{\tau * \sigma} s &\stackrel{\text{def}}{=} \langle \text{fst } t +_\tau \text{fst } s, \text{snd } t +_\sigma \text{snd } s \rangle \\
\underline{0}_{\tau \rightarrow \sigma} &\stackrel{\text{def}}{=} \lambda _ . \underline{0}_\sigma & t +_{\tau \rightarrow \sigma} s &\stackrel{\text{def}}{=} \lambda x . t \ x +_\sigma s \ x & \underline{0}_{\mathbf{LFun}(\tau, \sigma)} &\stackrel{\text{def}}{=} \lambda _ . \underline{0}_\sigma & t +_{\mathbf{LFun}(\tau, \sigma)} s &\stackrel{\text{def}}{=} \lambda x . t \ x +_\sigma s \ x \\
\underline{0}_{\mathbf{Copower}(\tau, \sigma)} &\stackrel{\text{def}}{=} [] & t +_{\mathbf{Copower}(\tau, \sigma)} s &\stackrel{\text{def}}{=} \text{fold } x :: \text{acc over } x \text{ in } t \text{ from } \text{acc} = s \\
\text{lid} &\stackrel{\text{def}}{=} \lambda x . x & t;_\ell s &\stackrel{\text{def}}{=} \lambda x . s \ (t \ x) & \text{lswap } t &\stackrel{\text{def}}{=} \lambda x . \lambda y . t \ y \ x & \text{leval}_t &\stackrel{\text{def}}{=} \lambda x . x \ t \\
\{(t, -)\} &\stackrel{\text{def}}{=} \lambda x . \langle t, x \rangle :: [] & \text{lcopowfold } t &\stackrel{\text{def}}{=} \lambda z . \text{fold } t \ (\text{fst } x) \ (\text{snd } x) + \text{acc over } x \text{ in } z \text{ from } \text{acc} = \underline{0} \\
\text{lfst} &\stackrel{\text{def}}{=} \lambda x . \text{fst } x & \text{lsnd} &\stackrel{\text{def}}{=} \lambda x . \text{snd } x & \text{lpair}(t, s) &\stackrel{\text{def}}{=} \lambda x . \langle t \ x, s \ x \rangle
\end{aligned}$$

¹¹Note that the implementation of $t +_{\mathbf{Copower}(\tau, \sigma)} s$ is merely list concatenation written using a fold.

Here, we write $\overline{\text{lop}}$ for the function that `lop` is intended to implement, $[]$ for the empty list, $t :: s$ for the list consisting of s with t prepended on the front, and `fold t over x in s from $acc = \text{init}$` for (right) folding an operation t over a list s , starting from init . The monoid structure $0_\tau, +_\tau$ can be defined by induction on the structure of types, by using, for example, type classes in Haskell or OCaml’s module system or (in any functional language with algebraic data types) by using reified types. Furthermore, the implementer of the AD library can determine which linear operations `lop` to include within the implementation of **LFun**. We expect these linear operations to include various forms of dense and sparse matrix-vector multiplication as well as code for computing Jacobian-vector and Jacobian-adjoint products for the operations `op` that avoids having to compute the full Jacobian. Another option is to simply include two linear operations `Dop` and `Dopt` for computing the derivative and transposed derivative of each operation `op`.

9.3 Maintaining Type Safety throughout the Compilation Pipe-Line in Our Reference Implementation

In a principled approach to building a define-then-run AD library, one would shield this implementation using the abstract data types **Copower** (τ, σ) and **LFun** (τ, σ) as we describe, both for reasons of type safety and because it conveys the intuition behind the algorithm and its correctness. By using such abstract data types in our Haskell implementation combined with GADTs and type families, we achieve a fully type-safe (well-scoped, well-typed De Bruijn) implementation of the source and target languages of Sections 3 and 4 with their semantics of Section 5 and statically type-checked code transformations of Section 7.

However, nothing stops library implementers from exposing the full implementation rather than working with abstract types. In fact, this seems to be the approach [43] have taken. A downside of that “exposed” approach is that the transformations then no longer respect equational reasoning principles. In our reference implementation, we include a compiler from the (linearly typed) target language to a less type-safe “concrete” target language (implementing Section 9.2 as a compilation step): essentially the source language extended with list (or array) types.¹² This demonstrates definitely that CHAD implements a notion of compile-time AD code transformation that takes in and spits out standard functional code without any custom semantics.

9.4 Compiling Away Copowers

As a final observation on this implementation, we would like to note that while the proposed implementation of copowers as lists is generally applicable, more efficient implementation strategies can often be achieved in practice. In fact, in unpublished follow-up work to this article led by Tom Smeding, we show that when we implement CHAD to operate on Accelerate [34], we can optimize away uses of copower types.

10 ADDING HIGHER-ORDER ARRAY PRIMITIVES

The aim of this article is to answer the foundational question of how to perform (reverse) AD at higher types. The problem of how to perform AD of evaluation and currying is highly challenging. For this reason, we have devoted this article to explaining a solution to that problem in detail, working with a toy language with ground types of black-box, sized arrays **realⁿ** with some

¹²To be exact, the to-concrete compilation step in the implementation does convert copowers and linear functions to lists and regular functions, but retains `_` and `+` primitives. We feel this brings a major increase in readability of the output. Actually implementing the inductive definitions for `_` and `+` is easy when reified types (singletons) are added to the implementation, which itself is an easy change if one modifies the `LT` type class.

first-order operations op . However, many of the interesting applications only arise once we can use higher-order array primitives such as **map** and **fold** on real^n .

Our definitions and correctness proofs extend to this setting with standard array processing primitives including **map**, **fold**, **filter**, **zipWith**, **permute** (aka **scatter**), **backpermute** (aka **gather**), **generate** (or **build**), and array indexing. We plan to discuss these primitives as well as CHAD applied to dynamically sized arrays in detail in an applied follow-up paper, which will focus on an implementation of CHAD to operate on Accelerate [34].

To illustrate the idea behind such an extension, we briefly discuss the case for **map** here and leave the rest to future work. Suppose that we add operations¹³

$$\frac{\Gamma, x : \text{real} \vdash t : \text{real} \quad \Gamma \vdash s : \text{real}^n}{\Gamma \vdash \text{map}(x.t, s) : \text{real}^n}$$

to the source language, to “map” functions over the black-box arrays. Then, supposing that we add to the target language primitives

$$\frac{\Gamma, x : \text{real}; v : \underline{\tau} * \text{real} \vdash t : \text{real} \quad \Gamma \vdash s : \text{real}^n \quad \Gamma; v : \underline{\tau} \vdash r : \text{real}^n}{\Gamma; v : \underline{\tau} \vdash D\text{map}(x.t, s, r) : \text{real}^n}$$

$$\frac{\Gamma, x : \text{real}; v : \text{real} \vdash t : \underline{\tau} * \text{real} \quad \Gamma \vdash s : \text{real}^n \quad \Gamma; v : \text{real}^n \vdash r : \underline{\tau}}{\Gamma; v : \text{real}^n \vdash D\text{map}^t(x.t, s, r) : \underline{\tau}},$$

we can define

$$\begin{aligned} \vec{D}_{\Gamma}(\text{map}(x.t, s)) &\stackrel{\text{def}}{=} \text{let } y = \lambda x. \vec{D}_{\Gamma, x}(t) \text{ in let } \langle z, z' \rangle = \vec{D}_{\Gamma}(s) \text{ in} \\ &\quad \langle \text{map}(x.\text{fst}(y x), z), \lambda v. D\text{map}(x.(\text{snd}(y x)) \bullet v, z, z' \bullet v) \rangle \\ \overleftarrow{D}_{\Gamma}(\text{map}(x.t, s)) &\stackrel{\text{def}}{=} \text{let } y = \lambda x. \overleftarrow{D}_{\Gamma, x}(t) \text{ in let } \langle z, z' \rangle = \overleftarrow{D}_{\Gamma}(s) \text{ in} \\ &\quad \langle \text{map}(x.\text{fst}(y x), z), \lambda v. D\text{map}^t(x.(\text{snd}(y x)) \bullet v, z, z' \bullet v) \rangle. \end{aligned}$$

In our practical API of Section 9.1, the required target language primitives correspond to

$$\begin{aligned} &\frac{\Gamma, x : \text{real} \vdash t : \text{LFun}(\underline{\tau} * \text{real}, \text{real}) \quad \Gamma \vdash s : \text{real}^n \quad \Gamma \vdash r : \text{LFun}(\underline{\tau}, \text{real}^n)}{\Gamma \vdash D\text{map}(x.t, s, r) : \text{LFun}(\underline{\tau}, \text{real}^n)} \\ &\frac{\Gamma, x : \text{real} \vdash t : \text{LFun}(\text{real}, \underline{\tau} * \text{real}) \quad \Gamma \vdash s : \text{real}^n \quad \Gamma \vdash r : \text{LFun}(\text{real}^n, \underline{\tau})}{\Gamma \vdash D\text{map}^t(x.t, s, r) : \text{LFun}(\text{real}^n, \underline{\tau})}. \end{aligned}$$

Extending Section 9.2, we can implement this API as

$$\begin{aligned} D\text{map}(x.t, s, r) &\stackrel{\text{def}}{=} \lambda y. \text{zipWith}((x, x').t \langle y, x' \rangle, s, r y) \\ D\text{map}^t(x.t, s, r) &\stackrel{\text{def}}{=} \lambda y. \text{let } zs = \text{zipWith}((x, x').t x', s, y) \text{ in sum}(\text{map}(w.\text{fst } w, zs)) + r \text{ map}(w.\text{snd } w, zs), \end{aligned}$$

¹³We prefer to work with this elementary formulation of maps rather than the usual higher-order formulation of

$$\frac{\Gamma \vdash t : \text{real} \rightarrow \text{real} \quad \Gamma \vdash s : \text{real}^n}{\Gamma \vdash \text{map}(t, s) : \text{real}^n},$$

because it makes sense in the wider context of languages without function types as well and because it simplifies its CHAD correctness proof. Note that both are equivalent in presence of function types: $\text{map}(t, s) = \text{map}(x.t x, s)$ and $\text{map}(x.t, s) = \text{map}(\lambda x. t, s)$.

where

$$\frac{\Gamma, x : \tau \vdash t : \sigma \quad \Gamma \vdash s : \tau^n}{\Gamma \vdash \text{map}(x.t, s) : \sigma^n} \quad \frac{\Gamma, x : \tau, x' : \sigma \vdash t : \rho \quad \Gamma \vdash s : \tau^n \quad \Gamma \vdash r : \sigma^n}{\Gamma \vdash \text{zipWith}((x, x').t, s, r) : \rho^n} \quad \frac{\Gamma \vdash t : \tau^n}{\Gamma \vdash \text{sum } t : \tau}$$

are the usual functional programming idiom for mapping a unary function over an array, zipping two arrays with a binary operation and for taking the sum of the elements in an array. Note that we assume that we have types τ^n for length- n arrays of elements of type τ , here, generalizing the arrays **real** ^{n} of elements of type **real**. We present a correctness proof for this implementation of the derivatives in Appendix A.

Applications frequently require AD of higher-order primitives such as differential and algebraic equation solvers, e.g., for use in pharmacological modelling in Stan [40]. Currently, derivatives of such primitives are derived using the calculus of variations (and implemented with define-by-run AD) [7, 18]. Our proof method provides a more lightweight and formal method for calculating, and establishing the correctness of, derivatives for such higher-order primitives. Indeed, most formalizations of the calculus of variations use infinite-dimensional vector spaces and are technically involved [26].

11 SCOPE OF CHAD AND FUTURE WORK

11.1 Memory Use of CHAD's Forward AD

CHAD has formulated reverse and forward AD to be precisely each other's categorical dual. The former first computes the primals in a forward pass and next the cotangents in a reverse pass. Dually, the latter first computes the primals in a forward pass and next the tangents in another forward pass. Seeing that the two forward passes in forward AD have identical control-flow, it can be advantageous to interleave them and simultaneously compute the primals and tangents. Such interleaving greatly reduces the memory consumption of the algorithm, as it means that all primals no longer need to be stored for most of the algorithm. We present such an interleaved formulation of forward AD in Reference [21].

While it is much more memory efficient, it has the conceptual downside of no longer being the mirror image of reverse AD. Furthermore, these interleaved formulations of forward AD work by operating on dual numbers. That is, they use an array-of-structs representation, by contrast with the struct-of-arrays representation used to pair primals with tangents in CHAD. Therefore, an SoA-to-AoS optimization is typically needed to make interleaved implementations of forward AD efficient [38].

Finally, we note that such interleaving techniques do not apply to reverse AD, as we need to have completed the forward primal pass before we can start the reverse cotangent pass, due to the dependency structure of the algorithm.

11.2 Applying CHAD to Richer Source Languages

The core observations that let us use CHAD for AD on a higher-order language were the following:

- (1) there is a class of categories with structure \mathcal{S} (in this case, Cartesian closure) such that the source language **Syn** that we want to perform AD on can be seen as the freely generated \mathcal{S} -category on the operations op ;
- (2) we identified structure \mathcal{T} that suffices for a **CMon**-enriched strictly indexed category $\mathcal{L} : C^{op} \rightarrow \mathbf{Cat}$ to ensure that $\Sigma_C \mathcal{L}$ and $\Sigma_C \mathcal{L}^{op}$ are \mathcal{S} -categories;
- (3) we gave a description $\mathbf{LSyn} : \mathbf{CSyn}^{op} \rightarrow \mathbf{Cat}$ of the freely generated **CMon**-enriched strictly indexed category with structure \mathcal{T} , on the Cartesian operations op in **CSyn** and linear

operations Dop and Dop^t in \mathbf{LSyn} ; we interpret these linear/non-linear language as the target language of our AD translations;

- (4) by the universal property of \mathbf{Syn} , we now obtained unique \mathcal{S} -homomorphic AD functors $\vec{D} : \mathbf{Syn} \rightarrow \Sigma_{\mathbf{CSyn}} \mathbf{LSyn}$ and $\tilde{D} : \mathbf{Syn} \rightarrow \Sigma_{\mathbf{CSyn}} \mathbf{LSyn}^{op}$ such that $\vec{D}(op) = (op, Dop)$ and $\tilde{D}(op) = (op, Dop^t)$, whose correctness proof follows immediately because of the well-known theory of subsconing for \mathcal{S} -categories.

CHAD applies equally to source languages with other choices of \mathcal{S} , provided that we can follow steps (1)–(4).

In particular, in Reference [30] it is shown how CHAD applies equally to languages with sum types and (co)inductive types (and tuple and function types). In that setting, the category \mathbf{LSyn} is a genuine strictly indexed category over \mathbf{CSyn} , to account for the fact that the (co)tangent space to a space of varying dimension depends on the chosen base point. That is, in its most principled formulation, the target language has (linear) dependent types. However, we can also work with a simply typed target language in this setting, at the cost of some extra type safety. In fact, our Haskell implementation already supports such a treatment of coproducts.

As discussed in Section 10, \mathcal{S} can also be chosen to include various operations for manipulating arrays, such as **map**, **fold**, **filter**, **zipWith**, **permute** (aka **scatter**), **backpermute** (aka **gather**), **generate** (aka **build**), and array indexing. We plan to describe this application of CHAD to array processing languages and we are implementing CHAD to operate on the Accelerate parallel array processing language.

In work in progress, we are applying CHAD to partial features such as real conditionals, iteration, recursion and recursive types. Our Haskell implementation of CHAD already supports real conditionals, iteration and recursion. The challenge in this setting is to understand the subtle interactions between the ω CPO-structure needed to model recursion and the commutative monoid structure that CHAD uses to accumulate (co)tangents.

11.3 CHAD for Other Dynamic Program Analyses

As noted in Reference [43], source-code transformation AD has a lot of similarities with other dynamic program analyses such as dynamic symbolic analysis and provenance analysis.

In fact, as the abstract perspective on CHAD given in Section 11.2 makes clear, CHAD is in no way tied to automatic differentiation. In many ways, it is much more general, and can best be seen as a framework for applying dynamic program analyses that accumulate data (either by going through the program forwards or backwards) in a commutative monoid to functional languages with expressive features. In fact, by varying the definitions of $\vec{D}(\mathcal{R})$ and $\tilde{D}(\mathcal{R})$ for the ground types \mathcal{R} and the definitions of $\vec{D}(op)$ and $\tilde{D}(op)$ for the primitive operations op (we do not even need to use $\vec{D}(\mathcal{R})_1 = \mathcal{R}$, $\tilde{D}(\mathcal{R})_1 = \mathcal{R}$, $\vec{D}(op)_1 = op$ or $\tilde{D}(op)_1 = op!$), we can completely change the nature of the analysis. In most cases, as long as a notion of correctness of the analysis can be phrased at the level of a denotational semantics, we conjecture that our subsconing techniques lead to straightforward correctness proofs of the analysis.

To give one more example application of such an analysis, beyond AD, dynamic symbolic analysis and provenance analysis, one can note that for a source language \mathbf{Syn} , generated from a base type \mathcal{R} that is a commutative semi-ring, we have a notion of algebraic (or formal) derivative of any polynomial $x_1 : \mathcal{R}, \dots, x_n : \mathcal{R} \vdash op(x_1, \dots, x_n) : \mathcal{R}$ [28]. CHAD can be used to extend and compute this notion of derivative for arbitrary functional programs generated from the polynomials op as basic operations. The particular case of such formal derivatives for the Boolean semi-ring $\mathcal{R} = \mathbb{B}$ are used in Reference [45] to feed into a gradient descent algorithm to learn

Boolean circuits. CHAD makes this method applicable to more general (higher-order) programs over (arrays of) Booleans.

12 RELATED WORK

This work is closely related to References [21] and [20], which introduced a similar semantic correctness proof for a dual-numbers version of forward mode AD and higher-order forward AD, using a subsconing construction. A major difference is that this article also phrases and proves correctness of reverse mode AD on a λ -calculus and relates reverse mode to forward mode AD. Using a syntactic logical relations proof instead, Reference [4] also proves correctness of forward mode AD. Again, it does not address reverse AD.

Reference [12] proposes a similar construction to that of Section 6, and it relates it to the differential λ -calculus. This article develops sophisticated axiomatics for semantic reverse differentiation. However, it neither relates the semantics to a source-code transformation, nor discusses differentiation of higher-order functions. Our construction of differentiation with a (biadditive) linear target language might remind the reader of differential linear logic [15]. In differential linear logic, (forward) differentiation is a first-class operation in a (biadditive) linear language. By contrast, in our treatment, differentiation is a meta-operation.

Importantly, Reference [16] describes and implements what are essentially our source-code transformations, though they were restricted to first-order functions and scalars. After completing this work, we realized that Reference [43] describes an extension of the reverse mode transformation to higher-order functions in a similar manner as we propose in this article, but without the linear or abstract types. Though that paper did not derive the algorithm or show its correctness, it does discuss important practical considerations for its implementation and offers a dependently typed variant of the algorithm based on typed closure conversion, inspired by Reference [37].

Next, there are various lines of work relating to correctness of reverse mode AD that we consider less similar to our work. For example, Reference [31] define and prove correct a formulation of reverse mode AD on a higher-order language that depends on a non-standard operational semantics, essentially a form of symbolic execution. Reference [2] does something similar for reverse mode AD on a first-order language extended with conditionals and iteration. Reference [8] defines a beautifully simple AD algorithm on a simply typed λ -calculus with linear negation (essentially, a more finely typed version of the continuation-based AD of Reference [21]) and proves it correct using operational techniques. Reference [33] extends this work to apply to recursion. Furthermore, they show with an impressive operational argument that this simple algorithm, surprisingly, corresponds to true reverse mode AD with the correct complexity under an operational semantics with a “linear factoring rule.” While this is a natural operational semantics for a linear λ -calculus, it is fundamentally different from normal call-by-value or call-by-name evaluation (under which the generated code has the wrong computational complexity). For this reason, we require a custom interpreter or compiler to use this reverse AD method in practice. Very recently, Reference [25] specified another purely functional reverse AD algorithm, which appears similar to, and, which we conjecture to be equivalent to an implementation of the techniques of Reference [33]. These formulations of reverse mode AD all depend on non-standard run-times and hence fall into the category of “define-by-run” formulations of reverse mode AD, for our purposes. Meanwhile, we are concerned with “define-then-run” formulations: source-code transformations producing differentiated code at compile-time that can then be optimized during compilation with existing compiler tool-chains (such as the Accelerate [11], Futhark [19], and TensorFlow [1] frameworks for generating highly performant GPU code). While we can compile such define-by-run transformations together with their interpreter to achieve a source-code transformation

(hence a sort of define-then-run transformation), the resulting code recursively traverses an AST, so does not obviously seem suitable for generating, via existing tool-chains, optimized machine code for the usual parallel hardware targets that we use as targets for AD, such as GPUs and TPUs.

Finally, there is a long history of work on reverse mode AD, though almost none of it applies the technique to higher-order functions. A notable exception is Reference [37], which gives an impressive source-code transformation implementation of reverse AD in Scheme. While very efficient, this implementation crucially uses mutation. Moreover, the transformation is complex and correctness is not considered. More recently, Reference [44] describes a much simpler implementation of a reverse AD code transformation, again very performant. However, the transformation is quite different from the one considered in this article as it relies on a combination of delimited continuations and mutable state. Correctness is not considered, perhaps because of the semantic complexities introduced by impurity.

Our work adds to the existing literature by presenting a novel, generally applicable method for compositional source-code transformation (forward and) reverse AD on expressive functional languages without a need for a non-standard runtime, by giving a method for compositional correctness proofs of such AD algorithms, and by observing that the CHAD method and its correctness proof are not limited to AD but apply generally to dynamic program analyses that accumulate data in a commutative monoid.

APPENDICES

A CHAD CORRECTNESS FOR HIGHER-ORDER OPERATIONS SUCH AS MAP

We extend the proofs of Lemmas 2.1 and 2.2 to apply to the `map`-constructs of Section 10.

A.1 The Semantics of `map` and Its Derivatives

First, we observe that

$$\begin{aligned} \llbracket \text{map}(x.t, s, r) \rrbracket : \llbracket \Gamma \rrbracket &\rightarrow \mathbb{R}^n \\ \gamma &\mapsto (\llbracket t \rrbracket(\gamma, \pi_1(\llbracket s \rrbracket(\gamma))), \dots, \llbracket t \rrbracket(\gamma, \pi_n(\llbracket s \rrbracket(\gamma)))) \end{aligned}$$

Similarly,

$$\begin{aligned} \llbracket D\text{map}(x.t, s, r) \rrbracket : \llbracket \Gamma \rrbracket &\rightarrow \llbracket \tau \rrbracket \multimap \mathbb{R}^n \\ \gamma &\mapsto v \mapsto (\llbracket t \rrbracket(\gamma, \pi_1(\llbracket s \rrbracket(\gamma)))(v, \pi_1(\llbracket r \rrbracket(\gamma)(v))), \dots, \llbracket t \rrbracket(\gamma, \pi_n(\llbracket s \rrbracket(\gamma)))(v, \pi_n(\llbracket r \rrbracket(\gamma)(v)))) \\ \llbracket D\text{map}^t(x.t, s) \rrbracket : \llbracket \Gamma \rrbracket &\rightarrow \mathbb{R}^n \multimap \llbracket \tau \rrbracket \\ \gamma &\mapsto v \mapsto \pi_1(\llbracket t \rrbracket(\gamma, \pi_1(\llbracket s \rrbracket(\gamma)))(\pi_1(v))) + \dots + \pi_1(\llbracket t \rrbracket(\gamma, \pi_n(\llbracket s \rrbracket(\gamma)))(\pi_n(v))) + \\ &\quad \llbracket r \rrbracket(\gamma)(\pi_2(\llbracket t \rrbracket(\gamma, \pi_1(\llbracket s \rrbracket(\gamma)))(\pi_1(v))), \dots, \pi_2(\llbracket t \rrbracket(\gamma, \pi_n(\llbracket s \rrbracket(\gamma)))(\pi_n(v)))) \end{aligned}$$

This implies that

$$\begin{aligned} \pi_1(\llbracket \vec{\mathcal{D}}_{\Gamma}(\text{map}(x.t, s)) \rrbracket(\gamma)) &= \left(\pi_1(\llbracket \vec{\mathcal{D}}_{\Gamma, x}(t) \rrbracket(\gamma, \pi_1(\pi_1(\llbracket \vec{\mathcal{D}}_{\Gamma}(s) \rrbracket(\gamma))))), \dots, \pi_1(\llbracket \vec{\mathcal{D}}_{\Gamma, x}(t) \rrbracket(\gamma, \pi_n(\pi_1(\llbracket \vec{\mathcal{D}}_{\Gamma}(s) \rrbracket(\gamma)))) \right) \\ \pi_2(\llbracket \vec{\mathcal{D}}_{\Gamma}(\text{map}(x.t, s)) \rrbracket(\gamma))(v) &= \left(\pi_2(\llbracket \vec{\mathcal{D}}_{\Gamma, x}(t) \rrbracket(\gamma, \pi_1(\pi_1(\llbracket \vec{\mathcal{D}}_{\Gamma}(s) \rrbracket(\gamma))))(v, \pi_1(\pi_2(\llbracket \vec{\mathcal{D}}_{\Gamma}(s) \rrbracket(\gamma))(v))), \dots, \right. \\ &\quad \left. \pi_2(\llbracket \vec{\mathcal{D}}_{\Gamma, x}(t) \rrbracket(\gamma, \pi_n(\pi_1(\llbracket \vec{\mathcal{D}}_{\Gamma}(s) \rrbracket(\gamma))))(v, \pi_n(\pi_2(\llbracket \vec{\mathcal{D}}_{\Gamma}(s) \rrbracket(\gamma))(v))) \right) \\ \pi_1(\llbracket \vec{\mathcal{D}}_{\Gamma}^{\dagger}(\text{map}(x.t, s)) \rrbracket(\gamma)) &= \left(\pi_1(\llbracket \vec{\mathcal{D}}_{\Gamma, x}^{\dagger}(t) \rrbracket(\gamma, \pi_1(\pi_1(\llbracket \vec{\mathcal{D}}_{\Gamma}^{\dagger}(s) \rrbracket(\gamma))))), \dots, \pi_1(\llbracket \vec{\mathcal{D}}_{\Gamma, x}^{\dagger}(t) \rrbracket(\gamma, \pi_n(\pi_1(\llbracket \vec{\mathcal{D}}_{\Gamma}^{\dagger}(s) \rrbracket(\gamma)))) \right) \\ \pi_2(\llbracket \vec{\mathcal{D}}_{\Gamma}^{\dagger}(\text{map}(x.t, s)) \rrbracket(\gamma))(v) &= \pi_1(\pi_2(\llbracket \vec{\mathcal{D}}_{\Gamma, x}^{\dagger}(t) \rrbracket(\gamma, \pi_1(\pi_1(\llbracket \vec{\mathcal{D}}_{\Gamma}^{\dagger}(s) \rrbracket(\gamma))))(\pi_1(v))) + \dots \\ &\quad + \pi_1(\pi_2(\llbracket \vec{\mathcal{D}}_{\Gamma, x}^{\dagger}(t) \rrbracket(\gamma, \pi_n(\pi_1(\llbracket \vec{\mathcal{D}}_{\Gamma}^{\dagger}(s) \rrbracket(\gamma))))(\pi_n(v))) \end{aligned}$$

$$\begin{aligned}
& + \pi_2(\llbracket \vec{\mathcal{D}}_{\Gamma}(s)(\gamma) \rrbracket)(\pi_2(\pi_2(\llbracket \vec{\mathcal{D}}_{\Gamma,x}(t) \rrbracket(\gamma, \pi_1(\pi_1(\llbracket \vec{\mathcal{D}}_{\Gamma}(s) \rrbracket(\gamma))))(\pi_1(v))), \dots, \\
& \quad \pi_2(\pi_2(\llbracket \vec{\mathcal{D}}_{\Gamma,x}(t) \rrbracket(\gamma, \pi_n(\pi_1(\llbracket \vec{\mathcal{D}}_{\Gamma}(s) \rrbracket(\gamma)))))(\pi_n(v))) \\
& = \sum_{i=1}^n \pi_1(\pi_2(\llbracket \vec{\mathcal{D}}_{\Gamma,x}(t) \rrbracket(\gamma, \pi_i(\pi_1(\llbracket \vec{\mathcal{D}}_{\Gamma}(s) \rrbracket(\gamma)))))(\pi_i(v))) \\
& + \pi_2(\llbracket \vec{\mathcal{D}}_{\Gamma}(s)(\gamma) \rrbracket)(0, \dots, 0, \pi_2(\pi_2(\llbracket \vec{\mathcal{D}}_{\Gamma,x}(t) \rrbracket(\gamma, \pi_1(\pi_1(\llbracket \vec{\mathcal{D}}_{\Gamma}(s) \rrbracket(\gamma)))))(\pi_1(v))), 0, \dots, 0),
\end{aligned}$$

where the last equation holds by linearity of $\pi_2(\llbracket \vec{\mathcal{D}}_{\Gamma}(s)(\gamma) \rrbracket)$.

A.2 Extending the Induction Proof of the Fundamental Lemma for Forward CHAD

First, we focus on extending the induction proof of the fundamental lemma for forward CHAD to apply to maps. Assume the induction hypothesis that t and s respect the logical relation. We show that $\mathbf{map}(x.t, s)$ does as well. (Where we assume all terms all well-typed.) Suppose that $(f, (g, h)) \in P_{\Gamma}$. We want to show that

$$(f', (g', h')) \stackrel{\text{def}}{=} (f; \llbracket \mathbf{map}(x.t, s) \rrbracket, (g; \llbracket \vec{\mathcal{D}}_{\Gamma}(\mathbf{map}(x.t, s)) \rrbracket; \pi_1, x \mapsto r \mapsto \pi_2(\llbracket \vec{\mathcal{D}}_{\Gamma}(\mathbf{map}(x.t, s)) \rrbracket(g(x)))(h(x)(r)))) \in P_{\text{real}^n}.$$

Note $f' = (f'_1, \dots, f'_n)$, $g' = (g'_1, \dots, g'_n)$ and $h'(x) = (h'_1(x), \dots, h'_n(x))$ that, as derivatives are computed componentwise, it is equivalent to show that $(f'_i, (g'_i, h'_i)) \in P_{\text{real}}$ for $i = 1, \dots, n$. That is, we need to show that

$$\begin{aligned}
& (x \mapsto \llbracket t \rrbracket(f(x), \pi_i(\llbracket s \rrbracket(f(x))))), (x \mapsto \pi_1(\llbracket \vec{\mathcal{D}}_{\Gamma,x}(t) \rrbracket(g(x), \pi_i(\pi_1(\llbracket \vec{\mathcal{D}}_{\Gamma}(s) \rrbracket(g(x)))))), \\
& \quad x \mapsto r \mapsto \pi_2(\llbracket \vec{\mathcal{D}}_{\Gamma,x}(t) \rrbracket(g(x), \pi_i(\pi_1(\llbracket \vec{\mathcal{D}}_{\Gamma}(s) \rrbracket(\gamma)))))(h(x)(r), \pi_i(\pi_2(\llbracket \vec{\mathcal{D}}_{\Gamma}(s) \rrbracket(g(x)))(h(x)(r)))) \in P_{\text{real}}.
\end{aligned}$$

As t respects the logical relation by our induction hypothesis, it is enough to show that

$$\begin{aligned}
& (x \mapsto (f(x), \pi_i(\llbracket s \rrbracket(f(x))))), (x \mapsto (g(x), \pi_i(\pi_1(\llbracket \vec{\mathcal{D}}_{\Gamma}(s) \rrbracket(g(x)))))) \\
& \quad x \mapsto r \mapsto (h(x)(r), \pi_i(\pi_2(\llbracket \vec{\mathcal{D}}_{\Gamma}(s) \rrbracket(g(x)))(h(x)(r)))) \in P_{\Gamma, x:\text{real}}.
\end{aligned}$$

Seeing that $(f, (g, h)) \in P_{\Gamma}$ by assumption, it is enough, by definition of $P_{\Gamma, x:\text{real}}$, to show that

$$\begin{aligned}
& (x \mapsto \pi_i(\llbracket s \rrbracket(f(x))), (x \mapsto \pi_i(\pi_1(\llbracket \vec{\mathcal{D}}_{\Gamma}(s) \rrbracket(g(x))))), \\
& \quad x \mapsto r \mapsto \pi_i(\pi_2(\llbracket \vec{\mathcal{D}}_{\Gamma}(s) \rrbracket(g(x)))(h(x)(r)))) \in P_{\text{real}}.
\end{aligned}$$

By definition of P_{real^n} , it is enough to show that

$$\begin{aligned}
& (x \mapsto \llbracket s \rrbracket(f(x)), (x \mapsto \pi_1(\llbracket \vec{\mathcal{D}}_{\Gamma}(s) \rrbracket(g(x))))), \\
& \quad x \mapsto r \mapsto \pi_2(\llbracket \vec{\mathcal{D}}_{\Gamma}(s) \rrbracket(g(x)))(h(x)(r))) \in P_{\text{real}}.
\end{aligned}$$

Seeing that s respects the logical relation by our induction hypothesis, it is enough to show that

$$\begin{aligned}
& (x \mapsto f(x), (x \mapsto g(x), \\
& \quad x \mapsto r \mapsto h(x)(r))) \in P_{\text{real}},
\end{aligned}$$

which is true by assumption.

A.3 Extending the Induction Proof of the Fundamental Lemma for Reverse CHAD

Next, we extend the fundamental lemma for reverse CHAD to apply to maps. Assume the induction hypothesis that t and s respect the logical relation. We show that $\mathbf{map}(x.t, s)$ does as well. (Where we assume all terms all well-typed.) Suppose that $(f, (g, h)) \in P_{\Gamma}$. We want to show that

$$(f', (g', h')) \stackrel{\text{def}}{=} (f; \llbracket \mathbf{map}(x.t, s) \rrbracket, (g; \llbracket \vec{\mathcal{D}}_{\Gamma}(\mathbf{map}(x.t, s)) \rrbracket; \pi_1, x \mapsto v \mapsto h(x)(\pi_2(\llbracket \vec{\mathcal{D}}_{\Gamma}(\mathbf{map}(x.t, s)) \rrbracket(g(x)))(v)))) \in P_{\text{real}^n}.$$

Observing that $(f', (g', h')) \in P_{\text{real}}^n$ are all of the form $f' = (f'_1, \dots, f'_n)$, $g' = (g'_1, \dots, g'_n)$ and $h'(x)(v) = h'_1(x)(\pi_1(v)) + \dots + h'_n(x)(\pi_n(v))$ where $(f'_i, (g'_i, h'_i)) \in P_{\text{real}}$ for $i = 1, \dots, n$, by basic multivariate calculus. That is, we need to show (by linearity of $h(x)$) that

$$\begin{aligned} & (x \mapsto \llbracket t \rrbracket(f(x), \pi_i(\llbracket s \rrbracket(f(x))))), (x \mapsto \pi_1(\llbracket \tilde{\mathcal{D}}_{\Gamma, x}^-(t) \rrbracket(g(x), \pi_i(\llbracket s \rrbracket(f(x))))), \\ & \quad x \mapsto v \mapsto h(x)(\pi_1(\pi_2(\llbracket \tilde{\mathcal{D}}_{\Gamma, x}^-(t) \rrbracket(g(x), \pi_i(\pi_1(\llbracket \tilde{\mathcal{D}}_{\Gamma}^-(s) \rrbracket(g(x)))))))(v)) \\ & \quad + \pi_2(\llbracket \tilde{\mathcal{D}}_{\Gamma}^-(s) \rrbracket(g(x)))(0, \dots, 0, \pi_2(\pi_2(\llbracket \tilde{\mathcal{D}}_{\Gamma, x}^-(t) \rrbracket(g(x), \pi_i(\pi_1(\llbracket \tilde{\mathcal{D}}_{\Gamma}^-(s) \rrbracket(g(x)))))))(v), 0, \dots, 0))) \in P_{\text{real}}. \end{aligned}$$

As t respects the logical relation by our induction hypothesis, it is enough to show that

$$\begin{aligned} & (x \mapsto (f(x), \pi_i(\llbracket s \rrbracket(f(x))))), (x \mapsto (g(x), \pi_i(\llbracket \tilde{\mathcal{D}}_{\Gamma}^-(s) \rrbracket(g(x))))), \\ & \quad x \mapsto v \mapsto h(x)(\pi_1(v) + \pi_2(\llbracket \tilde{\mathcal{D}}_{\Gamma}^-(s) \rrbracket(g(x)))(0, \dots, 0, \pi_2(v), 0, \dots, 0))) \in P_{\Gamma, x: \text{real}}. \end{aligned}$$

Seeing that $(f, (g, h)) \in P_{\Gamma}$ by assumption, we merely need to check the following, by definition of $P_{\Gamma, x: \text{real}}$:

$$\begin{aligned} & (x \mapsto \pi_i(\llbracket s \rrbracket(f(x))), (x \mapsto \pi_i(\llbracket \tilde{\mathcal{D}}_{\Gamma}^-(s) \rrbracket(g(x))), \\ & \quad x \mapsto v \mapsto h(x)(\pi_2(\llbracket \tilde{\mathcal{D}}_{\Gamma}^-(s) \rrbracket(g(x)))(0, \dots, 0, v, 0, \dots, 0))) \in P_{\text{real}}. \end{aligned}$$

By definition of P_{real}^n and linearity of $h(x)$, it is enough to show that

$$\begin{aligned} & (x \mapsto \llbracket s \rrbracket(f(x)), (x \mapsto \llbracket \tilde{\mathcal{D}}_{\Gamma}^-(s) \rrbracket(g(x))), \\ & \quad x \mapsto v \mapsto h(x)(\pi_2(\llbracket \tilde{\mathcal{D}}_{\Gamma}^-(s) \rrbracket(g(x)))(v))) \in P_{\text{real}}^n. \end{aligned}$$

Seeing that s respects the logical relation by our induction hypothesis, it is enough to show that

$$(f, (g, h)) \in P_{\Gamma},$$

which holds by assumption.

B TERM SIMPLIFICATIONS IN THE IMPLEMENTATION

Our implementation¹⁴ of the AD macros that is described in Section 9 includes a number of simplification rules on the concrete target language whose only purpose is to make the produced code more readable and clear (without changing its asymptotic runtime cost). The motivation for these rules is to be able to generate legible code when applying the AD macros to example programs. In this appendix, we list these simplification rules explicitly and show the implementation's output on the four example programs in Figures 1 and 2 under these simplification rules. We do this to illustrate that

- (1) the simplifications given here are evidently meaning-preserving, given the $\beta\eta$ -rules in Figures 4 and 6, and are standard rules that any optimizing compiler would apply;
- (2) the thus simplified output of the AD macros from Section 7 is indeed equivalent to the differentiated programs in Figures 1 and 2.

The simplification rules in question are given below in Table 1. In the implementation, these are (at the time of writing) implemented in the simplifier for the concrete target language.¹⁵

The last column in the table shows the justification for the simplification rule: “let substitution,” “ β pair,” “ η pair,” “lambda substitution,” and “ η lambda” refer to the corresponding rules in Figure 4. The equational rules are either from Figure 6 in the case of $t + \underline{0} = t$ and its symmetric variant, from

¹⁴As also mentioned in Section 9, the implementation is available at <https://github.com/VMatthijs/CHAD>.

¹⁵<https://github.com/VMatthijs/CHAD/blob/eedd6b12f224ed28ef9ca8650718d901c2b5e6a3/src/Concrete/Simplify.hs>.

Table 1. The Simplification Rules to Aid Legibility That Are Available in the CHAD Implementation in Haskell on the Concrete Target Language

Name	Rule	Justification
lamAppLet	$(\lambda x. e) a \rightsquigarrow \text{let } x = a \text{ in } e$	lambda subst., let subst.
letRotate	$\text{let } x = (\text{let } y = a \text{ in } b) \text{ in } e \rightsquigarrow \text{let } y = a \text{ in let } x = b \text{ in } e$	let substitution
letPairSplit	$\text{let } x = (a, b) \text{ in } e \rightsquigarrow \text{let } x_1 = a \text{ in let } x_2 = b \text{ in } e[\langle x_1, x_2 \rangle / x]$	let substitution
letInline	$\text{let } x = a \text{ in } e \rightsquigarrow e[a/x]$ (if a is cheap or used at most once in e)	let substitution
pairProj ₁	$\text{fst } \langle a, b \rangle \rightsquigarrow a$	β pair
pairProj ₂	$\text{snd } \langle a, b \rangle \rightsquigarrow b$	β pair
pairEta	$\langle \text{fst } a, \text{snd } a \rangle \rightsquigarrow a$	η pair
letProj ₁	$\text{fst } (\text{let } x = a \text{ in } e) \rightsquigarrow \text{let } x = a \text{ in fst } e$	let substitution
letProj ₂	$\text{snd } (\text{let } x = a \text{ in } e) \rightsquigarrow \text{let } x = a \text{ in snd } e$	let substitution
plusZero ₁	$\text{plus zero } a \rightsquigarrow a$	equational rule
plusZero ₂	$\text{plus } a \text{ zero} \rightsquigarrow a$	equational rule
plusPair	$\text{plus } \langle a, b \rangle \langle c, d \rangle \rightsquigarrow \langle \text{plus } a \ c, \text{plus } b \ d \rangle$	equational rule
plusLet ₁	$\text{plus } (\text{let } x = e \text{ in } a) \ b \rightsquigarrow \text{let } x = e \text{ in plus } a \ b$	let substitution
plusLet ₂	$\text{plus } a \ (\text{let } x = e \text{ in } b) \rightsquigarrow \text{let } x = e \text{ in plus } a \ b$	let substitution
algebra	$0 * x, 0 * x \rightsquigarrow 0$ (etc.)	basic algebra
letLamPairSplit	$\text{let } f = \lambda x. \langle a, b \rangle \text{ in } e \rightsquigarrow \text{let } f_1 = \lambda x. a \text{ in let } f_2 = \lambda x. b \text{ in } e[\langle f_1 \ x, f_2 \ x \rangle / f]$	η lambda, let subst.
mapPairSplit	$\text{map } (\lambda x. \langle b, c \rangle) a \rightsquigarrow \text{let } a' = a \text{ in } \langle \text{map } (\lambda x. b) \ a', \text{map } (\lambda x. c) \ a' \rangle$	equational rule
mapZero	$\text{map } (\lambda x. \text{zero}) a \rightsquigarrow \text{zero}$	equational rule
sumZip	$\text{sum } (\text{zip } a \ b) \rightsquigarrow \langle \text{sum } a, \text{sum } b \rangle$	equational rule
sumZero	$\text{sum zero} \rightsquigarrow \text{zero}$	equational rule
sumSingleton	$\text{sum } (\text{map } (\lambda x. [x]) \ e) \rightsquigarrow e$	equational rule

the type-based translation rules in Section 9.2 in the case of plus on pairs, or otherwise general laws that hold for **zero** (i.e., 0) and/or the array combinators in question.

Note that all of the rules preserve time complexity of the program through careful sharing of values with let-bindings. These let-bindings could only be argued to increase work if the value is only used once in the body of the **let** – but in that case, the “letInline” rule will eliminate the let-binding anyway.

B.1 First-order Example Programs

The output of our implementation for the forward derivative of Figure 1(a) and the reverse derivative of Figure 1(b) can be found below in Figure 9.

The simplification rules listed above have already been applied (otherwise the output would, indeed, be much less readable). The only change we made to the literal text output of the implementation is formatting and renaming of variables.

For both programs, we note that in the implementation, environments are encoded using snoc-lists: that is, the environment $\Gamma = x_1 : \text{real}, x_2 : \text{real}, x_3 : \text{real}$ is represented as $((\epsilon, x_1 : \text{real}), x_2 : \text{real}), x_3 : \text{real}$. Hence, for this Γ , $\vec{D}(\Gamma)_2$ would be represented as $((1*\text{real})*\text{real})*\text{real}$. This has an effect on the code in Figure 9, where in subfigure (a), the variable x' has type $1*\text{real}$ instead of **real**, which it had in Figure 1(c). Furthermore, the output in Figure 9(b), which is the cotangent (i.e., adjoint) of the environment, has type $((((1*\text{real})*\text{real})*\text{real})*\text{real}$, meaning that the “zero” term in the result has type 1 and is thus equal to $\langle \rangle$.

We hope that it is evident to the reader that these outputs are equivalent to the programs given in Figure 1(c) and (d).


```

let two = 2.0
    y  = two * x
    z  = x * y in
  <<(y, z), cos z>,
  λx'.
    let y' = two * snd x'
        z' = plus (x * y')
              (y * snd x') in
  <<(y', z')
  ,(0.0 - sin z) * z'>

let two = 2.0
    y  = x1 * x4 + two * x2
    w  = y * x3 + x4 in
  <sin w
  ,λv'.
    let w' = cos w * v'
        y' = x3 * w' in
  <<<<zero
    ,x4 * y'>
    ,two * y'>
    ,y * w'>
    ,plus w' (x1 * y')>>>

```

(a) The implementation's forward AD output on Fig. 1 (a). Compare this with Fig. 1 (c).

(b) The implementation's reverse AD output on Fig. 1 (b). Compare this with Fig. 1 (d).

Fig. 9. Output of our Haskell implementation when executed on the first-order example programs in Figure 1(a) and (b).

B.2 Second-order Example Program Figure 2(a)

The implementation's forward derivative of Figure 2(a) is shown below in Figure 10. This version contains a let-bound function “g” that does not occur in the code of Figure 2(c). However, inlining this function in the two places where it is used does not increase work, because due to pair projection and equational rules concerning “zero” and “plus,” only one half of the “plus” expression in “g” remains at both of the invocation sites of “g.” A version with “g” manually inlined and simplified using the stated rules is shown in Figure 11. (Our automatic simplifier is currently not smart enough to prove that inlining “g” does not increase work, and hence keeps it let-bound.)

```

let f' = λz. λd. plus (x * snd d) (z * snd (fst d))
    zs = vreplicate x
    g = λz. let dfun = f' z in
          (x * z + 1.0
          ,λdenv. plus (snd (fst (fst denv)) z)
                (dfun (zero , snd denv)))
in (vmap (λz. fst (g z)) zs
    ,λx'.
      let zs' = vreplicate (snd x')
          in plus
            (vzipWith (λz. λz'. snd (g z) (zero , z'))
              zs zs')
            (vmap (λz. snd (g z) (((x', λz. f' z (x', zero)), zs'), zero))
              zs))

```

Fig. 10. Output of our Haskell implementation of the forward AD macro when executed on Figure 2(a). The implementation writes operations on scalar arrays (**realⁿ**) with a “v” prefix.

First, we notice that the type of the variable x' here is **1*real** instead of **real** because of the snoc-list representation of environments, as was discussed in the previous subsection. Because of this, the expression **snd** x' in Figure 11 is equivalent to the expression x' in Figure 2(c). Knowing this, let us work out how the implementation got this output code, and how it is equivalent to the code in Figure 2(c).

```

let f' = λz. λv. plus (x * snd v) (z * snd (fst v))
    zs = vreplicate x in
(vmap (λz. x * z + 1.0) zs
,λx'.
    let zs' = vreplicate (snd x') in
    plus
      (vzipWith (λz. λz'. f' z (zero, z'))
              zs zs')
      (vmap (λz. f' z (x', zero)) zs))

```

Fig. 11. Manually simplified code from Figure 10, as described in the text.

For the purposes of this explanation, the most important component of the source of Figure 2(a) is its first line: “**let** $f = \lambda z. x * z + 1$ **in**...” Recall from the forward AD macro from Section 7.4:

$$\vec{\mathcal{D}}_{\Gamma}(\text{let } x = t \text{ in } s) \stackrel{\text{def}}{=} \text{let } \langle x, x' \rangle = \vec{\mathcal{D}}_{\Gamma}(t) \text{ in let } \langle y, y' \rangle = \vec{\mathcal{D}}_{\Gamma, x}(s) \text{ in } \langle y, \lambda v. y' \bullet \langle v, x' \bullet v \rangle \rangle.$$

Hence, the binding of t is transformed to a binding of $\vec{\mathcal{D}}_{\Gamma}(t)$, which is then used somehow in the body. Since t is the term “ $\lambda z. x * z + 1$ ” here, and since the macro rule for lambda abstraction is as follows:

$$\vec{\mathcal{D}}_{\Gamma}(\lambda x. t) \stackrel{\text{def}}{=} \text{let } y = \lambda x. \vec{\mathcal{D}}_{\Gamma, x}(t) \text{ in } \langle \lambda x. \text{let } \langle z, z' \rangle = y x \text{ in } \langle z, \lambda v. z' \bullet \langle \underline{0}, v \rangle \rangle, \lambda v. \lambda x. (\text{snd } (y x)) \bullet \langle v, \underline{0} \rangle \rangle,$$

we get the following result for $\vec{\mathcal{D}}_{\Gamma}(\lambda z. x * z + 1)$ with $\Gamma = \epsilon, x : \text{real}$:

$$\begin{aligned} \vec{\mathcal{D}}_{\epsilon, x}(\lambda z. x * z + 1) = & \text{let } y = \lambda z. \langle x * z + 1, \lambda \langle \langle \rangle, x' \rangle. x * z' + z * x' \rangle \\ & \text{in } \langle \lambda z. \text{let } \langle u, u' \rangle = y z \text{ in } \langle u, \lambda v. u' \bullet \langle \underline{0}, v \rangle \rangle, \lambda v. \lambda z. (\text{snd } (y z)) \bullet \langle v, \underline{0} \rangle \rangle. \end{aligned}$$

This is the term that appears on the right-hand side of a let-binding in the forward AD transformed version of the code from Figure 2(a). Inlining of y and some further simplification yields:

$$\vec{\mathcal{D}}_{\epsilon, x}(\lambda z. x * z + 1) = \langle \lambda z. \langle x * z + 1, \lambda v. x * v \rangle, \lambda v. \lambda z. z * \text{snd } v \rangle.$$

wherein we recognize f and f' from Figure 2(c).

The implementation proceeded differently in the simplification of $\vec{\mathcal{D}}_{\epsilon, x}(\lambda z. x * z + 1)$, namely by splitting the lambda bound to y using “letLamPairSplit.” The second of the resulting two lambda functions is $\lambda z. \lambda \langle \langle \rangle, x' \rangle. z'. x * z' + z * x'$, or by desugaring pattern matching, $\lambda z. \lambda v. x * \text{snd } v + z * \text{snd } (\text{fst } v)$. We recognise this expression as the right-hand side of the f' binding in Figure 11.

After seeing this, we leave it to the reader to show that the code in Figures 11 and 2(c) is equivalent under forward and reverse let-substitution.

B.3 Second-order Example Program Figure 2(b)

When the implementation performs reverse AD on the code in Figure 2(b) and simplifies the result using the simplification rules in Table 1, the result is the code that is shown below in Figure 12.

First, note the evalOp EScalProd. Since scalar multiplication is implemented as an operation (op), it does not take two separate arguments but rather a pair of the two scalars to multiply. Due to application of the “pairEta” simplification rule from Table 1, the argument to the multiplication was reduced from $\langle \text{fst } p, \text{snd } p \rangle$ to just “ p ,” preventing the pretty-printer from showing $\text{fst } p * \text{snd } p$; instead, this becomes a bare operation application to the argument “ p .” In the code in Figure 2(d), this lambda is the argument to the “map” in the cotangents block, meaning that “ p ” stands for $\langle x2i, y' \rangle$.

```

(vsum (vmap ( $\lambda x2i. x1 * x2i$ ) x2)
,  $\lambda w'$ .
  let ys' = vreplicate w' in
  (<<zero, sum (map ( $\lambda p. \text{evalOp } \text{EScalProd } p$ )
    (zip (toList x2) (toList ys')))))
, vzipWith ( $\lambda x2i. \lambda y'. x1 * y'$ ) x2 ys'))

```

Fig. 12. Output of our Haskell implementation of the reverse AD macro when executed on Figure 2(b). The implementation writes operations on scalar arrays (*realⁿ*) with a “v” prefix; list operations are written without a prefix. Contrast this to Figure 2, where all arrays are implemented as lists and hence there are no pure-array operations that would have been written using a “v” prefix.

In this example, a copower structure is created because of the usage of a function abstraction in the code to be differentiated using reverse AD. Here, this copower was interpreted using lists as described in Section 9.2. The “toList” function converts an *array* of scalars (i.e., a value of type *realⁿ*) to a *list* of scalars. The code in Figure 2(d) goes further and interprets all arrays as lists, effectively removing the distinction between arrays originating from arrays in the source program and lists originating from copower values. Together with recognizing the snoc-list representation of environments, it becomes sufficient to inline some let-bindings in Figure 2(d) to arrive at code equivalent to the code shown here.

ACKNOWLEDGMENTS

We are grateful to the anonymous reviewers who gave excellent comments on earlier versions of this article that prompted various much-needed rewrites.

REFERENCES

- [1] Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. 2016. Tensorflow: A system for large-scale machine learning. In *Proceedings of the 12th USENIX Symposium on Operating Systems Design and Implementation (OSDI'16)*. 265–283.
- [2] Martín Abadi and Gordon D. Plotkin. 2020. A simple differentiable programming language. In *Proceedings of the ACM SIGPLAN Symposium on Principles of Programming Languages (POPL'20)*. ACM.
- [3] Andrew Barber and Gordon Plotkin. 1996. *Dual Intuitionistic Linear Logic*. University of Edinburgh, Department of Computer Science, Laboratory for Foundations of Computer Science.
- [4] Gilles Barthe, Raphaëlle Crubillé, Ugo Dal Lago, and Francesco Gavazzo. 2020. On the versatility of open logical relations - continuity, automatic differentiation, and a containment theorem. In *Proceedings of the 29th European Symposium on Programming (ESOP'20) Held as Part of the European Joint Conferences on Theory and Practice of Software (ETAPS'20)*, Lecture Notes in Computer Science, Vol. 12075, Peter Müller (Ed.). Springer, 56–83. https://doi.org/10.1007/978-3-030-44914-8_3
- [5] Atilim Gunes Baydin and Barak A. Pearlmutter. 2018. Automatic differentiation in machine learning: a survey. *Journal of Machine Learning Research* 18, 1 (2018), 1–43.
- [6] P. Nick Benton. 1994. A mixed linear and non-linear logic: Proofs, terms and models. In *Proceedings of the International Workshop on Computer Science Logic*. Springer, 121–135.
- [7] Michael Betancourt, Charles C. Margossian, and Vianey Leos-Barajas. 2020. The discrete adjoint method: Efficient derivatives for functions of discrete sequences. arXiv:2002.00326. Retrieved from <https://arxiv.org/abs/2002.00326>.
- [8] Alois Brunel, Damiano Mazza, and Michele Pagani. 2020. Backpropagation in the simply typed lambda-calculus with linear negation. In *Proceedings of the ACM SIGPLAN Symposium on Principles of Programming Languages (POPL'20)*.
- [9] Bob Carpenter, Matthew D. Hoffman, Marcus Brubaker, Daniel Lee, Peter Li, and Michael Betancourt. 2015. The Stan math library: Reverse-mode automatic differentiation in C++. arXiv:1509.07164. Retrieved from <https://arxiv.org/abs/1509.07164>.
- [10] Manuel M. T. Chakravarty, Gabriele Keller, Sean Lee, Trevor L. McDonell, and Vinod Grover. 2011. Accelerating Haskell array codes with multicore GPUs. In *Proceedings of the 6th Workshop on Declarative Aspects of Multicore Programming*. 3–14.

- [11] Manuel M. T. Chakravarty, Gabriele Keller, Sean Lee, Trevor L. McDonell, and Vinod Grover. 2011. Accelerating haskell array codes with multicore GPUs. In *Proceedings of the 6th Workshop on Declarative Aspects of Multicore Programming (DAMP'11)*. ACM, New York, NY, 3–14. <https://doi.org/10.1145/1926354.1926358>
- [12] J. Robin B. Cockett, Geoff S. H. Cruttwell, Jonathan Gallagher, Jean-Simon Pacaud Lemay, Benjamin MacAdam, Gordon D. Plotkin, and Dorette Pronk. 2020. Reverse derivative categories. In *Proceedings of the Annual Conference on Computer Science Logic (CSL'20)*.
- [13] P.-L. Curien. 1986. Categorical combinators. *Inf. Contr.* 69, 1–3 (1986), 188–254.
- [14] Jeff Egger, Rasmus Ejlers Møgelberg, and Alex Simpson. 2009. Enriching an effect calculus with linear types. In *Proceedings of the International Workshop on Computer Science Logic*. Springer, 240–254.
- [15] Thomas Ehrhard. 2018. An introduction to differential linear logic: Proof-nets, models and antiderivatives. *Math. Struct. Comput. Sci.* 28, 7 (2018), 995–1060.
- [16] Conal Elliott. 2018. The simple essence of automatic differentiation. *Proc. ACM Program. Lang.* 2, ICFP (2018), 1–29.
- [17] Marcelo P. Fiore. 2007. Differential structure in models of multiplicative biadditive intuitionistic linear logic. In *Proceedings of the International Conference on Typed Lambda Calculi and Applications*. Springer, 163–177.
- [18] Ralf Hannemann-Tamas, Diego A. Munoz, and Wolfgang Marquardt. 2015. Adjoint sensitivity analysis for nonsmooth differential-algebraic equation systems. *SIAM J. Sci. Comput.* 37, 5 (2015), A2380–A2402.
- [19] Troels Henriksen, Niels G. W. Serup, Martin Elsmann, Fritz Henglein, and Cosmin E. Oancea. 2017. Futhark: Purely functional GPU-programming with nested parallelism and in-place array updates. In *Proceedings of the 38th ACM SIGPLAN Conference on Programming Language Design and Implementation*. 556–571.
- [20] Mathieu Huot, Sam Staton, and Matthijs Vákár. 2022. Higher order automatic differentiation of higher order functions. *Logical Methods in Computer Science* 18, 1 (2022), 41:1–41:34.
- [21] Mathieu Huot, Sam Staton, and Matthijs Vákár. 2020. Correctness of automatic differentiation via diffeologies and categorical gluing. In *Proceedings of the International Conference on Foundations of Software Science and Computation Structures (FoSSaCS'20)*.
- [22] Michael Innes. 2018. Don't unroll adjoint: Differentiating SSA-Form programs. arXiv:1810.07951. Retrieved from <https://arxiv.org/abs/1810.07951>.
- [23] Peter T. Johnstone. 2002. *Sketches of An Elephant: A Topos Theory Compendium*, Vol. 2. Oxford University Press.
- [24] Peter T. Johnstone, Stephen Lack, and P. Sobocinski. 2007. Quasitoposes, quasiadhesive categories and artin glueing. In *Proceedings of the Conference on Algebra and Coalgebra in Computer Science (CALCO'07)*.
- [25] Faustyna Krawiec, Neel Krishnaswami, Simon Peyton Jones, Tom Ellis, Andrew Fitzgibbon, and R. Eisenberg. 2022. Provably correct, asymptotically efficient, higher-order reverse-mode automatic differentiation. *Proc. ACM Program. Lang.* 6, POPL (2022), 1–30.
- [26] Andreas Kriegl and Peter W. Michor. 1997. *The Convenient Setting of Global Analysis*, Vol. 53. American Mathematical Soc.
- [27] Joachim Lambek and Philip J. Scott. 1988. *Introduction to Higher-order Categorical Logic*, Vol. 7. Cambridge University Press.
- [28] Serge Lang. 2002. *Algebra*. Springer, New York, NY.
- [29] Paul Blain Levy. 2012. *Call-by-push-value: A Functional/imperative Synthesis*, Vol. 2. Springer Science & Business Media.
- [30] Fernando Lucatelli Nunes and Matthijs Vákár. 2021. CHAD for expressive total languages. arXiv:2110.00446. Retrieved from <https://arxiv.org/abs/2110.00446>.
- [31] Carol Mak and Luke Ong. 2020. A differential-form pullback programming language for higher-order reverse-mode automatic differentiation. arxiv:2002.08241. Retrieved from <https://arxiv.org/abs/2002.08241>.
- [32] Charles C. Margossian. 2019. A review of automatic differentiation and its efficient implementation. *Data Min. Knowl. Discov.* 9, 4 (2019), e1305.
- [33] Damiano Mazza and Michele Pagani. 2021. Automatic differentiation in PCF. *Proc. ACM Program. Lang.* 5, POPL (2021), 1–27.
- [34] Trevor L. McDonell, Manuel M. T. Chakravarty, Gabriele Keller, and Ben Lippmeier. 2013. Optimising purely functional GPU programs. In *Proceedings of the 18th ACM SIGPLAN International Conference on Functional Programming (ICFP'13)*. ACM, New York, NY, 49–60. <https://doi.org/10.1145/2500365.2500595>
- [35] Paul-André Mellies. 2009. Categorical semantics of linear logic. *Panor. Synth.* 27 (2009), 15–215.
- [36] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. 2017. Automatic differentiation in pytorch. In *NIPS 2017 Workshop on Autodiff*. <https://openreview.net/forum?id=BJJsrnmcZ>.
- [37] Barak A. Pearlmutter and Jeffrey Mark Siskind. 2008. Reverse-mode AD in a functional framework: Lambda the ultimate backpropagator. *ACM Trans. Program. Lang. Syst.* 30, 2 (2008), 7.

- [38] Amir Shaikhha, Andrew Fitzgibbon, Dimitrios Vytiniotis, and Simon Peyton Jones. 2019. Efficient differentiable programming in a functional array-processing language. *Proc. ACM Program. Lang.* 3, ICFP (2019), 97.
- [39] Lau Skorstengaard. 2019. An introduction to logical relations. arXiv:1907.11133. <http://arxiv.org/abs/1907.11133>.
- [40] Periklis Tsiros, Frederic Y. Bois, Aristides Dokoumetzidis, Georgia Tsiliki, and Haralambos Sarimveis. 2019. Population pharmacokinetic reanalysis of a Diazepam PBPK model: A comparison of stan and GNU MCSim. *J. Pharmacokinet. Pharmacodynam.* 46, 2 (2019), 173–192.
- [41] Matthijs Vákár. 2017. In search of effectful dependent types. arXiv:1706.07997. Retrieved from <https://arxiv.org/abs/1706.07997>.
- [42] Matthijs Vákár. 2021. Reverse AD at higher types: Pure, principled and denotationally correct. In *Proceedings of the European Symposium on Programming (ESOP'21)*.
- [43] Dimitrios Vytiniotis, Dan Belov, Richard Wei, Gordon Plotkin, and Martin Abadi. 2019. The differentiable curry. In *NeurIPS Workshop Program Transformations*.
- [44] Fei Wang, Xilun Wu, Gregory Essertel, James Decker, and Tiark Rompf. 2019. Demystifying differentiable programming: Shift/reset the penultimate backpropagator. *Proc. ACM Program. Lang.* 3, ICFP (2019).
- [45] Paul W. Wilson and Fabio Zanasi. 2020. Reverse derivative ascent: A categorical approach to learning boolean circuits. In *Proceedings of the 3rd Annual International Applied Category Theory Conference (ACT'20)*, David I. Spivak and Jamie Vicary (Eds.). 247–260. <https://doi.org/10.4204/EPTCS.333.17>

Received March 2021; revised October 2021; accepted March 2022