

Categorical semantics of a simple differential programming language

Geoffrey Cruttwell

Mount Allison University, Sackville, Canada

gcruttwell@mta.ca

Jonathan Gallagher

Dalhousie University, Halifax, Canada

jonathan.gallagher@dal.ca

Dorette Pronk

Dalhousie University, Halifax, Canada

dorette.pronk@dal.ca

1 Introduction

With the increased interest in machine learning, and deep learning in particular (where one extracts progressively higher level features from data using multiple layers of processing), the use of automatic differentiation has become more wide-spread in computation. See for instance the surveys given in [32] and [4]. In fact, Facebook’s Chief AI scientist Yann LeCun has gone as far as famously exclaiming:

“Deep learning est mort. Vive Differentiable Programming! ...people are now building a new kind of software by assembling networks of parameterized functional blocks and by training them from examples using some form of gradient-based optimization.”¹

The point being that differentiation is no longer being viewed as merely a useful tool when creating software, but instead becoming viewed as a fundamental building block. This sort of ubiquity warrants a more in-depth study of automatic differentiation with a focus on treating it as a fundamental component.

There have been two recent developments to provide the theoretical support for this type of structure. In fact, the settings described above use two types of differentiation: the usual forward derivative to analyse the effect of changes in the data, as well as the reverse derivative to allow for error correction (i.e., training) through the efficient calculation of the gradients of functions. Thus, any theoretical approach must be able to deal with both types of differentiation. One approach is presented in [2], where Abadi and Plotkin provide a simple differential programming language with conditionals, recursive function definitions, and a notion of reverse-mode differentiation (from which forward differentiation can be derived) together with both a denotational and an operational semantics, and theorems showing that the two coincide. Another approach is given in [16], where the authors present reverse differential categories, a categorical setting for reverse differentiation. They also show how every reverse differential category gives rise to a (forward) derivative and a canonical “contextual linear dagger” operation. The converse is true as well: a category with a forward derivative (that is, a Cartesian differential category [8]) with a contextual linear dagger has a canonical reverse derivative.

In the present paper we bring these two approaches together. In particular, we show how an extension of reverse derivative categories models Abadi and Plotkin’s language, and describe how this categorical model allows one to consider potential improvements to the operational semantics of the language. To model Abadi and Plotkin’s language categorically, reverse derivative categories are not sufficient, due to their inability to handle partial functions and control structures. Thus, we need to add partiality to reverse differential categories. The standard categorical machinery to model partiality is restriction structure, which assigns to each map a partial identity map, subject to axioms as described in [14]. Combining this structure with reverse differential structure, we introduce *reverse differential restriction categories*. In

¹Facebook post on Jan 5, 2018

addition to the list of axioms **[RD.1 – 7]** given for reverse differential categories in [16], we require two additional axioms expressing how the restriction of the reverse derivative of a function is related to the restriction of the function itself and what the reverse derivative of a restriction of a function needs to be (cf. Definition 3.1). The results characterizing the relationship between differential and reverse differential categories in terms of a contextual linear dagger extend to the context of restriction categories. We also get for free that the reverse derivative preserves the order on the maps and preserves joins of maps, if they exist.

In Section 4 we show how Abadi and Plotkin’s language can be modelled in a reverse differential restriction category. We do this in two steps: at first we modify their language by omitting general recursion and instead only include while-loops. While-loops can be modelled in terms of recursion, but by separating them out we can see that source-transformation techniques (not discussed explicitly in Abadi and Plotkin but used in some commercial systems such as Theano [5], TensorFlow [1], and Tangent [45]) always hold in our semantics (see Section 4.3); source transformation techniques are not used for general recursion. We also note that in order to be able to push differentiation through the control structure, Abadi and Plotkin need that for each predicate the inverse images of *true* and *false* are both open sets. In the context of restriction categories we model this instead by providing two partially defined maps into the terminal object 1 for each predicate symbol (one for true and one for false) with the requirement that their restrictions do not overlap (cf. Section 4.2).

In the process of modelling Abadi and Plotkin’s language, we see that not all our axioms are needed. Specifically, Axioms **[RD.6]** and **[RD.7]** of a reverse differential restriction category (which deal with the behaviour of repeated reverse derivatives) are not strictly necessary to model Abadi and Plotkin’s language. However, in the final section of the paper, we show that if these axioms are present, changes can be made to the operational semantics to improve the efficiency and applicability of the simple differential programming language.

Abadi and Plotkin’s language represents an approach that makes the reverse derivative a language primitive in a functional language. Other approaches have been proposed to use reverse-mode accumulation for computing the derivative in a functional language. Given a function $\mathbb{R}^n \xrightarrow{f} \mathbb{R}^m$, Pearlmutter and Siskind [39] discuss how to compute the Jacobian matrix of f in a functional language by performing transformations on the function’s computational graph. This idea is similar to the symbolic differentiation of trace or tape terms in Abadi and Plotkin’s language. Elliot [26] shows how to view this sort of reverse-mode accumulation using continuations: when a function, written as a composition of simple operations, is written in continuation passing style, the reverse derivative of its computation graph corresponds to a sort of generalized derivative of the continuation. In [46], Wang et al extend Pearlmutter and Siskind’s work by showing that the move to continuations allows for getting around the *nonlocality* issue in the earlier work. Brunel et al [10] extend Wang’s work from the point of view of linear logic, and allow for additional analyses based on tracking the linearity of a variable. Abadi and Plotkin’s work represents a next step in this area by considering, in addition, control flow structures and general recursive functions.

This current work contrasts to work submitted to ACT 2019 on modelling differential programming using synthetic differential geometry (SDG) (see [34, 35] for an introduction to SDG). In the previous work a simple differential programming language featuring forward differentiation was introduced and an interpretation into a well-adapted model of SDG was given (see e.g. [23] for such models). The focus was on exploring what programming languages features might be able to exist soundly with differential programming. The current work develops the categorical semantics of Abadi and Plotkin’s language for reverse differentiation as well as the categorical semantics of source-transformations for their language. In particular we show that the operational semantics is modelled soundly by a denotational semantics into

our categorical framework. We will also see that using the axiomatic approach developed here leads to a sound exponential speedup when computing the reverse derivative of looping-phenomena.

2 Background: Relevant Categorical Structures

In this section, we briefly review some of the relevant structures from category theory which we will make use of.

2.1 Cartesian and reverse differential categories

The canonical category for differentiation is the category *Smooth* whose objects are the powers of the reals \mathbb{R} ($\mathbb{R}^0 = \{1\}, \mathbb{R}, \mathbb{R}^2, \mathbb{R}^3$, etc.) and whose maps are the smooth (infinitely differentiable) maps between them. To any map $f : A \rightarrow B$ in this category, there is an associated map

$$D[f] : A \times A \rightarrow B$$

whose value at $(x, v) \in A \times A$ is $J(f)(x) \cdot v$, the Jacobian of f at x , taken in the direction v . This map satisfies various rules; for example, the chain rule is equivalent to the statement that for any maps $f : A \rightarrow B, g : B \rightarrow C$,

$$D[fg] = \langle \pi_0 f, D[f] \rangle D[g].$$

(Note that we use the path-order for composition, so fg means “first f then g ”.) Many other familiar rules from calculus can be expressed via D ; for example, the symmetry of mixed partial derivatives can be expressed as a condition on $D^2[f] = D[D[f]]$.

Definition 2.1. ([8, Defn. 2.1.1]) A *Cartesian differential category* or *CDC* is a Cartesian left additive category ([8, Defn. 1.3.1]) which has, for any map $f : A \rightarrow B$, a map

$$D[f] : A \times A \rightarrow B$$

satisfying seven axioms [CD.1–7].

The formulation of CDCs and indeed the other flavours of categories with derivatives we will use have the intent that in $\langle a, v \rangle D[f]$, a is the point and v is the direction; this is in contrast to the original formulation of CDCs which had the point and direction swapped, and we chose the point-direction formulation because most of the literature follows this convention. This causes a change to axioms **CD.2,6,7**.

While *Smooth* is the canonical example, there are many others, including examples from infinite dimensional vector spaces, synthetic differential geometry, algebraic geometry, differential lambda calculus, etc: see [8, 28, 11, 19, 15].

In contrast, the reverse derivative, widely used in machine learning for its efficiency, is an operation which takes a smooth map $f : A \rightarrow B$ and produces a smooth map

$$R[f] : A \times B \rightarrow A$$

whose value at $(x, w) \in A \times B$ is $J(f)^T(x) \cdot w$, the transpose of the Jacobian of f at x , taken in the direction w .

There are two possible ways to categorically axiomatize the reverse derivative. One way is to start with a CDC and ask for a dagger structure (representing the transpose); one could then use the dagger

with the D from the CDC to define a reverse derivative R . However, there is some subtlety in this: the dagger structure is only present on the *linear* maps of the category, not on all the maps of the category. The other way is to axiomatize R directly, as was done in [16].

Definition 2.2. ([16, Defn. 13]) A **reverse differential category** or **RDC** is a Cartesian left additive category which has, for any map $f : A \rightarrow B$, a map

$$R[f] : A \times B \rightarrow A$$

satisfying seven axioms.

For example, in this formulation the chain rule is equivalent to **[RD.5]**, the rule that for any maps $f : A \rightarrow B$, $g : B \rightarrow C$,

$$R[fg] = \langle \pi_0, (f \times 1)R[g] \rangle R[f].$$

Moreover, there is something striking about a reverse differential structure: any RDC is automatically a CDC. If one applies the reverse derivative twice, 0's out a component and projects, the result is the forward derivative. That is, the following defines a (forward) differential structure from a reverse differential structure (see [16, Theorem 16]):

$$D[f] := A \times A \xrightarrow{\langle \pi_0, 0, \pi_1 \rangle} (A \times B) \times A \xrightarrow{R[R[f]]} A \times B \xrightarrow{\pi_1} B$$

Thus, while a “dagger on linear maps” is required to derive an RDC from a CDC, no such structure is required to go from an RDC to a CDC. In fact, one can show that a CDC with a “dagger on linear maps” is equivalent to an RDC: see Theorem 42 in [16].

For this reason, as well as the fact that the reverse derivative is of greater importance in machine learning, in this paper we take a reverse differential category to be the primary structure.

2.2 Restriction categories and differential restriction categories

Of course, to model a real-world programming language which involves non-terminating computations, we must also be able to handle partial functions. For this, we turn to restriction categories [14], which allow one to algebraically model categories whose maps may only be partially defined. Consider the category of sets and partial functions between them. To any map $f : A \rightarrow B$ in this category, there is an associated “partial identity” map $\bar{f} : A \rightarrow A$, which is defined to be the identity wherever f is defined, and undefined otherwise. This operation then has various properties such as $\bar{f}f = f$. This is then axiomatized:

Definition 2.3. ([14, Defn. 2.1.1]) A **restriction category** is a category which has for any map $f : A \rightarrow B$, a map $\bar{f} : A \rightarrow A$ satisfying various axioms.

In section 3, we will combine restriction structure with reverse differential structure to get the categorical structure we will use to model Abadi and Plotkin’s language.

Before we get to that, however, we will need to briefly review a few definitions from restriction category theory. It will also be helpful to consider the previously defined combination of restriction structure and (forward) differential structure.

A restriction category allows one to easily talk about when a map is “less than or equal to” a parallel map and when two parallel maps are “compatible”:

Definition 2.4. Suppose $f, g : A \rightarrow B$ are maps in a restriction category. Write $f \leq g$ if $\overline{f}g = f$, and write $f \sim g$ (and say “ f is compatible with g ”) if $\overline{f}g = \overline{g}f$.

That is, $f \leq g$ if g is defined wherever f is defined, and when restricted to f ’s domain of definition, g is equal to f ; $f \sim g$ if f and g are equal where they are both defined. One can show that \leq is a partial order on each hom-set; in fact, restriction categories are canonically partial order enriched by \leq .

Being able to “join” two compatible maps will be important when we define control structures such as “if” and “while”, as we will be able to discuss when maps are “disjoint”.

Definition 2.5. 1. If $f, g : A \rightarrow B$ and there is a least upper bound $f \vee g$ with respect to the partial order defined above, we call $f \vee g$ the **join** of f and g . Note that this implies that f and g are compatible.

2. The notion of join extends to families of maps that are pairwise compatible, and we write $\bigvee_i f_i$ to denote the join of the pairwise compatible family.

3. Say that a map $\emptyset : A \rightarrow B$ is **nowhere defined** if \emptyset is the minimum in the partial order.

4. Say that $f, g : A \rightarrow B$ are **disjoint** if $\overline{f}g$ is nowhere defined. Any two disjoint maps are compatible.

The formalization of disjoint joins in a restriction category was given in [18] as part of the story of formalizing Hoare semantics in a *classical* restriction category. Further analysis of joins in restriction categories was provided in [21]. Giles [31] used disjoint joins in connecting restriction categories to the semantics of reversible computing. Disjoint joins in partial map categories correspond to disjoint joins of monics, which often give a coproduct (e.g. as in coherent categories). One way to model iteration is to have a traced coproduct, and this can be directly expressed using disjoint joins: this approach was used in formalizing iteration in restriction categories and to build a partial combinatory algebra by iterating a step-function in [20, 13]. The formalization of iteration using disjoint joins was based on the work of Conway [22]. Another approach to formalizing the semantics of iterative processes in a category using algebraic formalizations was introduced in [25], refined in [7], and further developed categorically in [3].

Finally, it is worth noting that there has been previous work combining CDC structure with restriction structure [12]. The canonical example of such a category is the category of smooth *partial* maps between the \mathbb{R}^n ’s. The partiality acts in a compatible way with the derivative, as $D[f] : A \times A \rightarrow B$ is entirely defined in the second (vector) component: that is, the only partiality $D[f]$ has is from f itself. Thus, in a “differential restriction category”, one asks that $\overline{D[f]} = \overline{f} \times 1$. These are formulated on top of the notion of *cartesian left additive restriction category*: these are restriction categories with restriction products (which is a lax notion of product for restriction categories developed in [17]) and where each homset is a commutative monoid such that $x(f + g) = xf + xg$ and $0f \leq 0$, and finally projections fully preserve addition. The intuition comes from considering partial, smooth functions on open subsets of \mathbb{R}^n : not all smooth functions preserve addition, but smooth functions are addable under pointwise addition.

Definition 2.6. ([12, Defn. 3.18]) A **differential restriction category** is a Cartesian left additive restriction category, which has, for each map $f : A \rightarrow B$, a map

$$D[f] : A \times A \rightarrow B$$

satisfying various axioms², including **[DR.8]**: $\overline{D[f]} = \overline{f} \times 1$.

²There are nine equational axioms mirroring the axioms for reverse differential restriction categories given in the sequel.

3 Reverse differential restriction categories

We are now ready to define the new structure which we will use to model Abadi and Plotkin's language.

Definition 3.1. *A reverse differential restriction category or RDRC is a Cartesian left additive restriction category which has an operation on maps:*

$$\frac{A \xrightarrow{f} B}{A \times B \xrightarrow{R[f]} A}$$

such that

[RD.1] $R[f + g] = R[f] + R[g]$ and $R[0] = 0$;

[RD.2] for all a, b, c : $\langle a, b + c \rangle R[f] = \langle a, b \rangle R[f] + \langle a, c \rangle R[f]$ and $\langle a, 0 \rangle R[f] = \overline{a}f0$;

[RD.3] $R[\pi_j] = \pi_1 \iota_j$;

[RD.4] $R[\langle f, g \rangle] = (1 \times \pi_0)R[f] + (1 \times \pi_1)R[g]$;

[RD.5] $R[f g] = \langle \pi_0, \langle \pi_0 f, \pi_1 \rangle R[g] \rangle R[f]$;

[RD.6] $\langle 1 \times \pi_0, 0 \times \pi_1 \rangle (t_0 \times 1)R[R[R[f]]]\pi_1 = (1 \times \pi_1)R[f]$;

[RD.7] $(t_0 \times 1)R[R[(t_0 \times 1)R[R[f]]]\pi_1]\pi_1 = \text{ex}(t_0 \times 1)R[R[(t_0 \times 1)R[R[f]]]\pi_1]\pi_1$;

[RD.8] $\overline{R[f]} = \overline{f} \times 1$;

[RD.9] $R[\overline{f}] = (\overline{f} \times 1)\pi_1$.

As noted above, [RD.5] represents the chain rule, while [RD.8] says that the partiality of $R[f]$ is entirely determined by the partiality of f itself. [RD.9] says how to differentiate restriction idempotents. The other axioms are similar to those for an RDC; for an explanation of what they represent, see the discussion after Definition 13 in [16]. Also note that term logics have also been given to simplify reasoning in cartesian differential categories [8] and differential restriction categories [29]; a term logic for reverse differential restriction categories exists but will not be discussed further here.

Any Fermat theory [24] and more generally any Lawvere theory which is also a cartesian differential category can be given the structure of a reverse differential category; in these cases both the forward and reverse derivatives can be pushed down to sums and tuples of derivatives on maps $R \rightarrow R$, and here the forward and reverse derivative necessarily coincide. A restriction version of this example is given by considering a topological ring R that satisfies the axiom of determinacy (see [6]); the category with objects: powers of R , and C^∞ -maps that are smooth on restriction to an open set form a reverse differential restriction category. This meta-example includes the category Smooth_P of functions that are smooth on an open subset of \mathbb{R}^n . For an example whose objects are not of the form \mathbb{R}^n : the coKleisli category of the multiset comonad on the category of relations Rel is a cartesian differential category and the category of linear maps is Rel (see [9, 8] for details). As Rel is a compact closed self-dual category, and the derivative at a point is linear (hence a map in Rel), one can obtain a reverse derivative on the coKleisli category of the finite multiset comonad on Rel .

Just as with an RDC, we can derive a forward differential restriction structure from a reverse.

Theorem 3.2. *Every reverse differential restriction category \mathbb{X} is a differential restriction category with the derivative defined as previously (see also [16, Theorem 16]).*

Moreover, just as in [16, Theorem 42], one can prove that a DRC with a “contextual linear dagger” is equivalent to an RDRC; however, for space constraints we will not go into full details here. One must first describe fibrations for restriction categories: these were studied by Nester in [37]. One can give a version of the simple fibration for a restriction category as well as the dual of the simple fibration (this is remarkable – as the dual of a restriction category is not generally a restriction category). Importantly, maps in the simple fibration have their partiality concentrated in the context i.e. $\bar{f} = e \times 1$ where $e = \bar{e}$. A contextual dagger is an involution of fibrations $\text{Lin}(\mathbb{X})[\mathbb{X}] \rightarrow \text{Lin}(\mathbb{X})[\mathbb{X}]^*$ where $\text{Lin}(\mathbb{X})[\mathbb{X}]$ denotes a subfibration of the simple fibration consisting of linear maps in context, using the notion of fibration for restriction categories. From a reverse differential restriction category one obtains such an involution of fibrations from $(u, f) \mapsto (u, (\iota_0 \times 1)R[f]\pi_1)$, and the second component is sometimes written $f^{\dagger[I]}$ where I is the context object. There are a few subtleties that we will also not go further into here.

The reverse derivative automatically preserves the induced partial order (from the restriction structure) and joins, if they exist:

Proposition 3.3. *If \mathbb{X} is a reverse differential restriction category, then for any $f, g : A \rightarrow B$, $f \leq g$ implies $R[f] \leq R[g]$, and if \mathbb{X} has joins, then for any pairwise compatible family $\{f_i\}$, $R[\vee_i f_i] = \vee_i R[f_i]$.*

As we shall see, we will not strictly need [RD.6] and [RD.7] to model Abadi and Plotkin’s language; thus, we make the following definition:

Definition 3.4. *A **basic reverse differential restriction category** (or **basic RDRC**) is a structure satisfying all the requirements for an RDRC except [RD.6] and [RD.7].*

However, as we discuss in the final section, using axioms [RD.6] and [RD.7] allows one to consider improvements to the operational semantics of the language.

4 Interpretation of a simple differential language

We will make use of the language defined by Abadi and Plotkin [2]. We will make one modification up front. We will first consider the language without recursive function definitions and instead with while-loops (called SDPL); after showing the semantics works out, we will then add recursive definitions back in (called SDPL⁺). We remark that while the presentation of SDPL given in Plotkin and followed here is parametrized over a single generating type; however, we can add arbitrary generating types as long as those types have operations that provide the structure of a commutative monoid.

In [2], Abadi and Plotkin remarked that there are two approaches to differentiating over control structures: there are source transformations used in systems such as TensorFlow [1] and Theano [5] and there is the execution trace method used in systems such as Autograd [36] and PyTorch [38]. The source transformation method for dealing with derivatives of control structures defines a way to distribute the derivative into control structures; for example

$$\frac{\partial \text{if } B \text{ then } M \text{ else } N}{\partial x}$$

would be replaced by

$$\text{if } B \text{ then } \frac{\partial m}{\partial x} \text{ else } \frac{\partial n}{\partial x}$$

The execution trace allows defining a symbolic derivative on simpler terms with no control structures or derivatives, and then evaluating a term enough so that there are no control structures or derivatives present, allowing a symbolic trace through the derivative. This must be done at runtime – for example,

$\overline{\Gamma, x : A \vdash x : A}$	$\frac{r \in \mathbb{R}}{\Gamma \vdash r : \text{real}}$	$\frac{\Gamma \vdash m : \text{real} \quad \Gamma \vdash n : \text{real}}{\Gamma \vdash m + n : \text{real}}$
$\frac{\Gamma \vdash m : T \quad \text{op} : T \rightarrow U \in \Sigma}{\Gamma \vdash \text{op}(m) : U}$	$\frac{\Gamma \vdash m : T \quad \Gamma, x : T \vdash n : U}{\Gamma \vdash \text{let } x : T = m \text{ in } n : U}$	
$\overline{\Gamma \vdash * : 1}$	$\frac{\Gamma \vdash m : U \quad \Gamma \vdash n : T}{\Gamma \vdash (m, n)_{U, T} : U \times T}$	$\frac{\Gamma \vdash m : U \times T}{\Gamma \vdash \text{fst}_{U, T}(m) : U}$
$\frac{\Gamma \vdash m : U \times T}{\Gamma \vdash \text{snd}_{U, T}(m) : T}$	$\frac{\Gamma \vdash b \quad \Gamma \vdash m : T \quad \Gamma \vdash n : T}{\Gamma \vdash \text{if } b \text{ then } m \text{ else } n : T}$	
$\frac{p : U \vdash b \quad p : U \vdash f : U}{p : U \vdash \text{while } b \text{ do } f : U}$		
$\frac{\Gamma, x : U \vdash m : T \quad \Gamma \vdash a : U \quad \Gamma \vdash v : T}{\Gamma \vdash v.\text{rd}(x : U.m)(a) : U}$		
$\overline{\Gamma \vdash \text{true}}$	$\overline{\Gamma \vdash \text{false}}$	$\frac{\Gamma \vdash m : U \quad \text{pred} : U \in \text{Pred}}{\Gamma \vdash \text{pred}(m)}$

Table 1: Typing rules for SDPL

we need to know when differentiating over an if-then-else statement which branch was taken, and once this control structure is eliminated the derivative can be computed on the simpler resultant term. This has the advantage of making it simpler to adapt to derivatives over more subtle structures such as recursive function definitions. Since it's done at runtime, it can be performed by a source-transformation by a JIT compiler ensuring efficiency.

4.1 The core language SDPL

The types of SDPL are given by the following grammar:

$$\text{Ty} := \text{real} \mid 1 \mid \text{Ty} \times \text{Ty}$$

Powers are assumed to be left-associated so $\text{real}^{n+1} := \text{real}^n \times \text{real}$. To form the raw terms of SDPL we assume a countable supply of variables, a set of typed operation symbols Σ , and a set of typed predicate symbols Pred . The raw terms are then defined by the following grammar:

$$\begin{aligned} m := & x \mid r \ (r \in \mathbb{R}) \mid m + m \mid \text{op}(m) \ (\text{op} \in \Sigma) \mid \text{let } x : \text{Ty} = m \text{ in } m \\ & \mid * \mid (m, n)_{\text{Ty}, \text{Ty}} \mid \text{fst}_{\text{Ty}, \text{Ty}}(m) \mid \text{snd}_{\text{Ty}, \text{Ty}}(m) \mid \text{if } b \text{ then } m \text{ else } n \\ & \mid \text{while } b \text{ do } m \mid m.\text{rd}(x : \text{Ty}.m)(m) \end{aligned}$$

$$b := \text{pred}(m) \ (\text{pred} \in \text{Pred}) \mid \text{true} \mid \text{false}$$

Note that the typing rules will disallow inputs or outputs to come from boolean terms. This is to ensure that all typed terms are differentiable with respect to every argument. The typing rules for SDPL are given in Table 1. In the typing rules, Γ is assumed to be a list of typed variables $\Gamma = [x_i : A_i]_{i=1}^n$ where $A_i \in \text{Ty}$. Free variables are defined in the usual way; note that `let` expressions bind the variable x and

when forming the reverse differential term $v.\text{rd}(x : U.m)(a)$ the variable x is also bound. The reverse differential expression may be read as “the reverse differential of m with respect to x evaluated at the point a in the direction v .”

4.2 Categorical interpretation of SDPL

Let \mathbb{X} be a basic reverse differential restriction category with countable joins of disjoint maps. An **interpretation structure** for SDPL into \mathbb{X} is given by a tuple of structures:

$$(A \in \mathbb{X}_0, (1 \xrightarrow{a_r} A)_{r \in \mathbb{R}}, \llbracket \cdot \rrbracket, \llbracket \cdot \rrbracket_T, \llbracket \cdot \rrbracket_F)$$

and we extend such a structure to an interpretation of all the terms of SDPL as explained below. We must first interpret types, and to begin we need an object A from \mathbb{X} to carry our signatures. We also require that A has a point $1 \xrightarrow{a_r} A$ for each element $r \in \mathbb{R}$ (since we must interpret \mathbb{R} constants which are part of SDPL)³. With such an A we define an interpretation of types:

$$\llbracket 1 \rrbracket := 1 \quad \llbracket \text{real} \rrbracket := A \quad \llbracket T \times U \rrbracket := \llbracket T \rrbracket \times \llbracket U \rrbracket$$

We extend the interpretation to contexts:

$$\llbracket \cdot \rrbracket := 1 \quad \llbracket x : U \rrbracket := \llbracket U \rrbracket \quad \llbracket \Gamma, x : U \rrbracket := \llbracket \Gamma \rrbracket \times \llbracket U \rrbracket$$

We also require an interpretation of each operation symbol $\text{op} : T \rightarrow U \in \Sigma$ of the correct type: $\llbracket \text{op} \rrbracket : \llbracket T \rrbracket \rightarrow \llbracket U \rrbracket$. We additionally require two interpretations of each predicate symbol $\text{pred} : U \in \text{Pred}$: $\llbracket \text{pred} \rrbracket_T : \llbracket U \rrbracket \rightarrow 1$ and $\llbracket \text{pred} \rrbracket_F : \llbracket U \rrbracket \rightarrow 1$ such that $\llbracket \text{pred} \rrbracket_T \llbracket \text{pred} \rrbracket_F = \emptyset$. To summarize:

$$\Sigma(U, T) \xrightarrow{\llbracket \cdot \rrbracket} \mathbb{X}(\llbracket U \rrbracket, \llbracket T \rrbracket) \quad \text{Pred}(U) \xrightarrow{\llbracket \cdot \rrbracket_T, \llbracket \cdot \rrbracket_F} \mathbb{X}(\llbracket U \rrbracket, 1)$$

The intent for giving two interpretations of predicate symbols is that we must give an interpretation of the “true” part of the predicate and the “false” part. In [2] an interpretation of predicate symbols is given as maps $\llbracket U \rrbracket \rightarrow \{\text{true}, \text{false}\}$ with the property that the preimages of both true and false are open. This necessarily makes the interpretation of a predicate partial or trivial, and moreover it is equivalent to giving an interpretation of predicate symbols into disjoint open sets of $\llbracket U \rrbracket$, which is again equivalent to giving an interpretation into disjoint predicates on $\llbracket U \rrbracket$. A way around this non-standard interpretation of predicates is given by taking the manifold completion [33, 15] of the model, noting that $1 + 1$ is a manifold, and then requiring that we map into $1 + 1$ by an atlas morphism, which will necessarily yield two disjoint restriction idempotents on the domain. Another approach is to use the Heyting negation of the associated restriction idempotent, noting that this will always be disjoint from the starting map. These two approaches have interesting relationships with the approach we take, but their full development will not be pursued in this work.

We then extend the interpretation to all terms inductively. Most of these interpretations are standard; the more novel parts are the interpretations of if, while, and reverse derivatives.

Proj:

- $\llbracket x : U \vdash x : U \rrbracket := 1_{\llbracket U \rrbracket}$;

³It is not strictly necessary that SDPL contains a constant for every $r \in \mathbb{R}$ – as long as we include 0 we could only require constants that we actually use, such as the computable reals.

- $\llbracket \Gamma, x : U \vdash x : U \rrbracket := \llbracket \mathcal{Y} \rrbracket \times \llbracket U \rrbracket \xrightarrow{\pi_1} \llbracket U \rrbracket$;
- $\llbracket \Gamma, y : U \vdash x : T \rrbracket := \llbracket \mathcal{Y} \rrbracket \times \llbracket U \rrbracket \xrightarrow{\pi_0} \llbracket \mathcal{Y} \rrbracket \xrightarrow{\llbracket \Gamma \vdash x : T \rrbracket} \llbracket T \rrbracket$.

Real operations:

- We define $\llbracket \Gamma \vdash 0 : \text{real} \rrbracket := \llbracket \Gamma \rrbracket \xrightarrow{0} A$ and for the other elements $\llbracket \Gamma \vdash r : \text{real} \rrbracket := \llbracket \Gamma \rrbracket \xrightarrow{!} 1 \xrightarrow{a_r = \llbracket r \rrbracket} A = \llbracket \text{real} \rrbracket$
- $\llbracket \Gamma \vdash m + n : \text{real} \rrbracket := \llbracket \Gamma \rrbracket \xrightarrow{\llbracket \Gamma \vdash m : \text{real} \rrbracket + \llbracket \Gamma \vdash n : \text{real} \rrbracket} \llbracket \text{real} \rrbracket$

Operation terms: Given $\text{op} : T \rightarrow U \in \Sigma$

$$\llbracket \Gamma \vdash \text{op}(m) : U \rrbracket := \llbracket \Gamma \rrbracket \xrightarrow{\llbracket \Gamma \vdash m : T \rrbracket} \llbracket T \rrbracket \xrightarrow{\llbracket \text{op} \rrbracket} \llbracket U \rrbracket$$

Let:

$$\llbracket \Gamma \vdash \text{let } x : T = m \text{ in } n : U \rrbracket := \llbracket \Gamma \rrbracket \xrightarrow{\langle 1, \llbracket \Gamma \vdash m : T \rrbracket \rangle} \llbracket \Gamma \rrbracket \times \llbracket T \rrbracket \xrightarrow{\llbracket \Gamma, x : T \vdash n : U \rrbracket} \llbracket U \rrbracket$$

Product terms:

- $\llbracket \Gamma \vdash * : 1 \rrbracket := \llbracket \Gamma \rrbracket \xrightarrow{!} 1$
- $\llbracket \Gamma \vdash (m, n)_{A, B} : A \times B \rrbracket := \llbracket \Gamma \rrbracket \xrightarrow{\langle \llbracket \Gamma \vdash m : A \rrbracket, \llbracket \Gamma \vdash n : B \rrbracket \rangle} \llbracket A \rrbracket \times \llbracket B \rrbracket$
- $\llbracket \Gamma \vdash \text{fst}(m)_{A, B} : A \rrbracket := \llbracket \Gamma \rrbracket \xrightarrow{\llbracket \Gamma \vdash m : A \times B \rrbracket} \llbracket A \rrbracket \times \llbracket B \rrbracket \xrightarrow{\pi_0} \llbracket A \rrbracket$
- $\llbracket \Gamma \vdash \text{snd}(m)_{A, B} : B \rrbracket := \llbracket \Gamma \rrbracket \xrightarrow{\llbracket \Gamma \vdash m : A \times B \rrbracket} \llbracket A \rrbracket \times \llbracket B \rrbracket \xrightarrow{\pi_1} \llbracket B \rrbracket$

Control structures:

$$\begin{aligned} \llbracket \Gamma \vdash \text{if } b \text{ then } m \text{ else } n : U \rrbracket &:= \overline{\llbracket \Gamma \vdash b \rrbracket}_T \llbracket \Gamma \vdash m : U \rrbracket \vee \overline{\llbracket \Gamma \vdash b \rrbracket}_F \llbracket \Gamma \vdash n : U \rrbracket \\ \llbracket p : A \vdash \text{while } b \text{ dom} : A \rrbracket &:= \bigvee_{i=0}^{\infty} \left(\left(\overline{\llbracket p : A \vdash b \rrbracket}_T \llbracket p : A \vdash m : A \rrbracket \right)^i \overline{\llbracket p : A \vdash b \rrbracket}_F \right) \end{aligned}$$

Reverse derivatives:

$$\begin{aligned} \llbracket \Gamma \vdash v.\text{rd}(x : T.m)(a) : T \rrbracket \\ := \llbracket \Gamma \rrbracket \xrightarrow{\langle \langle 1, \llbracket \Gamma \vdash a : T \rrbracket \rangle, \llbracket \Gamma \vdash v : U \rrbracket \rangle} (\llbracket \Gamma \rrbracket \times \llbracket T \rrbracket) \times \llbracket U \rrbracket \xrightarrow{R[\llbracket \Gamma, x : T \vdash m : U \rrbracket] \pi_1} \llbracket T \rrbracket \end{aligned}$$

Boolean terms: $\llbracket \Gamma \vdash \text{true} \rrbracket_T := \llbracket \Gamma \rrbracket \xrightarrow{!} 1$ and $\llbracket \Gamma \vdash \text{true} \rrbracket_F := \emptyset$. Likewise, $\llbracket \Gamma \vdash \text{false} \rrbracket_T := \emptyset$ and $\llbracket \Gamma \vdash \text{false} \rrbracket_F := !$. Finally for any $\text{pred} \in \text{Pred}(A)$:

$$\llbracket \Gamma \vdash \text{pred}(m) \rrbracket_H := \llbracket \Gamma \rrbracket \xrightarrow{\llbracket m \rrbracket} \llbracket A \rrbracket \xrightarrow{\llbracket \text{pred} \rrbracket_H} 1$$

where H ranges over $\{T, F\}$.

For a brief explanation of the interpretation of while-loops, for $f : A \rightarrow A$ we set $f^0 = \text{id}$ and $f^{n+1} = f \circ f^n$. Then our interpretation says either the guard was false, or it was true and we executed m and then it was false, or it was true and we executed m and it was still true and we executed m again and then it was false, and so on. This yields

$$\llbracket \text{while } b \text{ dom} \rrbracket := \llbracket b \rrbracket_F \vee \llbracket b \rrbracket_T \llbracket m \rrbracket \llbracket b \rrbracket_F \vee \llbracket b \rrbracket_T \llbracket m \rrbracket \llbracket b \rrbracket_T \llbracket m \rrbracket_T \llbracket b \rrbracket_F \vee \dots$$

4.3 Categorical semantics of source code transformations

In this section we show that the interpretation above always soundly models source code transformations for differentiating if-then-else statements and while-loops.

Proposition 4.1. *In an interpretation structure on a basic RDRC, for any terms $\Gamma, x : U \vdash m : T$, $\Gamma, x : U \vdash n : T$, $\Gamma \vdash a : U$, and $\Gamma \vdash v : T$ and for any predicate $\Gamma, x : U \vdash B$ we have*

$$\begin{aligned} & \llbracket \Gamma \vdash v.\text{rd}(x : U.\text{if } b \text{ then } m \text{ else } n)(a) \rrbracket \\ &= \llbracket \Gamma \vdash \text{if}(let\ x = a\ \text{in } b)\ \text{then } v.\text{rd}(x : U.m)(a)\ \text{else } v.\text{rd}(x : U.n)(a) \rrbracket \end{aligned}$$

Corollary 4.2 (If-then-else transformation). *In an interpretation structure on a basic RDRC, we always have*

$$\begin{aligned} & \llbracket \Gamma, x : U \vdash v.\text{rd}(x : U.\text{if } b \text{ then } m \text{ else } n)(x) \rrbracket \\ &= \llbracket \Gamma, x : Y \vdash \text{if } b \text{ then } v.\text{rd}(x : U.m)(x)\ \text{else } v.\text{rd}(x : U.n)(x) \rrbracket \end{aligned}$$

Turning to iteration, if a while-loop terminates, then `while` b `do` f is f^n for some n . The forward derivative admits a tail recursive description:

$$D[f^{n+1}] = \langle \pi_0 f, D[f] \rangle D[f^n]$$

SDPL has two admissible operations: dagger and forward differentiation.

$$m^{\dagger[\Gamma]} := y.\text{rd}(x.m)(0) \quad \text{fd}(x.m)(a).v := \text{let } z = v \text{ in } (y.\text{rd}(x.m)(a))^{\dagger[\Gamma]}$$

where the y in $m^{\dagger[\Gamma]}$ is fresh. The recursive description of $D[f^n]$ is useful in proving the following:

Proposition 4.3 (Forward-differentiation for while-loops). *In an interpretation structure on a basic RDRC,*

1. For any $\Gamma, x : A \vdash m : B$, $\llbracket \text{fd}(x.m)(a).v \rrbracket = \langle \langle 1, \llbracket a \rrbracket \rrbracket, \llbracket v \rrbracket \rangle (1 \times \iota_1) D[\llbracket m \rrbracket]$
2. For any $x : A \vdash f : A$ we have $\llbracket \text{fd}(x.\text{while } b \text{ do } f)(a).v \rrbracket = \llbracket \text{let } x = a, y = v \text{ in } \text{snd}(\text{while } \pi_0 b \text{ do } (\pi_0 f, \text{fd}(x.f)(x).y)) \rrbracket$

On the other hand, the reverse derivative satisfies:

$$R[f^n](a, b) = R[f](a, R[f](f(a), R[f](f(f(a)), \dots, b)))$$

Which looks at first glance to be head recursive, and not like something that can be implemented by an iteration. However, with **[RD.6]**, we can do the following:

$$R[f^{n+1}] = D[f^{n+1}]^{\dagger[A]} = (T(f)^n D[f])^{\dagger[A]} \quad \text{where } T(f) = \langle \pi_0 f, D[f] \rangle$$

This is the basis of the following source transformation for while-loops.

Corollary 4.4 (Reverse-differentiation of while-loops). *In an interpretation structure on an RDRC, let $z : A \vdash f : A$ and $z : A \vdash b$; then we have*

$$\begin{aligned} & \llbracket v : A \vdash v.\text{rd}(x.\text{while } b \text{ do } f)(a) \rrbracket \\ &= \llbracket v : A \vdash (\text{let } x = a, y = v \text{ in } \text{snd}(\text{while } \pi_0 b \text{ do } (\pi_0 f, \text{fd}(x.f)(x).y)))^{\dagger[L]} \rrbracket \end{aligned}$$

where $\dagger[\cdot]$ denotes the dagger defined above with respect to the empty context.

4.4 Smooth recursive definitions

Until now we discussed the semantics of a fragment of the language described by [2]. We formally extended their language with while-loops to isolate their behaviour, but missed out on recursive function definitions. Given general recursion, one can implement loops using tail recursion. We now move to discuss their full language with recursive function definitions. This language will be called SDPL^+ . To give such an extension, we introduce two new raw terms

$$m := m \text{ as before} \mid f(m) \mid \text{letrec } f(x) := m \text{ in } n$$

In the above, when we form $f(a)$, the symbol f is taken to be a free function variable, and the term $\text{letrec } f(x) := m \text{ in } n$ binds the variable x in m and the function variable f in m and n . However, these function variables are of a different sort than ordinary variables because they have arity. That is $f(a)$ only makes sense if $a : B$ and f has arity $B \rightarrow C$, which we write as $f : B \rightarrow C$. Thus, our typing/term formation rules will have two sorts of contexts, one to record function names and the other for ordinary variables. Our terms in context then have the form $\Phi \mid \Gamma \vdash m : B$, and to update the rules from before, just add Φ to all the contexts. The two new rules are

$$\frac{\Phi, f : A \rightarrow B \mid \Gamma \vdash m : A}{\Phi, f : A \rightarrow B \mid \Gamma \vdash f(m) : B} \quad \frac{\Phi, f : A \rightarrow B \mid x : A \vdash m \quad \Phi, f : A \rightarrow B \mid \Gamma \vdash n : C}{\Phi \mid \Gamma \vdash \text{letrec } f(x) := m \text{ in } n : C}$$

We will now give the interpretation of recursive definitions and calls in a basic reverse differential join restriction category. But first, we will review a basic bit of intuition of recursive function theory in case the reader is unfamiliar.

We often write computable functions $A \xrightarrow{f} B$ as $f(n) := m$, but it is usually helpful to think of f as simply a name for the unnamed function $\lambda n.m$ and then write $f = \lambda n.m$. The idea is that as a computation f has an internal representation that uses the variable n somewhere. If f is recursive then that means that the symbol f also appears in m , and thus it is a function that depends on itself. To break this cycle we then abstract out the symbol f too. We write $\underline{f} := \lambda f.\lambda n.m$. This creates a function

$$\text{Fun}(A, B) \xrightarrow{\underline{f}} \text{Fun}(A, B)$$

\underline{f} takes an arbitrary computable function $A \xrightarrow{h} B$ and creates a function that uses h instead of f anywhere \underline{f} was used in the body m .

To give a quick example, consider the computable function $\text{fac}(n) := \text{if } n < 1 \text{ then } 1 \text{ else } n * \text{fac}(n - 1)$. Then

$$\underline{\text{fac}}(h)(n) := \text{if } n < 1 \text{ then } 1 \text{ else } n * h(n - 1)$$

The point is that this new function is not recursive. However, it is instructive to see what happens when we apply it to the function it represents. As an exercise, we leave it to the reader to prove that

$$\underline{\text{fac}}(\text{fac}) = \text{fac}$$

In other words, the recursive function fac is a *fixed point* of the functional $\underline{\text{fac}}$. It is also the best fixed point of fac in the sense that it is the least defined function that is a fixed point of $\underline{\text{fac}}$. This works in general, given any recursive function r it may be obtained as the least fixed point of \underline{r} .

To model least-fixed-point phenomena we will use the notion of a pointed directed complete partial order or (DCPPO) for short. The first use of DCPPOs to model recursive phenomena is due to Scott

[41] in giving models of the untyped λ -calculus. DCPPOs are used in the semantics of the functional programming language PCF in [40]. Abstract DCPPO-enriched categories of partial maps were used in modelling the semantics of the functional programming language FPC in [27]. The DCPPO structure on homsets of Smooth_P was used in [2] to provide a semantics of SDPL. The approach taken here generalizes [2] to an arbitrary basic reverse differential join restriction category, highlights the structural aspects of the interpretation, and uses the axioms of such a category to derive some simplifications to the operational behaviour. A connection of ω -CPPOs and restriction categories was introduced using the delay monad in [44].

Definition 4.5. Let (D, \leq) be a partial order. A subset $A \subseteq D$ is **directed** if A is nonempty and any two elements $f, g \in A$ have an upper bound in A ; i.e., there is an $h \in A$ with $f, g \leq h$. A partial order (D, \leq) is a **directed complete partial order** if every directed subset A has a supremum written $\bigvee_{a \in A} a \in D$. A directed complete partial order is **pointed** (DCPPO) if there is a supremum for the empty set, that is a minimal element $\emptyset \leq d$ for all $d \in D$.

By a **morphism of DCPPOs** $(P, \leq) \xrightarrow{g} (Q, \leq)$ we mean a function g on the underlying sets that is monotone and preserves suprema. We observe minimally that the category of DCPPOs is Cartesian closed.

Lemma 4.6. [42] Let (D, \leq) be a DCPPO. Then

1. Every morphism $D \xrightarrow{g} D$ has a least fixed point; i.e. a $u \in D$ such that $g(u) = u$.
2. For any other DCPPO (P, \leq) every morphism $P \times D \xrightarrow{g} P$ has a parametrized fixed point; i.e., a $P \xrightarrow{u} D$ such that

$$\begin{array}{ccc} P & \xrightarrow{u} & D \\ & \searrow \langle 1, u \rangle & \uparrow g \\ & & P \times D \end{array}$$

In other words, for each $x \in P$, $u(x)$ is a fixed point of $g(x, _)$. This parametrized fixed point is often denoted $\mu_y.g(_, y)$ and as the fixed point of $g(x, _)$ by $\mu_y.g(x, y)$.

Join restriction categories are DCPPO enriched:

Proposition 4.7. Let \mathbb{X} be a restriction category. Then with respect to the order enrichment of restriction categories:

1. \mathbb{X} is a join restriction category then the enrichment lies in DCPPOs.
2. If \mathbb{X} has joins and restriction products then those products are DCPPO enriched products; i.e., $\mathbb{X}(A, B \times C) \simeq \mathbb{X}(A, B) \times \mathbb{X}(A, C)$ qua an isomorphism of DCPPOs. Moreover, the “contraction operator”

$$\mathbb{X}(A, B) \xrightarrow{\Delta_A} \mathbb{X}(A, A \times B)$$

that sends $A \xrightarrow{f} B$ to $A \xrightarrow{\langle 1, f \rangle} A \times B$ is a morphism of DCPPOs.

3. If \mathbb{X} has joins and is a Cartesian left additive restriction category, then the addition on homsets

$$\mathbb{X}(A, B) \times \mathbb{X}(A, B) \xrightarrow{+} \mathbb{X}(A, B)$$

is a morphism of DCPPOs.

4. If \mathbb{X} is a reverse differential join restriction category then the operation of reverse differentiation

$$\mathbb{X}(A, B) \xrightarrow{R[_]} \mathbb{X}(A \times B, A)$$

is a morphism of DCPPOs.

Parts 2–4 of Proposition 4.7 implies that certain operations we will need to form from monotone and join preserving maps will again be monotone and join preserving.

To give the categorical semantics of SDPL^+ , we must extend the interpretation developed in section 4.2. To begin we first give the interpretation of function contexts. The idea being that a free function symbol could be any map of the correct type, the interpretation of function contexts is given as a product of homsets.

$$\llbracket \emptyset \rrbracket := 1 \quad \llbracket \Phi, f : A \rightarrow B \rrbracket := \llbracket \Phi \rrbracket \times \mathbb{X}(\llbracket A \rrbracket, \llbracket B \rrbracket)$$

The interpretation of a term in context $\Gamma \vdash m : B$ constructed a map $\llbracket \Gamma \rrbracket \xrightarrow{\llbracket m \rrbracket} \llbracket B \rrbracket$. With function contexts, the maps we build now depend on the morphism from $\mathbb{X}(\llbracket A \rrbracket, \llbracket B \rrbracket)$ to fill in the call to a function. That is, the interpretation is now a function

$$\llbracket \Phi \rrbracket \xrightarrow{\llbracket \Phi | \Gamma \vdash m : B \rrbracket} \mathbb{X}(\llbracket \Gamma \rrbracket, \llbracket B \rrbracket)$$

Now, we are building a function and to do so it suffices to build a map in $\mathbb{X}(\llbracket \Gamma \rrbracket, \llbracket B \rrbracket)$ for each element $\phi \in \llbracket \Phi \rrbracket$. We write $\llbracket m \rrbracket_\phi$ for the value of $\llbracket m \rrbracket$ at ϕ . The construction is by induction and for the terms from SDPL , the construction is exactly the same with the addition of a ϕ subscript decorating the terms appropriately. For example, $\llbracket \Phi | \Gamma \vdash \text{let } x = m \text{ in } n \rrbracket_\phi := \langle 1, \llbracket m \rrbracket_\phi \rangle \llbracket n \rrbracket_\phi$. However, we can build the interpretation entirely using external structure by induction as well.

For example, interpretation of $\text{let } x = m \text{ in } n$ may be given using the “contraction operator”, and this construction is element free.

$$\begin{array}{ccc} \llbracket \Phi \rrbracket \xrightarrow{\langle \llbracket m \rrbracket, \llbracket n \rrbracket \rangle} \mathbb{X}(\llbracket \Gamma \rrbracket, \llbracket A \rrbracket) \times \mathbb{X}(\llbracket \Gamma \rrbracket \times \llbracket A \rrbracket, \llbracket B \rrbracket) & \xrightarrow{\Delta_{\Gamma \times 1}} & \mathbb{X}(\llbracket \Gamma \rrbracket, \llbracket \Gamma \rrbracket \times \llbracket A \rrbracket) \times \mathbb{X}(\llbracket \Gamma \rrbracket \times \llbracket A \rrbracket, \llbracket B \rrbracket) \\ & \searrow \llbracket \text{let } x = m \text{ in } n \rrbracket & \downarrow \\ & & \mathbb{X}(\llbracket \Gamma \rrbracket, \llbracket B \rrbracket) \end{array}$$

We will leave it to the reader to construct the interpretation of $v.\text{rd}(x.m)(a)$ using a similar idea, as well as the reverse differential operator $\mathbb{X}(A, B) \xrightarrow{R[_]} \mathbb{X}(A \times B, A)$.

For SDPL^+ , we extend this to the two new terms. Given a function context $\Phi = (f_1, \dots, f_n)$ then for any $\phi \in \llbracket \Phi \rrbracket$ we have that $\phi = (\phi_1, \dots, \phi_n)$ are all maps in \mathbb{X} : if $f_i : A_i \rightarrow B_i$ then $\phi_i : \llbracket A_i \rrbracket \rightarrow \llbracket B_i \rrbracket$. We will write $\phi(f_i)$ to denote ϕ_i . We also make use of the “no-free-variable” assumption for recursive definitions; that is, in the type formation rule for recursive definitions $\text{letrec } f(x) := m \text{ in } n$, m must have at most a unique free variable, and it must be x .

Fun-Call: $\llbracket \Phi, f : A \rightarrow B | \Gamma \vdash f(m) : B \rrbracket_\phi := \llbracket \Gamma \rrbracket \xrightarrow{\llbracket m \rrbracket_\phi} \llbracket A \rrbracket \xrightarrow{\phi(f)} \llbracket B \rrbracket$

Rec-Def: First note that if we just translate a simple recursive function $\text{letrec } f(x) := m$, we see that x is a free variable and f is a free function variable in m . That is, we have $f : A \rightarrow B | x : A \vdash m : A$. Then note that the interpretation we are developing would interpret m as a function

$$\mathbb{X}(\llbracket A \rrbracket, \llbracket B \rrbracket) \xrightarrow{\llbracket m \rrbracket} \mathbb{X}(\llbracket A \rrbracket, \llbracket B \rrbracket)$$

This is exactly the sort of underlined function we looked at earlier: it takes each $h : \llbracket A \rrbracket \rightarrow \llbracket B \rrbracket$ in \mathbb{X} and uses it by the above translation of function calls, any where that f was used in m . Then by

Lemma 4.6, we may take its fixed point, μ . We then get a map $\llbracket A \rrbracket \xrightarrow{\mu} \llbracket B \rrbracket$ such that $\llbracket m \rrbracket \mu = \mu$ and is the least defined such map, giving us the interpretation of the recursive function, and we would write $\llbracket \text{letrec } f(x) := m \rrbracket = \mu$. More generally, in m the unique variable condition only applies to ordinary variables, but m could have multiple function variables. Then if we translate $\Phi, f : A \rightarrow B \mid x : A \vdash m : B$ we get a map

$$\llbracket \Phi \rrbracket \times \mathbb{X}(\llbracket A \rrbracket, \llbracket B \rrbracket) \xrightarrow{\llbracket m \rrbracket} \mathbb{X}(\llbracket A \rrbracket, \llbracket B \rrbracket)$$

We may then apply the second part of Lemma 4.6 and obtain a parametrized fixed point

$$\llbracket \Phi \rrbracket \xrightarrow{\mu_f \cdot \llbracket m \rrbracket_{(\cdot, f)}} \mathbb{X}(\llbracket A \rrbracket, \llbracket B \rrbracket)$$

Likewise if we translate $\Phi, f : A \rightarrow B \mid \Gamma \vdash n : C$, we get a map

$$\llbracket \Phi \rrbracket \times \mathbb{X}(\llbracket A \rrbracket, \llbracket B \rrbracket) \xrightarrow{\llbracket n \rrbracket} \mathbb{X}(\llbracket \Gamma \rrbracket, \llbracket C \rrbracket)$$

Then, finally, the interpretation of $\text{letrec } f(x) := m \text{ in } n$ is defined by the following diagram:

$$\begin{array}{ccc} \llbracket \Phi \rrbracket & \xrightarrow{\langle 1, \mu_f \cdot \llbracket m \rrbracket_{(\cdot, f)} \rangle} & \llbracket \Phi \rrbracket \times \mathbb{X}(\llbracket A \rrbracket, \llbracket B \rrbracket) \\ \llbracket \text{letrec } f(x) := m \text{ in } n \rrbracket & \searrow & \downarrow \llbracket n \rrbracket \\ & & \mathbb{X}(\llbracket \Gamma \rrbracket, \llbracket C \rrbracket) \end{array}$$

We may also define it componentwise as

$$\llbracket \text{letrec } f(x) := m \text{ in } n \rrbracket_{\phi} := \llbracket n \rrbracket_{(\phi, \mu_f \cdot \llbracket m \rrbracket_{(\phi, f)})}$$

Note that the above definition is only well-defined if we can prove that the interpretation $\llbracket \Phi \rrbracket \xrightarrow{\llbracket m \rrbracket} \mathbb{X}(\llbracket \Gamma \rrbracket, \llbracket B \rrbracket)$ always yields a monotone and join preserving function between the DCPPOs, so that in the last step the use of Lemma 4.6 is justified.

Proposition 4.8. *Let \mathbb{X} be a basic reverse differential join restriction category, with a specified interpretation structure for SDPL^+ . Then the interpretation of terms in context is always a monotone, join preserving function between the DCPPOs. In particular, the construction is well-defined.*

4.5 Operational semantics

The operational semantics used by [2] defined a sublanguage of the raw terms called trace terms. These are generated by the following grammar:

$$\text{tr} := x \mid r \ (r \in R) \mid \text{op}(\text{tr}) \mid \text{let } x = m \text{ in } n \mid * \mid (\text{tr}, \text{tr}) \mid \text{fst}(\text{tr}) \mid \text{snd}(\text{tr})$$

Abadi and Plotkin also defined a sublanguage of trace terms called values.

$$\begin{aligned} v &:= x \mid r \ (r \in R) \mid * \mid (v, v) \\ v_{\text{bool}} &:= \text{true} \mid \text{false} \end{aligned}$$

The operational semantics of a program then consists of two mutually inductively defined reductions: symbolic evaluation and ordinary evaluation – the former yields a trace term and the latter yields a value. Then the main idea is that to evaluate a term, when you hit a reverse differential, $v.\text{rd}(x.f)(a)$, you evaluate f symbolically, just enough to remove control structures and derivatives giving a trace term. And then this trace term is differentiated symbolically, yielding a trace term, and the evaluation continues.

Note that defining symbolic reverse differentiation does not require any evaluation functions. However, we do at this point require, as [2] did, that for each function symbol $\text{op} \in \Sigma(T, U)$ there is an associated a function symbol $\text{op}_R \in \Sigma(T \times U, T)$. The idea is that op_R is the reverse derivative of op . We will write $v.\text{op}_R(a)$ as notation for $\text{op}_R(a, v)$. Then define symbolic reverse differentiation $v.\mathcal{R}(x.f)(a)$ by induction over trace terms f and where v and a are values:

$$\begin{aligned} w.\mathcal{R}(x.y)(a) &= \begin{cases} w & x = y \\ 0 & x \neq y \end{cases} \\ w.\mathcal{R}(x.r)(a) &= 0 \quad r \in \mathbf{R} \\ w.\mathcal{R}(x.m+n)(a) &= w.\mathcal{R}(x.m)(a) + w.\mathcal{R}(x.n)(a) \\ w.\mathcal{R}(x.\text{op}(m))(a) &= \text{let } x = a, t = w.\text{op}_R(m) \text{ in } t.\text{rd}(m)(a) \quad t \text{ fresh} \\ w.\mathcal{R}(x.\text{let } y = d \text{ in } e)(a) &= \text{let } x = a, y = d \text{ in } w.\text{rd}(x.e)(a) \\ &\quad + (\text{let } t = w.\text{rd}(y.e)(y) \text{ in } t.\text{rd}(x.d)(a)) \quad t \text{ fresh} \\ w.\mathcal{R}(*) &= 0 \\ w.\mathcal{R}(x.(u, v))(a) &= \text{let } (y, z) = w \text{ in } y.\mathcal{R}(x.u)(a) + z.\mathcal{R}(x.v)(a) \\ w.\mathcal{R}(x.\text{fst}(m))(a) &= \text{let } x = a, t = m \text{ in } (w.0).\mathcal{R}(x.m)(a) \quad t \text{ fresh} \\ w.\mathcal{R}(x.\text{snd}(m))(a) &= \text{let } x = a, t = m \text{ in } (0, w).\mathcal{R}(x.m)(a) \quad t \text{ fresh} \end{aligned}$$

The let term is also the chain rule but for differentiating with respect to the two variable function $\Gamma, x, y \vdash n$, so that we get the usual rule $\partial n / \partial t = \partial n / \partial x \cdot \partial x / \partial t + \partial n / \partial y \cdot \partial y / \partial t$ appropriately reversed.

Also, for the projection rule, one might have expected just $(w, 0).\mathcal{R}(m)(a)$. Under interpretation we certainly get a term of the form $R[a\pi_0] = (\bar{a} \times \iota_0)R[a]$. However by **[RD.8]**, $(\bar{a} \times 1)R[a] = R[a]$. We will see below that if our evaluation satisfies a certain property, then the simpler translation is warranted.

Then as long as our interpretation always sends op_R to the reverse derivative of op , then symbolic and formal reverse differentiation agree under interpretation.

Proposition 4.9 (Symbolic differentiation correctness). *Suppose \mathbb{X} is a basic reverse differential join restriction category, and suppose that we have a fixed interpretation of SDPL into \mathbb{X} for which $\llbracket \text{op}_R \rrbracket = R[\llbracket \text{op} \rrbracket]$ then for all values a, v and for all traced terms m*

$$\llbracket v.\text{rd}(x.m)(a) \rrbracket = \llbracket v.\mathcal{R}(x.m)(a) \rrbracket$$

We have an analogous proposition for the interpretation of all SDPL^+ .

Proposition 4.10 (Symbolic differentiation correctness extended). *Suppose \mathbb{X} is a basic reverse differential join restriction category, and that we have a fixed interpretation of SDPL^+ into \mathbb{X} for which $\llbracket \text{op}_R \rrbracket = R[\llbracket \text{op} \rrbracket]$, then for all values a, v and for all traced terms m :*

$$\llbracket v.\text{rd}(x.m)(a) \rrbracket_\phi = \llbracket v.\mathcal{R}(x.m)(a) \rrbracket_\phi$$

We then define the operational semantics of SDPL exactly as done by Abadi and Plotkin [2]: an **operational structure** is given by $(\text{ev}, \text{bev}, _R)$ where

$$\text{ev}_{T,U} : \Sigma(T, U) \times v_T \rightarrow v_U \quad \text{bev}_T : \text{Pred}(T) \times v_T \rightarrow \text{val}_{\text{bool}}$$

are partial functions. We denote closed value terms v that have type Y as v_Y and the set of closed v_{bool} as val_{bool} (these sets are precisely those that require formation in an empty context $\vdash m : A$ and $\vdash b$). Further

$$_{-R} : \Sigma(T, U) \rightarrow \Sigma(T \times U, T) \quad \text{op} \mapsto \text{op}_R$$

With these three pieces one may define ordinary reduction \Rightarrow from terms to values and symbolic reduction \rightsquigarrow from terms to trace terms by induction; see [2] for details. For SDPL these reduction relations are formulated with respect to a value environment: this is a mapping of variable names to closed value terms. For SDPL^+ we also require a function environment: this is a mapping φ of function names to closures. A closure is a tuple (φ, f, x, m) where m has at most the free ordinary variable x , and additionally, all the free function variables in m except f are in the domain of φ . The idea is that closures are created when evaluating $\text{letrec } f(x) := m \text{ in } n$; if our current function environment is φ we extend it with (φ, f, x, m) and continue evaluating n – this way if n calls f then the definition of f can be looked up in the function environment, and any symbol that the body of f requires to operate will be there too.

5 Denotational Semantics

An interpretation structure $(A \in \mathbb{X}_0, (1 \xrightarrow{a_r} A)_{r \in R}, \llbracket _ \rrbracket, \llbracket _ \rrbracket_T, \llbracket _ \rrbracket_F)$ is a **differentially denotational interpretation structure** when

1. For all closed value terms of v we have that $1 \xrightarrow{\llbracket v \rrbracket} \llbracket A \rrbracket$ is a total point of $\llbracket A \rrbracket$;
2. For all $\text{op} \in \Sigma$ we have $R[\llbracket \text{op} \rrbracket] = \llbracket \text{op}_R \rrbracket$;
3. For all closed value terms $v \in v_A$ we have

$$\begin{array}{ccc} 1 \xrightarrow{\llbracket v \rrbracket} \llbracket A \rrbracket & & 1 \xrightarrow{\llbracket v \rrbracket} \llbracket A \rrbracket \\ \searrow \llbracket \text{ev}(\text{op}, v) \rrbracket & \downarrow \llbracket \text{op} \rrbracket & \searrow \llbracket \text{bev}(\text{pred}, v) \rrbracket_H \\ & \llbracket B \rrbracket & \downarrow \llbracket \text{op} \rrbracket_H \\ & & 1 \end{array}$$

where H is either T or F . In particular, both sides may be undefined, but they must be undefined simultaneously.

The idea behind showing that a denotational semantics captures a language's operational semantics is that if $m \Rightarrow v$ then $\llbracket m \rrbracket = \llbracket v \rrbracket$. However, the operational semantics for SDPL and SDPL^+ is defined with respect to value and function environments, and we have two operational relations. Interpreting a term m with free variables x_1, \dots, x_n in a value environment $\{x_i := v_i\}_{1 \leq i \leq n}$ is straightforward: since each v_i is a closed term, first interpret m as above: $\llbracket \Gamma \rrbracket \xrightarrow{\llbracket m \rrbracket} \llbracket B \rrbracket$, and then precompose with the point of $\llbracket \Gamma \rrbracket$ given by $1 \xrightarrow{\langle \llbracket v_i \rrbracket \rangle_{i \leq n}} \llbracket \Gamma \rrbracket$. Next we need the following lemma:

Lemma 5.1. *The interpretation of terms of SDPL^+ extends to allow the construction of an element of $\llbracket \Phi \rrbracket$ for each function environment φ whose domain is Φ .*

Note that for any trace term c it always fully evaluates. It requires no function context because it has no function symbols, and we have that for any value environment ρ , $c \Rightarrow v$ for some closed value term v .

The goal is then to prove the following theorem by mutual induction: for any term m , any value environment ρ , and function environment φ , we have that if $m \rightsquigarrow c$ then

$$\llbracket m \rrbracket = \llbracket c \rrbracket = \llbracket v \rrbracket.$$

6 Potential operational improvements

In this section we describe additional properties our categorical semantics has that may lead to a more refined operational semantics.

The compatibility between differentiation and restriction: **[RD.8,9]** state essentially that the definedness of the reverse derivative of a term is completely determined by the term itself. This is relevant to a more efficient semantics: the operational semantics used here has the property that when taking the reverse derivative over looping or recursive constructs, we first build a trace term, which turns out to be a (long) series of `let` expressions that describe the evolution of the state of the computation. We then symbolically differentiate these `let` expressions which always results in the creation of a sum of two expressions for each such `let` expression – and the number of `let` expressions created by recursion or looping is the number of times that the function recursed or the number of times the loop ran. Thus we quickly get wide trees of sums of symbolic terms that need to be evaluated. However, at each step of this process, one of these terms is of the form $v.\text{rd}(x.m)(a)$ where x does not occur freely in m , and hence can be proven to always evaluate to 0 if it evaluates to anything. Our semantics has the following property

Lemma 6.1. *For any term m in which x does not occur*

$$\llbracket v.\text{rd}(x.m)(a) \rrbracket = \overline{\langle 1, \langle \llbracket a \rrbracket, \llbracket v \rrbracket \rangle \rangle \llbracket m \rrbracket} 0$$

And moreover, the `let` expressions that get zeroed out have all their subterms occurring in the term that does not get zeroed out. We then have the following lemma

Lemma 6.2. *If we added the rule*

$$x \notin \text{fv}(e) \Rightarrow w.\mathcal{R}(x.\text{let } y = d \text{ in } e)(a) := \text{let } x = a, y = d, t = w.\text{rd}(y.e)(y) \text{ in } t.\text{rd}(x.d)(a)$$

Then Propositions 4.9 and 4.10 would still hold.

This gives an operational semantics where differentiating over looping constructs does not have a branching blowup, and hence experiences an exponential speedup.

Reverse differential restriction categories, as we have seen earlier, allow forming a forward derivative from the reverse derivative. They also allow forming a reverse derivative from that forward derivative. In a reverse differential restriction category, **[RD.6]** is equivalent to the requirement that the process of going from a reverse derivative to a forward derivative and then back to a reverse derivative gives exactly the starting reverse derivative.

Lemma 6.3. *For any map from $A \times B \xrightarrow{f} C$ define a map $A \times C \xrightarrow{f^\dagger[A] := (t_0 \times 1)R[f]\pi_1} B$. We always get a forward derivative as $D[f] := R[f]^\dagger[A]$. Then **[RD.6]** is equivalent to requiring that $D[f]^\dagger[A] = R[f]$.*

This kind of coherence for defining forward derivatives from their reverse could be useful in using the forward derivative and then converting back via daggering the result.

Lemma 6.4. *For any operational symbol op if the evaluation function used by the operational semantics satisfies*

$$\text{eval}(\text{snd}(op_{RRR}, (((a, 0), 0), (0, b)))) = \text{eval}(op_R, (a, b))$$

then this can be modelled in any reverse differential restriction category. Moreover, for every term m we have

$$\llbracket \text{rd}(x.\text{rd}(y.\text{rd}(z.m)(a).y)(b).x)(c).w \rrbracket = \overline{\llbracket b \rrbracket} \overline{\llbracket c \rrbracket} \llbracket \text{rd}(z.m)(a).w \rrbracket$$

Crucially for the above, we require **[RD.6]**.

An aspect of forward differentiation that is modelled in our semantics is that differentiating a differential with respect to its “direction” is just substitution. That is $\text{fd}(x.\text{fd}(y.m)(a).x)(b).v = \text{fd}(y.m)(a).v$ is modelled. This uses **[RD.6]**. More generally, we can modify the type system slightly to keep track of the arguments that a term is differentiated to by introducing another context, which we call a linearity context. Then the typing judgment for the reverse differential term would have two forms:

$$\frac{\Gamma, x : A \mid \Delta \vdash m : B}{\Gamma, a : A \mid \Delta, v : B \vdash \text{rdx}.m(a).v : A} \quad a, v \text{ fresh} \qquad \frac{\Gamma \mid \Delta, x : A \vdash m : B}{\Gamma, a : A \mid \Delta, v : B \vdash \text{rdx}.m(a).v : A} \quad a, v \text{ fresh}$$

And if we are forward differentiating with respect to a variable from the linearity context: i.e., if we form the forward derivative of a term with respect to a variable from the linearity context; i.e., if v was in the linearity context of a term m and we form $\text{fd}(v.m)(a).w$, then operational reduction

$$\text{fd}(v.m)(a).w \rightsquigarrow \text{let } v = w \text{ in } m$$

is modelled in our semantics. This means that we can completely avoid doing differentiation in some cases, at the cost of having to carry around more type information. There is a similar version of this rule for reverse derivatives and it has to do with “colet” expressions. In SDPL we can use the reverse derivative to create a term that substitutes linearly into the output variable of a term. We could use these “colet” expressions and allow for speedups of reverse derivatives as well. It might also be interesting to characterize these constructions in their own right. This approach also allows us to force **[RD.6]** into the operational semantics.

The axiom **[RD.7]**, dealing with the symmetry of mixed partial derivatives, may also have a role to play in simplifying the operational semantics. Some machine-learning algorithms use the Hessian of the error function to optimize backpropagation itself, allowing for both more efficient and effective training (for one example, see [30, 43]). These second derivatives are expected to satisfy a higher dimensional analog of the chain rule. In fact one might expect in general higher analogs of the chain rule to hold, which are sometimes called the Faa di Bruno formulae for higher chain rules on terms of the form $\partial^n(fg)$. These expected formulae will all hold in our semantics due to a result that shows **[CD.6,7]** are equivalent to having all the Faa di Bruno formulae [19]. These higher chain rule expansions can be used to determine a slightly different operational semantics for $\text{rd}(x.m)(a).v$ expressions, where the chain rule is maximally expanded first, and linearity reductions occur, and then symbolic differentiation is used. While it is unclear if this is more efficient, it would make things simpler, as it would guarantee that the operational semantics captured the higher chain rule formulae without having to make a requirement of the evaluation function on op_{RRR} .

References

- [1] Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, Manjunath Kudlur, Josh Levenberg, Rajat Monga, Sherry Moore, Derek G. Murray, Benoit Steiner, Paul Tucker, Vijay Vasudevan, Pete Warden, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. Tensorflow: a system for large-scale machine learning. *12th USENIX Symposium on Operating Systems Design and Implementation(OSDI 16)*, pages 265 – 283, 2016.
- [2] Martin Abadi and Gordon D. Plotkin. A simple differentiable programming language. *Proceedings of the ACM on Programming Languages*, 4:38:1 – 38:28, 2019. doi:<https://doi.org/10.1145/3371106>.
- [3] Jiří Adámek, Stefan Milius, and Jiří Velebil. Elgot algebras. *Electronic Notes Theoretical Computer Science*, 155:87–109, 2006. doi:10.1016/j.entcs.2005.11.053.

- [4] Atilim Günes Baydin, Barak A. Pearlmutter, Alexey Andreyevich Radul, and Jeffrey Mark Siskind. Automatic differentiation in machine learning: a survey. *Journal of Machine Learning Research*, 18(153):1 – 43, 2018.
- [5] James Bergstra, Olivier Breuleux, Frédéric Bastien, Pascal Lamblin, Razvan Pascanu, Guillaume Desjardins, Joseph Turian, David Warde-Farley, and Yoshua Bengio. Theano: A cpu and gpu math expression compiler. *Proceedings of the Python for scientific computing conference (SciPy)*, 4, 2010. doi:[10.25080/Majora-92bf1922-003](https://doi.org/10.25080/Majora-92bf1922-003).
- [6] W. Bertram, H. Glöckner, and K.-H. Neeb. Differential calculus over general base fields and rings. *Expositiones Mathematicae*, 22(3):213 – 282, 2004. doi:[https://doi.org/10.1016/s0723-0869\(04\)80006-9](https://doi.org/10.1016/s0723-0869(04)80006-9).
- [7] S. Bloom and Z. Esik. *Iteration theories*. Springer - EATCS Series, 1993. doi:https://doi.org/10.1007/978-3-642-78034-9_7.
- [8] R. Blute, R. Cockett, and R. Seely. Cartesian Differential Categories. *Theory and Applications of Categories*, 22:622–672, 2009.
- [9] R.F. Blute, J.R.B. Cockett, and R.A.G. Seely. Differential categories. *Mathematical structures in computer science*, 16(6):1049–1083, 2006.
- [10] Aloïs Brunel, Damiano Mazza, and Michele Pagani. Backpropagation in the simply typed lambda-calculus with linear negation. *Proc. ACM Program. Lang.*, 4(POPL), December 2019. doi:[10.1145/3371132](https://doi.org/10.1145/3371132).
- [11] J.R.B. Cockett and G.S.H. Cruttwell. Differential bundles and fibrations for tangent categories. *Cahiers de Topologie et Géométrie Différentielle Catégoriques*, LIX(1):10–92, 2018.
- [12] J.R.B. Cockett, G.S.H. Cruttwell, and J.D. Gallagher. Differential restriction categories. *Theory and Applications of Categories*, 25(21):537–613, 2011.
- [13] J.R.B. Cockett, P.J.W. Hofstra, and P. Hrubeš. Total maps of turing categories. *Electronic Notes in Theoretical Computer Science*, 308:129–146, 2014. Proceedings of the 30th Conference on the Mathematical Foundations of Programming Semantics (MFPS XXX). doi:<https://doi.org/10.1016/j.entcs.2014.10.008>.
- [14] J.R.B. Cockett and Stephen Lack. Restriction categories I: categories of partial maps. *Theoretical Computer Science*, 270(1):223 – 259, 2002. doi:[https://doi.org/10.1016/S0304-3975\(00\)00382-0](https://doi.org/10.1016/S0304-3975(00)00382-0).
- [15] R. Cockett and G. Cruttwell. Differential structure, tangent structure, and SDG. *Applied Categorical Structures*, 22:331–417, 2014. doi:<https://doi.org/10.1007/s10485-013-9312-0>.
- [16] R. Cockett, G. Cruttwell, J. Gallagher, J.-S. Lemay, B. MacAdam, G. Plotkin, and D. Pronk. Reverse derivative categories. *arxiv:1910.07065*, (18):1–25, 2019.
- [17] R. Cockett and S. Lack. Restriction categories III: colimits, partial limits and extensivity. *Mathematical Structures in Computer Science*, 17(4):775–817, 2007. doi:[10.1017/S0960129507006056](https://doi.org/10.1017/S0960129507006056).
- [18] R. Cockett and E. Manes. Boolean and classical restriction categories. *Mathematical Structures in Computer Science*, 19(2):357–416, 2009. doi:<https://doi.org/10.1017/s0960129509007543>.
- [19] R. Cockett and R. Seely. The Faá Di Bruno Construction. *Theory and Applications of Categories*, 25:294–425, 2011.
- [20] Robin Cockett, Joaquín Díaz-Boils, Jonathan Gallagher, and Pavel Hrubeš. Timed sets, functional complexity, and computability. *Electronic Notes in Theoretical Computer Science*, 286:117–137, 2012. Proceedings of the 28th Conference on the Mathematical Foundations of Programming Semantics (MFPS XXVIII). doi:<https://doi.org/10.1016/j.entcs.2012.08.009>.
- [21] Robin Cockett, Xiuzhan Guo, and Pieter Hofstra. Range categories II: Towards regularity. *Theory and applications of categories*, 26(18):453–500, 2012.
- [22] J. H. Conway. *Regular algebra and finite machines*. Chapman and Hall Mathematics Series, 1971.
- [23] E. Dubuc. Sur les modèles de la géométrie différentielle synthétique. *Cahiers de Topologie Géométrie Différentielle Catégoriques*, 20(1):231–279, 1979.
- [24] E. Dubuc and A. Kock. On 1-Form Classifiers. *Communications in Algebra*, 12(12):1471–1531, 1984. doi:<https://doi.org/10.1080/00927878408823064>.

- [25] Calvin C. Elgot. *Monadic Computation and Iterative Algebraic Theories*, pages 179–234. Springer New York, New York, NY, 1982. doi:https://doi.org/10.1007/978-1-4613-8177-8_6.
- [26] Conal Elliott. The simple essence of automatic differentiation. *Proceedings of the ACM on Programming Languages*, 2(ICFP):70, 2018. doi:<https://doi.org/10.1145/3236765>.
- [27] M. P. Fiore and G. D. Plotkin. An axiomatisation of computationally adequate domain theoretic models of fpc. In *Proceedings Ninth Annual IEEE Symposium on Logic in Computer Science*, pages 92–102, 1994. doi:<https://doi.org/10.1109/lics.1994.316083>.
- [28] J. Gallagher. *The differential lambda-calculus: syntax and semantics for differential geometry*. PhD thesis, University of Calgary, 2009.
- [29] J. Gallagher. What is a differential partial combinatory algebra? Master’s thesis, University of Calgary, 2011.
- [30] B. Ghorbani, S. Krishnan, and Y. Xiao. An investigation into neural net optimization via hessian eigenvalue density. *Proceedings of Machine Learning Research*, 97, 2019.
- [31] Brett Giles. *An investigation of some theoretical aspects of reversible computing*. PhD thesis, University of Calgary, 2014.
- [32] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016.
- [33] Marco Grandis. Cohesive categories and manifolds. *Annali di Matematica Pura ed Applicata*, 157(1):199–244, 1990. doi:<https://doi.org/10.1007/bf01765319>.
- [34] A. Kock. *Synthetic differential geometry*. Cambridge University Press, 1981. doi:<https://doi.org/10.1017/cbo9780511550812>.
- [35] R. Lavendhomme. *Basic Concepts of Synthetic Differential Geometry*. Kluwer Texts in Mathematical Sciences. Kluwer Academic Publishers, 1996. doi:<https://doi.org/10.1007/978-1-4757-4588-7>.
- [36] Dougal Maclaurin, David Duvenaud, and Ryan P. Adams. Autograd: effortless gradients in numpy. *ICML 2015 AutoML Workshop*, 238, 2015.
- [37] C. Nester. Turing categories and realizability. Master’s thesis, University of Calgary, 2017.
- [38] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, and Luca et al Antiga. Pytorch: An imperative style, high-performance deep learning library. *Advances in Neural Information Processing Systems*, pages 8024 – 8035, 2019.
- [39] Barak A. Pearlmutter and Jeffrey Mark Siskind. Reverse-mode ad in a functional framework: Lambda the ultimate backpropagator. *ACM Trans. Program. Lang. Syst.*, 30(2), March 2008. doi:[10.1145/1330017.1330018](https://doi.org/10.1145/1330017.1330018).
- [40] G.D. Plotkin. LCF considered as a programming language. *Theoretical Computer Science*, 5(3):223 – 255, 1977. doi:[https://doi.org/10.1016/0304-3975\(77\)90044-5](https://doi.org/10.1016/0304-3975(77)90044-5).
- [41] Dana S. Scott. A type-theoretical alternative to ISWIM, CUCH, OWHY. *Theoretical Computer Science*, 121(1–2):411–440, 1993. doi:[https://doi.org/10.1016/0304-3975\(93\)90095-b](https://doi.org/10.1016/0304-3975(93)90095-b).
- [42] Alfred Tarski. A lattice-theoretical fixpoint theorem and its applications. *Pacific Journal of Mathematics*, 5(2):285 – 309, 1955. doi:<https://doi.org/10.2140/pjm.1955.5.285>.
- [43] L. Tzu-Mao. *Differentiable Visual Computing*. PhD thesis, Massachusetts Institute of Technology, 2019.
- [44] Tarmo Uustalu and Niccolò Veltri. The delay monad and restriction categories. In *Theoretical Aspects of Computing – ICTAC 2017*, pages 32–50. Springer International Publishing, 2017. doi:https://doi.org/10.1007/978-3-319-67729-3_3.
- [45] Bart van Merriënboer, Dan Moldovan, and Alexander B. Wiltschko. Tangent: automatic differentiation using source-code transformation for dynamically typed array programming. *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, 3–8 December 2018, Montréal, Canada*, pages 6259 – 6268, 2018.

- [46] Fei Wang, Daniel Zheng, James Decker, Xilun Wu, Grégory M. Essertel, and Tiark Rompf. Demystifying differentiable programming: Shift/reset the penultimate backpropagator. *Proc. ACM Program. Lang.*, 3(ICFP), July 2019. doi:10.1145/3341700.