

Categorical Foundations of Gradient-Based Learning

G.S.H. CRUTTWELL, Mount Allison University, Canada

BRUNO GAVRANOVIĆ and NEIL GHANI, University of Strathclyde, UK

PAUL WILSON and FABIO ZANASI, University College London, UK

We propose a categorical semantics of gradient-based machine learning algorithms in terms of lenses, parametrised maps, and reverse derivative categories. This foundation provides a powerful explanatory and unifying framework: it encompasses a variety of gradient descent algorithms such as ADAM, AdaGrad, and Nesterov momentum, as well as a variety of loss functions such as MSE and Softmax cross-entropy, shedding new light on their similarities and differences. Our approach to gradient-based learning has examples generalising beyond the familiar continuous domains (modelled in categories of smooth maps) and can be realized in the discrete setting of boolean circuits. Finally, we demonstrate the practical significance of our framework with an implementation in Python.

Additional Key Words and Phrases: Category Theory, Semantics, Machine Learning, Lenses

1 INTRODUCTION

The last decade has witnessed a surge of interest in machine learning, fuelled by the numerous successes and applications that these methodologies have found in many fields of science and technology. As machine learning techniques become increasingly pervasive, algorithms and models become more sophisticated, posing a significant challenge both to the software developers and the users that need to interface, execute and maintain these systems. In spite of this rapidly evolving picture, the formal analysis of many learning algorithms mostly takes place at a heuristic level [Se-shia and Sadigh 2016], or using definitions that fail to provide a general and scalable framework for describing machine learning. Indeed, it is commonly acknowledged through academia, industry, policy makers and funding agencies that there is a pressing need for a unifying perspective, which can make this growing body of work more systematic, rigorous, transparent and accessible both for users and developers [Exp 2019; Olah 2015].

Consider, for example, one of the most common machine learning scenarios: supervised learning with a neural network. This technique trains the model towards a certain task, e.g. the recognition of patterns in a data set (*cf.* Figure 1). There are several different ways of implementing this scenario. Typically, at their core, there is a *gradient update* algorithm (often called the “optimiser”), depending on a given *loss function*, which updates in steps the parameters of the network, based on some *learning rate* controlling the “scaling” of the update. All of these components can vary independently in a supervised learning algorithm and a number of choices is available for loss maps (quadratic error, Softmax cross entropy, dot product, etc.) and optimisers (Adagrad [Duchi et al. 2011], Momentum [Polyak 1964], and Adam [Kingma and Ba 2015], etc.).

This scenario highlights several questions: is there a uniform mathematical language capturing the different components of the learning process? Can we develop a unifying picture of the various optimisation techniques, allowing for their comparative analysis? Moreover, it should be noted that supervised learning is not limited to neural networks. For example, supervised learning is surprisingly applicable to the discrete setting of boolean circuits [Wilson and Zanasi 2020]

Authors' addresses: G.S.H. Cruttwell, Department of Mathematics and Computer Science, Mount Allison University, Canada; Bruno Gavranović; Neil Ghani, University of Strathclyde, UK, bruno@brunogavranovic.com; Paul Wilson, paul@statusfailed.com; Fabio Zanasi, University College London, UK.

2018. 2475-1421/2018/1-ART1 \$15.00

<https://doi.org/>




Fig. 1. An informal illustration of gradient-based learning. This neural network is trained to distinguish different kinds of animals in the input image. Given an input X , the network predicts an output Y , which is compared by a ‘loss map’ with what would be the correct answer (‘label’). The loss map returns a real value expressing the error of the prediction; this information, together with the *learning rate* (a weight controlling how much the model should be changed in response to error) is used by an *optimiser*, which computes by gradient-descent the update of the parameters of the network, with the aim of improving its accuracy. The neural network, the loss map, the optimiser and the learning rate are all components of a supervised learning system, and can vary independently of one another.

where continuous functions are replaced by boolean-valued functions. Can we identify an abstract perspective encompassing both the real-valued and the boolean case? In a nutshell, this paper seeks to answer the question:


what are the fundamental mathematical structures underpinning gradient-based learning?

Our approach to this question stems from the identification of three fundamental aspects of the gradient-descent learning process:

- (I) computation is **parameterised**, e.g. in the most simple case we are given a function $f : P \times X \rightarrow Y$ and learning consists of finding a parameter $p : P$ such that $f(p, -)$ is the best function according to some criteria. More generally, the weights on the internal nodes of a neural network are a parameter which the learning is seeking to optimize. Parameters also arise elsewhere, e.g. in the loss function (see later).
- (II) information flows **bidirectionally**: in the forward direction the computation turns inputs via a sequence of *layers* into predicted outputs, and then into a loss value; in the reverse direction, backpropagation is used to propagate the changes *backwards* through the layers, and then turn them into parameter updates.
- (III) the basis of parameter update via gradient descent is **differentiation** e.g. in the simple case we differentiate the function mapping a parameter to its associated loss to reduce that loss.

We model bidirectionality via lenses [Bohannon et al. 2008; Clarke et al. 2020; Hedges 2018] and based upon the above three insights, we propose the notion of **parametric lens** as the fundamental semantic structure of learning. In a nutshell, a parametric lens is a process with three kinds of interfaces: inputs, outputs, and parameters. On each interface, information flows both ways, i.e. computations are bidirectional. These data are best explained with our graphical representation of parametric lenses, with inputs A, A' , outputs B, B' , parameters P, P' , and arrows indicating information flow (below left). The graphical notation also makes evident that parametric lenses are

open systems, which may be composed along their interfaces (below center and right).



This pictorial formalism is not just an intuitive sketch: as we will show, it can be understood as a completely formal (graphical) syntax using the formalism of *string diagrams* [Selinger 2010], in a way similar to how other computational phenomena have been recently analysed e.g. in quantum theory [Coecke and Kissinger 2017], control theory [Baez and Erbele 2014; Bonchi et al. 2017], and digital circuit theory [Ghica et al. 2017].

It is intuitively clear how parametric lens express aspects (I) and (II) above, whereas (III) will be achieved by studying them in a space of ‘differentiable objects’ (in a sense that will be made precise). The main technical contribution of our paper is showing how the various ingredients involved in learning (the model, the optimiser, the error map and the learning rate) can be uniformly understood as being built from parametric lenses.




Fig. 2. The parametric lens that captures the learning process informally sketched in Figure 1. Note each component is a lens itself, whose composition yields the interactions described in Figure 1. Defining this picture formally will be the subject of Sections 3-4. Also, see Appendix B for an animation of lens computation within this structure.

We will use *category theory* as the formal language to develop our notion of parametric lenses, and make Figure 2 mathematically precise. The categorical perspective brings several advantages, which are well-known, established principles in programming language semantics [Abramsky and Coecke 2004; Selinger 2001; Turi and Plotkin 1997]. Three of them are particularly important to our contribution, as they constitute distinctive advantages of our semantic foundations:

Abstraction Our approach studies which categorical structures are sufficient to perform gradient-based learning. This analysis abstracts away from the standard case of neural networks in several different ways: as we will see, it encompasses other models (namely Boolean circuits), different kinds of optimisers (including Adagrad, Adam, Nesterov momentum), and error maps (including quadratic and softmax cross entropy loss). These can be all understood as parametric lenses, and different forms of learning result from their interaction.

Uniformity As seen in Figure 1, learning involves ingredients that are seemingly quite different: a model, an optimiser, a loss map, etc. We will show how all these notions may be seen as instances of the categorical definition of a parametric lens, thus yielding a remarkably

uniform description of the learning process, and supporting our claim of parametric lenses being a fundamental semantic structure of learning.

Compositionality The use of categorical structures to describe computation naturally enables *compositional reasoning* whereby complex systems are analysed in terms of smaller, and hence easier to understand, components. Compositionality is a fundamental tenet of programming language semantics; in the last few years, it has found application in the study of diverse kinds of computational models, across different fields— see e.g. [Bonchi et al. 2017; Coecke and Kissinger 2017; Ghani et al. 2016; Spivak 2010]. As made evident by Figure 2, our approach models a neural network as a parametric lens, resulting from the *composition* of simpler parametric lenses, capturing the different ingredients involved in the learning process. Moreover, as all the simpler parametric lenses are themselves composable, one may engineer a different learning process by simply plugging a new lens on the left or right of existing ones. This means that one can glue together smaller and relatively simple networks to create larger and more sophisticated neural networks.

We now give a synopsis of our contributions:

- In Section 2, we introduce the tools necessary to define our notion of **parametric lens**. First, in Section 2.1, we introduce a notion of parametrisation for categories, which amounts to a functor $\text{Para}(-)$ turning a category C into one $\text{Para}(C)$ of ‘parametrised C -maps’. Second, we recall *lenses* (Section 2.2). In a nutshell, a lens is a categorical morphism equipped with operations to view and update values in a certain data structure. Lenses play a prominent role in functional programming [Steckermeier 2015], as well as in the foundations of database theory [Johnson et al. 2012] and more recently game theory [Ghani et al. 2016]. Considering lenses in C simply amounts to the application of a functorial construction $\text{Lens}(-)$, yielding $\text{Lens}(C)$. Finally, we recall the notion of a *cartesian reverse differential category* (CRDC): a categorical structure axiomatising the notion of differentiation [Cockett et al. 2019] (Section 2.4). We wrap up in Section 2.3, by combining these ingredients into the notion of parametric lens, formally defined as a morphism in $\text{Para}(\text{Lens}(C))$ for a CRDC C . In terms of our desiderata (I)-(III) above, note that $\text{Para}(-)$ accounts for (I), $\text{Lens}(-)$ accounts for (II), and the CRDC structure accounts for (III).
- As seen in Figure 1, in the learning process there are many components at work: the model, the optimiser, the loss map, the learning rate, etc.. In Section 4, we show how the previously introduced notion of parametric lens provides a uniform characterisation for such components. Moreover, for each of them, we show how different variations appearing in the literature become instances of our abstract characterisation. The plan is as follows:
 - In Section 3.1, we show how the combinatorial **model** subject of the training can be seen as a parametric lens. The conditions we provide are met by the ‘standard’ case of neural networks, but also enables the study of learning for other classes of models. In particular, another instance are Boolean circuits: learning of these structures is relevant to binarisation [Courbariaux et al. [n.d.]] and it has been explored recently using a categorical approach [Wilson and Zanasi 2020], which turns out to be a particular case of our framework.
 - In Section 3.2, we show how the **loss maps** associated with training are also parametric lenses. We also show how our approach covers the cases of quadratic error, Boolean error, Softmax cross entropy, but also the ‘dot product loss’ associated with the phenomenon of deep dreaming [Dosovitskiy and Brox 2015; Mahendran and Vedaldi 2014; Nguyen et al. 2014; Simonyan et al. 2014].

- In Section 3.3, we model the **learning rate** as a parametric lens. This analysis also allows us to contrast how learning rate is handled in the ‘real-valued’ case of neural networks with respect to the ‘Boolean-valued’ case of Boolean circuits.
- In Section 3.4, we show how **optimisers** can be modelled as ‘reparametrisations’ of models as parametric lenses. As case studies, in addition to basic gradient update, we consider the stateful variants: Momentum [Polyak 1964], Nesterov Momentum [Sutskever et al. 2013], Adagrad [Duchi et al. 2011], and Adam (Adaptive Moment Estimation) [Kingma and Ba 2015]. Also, on Boolean circuits, we show how the reverse derivative ascent of [Wilson and Zanasi 2020] can be also regarded in such way.
- In Section 4, we study how the composition of the lenses defined in Section 3 yields a description of different kinds of learning processes.
 - Section 4.1 is dedicated to modelling supervised **learning of parameters**, in the way described in Figure 1. This amounts essentially to study the composite of lenses expressed in Figure 2, for different choices of the various components. In particular we show (i) quadratic loss with basic gradient descent, (ii) softmax cross entropy loss with basic gradient descent, (iii) quadratic loss with Nesterov momentum, and (iv) learning in Boolean circuits with XOR loss and basic gradient ascent.
 - In order to showcase the flexibility of our approach, in Section 4.2 we depart from our ‘core’ case study of parameter learning, and turn attention to supervised **learning of inputs**. The idea behind this technique, sometimes called **deep dreaming**, is that, instead of the network parameters, one updates the inputs, in order to elicit a particular interpretation [Dosovitskiy and Brox 2015; Mahendran and Vedaldi 2014; Nguyen et al. 2014; Simonyan et al. 2014]. Deep dreaming can be easily expressed within our approach, with a different rearrangement of the various parametric lenses involved in the learning process, see Figure 7 below. The abstract viewpoint of categorical semantics provides a mathematically precise and visually captivating description of the differences between the usual parameter learning process and deep dreaming.
- In Section 5 we describe a proof-of-concept Python **implementation**¹ based on the theory developed in this paper. This code is intended to show more concretely the payoff of our approach. Model architectures, as well as the various components participating in the learning process, are now expressed in a uniform, principled mathematical language, in terms of lenses. As a result, computing network gradients is greatly simplified, as it amounts to lens composition. Moreover, the modularity of this approach allows one to more easily tune the various parameters of training. We give a demonstration of our library via a number of experiments, and prove correctness by achieving accuracy on par with an equivalent model in Keras, a mainstream deep learning framework [Chollet et al. 2015]. In particular, we create a working non-trivial neural network model for the MNIST image-classification problem [Lecun et al. 1998].
- Finally, in Sections 6 and 7, we discuss related and future work.

2 CATEGORICAL TOOLKIT

In this section, we describe the three categorical components of our framework, each corresponding to an aspect of gradient-based learning. In Section 2.1 we review the **Para** construction which builds a category of parametrised maps from a monoidal category, and describe its graphical language. In Section 2.2, we review the **Lens** construction, which builds a category of “bidirectional” maps out of a Cartesian category, and describe its graphical language. In Section 2.3, we look at what happens

¹Available at <https://github.com/statusfailed/numeric-optics-python/>.

when we combine these two constructions, and the resulting graphical language of “parametric lenses”. In Section 2.4, we review Cartesian reverse differential categories, a setting for categories equipped with an abstract gradient operator, and how their structure relates to categories of lenses and parameteric lenses.

Then, in the following sections, we will see how these components fit together, allowing us to describe parametrised models and the algorithms used to train them.

2.1 Parametrized Maps

In supervised learning one is typically interested in approximating a function

$$g : \mathbb{R}^n \rightarrow \mathbb{R}^m$$

for some n and m . To do this, one begins by building a neural network, which is a smooth map

$$f : \mathbb{R}^P \times \mathbb{R}^n \rightarrow \mathbb{R}^m$$

where \mathbb{R}^P is the set of possible weights of that neural network. Then one looks for a value of $p \in \mathbb{R}^P$ such that the function $f(p, -) : \mathbb{R}^n \rightarrow \mathbb{R}^m$ closely approximates g . The first thing we need to do is formalize these types of maps categorically, and this is done via the **Para** construction [Capucci et al. 2021; Fong et al. 2017; Gavranovic 2019].

Definition 2.1 (Parametrised category). If C is a strict² symmetric monoidal category (with monoidal product \otimes and monoidal unit I) then we define a category **Para**(C) with

- objects those of C
- a map from A to B in **Para**(C) is a pair (P, f) where P is an object of C and $f : P \otimes A \rightarrow B$
- the identity on A is the pair $(I, 1_A)$ (since \otimes is strict monoidal, $I \otimes A = A$)
- the composite of $(P, f) : A \rightarrow B$ with $(P', f') : B \rightarrow C$ is the pair

$$(P' \otimes P, (1 \otimes f); g).$$

Example 2.2. Our primary example for the above construction is the category **Smooth** whose objects are natural numbers with a map $f : n \rightarrow m$ a smooth map from \mathbb{R}^n to \mathbb{R}^m . As described above, the category **Para**(**Smooth**) can be thought of as a category of neural networks: a map in this category from n to m consists of a choice of p and a map $f : \mathbb{R}^P \times \mathbb{R}^n \rightarrow \mathbb{R}^m$ with \mathbb{R}^P representing the set of possible weights of the neural network.

As anticipated in the introduction, we represent the morphisms of **Para**(C) graphically using the formalism of *string diagrams*—see [Selinger 2010] for a general overview. As we will see in the next sections, the interplay of the various components at work in the learning process becomes much clearer once represented with this pictorial notation.

In fact, we will mildly massage the traditional notation for string diagrams, which would represent a morphism in **Para**(C) from A to B as below left.



Note the standard notation does not emphasise the special role played by P , which is part of the data of the morphism itself. Parameters and data in machine learning have different semantics: by separating them on two different axes, we obtain a graphical language which is more closely tied to these semantics. Thus, we will use a slightly different convention for **Para**(C), writing a morphism $(P, f) : A \rightarrow B$ as on the right above. Incidentally, this clarifies why composition in **Para**(C) is

²One can also define **Para**(C) in the case when C is non-strict; however, the result would be not a category but a bicategory.

defined the way it is: the composite of $(P, f) : A \rightarrow B$ with $(P', f') : B \rightarrow C$ is simply given by hooking up the B wires:

$$\begin{array}{c}
 P \qquad P' \\
 | \qquad | \\
 \boxed{f} \text{---} B \text{---} \boxed{f'} \\
 A \qquad \qquad C
 \end{array} \quad (1)$$

This notation also yields a neat visualisation of “reparameterisation”, as defined below.

Definition 2.3. A **reparameterisation** of $(P, f) : A \rightarrow B$ in $\mathbf{Para}(C)$ by a map $\alpha : Q \rightarrow P$ (below left) is the $\mathbf{Para}(C)$ map $(Q, (\alpha \otimes 1_A); f) : A \rightarrow B$ (represented below right).

$$\begin{array}{ccc}
 \begin{array}{c} P \\ | \\ \boxed{f} \\ A \rightarrow \quad B \end{array} & & \begin{array}{c} Q \\ | \\ \boxed{\alpha} \\ | \\ P \\ | \\ \boxed{f} \\ A \rightarrow \quad B \end{array}
 \end{array} \quad (2)$$

Intuitively, reparameterisation changes the parameter space of $(P, f) : A \rightarrow B$ to some other object Q , via some map $\alpha : Q \rightarrow P$. We shall see later that gradient descent and its many variants can naturally be viewed as reparameterisations.

Note coherence rules in combining operations (1) and (2) just work as expected, as these diagrams can be ultimately ‘compiled’ down to string diagrams for monoidal categories. For example, given maps $(P, f) : A \rightarrow B$, $(Q, g) : B \rightarrow C$ with reparameterisations $\alpha : P' \rightarrow P$, $\beta : Q' \rightarrow Q$, one could either first reparameterise f and g separately and then compose the results (below left), or compose first then reparameterise jointly (below right):

$$\begin{array}{ccc}
 \begin{array}{c} P' \\ | \\ \boxed{\alpha} \\ | \\ P \\ | \\ \boxed{f} \\ A \rightarrow \quad B \end{array} & \circledast & \begin{array}{c} Q' \\ | \\ \boxed{\beta} \\ | \\ Q \\ | \\ \boxed{g} \\ B \rightarrow \quad C \end{array} \\
 & & \begin{array}{c} P' \qquad Q' \\ | \qquad | \\ \boxed{\alpha} \qquad \boxed{\beta} \\ | \qquad | \\ P \qquad Q \\ | \qquad | \\ \boxed{f} \text{---} \boxed{g} \\ A \rightarrow \quad B \rightarrow \quad C \end{array}
 \end{array} \quad (3)$$

As expected, translating these two operations into string diagrams for monoidal categories yield equivalent representations of the same morphism.

$$\begin{array}{ccc}
 \begin{array}{c} Q' \text{---} \boxed{\beta} \text{---} Q \\ | \qquad | \\ P' \text{---} \boxed{\alpha} \text{---} P \\ | \qquad | \\ A \text{---} \boxed{f} \text{---} B \end{array} & \rightarrow & \begin{array}{c} Q' \text{---} \boxed{\beta} \text{---} Q \\ | \qquad | \\ P' \text{---} \boxed{\alpha} \text{---} P \\ | \qquad | \\ A \text{---} \boxed{f} \text{---} B \end{array} \\
 & & \begin{array}{c} \boxed{g} \text{---} C \\ | \\ B \end{array}
 \end{array} \quad (4)$$

REMARK 2.1. *There is a 2-categorical perspective on $\mathbf{Para}(C)$, which we glossed over in this paper for the sake of simplicity. In particular, the reparameterisations described above can also be seen as equipping $\mathbf{Para}(C)$ with 2-cells, giving a 2-categorical structure on $\mathbf{Para}(C)$. This is also coherent with respect to base change: if C and D are strict symmetric monoidal categories, and $F : C \rightarrow D$ a lax symmetric monoidal functor, then there is an induced 2-functor $\mathbf{Para}(F) : \mathbf{Para}(C) \rightarrow \mathbf{Para}(D)$ which agrees with F on objects. This 2-functor is straightforward: for a 1-cell $(P, f) : A \rightarrow B$, it applies F to P and f and uses the (lax) comparison to get a map of the correct type. We will see how this base change becomes important when performing backpropagation on parameterised maps (Eq. 8)*

Lastly, we mention that $\mathbf{Para}(C)$ inherits the symmetric monoidal structure from C and that the induced 2-functor $\mathbf{Para}(F)$ respects that structure. This will allow us to compose neural networks not only in series, but also in parallel. For more detail on alternative viewpoints on the \mathbf{Para} construction, including how it can be viewed as the Grothendieck construction of a certain indexed category, see [Capucci et al. 2021].

2.2 Lenses

We next consider a very different categorical construction. In machine learning (or even learning in general) it is fundamental that information flows both forwards and backwards: the ‘forward’ flow corresponds to a model’s predictions, and the ‘backwards’ flow to *corrections* to the model. The category of lenses is the ideal setting to capture this type of structure, as it is a category consisting of maps with both a “forward” and a “backward” part.


Definition 2.4. For any Cartesian category C , the category of lenses³ in C , $\mathbf{Lens}(C)$, is the category with the following data:

- Its objects are pairs (A, A') , where both A and A' are objects in C
- A map from (A, A') to (B, B') consists of a pair (f, f^*) where $f : A \rightarrow B$ (called the **get** or **forward** part of the lens) and $f^* : A \times B' \rightarrow A'$ (called the **put** or **backwards** part of the lens)
- the composite of $(f, f^*) : (A, A') \rightarrow (B, B')$ and $(g, g^*) : (B, B') \rightarrow (C, C')$ is given by get $f;g$ and put


$$\langle \pi_0, \langle \pi_0; f, \pi_1 \rangle; g^* \rangle; f^*$$

- The identity on (A, A') is the pair $(1_A, \pi_1)$.

It is much easier to visualize the morphisms of $\mathbf{Lens}(C)$ and their composites with a graphical calculus as described in [Boisseau 2020, Thm. 23]. In this language, a morphism $(f, f^*) : (A, A') \rightarrow (B, B')$ is written as




where \langle is the string diagram (‘built-in’ in any cartesian category) duplicating the value A . It is clear in this language how to describe the composite of $(f, f^*) : (A, A') \rightarrow (B, B')$ and $(g, g^*) : (B, B') \rightarrow (C, C')$: simply join the B/B' wires together to get the composite lens



³These lenses are often called *bimorphic* [Hedges 2018], in contrast to *simple* lenses whose objects are all of the form (A, A) : the forward and backward types of simple lenses are the same.

and the formula for the composite in terms of equations (as described above) follows from this.

We will often write lenses without the inside wires exposed, thinking of the entire lens as a black-box:



Note $\mathbf{Lens}(C)$ is a monoidal category, with $(A, A') \otimes (B, B')$ defined as $(A \times B, A' \times B')$. However, in general $\mathbf{Lens}(C)$ is not itself Cartesian. This is easy to see when looking at even a terminal object: if T is a terminal object in C , then in general (T, T) will not be a terminal object in $\mathbf{Lens}(C)$ – if it was, there would be a unique lens $(!_A, !_A) : (A, A') \rightarrow (T, T)$ whose put part would need to be a (unique) map $A \times T \rightarrow A'$, but in general there are many such maps.

2.3 Parametric Lenses

The fundamental category where supervised learning takes place is the composite of the two constructions in the previous sections. As noted in the previous section, for a Cartesian category C , $\mathbf{Lens}(C)$ is monoidal, and so we can form the the category $\mathbf{Para}(\mathbf{Lens}(C))$, which we shall call the category of **parametric lenses** of C . The definition of this category follows automatically from definitions of \mathbf{Para} and $\mathbf{Lens}(C)$:

Definition 2.5. The category $\mathbf{Para}(\mathbf{Lens}(C))$ of parametric lenses on C is defined as follows.

- An object is a pair of objects (A, A') from C
- A morphism from (A, A') to (B, B') , called a parametric lens⁴, is a choice of parameter pair (P, P') and a lens


$$(f, f^*) : (A, A') \times (P, P') \rightarrow (B, B')$$

so that $f : P \times A \rightarrow B$ and $f^* : P \times A \times B' \rightarrow P' \times A'$

By the previous two sections, we get a graphical language for $\mathbf{Para}(\mathbf{Lens}(C))$ which uses the graphical language for $\mathbf{Lens}(C)$ from section 2.2 as a base, then augments it with parameters as described in section 2.1. Thus a morphism of $\mathbf{Para}(\mathbf{Lens}(C))$ from (A, A') to (B, B') is a box with input/output of (A, A') on the left, input/output of (B, B') on the right, and input/output of (P, P') (the parameter space) on top:



Composition is again quite natural in this formulation: given one box with input/output wires (B, B') on the right and another box with input/output wires (B, B') on the left, one simply hooks up those input/output wires to get the desired composite:



⁴In [Fong et al. 2017], these are called *learners*. However, in this paper we study them in a much broader light; see section 6.

A reparameterisation in $\text{Para}(\text{Lens}(C))$ is depicted graphically as drawing a box on top of the (P, P') wires:



Given a generic morphism f in $\text{Para}(\text{Lens}(C))$ as depicted in (5), one can see how it is possible to “learn” new values from f : it takes as input an input A , a parameter P , and a change B' , and outputs a change in A , a value of B , and a change P' . This last element is the key component for supervised learning: intuitively, it says how to change the parameter values to get the neural network closer to the true value of the desired function.

The question, then, is how one is to define such a parametric lens given nothing more than a neural network, i.e., a parameterized map $(P, f) : A \rightarrow B$. This is precisely what the gradient operation provides, and its generality to categories is explored in the next subsection.

2.4 Cartesian Reverse Differential Categories

Fundamental to all gradient-based learning is, of course, the gradient operation. In most cases this gradient operation is performed in the category of smooth maps between Euclidean spaces. However, recent work [Wilson and Zanasi 2020] has shown that gradient-based learning can also work well in other categories; for example, in a category of boolean circuits. Thus, to encompass these examples in a single framework, it is helpful to work in a category with an abstract gradient operation. Specifically, we will work in a Cartesian reverse differential category (first defined in [Cockett et al. 2019]), a category in which every map has an associated reverse derivative.

Definition 2.6.

- [Cockett et al. 2019, Defn. 1] A **Cartesian left additive category** consists of a category C with chosen finite products (including a terminal object), and an addition operation and zero morphism in each homset, satisfying various axioms.
- [Cockett et al. 2019, Defn. 13] A **Cartesian reverse differential category** (CRDC) consists of a Cartesian left additive category C , together with an operation which provides, for each map $f : A \rightarrow B$ in C , a map


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

satisfying seven axioms (for full details, see the appendix).


Why are reverse derivatives helpful for learning? For $f : A \rightarrow B$, the pair $(f, R[f])$ form a lens from (A, A) to (B, B) , with $R[f]$ acting as backwards map. Thus having a reverse derivative already provides a way to turn an ordinary map in the category into which one can pass information backwards, that is, a map which can “learn”.

Note assigning type $A \times B \rightarrow A$ to $R[f]$ hides some relevant information: B -values in the domain and A -values in the codomain of $R[f]$ do not play the same role as values of the same types in $f : A \rightarrow B$: in $R[f]$, they really take in a tangent vector at B and output a tangent vector at A (cf. the definition of $R[f]$ in Smooth, Example 2.8 below). To emphasise this, we will type $R[f]$ as a map $A \times B' \rightarrow A'$ (even though in reality $A = A'$ and $B = B'$), thus meaning that $(f, R[f])$ is actually a lens from (A, A') to (B, B') . This typing distinction will be helpful later on, when we want to add additional components to our learning algorithms.

Graphically, then, we represent the pair $(f, R[f])$ as a lens:



This point of view also makes clear the usefulness of the reverse chain rule (axiom [RD.5] from [Cockett et al. 2019, Defn. 13]): it tells us that the operation which takes a map f and produces the lens $(f, R[f])$ preserves composition: that is,



Combined with axiom [RD.3] for a CRDC, this justifies the following fact, which we record for later use.

PROPOSITION 2.7. [Cockett et al. 2019, Prop. 31] *If \mathcal{C} is a CRDC, there is a functor $\mathbf{R}: \mathcal{C} \rightarrow \mathbf{Lens}(\mathcal{C})$ which on objects sends A to the pair (A, A) and on maps sends $f: A \rightarrow B$ to the pair $(f, R[f])$.*

The following two examples of CRDCs will serve as the basis for the learning scenarios of the upcoming sections.

Example 2.8. The category Smooth has as objects natural numbers and maps $n \rightarrow m$ are m -tuples of smooth maps $f: \mathbb{R}^n \rightarrow \mathbb{R}$. Smooth is Cartesian with product given by addition. Smooth is a Cartesian reverse differential category: given a smooth map $f: \mathbb{R}^n \rightarrow \mathbb{R}^m$, the map

$$R[f]: \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^n$$

sends a pair (x, v) to $J[f]^T(x) \cdot v$: the transpose of the Jacobian of f at x in the direction v .

For example, if $f: \mathbb{R}^2 \rightarrow \mathbb{R}^3$ is defined as $f(x_1, x_2) := (x_1^3 + 2x_1x_2, x_2, \sin(x_1))$, then $R[f]:$

$$\mathbb{R}^2 \times \mathbb{R}^3 \rightarrow \mathbb{R}^2 \text{ is given by } (x, v) \mapsto \begin{bmatrix} 3x_1 + 2x_2 & 0 & \cos(x_1) \\ 2x_1 & 1 & 0 \end{bmatrix} \cdot \begin{bmatrix} v_1 \\ v_2 \\ v_3 \end{bmatrix}. \text{ Using the reverse derivative}$$

(as opposed to the forward derivative⁵) is well-known to be much more computationally efficient for functions $f: \mathbb{R}^n \rightarrow \mathbb{R}^m$ when $m \ll n$ (for example, see [Griewank and Walther 2008]), as is the case in most supervised learning situations (where often $m = 1$).

Example 2.9. Another RDC is the PROP $\text{POLY}_{\mathbb{Z}_2}$ [Cockett et al. 2019, Example 14], whose morphisms $f: A \rightarrow B$ are B -tuples of polynomials $\mathbb{Z}_2[x_1 \dots x_A]$. When presented by generators and relations these morphisms can be viewed as a syntax for boolean circuits, with parametric lenses for such circuits (and their reverse derivative) described in [Wilson and Zanasi 2020].

⁵Forward derivatives can analogously be modelled with Cartesian (forward) differential categories [Cruttwell 2012]. Given a map $f: A \rightarrow B$ the forward derivative $D[f]: A \times A \rightarrow B$ uses the Jacobian of f , instead of its transpose.

3 COMPONENTS OF LEARNING AS PARAMETRIC LENSES

As seen in the introduction, in the learning process there are many components at work: a model, an optimiser, a loss map, a learning rate, etc. In this section we show how each such component can be understood as a parametric lens. Moreover, for each component, we show how our framework encompasses several variations of the gradient-descent algorithms, thus offering a unifying perspective on many different approaches that appear in the literature.


3.1 Models as Parametric Lenses

We begin by characterising the models used for training as a parametric lenses. In essence, our approach identifies a set of abstract requirements necessary to perform training by gradient descent, which covers the case studies that we will consider in the next sections.

The leading intuition is that a suitable model is a parametrised map, equipped with a reverse derivative operator. Using the formal developments of Section 2, this amounts to assuming that a model is a morphism in $\mathbf{Para}(C)$, for a CRDC C . In order to visualise such morphism as a parametric lens, it then suffices to apply \mathbf{R} from Proposition 2.7 under $\mathbf{Para}(-)$, yielding a functor

$$\mathbf{Para}(\mathbf{R}) : \mathbf{Para}(C) \rightarrow \mathbf{Para}(\mathbf{Lens}(C)) \quad (8)$$

Pictorially, $\mathbf{Para}(\mathbf{R})$ send a map as on the left below to a parametric lens as on the right.



Example 3.1 (Neural networks). As noted previously, to learn a function of type $\mathbb{R}^n \rightarrow \mathbb{R}^m$, one constructs a neural network, which can be seen as a function of type $\mathbb{R}^p \times \mathbb{R}^n \rightarrow \mathbb{R}^m$ where \mathbb{R}^p is the space of parameters of the neural network. As seen in Example 2.2, this is a map in the category $\mathbf{Para}(\mathbf{Smooth})$ of type $\mathbb{R}^n \rightarrow \mathbb{R}^m$ with parameter space \mathbb{R}^p . Then one can apply the functor in (8) to present a neural network together with its reverse derivative operator as a parameterised lens, i.e. a morphism in $\mathbf{Para}(\mathbf{Lens}(\mathbf{Smooth}))$.

Example 3.2 (Boolean circuits). For learning of Boolean circuits as described in [Wilson and Zanasi 2020], almost everything is the same as the previous example except that the base category is changed to polynomials over \mathbb{Z}_2 , $\mathbf{POLY}_{\mathbb{Z}_2}$. To learn a function of type $\mathbb{Z}_2^n \rightarrow \mathbb{Z}_2^m$, one constructs some map of type $\mathbb{Z}^p \times \mathbb{Z}^n \rightarrow \mathbb{Z}^m$, which one can view as a map in the category $\mathbf{Para}(\mathbf{POLY}_{\mathbb{Z}_2})$ of type $\mathbb{Z}^n \rightarrow \mathbb{Z}^m$ with parameter space \mathbb{Z}^p , and again applying the functor in (8) for $\mathbf{POLY}_{\mathbb{Z}_2}$ yields a parameterised lens.


Note a model/parametric lens f can take as inputs an element of A , an element of B' (a change in B) and a parameter P and outputs an element of B , a change in A , and a change in P . This is not yet sufficient to do machine learning! When we perform learning, we want to input a parameter P and a pair $A \times B$ and receive a new parameter P . Instead, f expects a change in B (not an element of B) and outputs a change in P (not an element of P). Deep dreaming, on the other hand, wants to return an element of A (not a change in A). Thus, to do machine learning (or deep dreaming) we need to add additional components to f ; we will consider these additional components in the next sections.

3.2 Loss Maps as Parametric Lenses

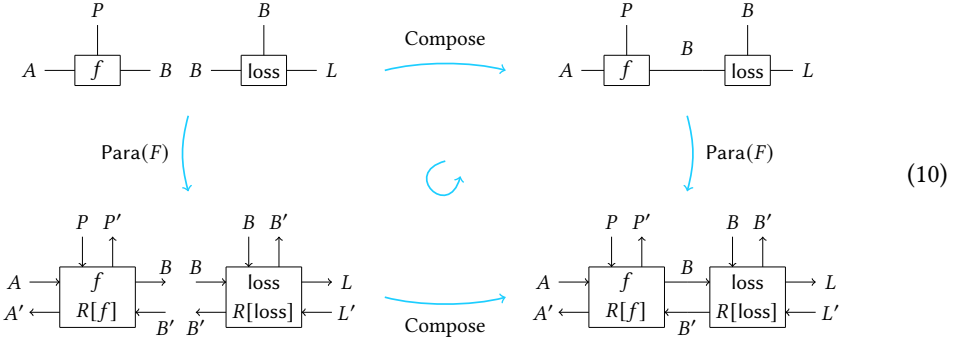
Another key component of any learning algorithm is the choice of loss map. This gives a measurement of how far the current output of the model is from the desired output. In standard learning in Smooth, this loss map is viewed as a map of type $B \times B \rightarrow \mathbb{R}$. However, in our setup, this is naturally viewed as a parameterised map from B to \mathbb{R} with parameter space B .⁶ We also generalize the codomain to an arbitrary object L .

Definition 3.3. A **loss map on B** consists of a $\text{Para}(C)$ map $(\text{loss}, B) : B \rightarrow L$ for some object L .

We can then compose such a map with a neural network $(N, P) : A \rightarrow B$ to get the composite



Note we can apply Para to either composite NL or N and L individually, giving a parametric lens



This is getting closer to the parametric lens we want: it can now receive inputs of type B . However, this is at the cost of now needing an input to L' ; we consider how to handle this in the next section.

Example 3.4 (Quadratic error). In Smooth, the standard loss function on \mathbb{R}^b is quadratic error: it uses $L = \mathbb{R}$ and has parameterised map $e : \mathbb{R}^b \times \mathbb{R}^b \rightarrow \mathbb{R}$ given by

$$e(b_t, b_p) = \frac{1}{2} \sum_{i=1}^b ((b_p)_i - (b_t)_i)^2.$$

where we think of b_t as the “true” value and b_p the predicted value. This has reverse derivative $R[e] : \mathbb{R}^b \times \mathbb{R}^b \times \mathbb{R} \rightarrow \mathbb{R}^b \times \mathbb{R}^b$ given by⁷ $R[e](b_t, b_p, \alpha) = \alpha \cdot (b_p - b_t, b_t - b_p)$.

Example 3.5 (Boolean error). In $\text{POLY}_{\mathbb{Z}_2}$, the loss function on \mathbb{Z}^b which is implicitly used in [Wilson and Zanasi 2020] is a bit different: it uses $L = \mathbb{Z}^b$ and has parameterised map $e : \mathbb{Z}^b \times \mathbb{Z}^b \rightarrow \mathbb{Z}^b$ given by

$$e(b_t, b_p) = b_t + b_p.$$

(Note that this is $+$ in \mathbb{Z}_2 ; equivalently this is given by XOR.) Its reverse derivative is of type $R[e] : \mathbb{Z}^b \times \mathbb{Z}^b \times \mathbb{Z}^b \rightarrow \mathbb{Z}^b \times \mathbb{Z}^b$ given by $R[e](b_t, b_p, \alpha) = (\alpha, \alpha)$.

⁶Here the loss map has its parameter space equal to its input space. However, putting loss maps on the same footing as models lends itself to further generalizations where the parameter space is different, and where the loss map can itself be learned. See Generative Adversarial Networks, [Capucci et al. 2021, Figure 7].

⁷The last argument in the reverse derivative implementation is suggestively named α to evoke the idea of *learning rate*, described in the next subsection.

Example 3.6 (Softmax cross entropy). The Softmax cross entropy loss is a \mathbb{R}^b -parameterized map $\mathbb{R}^b \rightarrow \mathbb{R}$ defined by

$$e(b_t, b_p) = \sum_{i:N} (b_t)_i ((b_p)_i - \log(\text{Softmax}(b_p)_i))$$

where $\text{Softmax}(b_p) = \frac{\exp((b_p)_i)}{\sum_{j:N} \exp((b_p)_j)}$ is defined componentwise for $i : N$.

We note that, although q needs to be a probability distribution, at the moment there is no need to ponder the question of interaction of probability distributions with the reverse derivative framework: one can simply consider q as the image of some logits under the Softmax function.

Example 3.7 (Dot product). In Deep Dreaming (Section 4.2) we often want to focus only on a particular element of the network output \mathbb{R}^b . This is done by supplying a one-hot vector b_t as the ground truth to the loss function which takes two vectors and computes their dot product:

$$e(b_t, b_p) = b_t \cdot b_p$$


The reverse derivative of the dot product has the type $\mathbb{R}^b \times \mathbb{R}^b \times \mathbb{R} \rightarrow \mathbb{R}^b \times \mathbb{R}^b$ and the following implementation

$$R[e](b_t, b_p, \alpha) = (\alpha \cdot b_p, \alpha \cdot b_t)$$

If the ground truth vector y is a one-hot vector (active at the i -th element), then the dot product essentially performs masking of all inputs except the i -th one.

3.3 Learning Rates as Parametric Lenses

With a loss function, we are getting closer to our goal of having a parameterised lens which represents a learning process. We have the following parametrised lens:



In this section we focus on the right side of the diagram. There is an output of L and an input of L' to the diagram. This is precisely the place where gradient-based learning algorithms input a *learning rate*.

Definition 3.8. A **learning rate** α on L consists of a lens from (L, L') to $(1, 1)$ where 1 is a terminal object in \mathcal{C} .

Note that the get component of such a lens must be the unique map to 1 , while the put component is a map $L \times 1 \rightarrow L'$; that is, simply a map $\alpha^* : L \rightarrow L'$. Moreover, we can view α as a **Para(Lens(C))** map from $(L, L') \rightarrow (1, 1)$ (with trivial parameter space). We write such a morphism as a cap, and compose it with the parameterised map above to get




Fig. 3. Model composed with a loss function and a learning rate

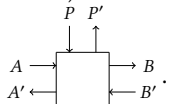
Example 3.9. In standard supervised learning in Smooth, one fixes some $\epsilon > 0$ as a learning rate, and this is used to define α : α is simply constantly $-\epsilon$, ie., $\alpha(l) = -\epsilon$ for any $l \in L$.

Example 3.10. In supervised learning in $\text{POLY}_{\mathbb{Z}_2}$, the standard learning rate is quite different: for a given L it is defined as the identity function, $\alpha(l) = l$.

Other learning rate morphisms are possible as well: for example, one could fix some $\epsilon > 0$ and define a learning rate in Smooth by $\alpha(l) = -\epsilon \cdot l$. Such a learning rate would take into account how far away the network is from its desired goal and adjust its learning rate accordingly.

3.4 Optimisers as Reparameterisations

In the previous sections we have seen how to incorporate the loss map and learning rate into our formalism. In this section we consider how to implement gradient descent (and its variants) into the

picture. Recall that we are writing our model graphically as the parameterised lens 

Note that this diagram outputs a P' , which represents a *change* in the parameter space. But we would like to receive not just the requested change in the parameter, but the new parameter itself. Thus, we need to add a box above the P/P' wires in the image above to get something whose input and output are both P . Recall that a box in the graphical language is a lens; thus, we are asking for a lens of type $(P, P) \rightarrow (P, P')$. This is precisely what gradient descent accomplishes.

Definition 3.11. In any CRDC C we can define gradient update as a map G in $\text{Lens}(C)$ from (P, P) to (P, P') consisting of

$$(G, G^*) : (P, P) \rightarrow (P, P')$$

where $G(p) = p$ and $G^*(p, p') = p + p'$.

Note that gradient descent is not typically seen as a lens - but it precisely fits this way into the picture we are creating!

Gradient descent allows one to receive the requested change in parameter and implement that change by adding that value to the current parameter. We attach this lens, seen as a reparameterisation, to the top of the diagram above, giving us Figure 4 (left).




Fig. 4. Model reparameterised by basic gradient descent (left) and a generic stateful optimiser (right).

Example 3.12 (Gradient update in Smooth). In Smooth, the gradient descent reparameterisation will take the output from P' and add it to the current value of P to get a new value of P .

Example 3.13 (Gradient update in Boolean circuits). In the CRDC $\text{POLY}_{\mathbb{Z}_2}$, the gradient descent reparameterisation will again take the output from P' and add it to the current value of P to get a new value of P ; however, since $+$ in \mathbb{Z}_2 is the same as XOR, this can be also be seen as taking the XOR of the current parameter and the requested change; this is exactly how this algorithm is implemented in [Wilson and Zanasi 2020].

Moreover, other variants of gradient descent also fit naturally into this framework by allowing for additional input/output data with P . In particular, many important variants of gradient descent keep track of the history of previous updates and use that to inform the next one. This is easy to model in our setup: instead of asking for a lens from (P, P) to (P, P') , we ask instead for a lens from $(S \times P, S \times P)$ to (P, P') where S is some other object which holds a “state”.

Definition 3.14. A **stateful parameter update** consists of a choice of object S (the **state** object) and a lens $U : (S \times P, S \times P) \rightarrow (P, P')$.

Again, we view this optimizer as a reparameterisation and attach it above the P/P' wires to the image from the previous section, giving us Figure 4 (right). Let us consider how several well-known optimizers can be implemented in this way.

Example 3.15 (Momentum). In the momentum variant of gradient descent, one keeps track of the previous change and uses this to inform how the current parameter should be changed. Thus, in this case, we set $S = P$, fix some $\gamma > 0$, and define the **momentum** lens

$$(U, U^*) : (P \times P, P \times P) \rightarrow (P, P')$$

by $U(s, p) = p$ and $U^*(s, p, p') = (s', p + s')$, where $s' = -\gamma s + p'$. Note momentum recovers gradient descent when $\gamma = 0$.

In both standard gradient descent and momentum, our lens representation has trivial get/forward part. Thus it is reasonable to wonder whether this formulation is really capturing the essence of what is going on. However, as soon as we move to more complicated variants, having non-trivial forward part of the lens is important, and Nesterov momentum is a key example of this.

Example 3.16 (Nesterov momentum). In Nesterov momentum, one makes a modification to the input of the derivative by the previous update. We can precisely capture this by using a small variation of the lens in the previous example. Again, we set $S = P$, fix some $\gamma > 0$, and define the **Nesterov momentum** lens

$$(U, U^*) : (P \times P, P \times P) \rightarrow (P, P')$$

by $U(s, p) = p + \gamma s$ and U^* as in the previous example.

Example 3.17 (Adagrad). Given any fixed $\epsilon > 0$ and $\delta \sim 10^{-7}$, Adagrad [Duchi et al. 2011] is given by $S = P$, with the lens whose get part is $(g, p) \mapsto p$. The put is $(g, p, p') \mapsto (g', p + \frac{\epsilon}{\delta + \sqrt{g}} \odot p')$ where $g' = g + p' \odot p'$ and \odot is the elementwise (Hadamard) product.

Unlike with other optimization algorithms where the learning rate is the same for all parameters, Adagrad divides the learning rate of each individual parameter with the square root of the past accumulated gradients.

Example 3.18 (Adam). Adaptive Moment Estimation (Adam) [Kingma and Ba 2015] is another method that computes adaptive learning rates for each parameter by storing exponentially decaying average of past gradients (m) and past squared gradients (v). Fixed $\beta_1, \beta_2 \in [0, 1)$, $\epsilon > 0$, and $\delta \sim 10^{-8}$, Adam is given by $S = P \times P$, with the lens whose get part is $(m, v, p) \mapsto p$ and whose put part is

$$\text{put}(m, v, p, p') = (\widehat{m}', \widehat{v}', p + \frac{\epsilon}{\delta + \sqrt{\widehat{v}'}} \odot \widehat{m}')$$

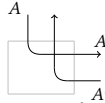
where $m' = \beta_1 m + (1 - \beta_1)p'$, $v' = \beta_2 v + (1 - \beta_2)p'^2$, and $\widehat{m}' = \frac{m'}{1 - \beta_1}$, $\widehat{v}' = \frac{v'}{1 - \beta_2}$.

Note that, so far, optimisers/reparameterisations have been added to the P/P' wires, in order to change the model's parameters. We will see in section 4.2 how we can also attach them to the A/A' wires instead, giving *deep dreaming*.

4 LEARNING WITH PARAMETRIC LENSES

In the previous section we have seen how all the components of learning can be modeled as parametric lenses. We now study how all these components can be put together to form supervised learning systems. In addition to studying the most common examples of supervised learning: systems that learn *parameters*, we also study different kinds systems: those that learn their *inputs*. This is a technique commonly known as deep dreaming, and we present it as a natural counterpart of supervised learning of parameters.

Before we describe these systems, it will be convenient to represent all the inputs and outputs of our parametric lenses as parameters. In Figure 3, we see the P/P' and B/B' inputs and outputs as parameters; however, the A/A' wires are not. To view the A/A' inputs as parameters, we compose that system with the parameterised lens η we now define. The parameterised lens η has the type $(1, 1) \rightarrow (A, A')$ with parameter space (A, A') defined by $(\text{get}_\eta = 1_A, \text{put}_\eta = \pi_1)$ and can be depicted

graphically as . Composing η with the rest of the learning system in Figure 3 gives us the closed parametric lens in Figure 5.




Fig. 5. Closed parametric lens whose all inputs and outputs are now vertical wires.

This composite is now a map in $\text{Para}(\text{Lens}(C))$ from $(1, 1)$ to $(1, 1)$; all its inputs and outputs are now vertical wires, i.e., parameters. Unpacking it further, this is a lens of type $(A \times P \times B, A' \times P' \times B') \rightarrow (1, 1)$ whose get map is the terminal map, and whose put map is of the type $A \times P \times B \rightarrow A' \times P' \times B'$. It can be unpacked as the composite

$$\begin{aligned} \text{put}(a, p, b_t) &= (a', p', b'_t) \quad \text{where} \quad b_p = f(p, a) \\ (b'_t, b'_p) &= R[\text{loss}](b_t, b_p, \alpha(\text{loss}(b_t, b_p))) \\ (p', a') &= R[f](p, a, b'_p) \end{aligned}$$

In the next two sections we consider further additions to the image above which correspond to different types of supervised learning: supervised learning of parameters and supervised learning of inputs.

4.1 Supervised Learning of Parameters

The most common type of learning that is performed on the image in Figure 5 is supervised learning of *parameters*. This is done by reparameterising the image (Def. 2.3) in the following manner. The parameter ports are reparameterised by one of the (potentially stateful) optimisers described in the previous section, while the backward wires A' of inputs and B' of outputs are discarded. This finally gives us the complete picture of a system which learns the parameters in a supervised manner (Figure 6).




Fig. 6. Closed parametric lens whose parameters are being learned. An animation of this supervised learning system is available in Appendix B.

Fixing a particular optimiser $(U, U^*) : (S \times P, S \times P) \rightarrow (P, P')$ we again unpack the entire construction. This is a map in $\text{Para}(\text{Lens}(C))$ from $(1, 1)$ to $(1, 1)$ whose parameter space is $(A \times S \times P \times B, S \times P)$. In other words, this is a lens of type $(A \times S \times P \times B, S \times P) \rightarrow (1, 1)$ whose get component is identity. Its put map has the type $A \times S \times P \times B \rightarrow S \times P$ and unpacks to

$$\begin{aligned} \text{put}(a, s, p, b_t) = U^*(s, p, p') \quad \text{where} \quad & \bar{p} = U(s, p) \\ & b_p = f(\bar{p}, a) \\ (b'_t, b'_p) = R[\text{loss}](b_t, b_p, \alpha(\text{loss}(b_t, b_p))) \\ (p', a') = R[f](\bar{p}, a, b'_p) \end{aligned}$$

While this formulation might seem daunting, we note that it just explicitly specifies the computation performed by a supervised learning system. The variable \bar{p} represents the parameter supplied to the network by the stateful gradient update rule (in many cases this is equal to p); b_p represents the prediction of the network (contrast this with b_t which represents the ground truth from the dataset), Variables with a tick ' represent changes: b'_p and b'_t are the changes on predictions and true values respectively, while p' and a' are changes on the parameters and inputs. Furthermore, this arises automatically out of the rule for lens composition (Figure 2.2); what we needed to specify is just the lenses themselves.

We justify and illustrate our approach on a series of case studies drawn from the machine learning literature, showing how in each case the parameters of our framework (in particular, loss functions and gradient updates) instantiate to familiar concepts. This presentation has the advantage of treating all these case studies uniformly in terms of our basic constructs, highlighting their similarities and differences.

We start in Smooth, fixing some parameterised map $(\mathbb{R}^P, f) : \text{Para}(\text{Smooth})(\mathbb{R}^a, \mathbb{R}^b)$ and the constant *negative* learning rate $\alpha : \mathbb{R}$ (Example 3.9). We then vary the loss function and the gradient update, seeing how the put map above reduces to many of the known cases in the literature.

Example 4.1 (Quadratic error, basic gradient descent). Fix the quadratic error (Example 3.4) as the loss map and basic gradient update (Example 3.12). Then the aforementioned put map simplifies. Since there is no state, its type reduces to $A \times P \times B \rightarrow P$, and its implementation to: $\text{put}(a, p, b_t) = p + p'$, where $(p', a') = R[f](p, a, \alpha \cdot (f(p, a) - b_t))$.

Note that α here is simply a constant, and due to the linearity of the reverse derivative (Def A.5), we can slide the α from the costate into the gradient descent lens. Rewriting this update, and

performing this sliding we obtain a closed form update step

$$\text{put}(a, p, b_t) = p + \alpha \cdot (R[f](p, a, f(p, a) - b_t); \pi_0)$$

where the negative *descent* component of gradient descent is here contained in the choice of the negative constant α .

This example gives us a variety of *regression* algorithms solved iteratively by gradient descent: it embeds some parameterised map $(\mathbb{R}^p, f) : \mathbf{Para}(\text{Smooth})(\mathbb{R}^a, \mathbb{R}^b)$ into the system which performs regression on input data - where a denotes the input to the model and b_t denotes the ground truth. If the corresponding map f is linear and $m = 1$, we recover simple linear regression with gradient descent. If the codomain is additionally multi-dimensional, i.e. we're predicting multiple scalars, then we recover multivariate linear regression. Likewise, we can model a multi-layer perceptron or even more complex neural network architectures performing supervised learning of parameters simply by changing the underlying parameterised map.

Example 4.2 (Softmax cross entropy, basic gradient descent). Fix Softmax cross entropy (Example 3.6) as the loss map and basic gradient update (Example 3.12). Again the put map simplifies. The type reduces to $A \times P \times B \rightarrow P$ and the implementation to

$$\text{put}(a, p, b_t) = p + p'$$

where $(p', a') = R[f](\bar{p}, a, \alpha \cdot (\text{Softmax}(f(p, a)) - b_t))$. The same rewriting performed on the previous example can be done here.

This example recovers *logistic regression*, e.g. classification.

Example 4.3 (Mean squared error, Nesterov Momentum). Fix the quadratic error (Example 3.4) as the loss map and Nesterov momentum (Example 3.16) as the gradient update. This time the put map doesn't have a simplified type, it is still $A \times S \times P \times B \rightarrow S \times P$. The implementation of put reduces to

$$\begin{aligned} \text{put}(a, s, p, b_t) = (s', p + s') \quad \text{where} \quad & \bar{p} = p + \gamma s \\ & (p', a') = R[f](\bar{p}, a, \alpha \cdot (f(\bar{p}, a) - b_t)) \\ & s' = -\gamma s + p' \end{aligned}$$

This example with Nesterov momentum differs in two key points from all the other ones: i) the optimiser is stateful, and ii) its get map is not trivial. While many other optimisers are stateful, the non-triviality of the get map here showcases the importance of lenses. They allow us to make precise the notion of computing a "lookahead" value for Nesterov momentum, something that is in practice usually handled in ad-hoc ways. Here, the algebra of lens composition handles this case naturally by using the get map, a seemingly trivial, unused piece of data for previous optimisers.

We finish off these examples by moving to a different base category $\text{POLY}_{\mathbb{Z}_2}$. This example shows that our framework describes learning in not just continuous, but discrete settings too. Again, we fix a parameterised map $(\mathbb{Z}^p, f) : \text{POLY}_{\mathbb{Z}_2}(\mathbb{R}^a, \mathbb{R}^b)$ but this time we fix the identity learning rate (Example 3.10), instead of a constant one.


Example 4.4 (Basic learning in Boolean circuits). Fix XOR as the loss map (Example 3.5) and the basic gradient update (Example 3.13). The put map again simplifies. The type reduces to $A \times P \times B \rightarrow P$ and the implementation to $\text{put}(a, p, b_t) = p + p'$ where $(p', a') = R[f](p, a, f(p, a) + b_t)$.

A sketch of learning iteration. Having described a number of examples in supervised learning, we outline how to model learning iteration in our framework. Recall the aforementioned put map

whose type is $A \times P \times B \rightarrow P$ (for simplicity here modelled without state S). This map takes an input-output pair (a_0, b_0) , the current parameter p_i and produces an updated parameter p_{i+1} . At the next time step, it takes a potentially different input-output pair (a_1, b_1) , the updated parameter p_{i+1} and produces p_{i+2} . This process is then repeated. We can model this iteration as a composition of the put map with itself, as a composite $(A \times \text{put} \times B)$; put whose type is $A \times A \times P \times B \times B \rightarrow P$. This map takes two input-output pairs $A \times B$, a parameter and produces a new parameter by processing these datapoints in sequence. One can see how this process can be iterated any number of times, and even represented as a string diagram.

But we note that with a slight reformulation of the put map, it is possible to obtain a conceptually much simpler definition. The key insight lies in seeing that the map $\text{put} : A \times P \times B \rightarrow P$ is essentially an endo-map $P \rightarrow P$ with some extra inputs $A \times B$; it's a parameterised map!

In other words, we can recast the put map as a parameterised map $(A \times B, \text{put}) : \text{Para}(C)(P, P)$. Since it is an endo-map, it can be composed with itself. The resulting composite is too an endo-map, which now takes two "parameters": input-output pair at the time step 0 and time step 1. This process can then be repeated, with **Para** composition automatically taking care of the algebra of iteration.



This reformulation captures the essence of parameter iteration: one can think of it as a trajectory $p_i, p_{i+1}, p_{i+2}, \dots$ through the parameter space; but it is a *trajectory parameterised by the dataset*. With different datasets the algorithm will take a different path through this space and learn different things.

4.2 Deep Dreaming: Supervised Learning of Inputs

We have seen that attaching gradient descent to the parameter port of the parametric lens as a reparameterisation allows us to learn the parameters in a supervised way. In this section we describe how attaching the gradient descent lens to the *input port* provides us with a way to enhance an input image to elicit a particular interpretation. This is the idea behind the technique called Deep Dreaming, appearing in the literature in many forms [Dosovitskiy and Brox 2015; Mahendran and Vedaldi 2014; Nguyen et al. 2014; Simonyan et al. 2014].




Fig. 7. Deep dreaming: supervised learning of inputs

Deep dreaming is a technique which uses the parameters p of some trained classifier network to iteratively dream up, or amplify some features of a class b on a chosen input a . For example, if

we start with an image of a landscape a_0 , a label b of a “cat” and a parameter p of a sufficiently well-trained classifier, we can start performing “learning” as usual: computing the predicted class for the landscape a_0 for the network with parameters p , and then computing the distance between the prediction and our label of a cat b . When performing backpropagation, the respective changes computed for each layer tell us how the activations of that layer should have been changed to be more “cat” like. This includes the first (input) layer of the landscape a_0 . Usually, we discard these changes and apply gradient update to the parameters. In deep dreaming we *discard the parameters* and *apply gradient update to the input* (Figure 7). Gradient update here takes these changes and computes a new image a_1 which is the same image of the landscape, but changed slightly so to look more like whatever the network thinks a cat looks like. This is the essence of deep dreaming, where iteration of this process allows networks to dream up features and shapes on a particular chosen image [Goo 2015].

Just like in the previous subsection, we can write this deep dreaming system as a map in $\text{Para}(\text{Lens}(C))$ from $(1, 1)$ to $(1, 1)$ whose parameter space is $(S \times A \times P \times B, S \times A)$. In other words, this is a lens of type $(S \times A \times P \times B, S \times A) \rightarrow (1, 1)$ whose get map is trivial. Its put map has the type $S \times A \times P \times B \rightarrow S \times A$ and unpacks to

$$\begin{aligned} \text{put}(s, a, p, b_t) = U^*(s, a, a') \quad \text{where} \quad & \bar{a} = U(s, a) \\ & b_p = f(p, \bar{a}) \\ & (b'_t, b'_p) = R[\text{loss}](b_t, b_p, \alpha(\text{loss}(b_t, b_p))) \\ & (p', a') = R[f](p, \bar{a}, b'_p) \end{aligned}$$

We note that deep dreaming is usually presented without any loss function as a maximisation of a particular activation in the last layer of the network output [Simonyan et al. 2014, Section 2.]. This maximisation is done with gradient ascent, as opposed to gradient descent. However, this is just a special case of our framework where the loss function is the dot product (Example 3.7). The choice of the particular activation is encoded as a one-hot vector, and the loss function in that case essentially masks the network output, leaving active only the particular chosen activation. The final component is the gradient *ascent*: this is simply recovered by choosing a positive, instead of a negative learning rate [Simonyan et al. 2014]. We explicitly unpack this in the following example.

Example 4.5 (Deep dreaming, dot product loss, gradient descent). Fix the base category to Smooth and a parameterised map $(\mathbb{R}^p, f) : \text{Para}(\text{Smooth})(\mathbb{R}^a, \mathbb{R}^b)$. Fix the dot product loss (Example 3.7), basic gradient descent (Example 3.12), and a *positive* learning rate $\alpha : \mathbb{R}$. Then the above put map simplifies. Since there is no state, its type reduces to $A \times P \times B \rightarrow A$ and its implementation to

$$\text{put}(a, p, b_t) = a + a' \quad \text{where} \quad (p', a') = R[f](p, a, \alpha \cdot b_t).$$

Like in Example 4.1, this update can be rewritten as

$$\text{put}(a, p, b_t) = a + \alpha \cdot (R[f](p, a, b_t); \pi_1)$$

making a few things apparent. This update does not depend on the prediction $f(p, a)$: no matter what the network has predicted, the goal is always to maximize particular activations. Which activations? The ones chosen by b_t . When b_t is a one-hot vector, this picks out the activation of just one class to maximize, which is often done in practice.

While we present only the most basic image, there is plenty of room left for exploration. The work of [Simonyan et al. 2014, Section 2.] adds an extra regularization term to the image. In general, the neural network f is sometimes changed to copy a number of internal activations which are then exposed on the output layer. Maximizing all these activations often produces more visually appealing results. In the literature we did not find an example which uses the Softmax-cross entropy

(Example 3.6) as a loss function in deep dreaming, which seems like the more natural choice in this setting. Furthermore, while deep dreaming is commonly done with basic gradient descent, there is nothing preventing one from doing deep dreaming with any of the optimizer lenses discussed in the previous section, or even doing deep dreaming in the context of Boolean circuits. Lastly, learning iteration which was described in at the end of previous subsection can be modelled here in an analogous way.

5 IMPLEMENTATION

We provide a proof-of-concept implementation as a Python library.⁸ We demonstrate the correctness of our library empirically using a number of experiments implemented both in our library and in Keras[Chollet et al. 2015], a popular framework for deep learning. For example, one experiment is a model for the MNIST image classification problem[Lecun et al. 1998]: we implement the same model in both frameworks and achieve comparable accuracy.

Our implementation also demonstrates the advantages of our approach. Firstly, computing the gradients of the network is greatly simplified through the use of lens composition. Secondly, model architectures can be expressed in a principled, mathematical language; as morphisms of a monoidal category. Finally, the modularity of our approach makes it easy to see how various aspects of training can be modified: for example, one can define a new optimization algorithm simply by defining an appropriate lens. We now give a brief sketch of our implementation.


5.1 Constructing a Model with Lens and Para

We model a lens (f, f^*) in our library with the Lens class, which consists of a pair of maps fwd and rev corresponding to f and f^* , respectively. For example, we write the identity lens $(1_A, \pi_2)$ as follows:

```
identity = Lens(lambda x: x, lambda x_dy: x_dy[1])
```

The composition (in diagrammatic order) of Lens values f and g is written $f \gg g$, and monoidal composition as $f @ g$. Similarly, the type of Para maps is modeled by the Para class, with composition and monoidal product written the same way. Our library provides several primitive Lens and Para values.

Let us now see how to construct a single layer neural network from the composition of such primitives. Diagrammatically, we wish to construct the following model, representing a single ‘dense’ layer of a neural network:



Here, the parameters of `linear` are the coefficients of a $b \times a$ matrix, and the underlying lens has as its forward map the function $(M, x) \rightarrow M \cdot x$, where M is the $b \times a$ matrix whose coefficients are the $\mathbb{R}^{b \times a}$ parameters, and $x \in \mathbb{R}^a$ is the input vector. The `bias` map is even simpler: the forward map of the underlying lens is simply pointwise addition of inputs and parameters: $(b, x) \rightarrow b + x$. Finally, the `activation` map simply applies a nonlinear function (e.g., sigmoid) to the input, and thus has the trivial (unit) parameter space. The representation of this composition in code is straightforward: we can simply compose the three primitive Para maps as in (12):

⁸Full usage examples, source code, and experiments using our proof-of-concept can be found at <https://github.com/statusfailed/numeric-optics-python/>.


```
def dense(a, b, activation):
    return linear(a, b) >> bias(b) >> activation
```


Note that by constructing model architectures in this way, the computation of reverse derivatives is greatly simplified: we obtain the reverse derivative ‘for free’ as the put map of the model. Furthermore, adding new primitives is also simplified: the user need simply provide a function and its reverse derivative in the form of a Para map. Finally, notice also that our approach is truly compositional: we can define a hidden layer neural network with n hidden units simply by composing two dense layers, as follows:

```
dense(a, n, activation) >> dense(n, b, activation)
```

5.2 Learning

Now that we have constructed a model, we also need to use it to *learn* from data. Concretely, we will construct a full parametric lens as in Figure 6 then extract its put map to iterate over the dataset.

By way of example, let us see how to construct the following parametric lens, representing basic gradient descent over a single layer neural network with a fixed learning rate:



This morphism is constructed essentially as below, where $\text{apply_update}(\alpha, f)$ represents the ‘vertical stacking’ of α atop f :

```
apply_update(basic_update, dense) >> loss >> learning_rate( $\epsilon$ )
```

Now, given the parametric lens of (13), one can construct a morphism $\text{step} : B \times P \times A \rightarrow P$ which is simply the put map of the lens. Training the model then consists of iterating the step function over dataset examples $(x, y) \in A \times B$ to optimise some initial choice of parameters $\theta_0 \in P$, by letting $\theta_{i+1} = \text{step}(y_i, \theta_i, x_i)$.

Note that our library also provides a utility function to construct step from its various pieces:

```
step = supervised_step(model, update, loss, learning_rate)
```

For an end-to-end example of model training and iteration, we refer the interested reader to the experiments accompanying the code: <https://github.com/statusfailed/numeric-optics-python/>.

6 RELATED WORK

The work [Fong et al. 2017] is closely related to ours, in that it provides an abstract categorical model of back-propagation. However, it differs in a number of key aspects. We give a complete lens-theoretic explanation of *what* is back-propagated via (i) the use of CRDCs to model gradients; and (ii) the Para construction to model parameterized functions and parameter update. We thus can go well beyond [Fong et al. 2017] in terms of examples - their example of smooth functions and basic gradient descent is covered in our subsection 4.1.

We also explain some of the constructions of [Fong et al. 2017] in a more structured way. For example, rather than considering the category **Learn** of [Fong et al. 2017] as primitive, here we construct it as a composite of two more basic constructions (the **Para** and **Lens** constructions). The flexibility could be used, for example, to compositionally replace **Para** with a variant allowing parameters to come from a different category, or lenses with the category of optics [Riley 2018] enabling us to model things such as control flow using prisms.

One more important thing is related to functoriality. We use a functor to augment a parameterised map with its backward pass, just like [Fong et al. 2017]. However, they additionally augmented this map with a loss map and gradient descent using a functor as well. This added extra conditions on the partial derivatives of the loss function: it needed to be invertible in the 2nd variable. This constraint was not justified in [Fong et al. 2017], nor is it a constraint that appears in machine learning practice. This led us to reexamine their constructions, coming up with our reformulation that does not require it. While loss maps and optimisers are mentioned in [Fong et al. 2017] as parts of the aforementioned functor, here they are extracted out and play a key role: loss maps are parameterised lenses and optimisers are reparameterisations. Thus, in this paper we instead use **Para**-composition to add the loss map to the model, and **Para** 2-cells to add optimisers. The mentioned inverse of the partial derivative of the loss map in the 2nd variable was also hypothesised to be relevant to deep dreaming. In our paper we have given a complete picture of deep dreaming systems, showing it is gradient update which is used to dream up pictures.

We also correct a small issue in Theorem III.2 of [Fong et al. 2017]. There, the morphisms of **Learn** were defined up to an equivalence (pg. 4 of [Fong et al. 2017]) but, unfortunately, the functor defined in Theorem III.2 does not respect this equivalence relation. Our approach instead uses 2-cells which comes from the universal property of **Para** — a 2-cell from $(P, f) : A \rightarrow B$ to $(Q, g) : A \rightarrow B$ is a lens, and hence has two components: a map $\alpha : Q \rightarrow P$ and $\alpha^* : Q \times P \rightarrow Q$. By comparison, we can see the equivalence relation of [Fong et al. 2017] as being induced by map $\alpha : Q \rightarrow P$, and not a lens. Our approach highlights the importance of the 2-categorical structure of learners. In addition, it does not treat the functor $\mathbf{Para}(C) \rightarrow \mathbf{Learn}$ as a primitive. In our case, this functor has the type $\mathbf{Para}(C) \rightarrow \mathbf{Para}(\mathbf{Lens}(C))$ and arises from applying **Para** to a canonical functor $C \rightarrow \mathbf{Lens}(C)$ existing for *any* reverse derivative category, not just **Smooth**. Lastly, in our paper we have taken the advantage of the graphical calculus for **Para**, redrawing many diagrams appearing in [Fong et al. 2017] in a structured way.

Other than [Fong et al. 2017], there are a few more relevant papers. The work of ([Dalrymple 2019]) contains a sketch of some of the ideas this paper evolved from. They are based on the interplay of optics with parameterisation, albeit framed in the setting diffeological spaces, and requiring cartesian and local cartesian closed structure on the base category. Lenses and Learners are studied in the eponymous work of [Fong and Johnson 2019] which observes that lenses are parameterised learners. They do not explore any of the relevant **Para** or CRDC structure, but make the distinction between *symmetric* and *asymmetric lenses*, studying how they are related to learners defined in [Fong et al. 2017]. A lens-like implementation of automatic differentiation is the focus of [Elliott 2018], but learning algorithms aren't studied. A relationship between category-theoretic perspective on probabilistic modeling and gradient-based optimisation is studied in [Shiebler 2020] which also studies a variant of the **Para** construction. Usage of Cartesian differential categories to study learning is found in [Sprunger and Katsumata 2019]. They extend the differential operator to work on stateful maps, but do not study lenses, parameterisation nor update maps. The work of [Gavranovic 2019] studies deep learning in the context of Cycle-consistent Generative Adversarial Networks [Zhu et al. 2017] and formalises it via free and quotient categories, making parallels to the categorical formulations of database theory [Spivak 2010]. They do use the **Para** construction, but do not relate it to lenses nor reverse derivative categories.

Lastly, the concept of parameterised lenses has started appearing in recent formulations of categorical game theory and cybernetics [Capucci et al. 2021; Capucci et al. 2021]. The work of [Capucci et al. 2021] generalises the study of parameterised lenses into parameterised optics and connects it to game theoretic concepts such as Nash equilibria. A general survey of category theoretic approaches to machine learning, covering many of the above papers in detail, can be found in [Shiebler et al. 2021].

7 CONCLUSIONS AND FUTURE DIRECTIONS

We have given a categorical foundation of gradient-based learning algorithms which achieves a number of important goals. The foundation is principled and mathematically clean, based on the fundamental idea of a *parameterised lens*. The foundation covers a wide variety of examples: it covers different optimisers and loss maps in gradient-based learning, it covers different settings where gradient-based learning happens (smooth functions vs. boolean circuits) and it covers both learning of parameters and learning of inputs (deep dreaming). Finally, the foundation is more than a mere abstraction: we have also shown how it can be used to give a practical implementation of learning, as discussed in section 5.

There are a number of important directions which are possible to explore because of this work. One of the most exciting ones is the extension to more complex neural network architectures. Our formulation of the loss map as a parameterised lens should pave the way for Generative Adversarial Networks [Goodfellow et al. 2014], an exciting new architecture whose loss map can be said to be *learned* in tandem with the base network. In all our settings we have fixed an optimiser beforehand. The work of [Andrychowicz et al. 2016] describes a *meta-learning* approach which sees the optimiser as a neural network whose parameters and gradient update rule can be learned. This is an exciting prospect since one can model optimisers as parameterised lenses; and our framework covers learning with parameterised lenses. Recurrent neural networks are another example of a more complex architecture, which has already been studied in the context of differential categories in [Sprunger and Katsumata 2019]. When it comes to architectures, future work includes modelling some classical systems as well, such as the Support Vector Machines [Cortes and Vapnik 1995], which should be possible with the usage of loss maps such as Hinge loss.

We have not made use of the full power of the CRDC axioms; in particular, we did not explicitly need axioms RD.6 or RD.7, which deal with the behaviour of higher-order derivatives. However, some supervised learning algorithms do use the higher-order derivatives (for example, the Hessian) for additional optimisations; as such, future work includes exploring how to use those axioms to capture these optimisations. Taking this idea in a different direction, one can see that much of our work can be applied to any functor of the form $F : C \rightarrow \mathbf{Lens}(C)$ - F does not necessarily have to be of the form $f \mapsto (f, R[f])$ for a CRDC R . Moreover, by working with more generalised forms of the lens category (such as dependent lenses), we may be able to capture ideas related to supervised learning on manifolds. And, of course, we can vary the parameter space to endow it with different structure from the functions we wish to learn. In this vein, we wish to use fibrations/dependent types to model the use of tangent bundles: this would foster the extension of the *correct by construction* paradigm to machine learning, and thereby addressing the widely acknowledged problem of trusted machine learning. The possibilities are made much easier by the compositional nature of our framework. Another key topic for future work is to link gradient-based learning with game theory. At a high level, the former takes little incremental steps to achieve an equilibrium while the later aims to do so in one fell swoop. Formalising this intuition is possible with our lens-based framework and the lens-based framework for game theory [Ghani et al. 2016]. Finally, because our framework is quite general, in future work we plan to consider further modifications and additions to encompass

non-supervised, probabilistic and non-gradient based learning. This includes genetic algorithms and reinforcement learning.

REFERENCES

2015. Inceptionism: Going Deeper into Neural Networks. (2015). <https://ai.googleblog.com/2015/06/inceptionism-going-deeper-into-neural.html>
2019. Explainable AI: the Basics - Policy Briefing. [royalsociety.org/ai-interpretability](https://royalsocietypublishing.org/journal/rsos/190101)
- S. Abramsky and B. Coecke. 2004. A categorical semantics of quantum protocols. In *Proceedings of the 19th Annual IEEE Symposium on Logic in Computer Science, 2004*. 415–425. <https://doi.org/10.1109/LICS.2004.1319636>
- Marcin Andrychowicz, Misha Denil, Sergio Gomez, Matthew W. Hoffman, David Pfau, Tom Schaul, Brendan Shillingford, and Nando de Freitas. 2016. Learning to learn by gradient descent by gradient descent. *arXiv e-prints*, Article arXiv:1606.04474 (June 2016), arXiv:1606.04474 pages. arXiv:1606.04474 [cs.NE]
- John C. Baez and Jason Erbele. 2014. Categories in Control. *arXiv e-prints*, Article arXiv:1405.6881 (May 2014), arXiv:1405.6881 pages. arXiv:1405.6881 [math.CT]
- R. Blute, R. Cockett, and R. Seely. 2009. Cartesian Differential Categories. *Theory and Applications of Categories* 22 (2009), 622–672.
- Aaron Bohannon, J. Nathan Foster, Benjamin C. Pierce, Alexandre Pilkiewicz, and Alan Schmitt. 2008. Boomerang: Resourceful Lenses for String Data. *SIGPLAN Not.* 43, 1 (Jan. 2008), 407–419. <https://doi.org/10.1145/1328897.1328487>
- Guillaume Boisseau. 2020. String Diagrams for Optics. *arXiv e-prints*, Article arXiv:2002.11480 (Feb. 2020), arXiv:2002.11480 pages. arXiv:2002.11480 [math.CT]
- Filippo Bonchi, Pawel Sobocinski, and Fabio Zanasi. 2017. The Calculus of Signal Flow Diagrams I: Linear relations on streams. *Inf. Comput.* 252 (2017), 2–29. <https://doi.org/10.1016/j.ic.2016.03.002>
- Matteo Capucci, Bruno Gavranović, Jules Hedges, and E. F. Rischel. 2021. Towards foundations of categorical cybernetics.
- Matteo Capucci, Neil Ghani, Jérémy Ledent, and Fredrik Nordvall Forsberg. 2021. Translating Extensive Form Games to Open Games with Agency. *arXiv e-prints*, Article arXiv:2105.06763 (May 2021), arXiv:2105.06763 pages. arXiv:2105.06763 [cs.GT]
- Francois Chollet et al. 2015. *Keras*. <https://github.com/fchollet/keras>
- Bryce Clarke, Derek Elkins, Jeremy Gibbons, Fosco Loregian, Bartosz Milewski, Emily Pillmore, and Mario Román. 2020. Profunctor optics, a categorical update. *arXiv e-prints*, Article arXiv:2001.07488 (Jan. 2020), arXiv:2001.07488 pages. arXiv:2001.07488 [cs.PL]
- J. Robin B. Cockett, Geoff S. H. Cruttwell, Jonathan Gallagher, Jean-Simon Pacaud Lemay, Benjamin MacAdam, Gordon D. Plotkin, and Dorette Pronk. 2019. Reverse derivative categories. In *CSL*.
- Bob Coecke and Aleks Kissinger. 2017. *Picturing Quantum Processes: A First Course in Quantum Theory and Diagrammatic Reasoning*. Cambridge University Press. <https://doi.org/10.1017/9781316219317>
- Corinna Cortes and Vladimir Vapnik. 1995. Support-vector networks. *Machine learning* 20, 3 (1995), 273–297.
- Matthieu Courbariaux, Yoshua Bengio, and Jean-Pierre David. [n.d.]. BinaryConnect: Training Deep Neural Networks with binary weights during propagations. *arXiv:1511.00363 [cs]* ([n. d.]).
- G. S. H. Cruttwell. 2012. Cartesian differential categories revisited. *arXiv e-prints*, Article arXiv:1208.4070 (Aug. 2012), arXiv:1208.4070 pages. arXiv:1208.4070 [math.CT]
- David Dalrymple. 2019. Dioptics: a Common Generalization of Open Games and Gradient-Based Learners. *SYCO7* (2019). <https://research.protocol.ai/publications/dioptics-a-common-generalization-of-open-games-and-gradient-based-learners/dalrymple2019.pdf>
- Alexey Dosovitskiy and Thomas Brox. 2015. Inverting Convolutional Networks with Convolutional Networks. *CoRR abs/1506.02753* (2015). arXiv:1506.02753 <http://arxiv.org/abs/1506.02753>
- John Duchi, Elad Hazan, and Yoram Singer. 2011. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research* 12, Jul (2011), 2121–2159.
- Conal Elliott. 2018. The simple essence of automatic differentiation (Differentiable functional programming made easy). *CoRR abs/1804.00746* (2018). arXiv:1804.00746 <http://arxiv.org/abs/1804.00746>
- Brendan Fong and Michael Johnson. 2019. Lenses and Learners. *CoRR abs/1903.03671* (2019). arXiv:1903.03671 <http://arxiv.org/abs/1903.03671>
- Brendan Fong, David I. Spivak, and Rémy Tuyéras. 2017. Backprop as Functor: A compositional perspective on supervised learning. *CoRR abs/1711.10455* (2017). arXiv:1711.10455 <http://arxiv.org/abs/1711.10455>
- Bruno Gavranovic. 2019. Compositional Deep Learning. *ArXiv abs/1907.08292* (2019).
- Neil Ghani, Jules Hedges, Viktor Winschel, and Philipp Zahn. 2016. Compositional game theory. *arXiv e-prints*, Article arXiv:1603.04641 (March 2016), arXiv:1603.04641 pages. arXiv:1603.04641 [cs.GT]
- Dan R. Ghica, Achim Jung, and Aliaume Lopez. 2017. Diagrammatic Semantics for Digital Circuits. *arXiv e-prints*, Article arXiv:1703.10247 (March 2017), arXiv:1703.10247 pages. arXiv:1703.10247 [cs.PL]

- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative Adversarial Nets. In *Advances in Neural Information Processing Systems 27*, Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger (Eds.). Curran Associates, Inc., 2672–2680. <http://papers.nips.cc/paper/5423-generative-adversarial-nets.pdf>
- Andreas Griewank and Andrea Walther. 2008. *Evaluating derivatives: principles and techniques of algorithmic differentiation*. Society for Industrial and Applied Mathematics.
- Jules Hedges. 2018. Limits of bimorphic lenses. *arXiv e-prints*, Article arXiv:1808.05545 (Aug. 2018), arXiv:1808.05545 pages. arXiv:1808.05545 [math.CT]
- Michael Johnson, Robert Rosebrugh, and R.J. Wood. 2012. Lenses, fibrations and universal translations. *Mathematical structures in computer science* 22 (2012), 25–42.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, Yoshua Bengio and Yann LeCun (Eds.). <http://arxiv.org/abs/1412.6980>
- Yann Lecun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. 1998. Gradient-Based Learning Applied to Document Recognition. In *Proceedings of the IEEE*. 2278–2324. <https://doi.org/10.1109/5.726791>
- Aravindh Mahendran and Andrea Vedaldi. 2014. Understanding Deep Image Representations by Inverting Them. *CoRR abs/1412.0035* (2014). arXiv:1412.0035 <http://arxiv.org/abs/1412.0035>
- Anh Mai Nguyen, Jason Yosinski, and Jeff Clune. 2014. Deep Neural Networks are Easily Fooled: High Confidence Predictions for Unrecognizable Images. *CoRR abs/1412.1897* (2014). arXiv:1412.1897 <http://arxiv.org/abs/1412.1897>
- Christopher Olah. 2015. Neural Networks, Types, and Functional Programming. (2015). <http://colah.github.io/posts/2015-09-NN-Types-FP/>
- B.T. Polyak. 1964. Some methods of speeding up the convergence of iteration methods. *U. S. S. R. Comput. Math. and Math. Phys.* 4, 5 (1964), 1 – 17. [https://doi.org/10.1016/0041-5553\(64\)90137-5](https://doi.org/10.1016/0041-5553(64)90137-5)
- Mitchell Riley. 2018. Categories of Optics. arXiv:1809.00738 [math.CT]
- Peter Selinger. 2001. Control categories and duality: on the categorical semantics of the lambda-mu calculus. *Mathematical Structures in Computer Science* 11, 02 (4 2001), 207–260. <https://doi.org/null>
- P. Selinger. 2010. A Survey of Graphical Languages for Monoidal Categories. *Lecture Notes in Physics* (2010), 289–355. https://doi.org/10.1007/978-3-642-12821-9_4
- Sanjit A. Seshia and Dorsa Sadigh. 2016. Towards Verified Artificial Intelligence. *CoRR abs/1606.08514* (2016). arXiv:1606.08514 <http://arxiv.org/abs/1606.08514>
- Dan Shiebler. 2020. Categorical Stochastic Processes and Likelihood. *arXiv e-prints*, Article arXiv:2005.04735 (May 2020), arXiv:2005.04735 pages. arXiv:2005.04735 [cs.AI]
- Dan Shiebler, Bruno Gavranović, and Paul Wilson. 2021. Category Theory in Machine Learning. *arXiv e-prints*, Article arXiv:2106.07032 (June 2021), arXiv:2106.07032 pages. arXiv:2106.07032 [cs.LG]
- Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. 2014. Deep Inside Convolutional Networks: Visualising Image Classification Models and Saliency Maps. arXiv:1312.6034 [cs.CV]
- David I. Spivak. 2010. Functorial Data Migration. *CoRR abs/1009.1166* (2010). arXiv:1009.1166 <http://arxiv.org/abs/1009.1166>
- David Sprunger and Shin-ya Katsumata. 2019. Differentiable Causal Computations via Delayed Trace. *CoRR abs/1903.01093* (2019). arXiv:1903.01093 <http://arxiv.org/abs/1903.01093>
- Albert Steckermeier. 2015. Lenses in Functional Programming. *Preprint* (2015).
- Ilya Sutskever, James Martens, George Dahl, and Geoffrey Hinton. 2013. On the importance of initialization and momentum in deep learning. In *Proceedings of the 30th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 28)*, Sanjoy Dasgupta and David McAllester (Eds.). PMLR, Atlanta, Georgia, USA, 1139–1147. <http://proceedings.mlr.press/v28/sutskever13.html>
- D. Turi and G. Plotkin. 1997. Towards a mathematical operational semantics. In *Proceedings of Twelfth Annual IEEE Symposium on Logic in Computer Science*. 280–291. <https://doi.org/10.1109/LICS.1997.614955>
- Paul Wilson and Fabio Zanasi. 2020. Reverse Derivative Ascent: A Categorical Approach to Learning Boolean Circuits. *EPTCS* (2020). <https://cgi.cse.unsw.edu.au/~eptcs/paper.cgi?ACT2020:31>
- Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. 2017. Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networks. *arXiv e-prints*, Article arXiv:1703.10593 (March 2017), arXiv:1703.10593 pages. arXiv:1703.10593 [cs.CV]

A BACKGROUND ON CARTESIAN REVERSE DIFFERENTIAL CATEGORIES

Here we briefly review the definitions of Cartesian left additive category (CLAC), Cartesian reverse differential category (CRDC) and additive and linear maps in these categories. Note that in this appendix we follow the convention of [Cockett et al. 2019] and write composition in diagrammatic order by juxtaposition of terms (rather than a semicolon) to shorten the form of many of the expressions.

Definition A.1. A category C is said to be **Cartesian** when there are chosen binary products \times , with projection maps π_i and pairing operation $\langle -, - \rangle$, and a chosen terminal object T , with unique maps $!$ to the terminal object.

Definition A.2. A **left additive category** [Blute et al. 2009, Definition 1.1.1] (CLAC) is a category C such that each hom-set has commutative monoid structure, with addition operation $+$ and zero maps 0 , such that composition on the left preserves the additive structure: for any appropriate f, g, h , $f(g + h) = fg + fh$ and $f0 = 0$.

Definition A.3. A map $h : X \rightarrow Y$ in a CLAC is **additive** if it has the property that it preserves additive structure by composition on the right: for any maps $x, y : Z \rightarrow X$, $(x + y); h = x; h + y; h$, and $0; h = 0$.

Definition A.4. A **Cartesian left additive category** [Blute et al. 2009, Definition 1.2.1] is a left additive category C which is Cartesian and such that all projection maps π_i are additive.

The central definition of [Cockett et al. 2019] is the following:

Definition A.5. A **Cartesian reverse differential category** (CRDC) is a Cartesian left additive category C which has, for each map $f : A \rightarrow B$ in C , a map

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

satisfying seven axioms:

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

[RD.2] $\langle a, b + c \rangle R[f] = \langle a, b \rangle R[f] + \langle a, c \rangle R[f]$ and $\langle a, 0 \rangle R[f] = 0$.

[RD.3] $R[1] = \pi_1$, $R[\pi_0] = \pi_1 \iota_0$, and $R[\pi_1] = \pi_1 \iota_1$. **[RD.4]** For a tupling of maps f and g , the following equality holds:

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

And if $!_A : A \rightarrow T$ is the unique map to the terminal object, $R[!_A] = 0$.

[RD.5] For composable maps f and g ,

$$R[fg] = \langle \pi_0, \pi_0 f, \pi_1 \rangle (1 \times R[g]) R[f]$$

[RD.6]

$$\langle 1 \times \pi_0, 0 \times \pi_1 \rangle (\iota_0 \times 1) R[R[f]] \pi_1 = (1 \times \pi_1) R[f].$$

[RD.7]

$$(\iota_0 \times 1); R[R[(\iota_0 \times 1)R[R[f]]\pi_1]]; \pi_1 = \text{ex}; (\iota_0 \times 1)R[R[(\iota_0 \times 1)R[R[f]]\pi_1]]\pi_1$$

(where ex is the map that exchanges the middle two variables).

As discussed in [Cockett et al. 2019], these axioms correspond to familiar properties of the reverse derivative:

- **[RD.1]** says that differentiation preserves addition of maps, while **[RD.2]** says that differentiation is additive in its vector variable.

- [RD.3] and [RD.4] handle the derivatives of identities, projections, and tuples.
- [RD.5] is the (reverse) chain rule.
- [RD.6] says that the reverse derivative is linear in its vector variable.
- [RD.7] expresses the independence of order of mixed partial derivatives.

B A VISUALISATION OF SUPERVISED PARAMETER LEARNING

We animated the learning system described in subsection subsection 4.1. Visualising the following animations require a fully-featured PDF reader.⁹ Alternatively, it is available at the following [hyperlink](#).

Fig. 8. Animation of supervised learning with basic gradient update

⁹<https://acrobat.adobe.com/uk/en/acrobat/pdf-reader.html>