








Categorical Foundations of Gradient-Based Learning

Geoffrey S. H. Cruttwell¹ (✉) , Bruno Gavranović² (✉) , Neil Ghani² (✉) ,
Paul Wilson⁴ (✉) , and Fabio Zanasi⁴ (✉) 

¹ Mount Allison University, Canada

² University of Strathclyde, United Kingdom

³ University College London

Abstract. We propose a categorical semantics of gradient-based machine learning algorithms in terms of lenses, parametric 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.

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 [41], 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 [2, 36].

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 [20], Momentum [37], and Adam [32], etc.).




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.

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 [50] 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 **parametric**, e.g. in the simplest 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. Specifically, 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 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 [6, 12, 29] 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* [39], in a way similar to how other computational phenomena have been recently analysed e.g. in quantum theory [14], control theory [5, 8], and digital circuit theory [26].

It is intuitively clear how parametric lenses 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.

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 [3, 40, 49]. Three of them are particularly




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.

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. [8, 14, 25, 45]. 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 parametric categories, which amounts to a functor $\mathbf{Para}(-)$ turning a category \mathcal{C} into one $\mathbf{Para}(\mathcal{C})$ of ‘parametric \mathcal{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 [47], as well as in the foundations of database theory [31] and more recently game theory [25]. Considering lenses in \mathcal{C} simply amounts to the application of a functorial construction $\mathbf{Lens}(-)$, yielding $\mathbf{Lens}(\mathcal{C})$. Finally, we recall the notion of a *cartesian reverse differential category* (CRDC): a categorical structure axiomatising the notion of differentiation [13] (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 $\mathbf{Para}(\mathbf{Lens}(\mathcal{C}))$ for a CRDC \mathcal{C} . In terms of our desiderata (I)-(III) above, note that $\mathbf{Para}(-)$ accounts for (I), $\mathbf{Lens}(-)$ 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 3, we show how the 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 [16] and it has been explored recently using a categorical approach [50], 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. 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 [19, 34, 35, 44].
 - 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 ‘reparameterisations’ of models as parametric lenses. As case studies, in addition to basic gradient update, we consider the stateful variants: Momentum [37], Nesterov Momentum [48], Adagrad [20], and Adam (Adaptive Moment Estimation) [32]. Also, on Boolean circuits, we show how the reverse derivative ascent of [50] 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 of the composite of lenses expressed in Figure 2, for different choices of the various components. In particular we look at (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**, also called **deep dreaming** — the idea behind this technique is that, instead of the network parameters, one updates the inputs, in order to elicit a particular interpretation [19, 34, 35, 44]. Deep dreaming can be easily expressed within our approach, with a different rearrangement of the parametric lenses involved in the learning process, see (8) 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**, available at [17], 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 show 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 [11]. In particular, we create a working non-trivial neural network model for the MNIST image-classification problem [33].
- 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: (I) the **Para** construction (Section 2.1), which builds a category of parametric maps, (II) the **Lens** construction, which builds a category of “bidirectional” maps (Section 2.2), and (III) the combination of these two constructions into the notion of “parametric lenses” (Section 2.3). Finally (IV) we recall Cartesian reverse differential categories — categories equipped with an abstract gradient operator.

Notation We shall use $f;g$ for sequential composition of morphisms $f: A \rightarrow B$ and $g: B \rightarrow C$ in a category, 1_A for the identity morphism on A , and I for the unit object of a symmetric monoidal category.

2.1 Parametric 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 $q \in \mathbb{R}^p$ such that the function $f(q, -) : \mathbb{R}^n \rightarrow \mathbb{R}^m$ closely approximates g . We formalise these maps categorically via the **Para** construction [9, 23, 24, 30].


Definition 1 (Parametric category). Let $(\mathcal{C}, \otimes, I)$ be a strict⁴ symmetric monoidal category. We define a category **Para**(\mathcal{C}) with objects those of \mathcal{C} , and a map from A to B a pair (P, f) , with P an object of \mathcal{C} and $f : P \otimes A \rightarrow B$. The composite of maps $(P, f) : A \rightarrow B$ and $(P', f') : B \rightarrow C$ is the pair $(P' \otimes P, (1_{P'} \otimes f); f')$. The identity on A is the pair $(I, 1_A)$.

Example 1. Take the category **Smooth** whose objects are natural numbers and whose morphisms $f : n \rightarrow m$ are smooth maps 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 we will see in the next sections, the interplay of the various components at work in the learning process becomes much clearer once represented the morphisms of **Para**(\mathcal{C}) using the pictorial formalism of *string diagrams*, which we now recall. In fact, we will mildly massage the traditional notation for string diagrams (below left), by representing a morphism $f : A \rightarrow B$ in **Para**(\mathcal{C}) as below right.



This is to emphasise the special role played by P , reflecting the fact that in machine learning data and parameters have different semantics. String diagrammatic notations also allows to neatly represent composition of maps $(P, f) : A \rightarrow B$ and $(P', f') : B \rightarrow C$ (below left), and “reparameterisation” of $(P, f) : A \rightarrow B$ by a map $\alpha : Q \rightarrow P$ (below right), yielding a new map $(Q, (\alpha \otimes 1_A); f) : A \rightarrow B$.



⁴ One can also define **Para**(\mathcal{C}) in the case when \mathcal{C} is non-strict; however, the result would be not a category but a bicategory.

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 the two operations in (2) just work as expected, as these diagrams can be ultimately ‘compiled’ down to string diagrams for monoidal categories.

2.2 Lenses


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. For any Cartesian category \mathcal{C} , the category of (bimorphic) lenses in \mathcal{C} , $\mathbf{Lens}(\mathcal{C})$, is the category with the following data. Objects are pairs (A, A') of objects in \mathcal{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 $(1_A, \pi_1)$.

The embedding of $\mathbf{Lens}(\mathcal{C})$ into the category of Tambara modules over \mathcal{C} (see [7, Thm. 23]) provides a rich string diagrammatic language, in which lenses may be represented with forward/backward wires indicating the information flow. In this language, a morphism $(f, f^*) : (A, A') \rightarrow (B, B')$ is written as below left, which can be ‘expanded’ as below right.



It is clear in this language how to describe the composite of $(f, f^*) : (A, A') \rightarrow (B, B')$ and $(g, g^*) : (B, B') \rightarrow (C, C')$:



2.3 Parametric Lenses

The fundamental category where supervised learning takes place is the composite $\mathbf{Para}(\mathbf{Lens}(\mathcal{C}))$ of the two constructions in the previous sections:

Definition 3. *The category $\mathbf{Para}(\mathbf{Lens}(\mathcal{C}))$ of **parametric lenses** on \mathcal{C} has as objects pairs (A, A') of objects from \mathcal{C} . A morphism from (A, A') to (B, B') , called a **parametric lens**⁵, is a choice of parameter pair (P, P') and a lens $(f, f^*) : (P, P') \times (A, A') \rightarrow (B, B')$ so that $f : P \times A \rightarrow B$ and $f^* : P \times A \times B' \rightarrow P' \times A'$*

String diagrams for parametric lenses are built by simply composing the graphical languages of the previous two sections — see (1), where respectively a morphism, a composition of morphisms, and a reparameterisation are depicted.

Given a generic morphism in $\mathbf{Para}(\mathbf{Lens}(\mathcal{C}))$ as depicted in (1) on the left, 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, ie., a parametric map $(P, f) : A \rightarrow B$. This is precisely what the gradient operation provides, and its generalization to categories is explored in the next subsection.

2.4 Cartesian Reverse Differential Categories

Fundamental to all types of 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 [50] 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, we will work in a category with an abstract gradient operation.

Definition 4. *A **Cartesian left additive category** [13, Defn. 1] consists of a category \mathcal{C} with chosen finite products (including a terminal object), and an addition operation and zero morphism in each homset, satisfying various axioms. A **Cartesian reverse differential category** (CRDC) [13, Defn. 13] consists of a Cartesian left additive category \mathcal{C} , together with an operation which provides, for each map $f : A \rightarrow B$ in \mathcal{C} , a map $R[f] : A \times B \rightarrow A$ satisfying various axioms.*

For $f : A \rightarrow B$, the pair $(f, R[f])$ forms a lens from (A, A) to (B, B) . We will pursue the idea that $R[f]$ acts as backwards map, thus giving a means to “learn” f .

⁵ In [23], these are called *learners*. However, in this paper we study them in a much broader light; see Section 6.

Note that 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 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.

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

Example 2. The category **Smooth** (Example 1) is Cartesian with product given by addition, and it is also 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

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

the reverse derivative (as opposed to the forward derivative) 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 [28]), as is the case in most supervised learning situations (where often $m = 1$).

Example 3. Another CRDC is the symmetric monoidal category $\text{POLY}_{\mathbb{Z}_2}$ [13, Example 14] with objects the natural numbers and morphisms $f: A \rightarrow B$ the 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 [50].


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 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 parametric 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 $\mathbf{Para}(\mathcal{C})$, for a CRDC \mathcal{C} . In order to visualise such morphism as a parametric lens, it then suffices to apply under $\mathbf{Para}(-)$ the canonical morphism $\mathbf{R}: \mathcal{C} \rightarrow \mathbf{Lens}(\mathcal{C})$ (which exists for any CRDC \mathcal{C} , see [13, Prop. 31]), mapping f to $(f, R[f])$. This yields a functor $\mathbf{Para}(\mathbf{R}): \mathbf{Para}(\mathcal{C}) \rightarrow \mathbf{Para}(\mathbf{Lens}(\mathcal{C}))$, pictorially defined as



Example 4 (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 1, 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 (4) to present a neural network together with its reverse derivative operator as a parametric lens, i.e. a morphism in $\mathbf{Para}(\mathbf{Lens}(\mathbf{Smooth}))$.

Example 5 (Boolean circuits). For learning of Boolean circuits as described in [50], the recipe is the same as in Example 4, except that the base category is $\mathbf{POLY}_{\mathbb{Z}_2}$ (see Example 3). The important observation here is that $\mathbf{POLY}_{\mathbb{Z}_2}$ is a CRDC, see [13, 50], and thus we can apply the functor in (4).

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 \mathbf{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

parametric map from B to \mathbb{R} with parameter space B .⁶ We also generalize the codomain to an arbitrary object L .

Definition 5. A *loss map on B* consists of a parametric map $(B, \text{loss}) : \text{Para}(\mathcal{C})(B, L)$ for some object L .

Note that we can precompose a loss map $(B, \text{loss}) : B \rightarrow L$ with a neural network $(P, f) : A \rightarrow B$ (below left), and apply the functor in (4) (with $\mathcal{C} = \text{Smooth}$) to obtain the parametric lens below right.

$$\begin{array}{c}
 \begin{array}{ccc}
 & P & B \\
 & \downarrow & \downarrow \\
 A & \xrightarrow{f} & B \\
 & & \downarrow \\
 & & \text{loss} \\
 & & \downarrow \\
 & & L
 \end{array}
 & \mapsto &
 \begin{array}{ccccc}
 & P & P' & B & B' \\
 & \downarrow & \uparrow & \downarrow & \uparrow \\
 A & \xrightarrow{f} & B & \xrightarrow{\text{loss}} & L \\
 A' & \xleftarrow{R[f]} & B' & \xleftarrow{R[\text{loss}]} & L'
 \end{array}
 \end{array} \quad (5)$$

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 6 (Quadratic error). In Smooth , the standard loss function on \mathbb{R}^b is quadratic error: it uses $L = \mathbb{R}$ and has parametric 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 by $R[e](b_t, b_p, \alpha) = \alpha \cdot (b_p - b_t, b_t - b_p)$ — note α suggests the idea of *learning rate*, which we will explore in Section 3.3.

Example 7 (Boolean error). In $\text{POLY}_{\mathbb{Z}_2}$, the loss function on \mathbb{Z}^b which is implicitly used in [50] is a bit different: it uses $L = \mathbb{Z}^b$ and has parametric 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 $R[e](b_t, b_p, \alpha) = (\alpha, \alpha)$.

Example 8 (Softmax cross entropy). The Softmax cross entropy loss is a \mathbb{R}^b -parametric map $\mathbb{R}^b \rightarrow \mathbb{R}$ defined by $e(b_t, b_p) = \sum_{i=1}^b (b_t)_i ((b_p)_i - \log(\text{Softmax}(b_p)_i))$ where $\text{Softmax}(b_p) = \frac{\exp((b_p)_i)}{\sum_{j=1}^b \exp((b_p)_j)}$ is defined componentwise for each class i .

We note that, although b_t 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 b_t as the image of some logits under the Softmax function.

⁶ Here 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, [9, Figure 7.].


Example 9 (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 $e(b_t, b_p) = b_t \cdot b_p$ which computes the dot product of two vectors. If the ground truth vector y is a one-hot vector (active at the i -th element), then the dot product performs masking of all inputs except the i -th one. Note the reverse derivative $R[e]: \mathbb{R}^b \times \mathbb{R}^b \times \mathbb{R} \rightarrow \mathbb{R}^b \times \mathbb{R}^b$ of the dot product is defined as $R[e](b_t, b_p, \alpha) = (\alpha \cdot b_p, \alpha \cdot b_t)$.

3.3 Learning Rates as Parametric Lenses

After models and loss maps, another ingredient of the learning process are *learning rates*, which we formalise as follows.

Definition 6. 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 the learning rate 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 $\alpha^*: L \rightarrow L'$. Thus we can view α as a parametric lens from $(L, L') \rightarrow (1, 1)$ (with trivial parameter space) and compose it in $\mathbf{Para}(\mathbf{Lens}(\mathcal{C}))$ with a model and a loss map (cf. (5)) to get



Example 10. In standard supervised learning in \mathbf{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 11. In supervised learning in $\mathbf{POLY}_{\mathbb{Z}_2}$, 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 \mathbf{Smooth} by $\alpha(l) = -\epsilon \cdot l$. Such a choice would take into account how far away the network is from its desired goal and adjust the learning rate accordingly.

3.4 Optimisers as Reparameterisations

In this section we consider how to implement gradient descent (and its variants) into our framework. To this aim, note that the parametric lens $(f, R[f])$ representing our model (see (4)) outputs a P' , which represents a *change* in the parameter space. Now, we would like to receive not just the requested change in the parameter, but the new parameter itself. This is precisely what gradient descent accomplishes, when formalised as a lens.

Definition 7. In any CRDC \mathcal{C} we can define gradient update as a map G in $\mathbf{Lens}(\mathcal{C})$ from (P, P) to (P, P') consisting of $(G, G^*) : (P, P) \rightarrow (P, P')$, where $G(p) = p$ and $G^*(p, p') = p + p'$ ⁷.

Intuitively, such a lens allows one to receive the requested change in parameter and implement that change by adding that value to the current parameter. By its type, we can now “plug” the gradient descent lens $G : (P, P) \rightarrow (P, P')$ above the model $(f, R[f])$ in (4) — formally, this is accomplished as a *reparameterisation* of the parametric morphism $(f, R[f])$, cf. Section 2.1. This gives us Figure 3 (left).




Fig. 3: Model reparameterised by basic gradient descent (left) and a generic stateful optimiser (right).

Example 12 (Gradient update in Smooth). In \mathbf{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 13 (Gradient update in Boolean circuits). In the CRDC $\mathbf{POLY}_{\mathbb{Z}_2}$, 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 also be seen as taking the XOR of the current parameter and the requested change; this is exactly how this algorithm is implemented in [50].

Other variants of gradient descent also fit naturally into this framework by allowing for additional input/output data with P . In particular, many of them 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 $(P, P) \rightarrow (P, P')$, we ask instead for a lens $(S \times P, S \times P) \rightarrow (P, P')$ where S is some “state” object.

⁷ Note that as in the discussion in Section 2.4, we are implicitly assuming that $P = P'$; we have merely notated them differently to emphasize the different “roles” they play (the first P can be thought of as “points”, the second as “vectors”)

Definition 8. 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 optimiser as a reparameterisation which may be “plugged in” a model as in Figure 3 (right). Let us now consider how several well-known optimisers can be implemented in this way.

Example 14 (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 part. However, as soon as we move to more complicated variants, this is not anymore the case, as for instance in Nesterov momentum below.

Example 15 (Nesterov momentum). In Nesterov momentum, one uses the momentum from previous updates to tweak the input parameter supplied to the network. 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 16 (Adagrad). Given any fixed $\epsilon > 0$ and $\delta \sim 10^{-7}$, Adagrad [20] 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 17 (Adam). Adaptive Moment Estimation (Adam) [32] 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). For 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 $\text{put}(m, v, p, p') = (\hat{m}', \hat{v}', p + \frac{\epsilon}{\delta + \sqrt{\hat{v}'}} \odot \hat{m}')$ where $m' = \beta_1 m + (1 - \beta_1)p'$, $v' = \beta_2 v + (1 - \beta_2)p'^2$, and $\hat{m}' = \frac{m'}{1 - \beta_1^t}$, $\hat{v}' = \frac{v'}{1 - \beta_2^t}$.

Note that, so far, optimisers/reparameterisations have been added to the P/P' wires. In order to change the model’s parameters (Fig. 3). In Section 4.2 we will study them on 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 (6), 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 parametric lens η we now define. The parametric 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 (6) gives us the closed parametric lens




This composite is now a map in $\mathbf{Para}(\mathbf{Lens}(\mathcal{C}))$ 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 $\text{put}(a, p, b_t) = (a', p', b'_t)$, where

$$b_p = f(p, a) \quad (b'_t, b'_p) = R[\text{loss}](b_t, b_p, \alpha(\text{loss}(b_t, b_p))) \quad (p', a') = R[f](p, a, b'_p).$$

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

4.1 Supervised Learning of Parameters

The most common type of learning performed on (7) is supervised learning of *parameters*. This is done by reparameterising (*cf.* Section 2.1) the image in the following manner. The parameter ports are reparameterised by one of the (possibly stateful) optimisers described in the previous section, while the backward wires A' of inputs and B' of outputs are discarded. This finally yields the complete picture of a system which learns the parameters in a supervised manner:



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 $\mathbf{Para}(\mathbf{Lens}(\mathcal{C}))$ 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 the terminal map. Its `put` map has the type $A \times S \times P \times B \rightarrow S \times P$ and unpacks to $\text{put}(a, s, p, b_t) = U^*(s, p, p')$, where

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

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 (3); 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 literature. This presentation has the advantage of treating all these instances uniformly in terms of basic constructs, highlighting their similarities and differences. First, we fix some parametric map $(\mathbb{R}^p, f) : \mathbf{Para}(\mathbf{Smooth})(\mathbb{R}^a, \mathbb{R}^b)$ in \mathbf{Smooth} and the constant *negative* learning rate $\alpha : \mathbb{R}$ (Example 10). 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 18 (Quadratic error, basic gradient descent). Fix the quadratic error (Example 6) as the loss map and basic gradient update (Example 12). Then the aforementioned `put` map simplifies. Since there is no state, its type reduces to $A \times P \times B \rightarrow P$, and we have $\text{put}(a, p, b_t) = p + p'$, 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 4), we can slide the α from the costate into the basic gradient update 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 parametric map $(\mathbb{R}^p, f): \mathbb{R}^a \rightarrow \mathbb{R}^b$ 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 f is linear and $b = 1$, we recover simple linear regression with gradient descent. If the codomain is multi-dimensional, i.e. we are 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 parametric map.

Example 19 (Softmax cross entropy, basic gradient descent). Fix Softmax cross entropy (Example 8) as the loss map and basic gradient update (Example 12). Again the **put** map simplifies. The type reduces to $A \times P \times B \rightarrow P$ and we have $\text{put}(a, p, b_t) = p + p'$ 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 20 (Mean squared error, Nesterov Momentum). Fix the quadratic error (Example 6) as the loss map and Nesterov momentum (Example 15) as the gradient update. This time the **put** map $A \times S \times P \times B \rightarrow S \times P$ does not have a simplified type. The implementation of **put** reduces to $\text{put}(a, s, p, b_t) = (s', p + s')$, where $\bar{p} = p + \gamma s$, $(p', a') = R[f](\bar{p}, a, \alpha \cdot (f(\bar{p}, a) - b_t))$, and $s' = -\gamma s + p'$.

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.

Our last example, using a different base category $\text{POLY}_{\mathbb{Z}_2}$, shows that our framework captures learning in not just continuous, but discrete settings too. Again, we fix a parametric map $(\mathbb{Z}^p, f) : \text{POLY}_{\mathbb{Z}_2}(\mathbb{Z}^a, \mathbb{Z}^b)$ but this time we fix the identity learning rate (Example 11), instead of a constant one.


Example 21 (Basic learning in Boolean circuits). Fix XOR as the loss map (Example 7) and the basic gradient update (Example 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 parametric map!


In other words, we can recast the **put** map as a parametric map $(A \times B, \text{put}) : \mathbf{Para}(\mathcal{C})(P, P)$. Being an endo-map, it can be composed with itself. The resulting composite is an endo-map taking 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 reparameterising the parameter port with gradient descent allows us to capture supervised parameter learning. In this section we describe how reparameterising 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 [19, 34, 35, 44].



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* (see (8)). 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 [1].

Just like in the previous subsection, we can write this deep dreaming system as a map in $\mathbf{Para}(\mathbf{Lens}(\mathcal{C}))$ 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 $\mathbf{put}(s, a, p, b_t) = U^*(s, a, a')$, where $\bar{a} = U(s, a)$, $b_p = f(p, \bar{a})$, $(b'_t, b'_p) = R[\mathbf{loss}](b_t, b_p, \alpha(\mathbf{loss}(b_t, b_p)))$, and $(p', a') = R[f](p, \bar{a}, b'_p)$.

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 [44, 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 9). 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 [44]. We explicitly unpack this in the following example.

Example 22 (Deep dreaming, dot product loss, basic gradient update). Fix \mathbf{Smooth} as base category, a parametric map $(\mathbb{R}^p, f) : \mathbf{Para}(\mathbf{Smooth})(\mathbb{R}^a, \mathbb{R}^b)$, the dot product loss (Example 9), basic gradient update (Example 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 $\mathbf{put}(a, p, b_t) = a + a'$, where $(p', a') = R[f](p, a, \alpha \cdot b_t)$. Like in Example 18, this update can be rewritten as $\mathbf{put}(a, p, b_t) = a + \alpha \cdot (R[f](p, a, b_t); \pi_1)$, 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 [44, 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 8) as a loss function in deep dreaming, which seems like the more natural choice in this setting. Furthermore, while deep dreaming commonly uses basic gradient descent, there is nothing preventing the use of any of the optimiser 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 — full usage examples, source code, and experiments can be found at [17]. We demonstrate the correctness of our library empirically using a number of experiments implemented both in our library and in Keras [11], a popular framework for deep learning. For example, one experiment is a model for the MNIST image classification problem [33]: we implement the same model in both frameworks and achieve comparable accuracy. Note that despite similarities between the user interfaces of our library and of Keras, a model in our framework is constructed as a composition of parametric lenses. This is fundamentally different to the approach taken by Keras and other existing libraries, and highlights how our proposed algebraic structures naturally guide programming practice

In summary, our implementation 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 >> 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 (9):

```
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 2 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 `apply_update(α , f)` represents the ‘vertical stacking’ of α atop f :

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

Now, given the parametric lens of (10), one can construct a morphism `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 [17].

6 Related Work

The work [23] is closely related to ours, in that it provides an abstract categorical model of backpropagation. 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 parametric functions and parameter update. We thus can go well beyond [23] 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 [23] in a more structured way. For example, rather than considering the category **Learn** of [23] 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 [38] enabling us to model things such as control flow using prisms.

One more relevant aspect is functoriality. We use a functor to augment a parametric map with its backward pass, just like [23]. 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 [23], 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 [23] as parts of the aforementioned functor, here they are extracted out and play a key role: loss maps are parametric 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. We have investigated this possibility thoroughly in our paper, showing

it is gradient update which is used to dream up pictures. We also correct a small issue in Theorem III.2 of [23]. There, the morphisms of **Learn** were defined up to an equivalence (pg. 4 of [23]) 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 [23] 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 $\mathbf{Para}(\mathcal{C}) \rightarrow \mathbf{Learn}$ as a primitive. In our case, this functor has the type $\mathbf{Para}(\mathcal{C}) \rightarrow \mathbf{Para}(\mathbf{Lens}(\mathcal{C}))$ and arises from applying **Para** to a canonical functor $\mathcal{C} \rightarrow \mathbf{Lens}(\mathcal{C})$ existing for *any* reverse derivative category, not just **Smooth**. Lastly, in our paper we took advantage of the graphical calculus for **Para**, redrawing many diagrams appearing in [23] in a structured way.

Other than [23], there are a few more relevant papers. The work of [18] 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 of diffeological spaces, and requiring cartesian and local cartesian closed structure on the base category. Lenses and Learners are studied in the eponymous work of [22] which observes that learners are parametric lenses. 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 [23]. A lens-like implementation of automatic differentiation is the focus of [21], but learning algorithms aren't studied. A relationship between category-theoretic perspective on probabilistic modeling and gradient-based optimisation is studied in [42] which also studies a variant of the **Para** construction. Usage of Cartesian differential categories to study learning is found in [46]. They extend the differential operator to work on stateful maps, but do not study lenses, parameterisation nor update maps. The work of [24] studies deep learning in the context of Cycle-consistent Generative Adversarial Networks [51] and formalises it via free and quotient categories, making parallels to the categorical formulations of database theory [45]. They do use the **Para** construction, but do not relate it to lenses nor reverse derivative categories. A general survey of category theoretic approaches to machine learning, covering many of the above papers, can be found in [43]. Lastly, the concept of parametric lenses has started appearing in recent formulations of categorical game theory and cybernetics [9, 10]. The work of [9] generalises the study of parametric lenses into parametric optics and connects it to game theoretic concepts such as Nash equilibria.

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 *parametric lens*. The foundation covers a wide variety of examples: different optimisers and loss maps

in gradient-based learning, different settings where gradient-based learning happens (smooth functions vs. boolean circuits), and 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 parametric lens should pave the way for Generative Adversarial Networks [27], 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 [4] 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 parametric lenses; and our framework covers learning with parametric lenses. Recurrent neural networks are another example of a more complex architecture, which has already been studied in the context of differential categories in [46]. When it comes to architectures, future work includes modelling some classical systems as well, such as the Support Vector Machines [15], which should be possible with the usage of loss maps such as Hinge loss.

Future work also includes using the full power of CRDC axioms. In particular, axioms RD.6 or RD.7, which deal with the behaviour of higher-order derivatives, were not exploited in our work, but they should play a role in modelling some supervised learning algorithms using higher-order derivatives (for example, the Hessian) for additional 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 $F : \mathcal{C} \rightarrow \mathbf{Lens}(\mathcal{C})$ - 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 [25]. 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.

Acknowledgements Fabio Zanasi acknowledges support from EPSRC EP/V002376/1. Geoff Cruttwell acknowledges support from NSERC.

References

1. Inceptionism: Going deeper into neural networks (2015), <https://ai.googleblog.com/2015/06/inceptionism-going-deeper-into-neural.html>
2. Explainable AI: the basics - policy briefing (2019), royalsociety.org/ai-interpretability
3. Abramsky, S., Coecke, B.: A categorical semantics of quantum protocols. In: Proceedings of the 19th Annual IEEE Symposium on Logic in Computer Science, 2004. pp. 415–425 (2004). <https://doi.org/10.1109/LICS.2004.1319636>
4. Andrychowicz, M., Denil, M., Gomez, S., Hoffman, M.W., Pfau, D., Schaul, T., Shillingford, B., de Freitas, N.: Learning to learn by gradient descent by gradient descent. In: 30th Conference on Neural Information Processing Systems (NIPS) (2016)
5. Baez, J.C., Erbele, J.: Categories in Control. *Theory and Applications of Categories* **30**(24), 836–881 (2015)
6. Bohannon, A., Foster, J.N., Pierce, B.C., Pilkiewicz, A., Schmitt, A.: Boomerang: Resourceful lenses for string data. *SIGPLAN Not.* **43**(1), 407–419 (Jan 2008). <https://doi.org/10.1145/1328897.1328487>
7. Boisseau, G.: String Diagrams for Optics. arXiv:2002.11480 (2020)
8. Bonchi, F., Sobocinski, P., Zanasi, F.: The calculus of signal flow diagrams I: linear relations on streams. *Inf. Comput.* **252**, 2–29 (2017). <https://doi.org/10.1016/j.ic.2016.03.002>, <https://doi.org/10.1016/j.ic.2016.03.002>
9. Capucci, M., Gavranović, B., Hedges, J., Rischel, E.F.: Towards foundations of categorical cybernetics. arXiv:2105.06332 (2021)
10. Capucci, M., Ghani, N., Ledent, J., Nordvall Forsberg, F.: Translating Extensive Form Games to Open Games with Agency. arXiv:2105.06763 (2021)
11. Chollet, F., et al.: Keras (2015), <https://github.com/fchollet/keras>
12. Clarke, B., Elkins, D., Gibbons, J., Loregian, F., Milewski, B., Pillmore, E., Román, M.: Profunctor optics, a categorical update. arXiv:2001.07488 (2020)
13. Cockett, J.R.B., Cruttwell, G.S.H., Gallagher, J., Lemay, J.S.P., MacAdam, B., Plotkin, G.D., Pronk, D.: Reverse derivative categories. In: Proceedings of the 28th Computer Science Logic (CSL) conference (2020)
14. Coecke, B., Kissinger, A.: *Picturing Quantum Processes: A First Course in Quantum Theory and Diagrammatic Reasoning*. Cambridge University Press (2017). <https://doi.org/10.1017/9781316219317>
15. Cortes, C., Vapnik, V.: Support-vector networks. *Machine learning* **20**(3), 273–297 (1995)
16. Courbariaux, M., Bengio, Y., David, J.P.: BinaryConnect: Training Deep Neural Networks with binary weights during propagations. arXiv:1511.00363
17. CRCoauthors, A.: Numeric Optics: A python library for constructing and training neural networks based on lenses and reverse derivatives. <https://github.com/anonymous-c0de/esop-2022>
18. Dalrymple, D.: 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>
19. Dosovitskiy, A., Brox, T.: Inverting convolutional networks with convolutional networks. arXiv:1506.02753 (2015)

20. Duchi, J., Hazan, E., Singer, Y.: Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research* **12**(Jul), 2121–2159 (2011)
21. Elliott, C.: The simple essence of automatic differentiation (differentiable functional programming made easy). arXiv:1804.00746 (2018)
22. Fong, B., Johnson, M.: Lenses and learners. In: *Proceedings of the 8th International Workshop on Bidirectional Transformations (Bx@PLW)* (2019)
23. Fong, B., Spivak, D.I., Tuyéras, R.: Backprop as functor: A compositional perspective on supervised learning. In: *Proceedings of the Thirty fourth Annual IEEE Symposium on Logic in Computer Science (LICS 2019)*. pp. 1–13. IEEE Computer Society Press (June 2019)
24. Gavranovic, B.: Compositional deep learning. arXiv:1907.08292 (2019)
25. Ghani, N., Hedges, J., Winschel, V., Zahn, P.: Compositional game theory. In: *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science*. p. 472–481. LICS '18 (2018). <https://doi.org/10.1145/3209108.3209165>
26. Ghica, D.R., Jung, A., Lopez, A.: Diagrammatic Semantics for Digital Circuits. arXiv:1703.10247 (2017)
27. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N.D., Weinberger, K.Q. (eds.) *Advances in Neural Information Processing Systems 27*, pp. 2672–2680 (2014), <http://papers.nips.cc/paper/5423-generative-adversarial-nets.pdf>
28. Griewank, A., Walther, A.: *Evaluating derivatives: principles and techniques of algorithmic differentiation*. Society for Industrial and Applied Mathematics (2008)
29. Hedges, J.: Limits of bimorphic lenses. arXiv:1808.05545 (2018)
30. Hermida, C., Tennent, R.D.: Monoidal indeterminates and categories of possible worlds. *Theor. Comput. Sci.* **430**, 3–22 (Apr 2012). <https://doi.org/10.1016/j.tcs.2012.01.001>
31. Johnson, M., Rosebrugh, R., Wood, R.: Lenses, fibrations and universal translations. *Mathematical structures in computer science* **22**, 25–42 (2012)
32. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: Bengio, Y., LeCun, Y. (eds.) *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings* (2015), <http://arxiv.org/abs/1412.6980>
33. Lecun, Y., Bottou, L., Bengio, Y., Haffner, P.: Gradient-based learning applied to document recognition. In: *Proceedings of the IEEE*. pp. 2278–2324 (1998). <https://doi.org/10.1109/5.726791>
34. Mahendran, A., Vedaldi, A.: Understanding deep image representations by inverting them. arXiv:1412.0035 (2014)
35. Nguyen, A.M., Yosinski, J., Clune, J.: Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. arXiv:1412.1897 (2014)
36. Olah, C.: Neural networks, types, and functional programming (2015), <http://colah.github.io/posts/2015-09-NN-Types-FP/>
37. Polyak, B.: Some methods of speeding up the convergence of iteration methods. *USSR Computational Mathematics and Mathematical Physics* **4**(5), 1 – 17 (1964). [https://doi.org/https://doi.org/10.1016/0041-5553\(64\)90137-5](https://doi.org/https://doi.org/10.1016/0041-5553(64)90137-5), <http://www.sciencedirect.com/science/article/pii/0041555364901375>
38. Riley, M.: Categories of optics. arXiv:1809.00738 (2018)
39. Selinger, P.: A survey of graphical languages for monoidal categories. *Lecture Notes in Physics* p. 289–355 (2010)

40. Selinger, P.: Control categories and duality: on the categorical semantics of the lambda-mu calculus. *Mathematical Structures in Computer Science* **11**(02), 207–260 (4 2001). <https://doi.org/null>, http://journals.cambridge.org/article_S096012950000311X
41. Seshia, S.A., Sadigh, D.: Towards verified artificial intelligence. *CoRR abs/1606.08514* (2016), <http://arxiv.org/abs/1606.08514>
42. Shiebler, D.: Categorical Stochastic Processes and Likelihood. *Compositionality* **3**(1) (2021)
43. Shiebler, D., Gavranović, B., Wilson, P.: Category Theory in Machine Learning. *arXiv:2106.07032* (2021)
44. Simonyan, K., Vedaldi, A., Zisserman, A.: Deep inside convolutional networks: Visualising image classification models and saliency maps. *arXiv:1312.6034* (2014)
45. Spivak, D.I.: Functorial data migration. *arXiv:1009.1166* (2010)
46. Sprunger, D., Katsumata, S.y.: Differentiable causal computations via delayed trace. In: *Proceedings of the 34th Annual ACM/IEEE Symposium on Logic in Computer Science. LICS '19*, IEEE Press (2019)
47. Steckermeier, A.: Lenses in functional programming. Preprint, available at <https://sinusoid.es/misc/lager/lenses.pdf> (2015)
48. Sutskever, I., Martens, J., Dahl, G., Hinton, G.: On the importance of initialization and momentum in deep learning. In: Dasgupta, S., McAllester, D. (eds.) *Proceedings of the 30th International Conference on Machine Learning*. vol. 28, pp. 1139–1147 (2013), <http://proceedings.mlr.press/v28/sutskever13.html>
49. Turi, D., Plotkin, G.: Towards a mathematical operational semantics. In: *Proceedings of Twelfth Annual IEEE Symposium on Logic in Computer Science*. pp. 280–291 (1997). <https://doi.org/10.1109/LICS.1997.614955>
50. Wilson, P., Zanasi, F.: Reverse derivative ascent: A categorical approach to learning boolean circuits. In: *Proceedings of Applied Category Theory (ACT)* (2020), <https://cgi.cse.unsw.edu.au/~eptcs/paper.cgi?ACT2020:31>
51. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networks. *arXiv:1703.10593* (2017)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

