

# Beyond Backpropagation: Bilevel Optimization Through Implicit Differentiation and Equilibrium Propagation

**Nicolas Zucchet**

*nzucchet@inf.ethz.ch*

*Department of Computer Science, ETH Zurich 8092, Zurich, Switzerland*

**João Sacramento**

*rjoao@ethz.ch*

*Institute of Neuroinformatics, University of Zurich and ETH Zurich 8092, Zurich, Switzerland*

**This review examines gradient-based techniques to solve bilevel optimization problems. Bilevel optimization extends the loss minimization framework underlying statistical learning to systems that are implicitly defined through a quantity they minimize. This characterization can be applied to neural networks, optimizers, algorithmic solvers, and even physical systems and allows for greater modeling flexibility compared to the usual explicit definition of such systems. We focus on solving learning problems of this kind through gradient descent, leveraging the toolbox of implicit differentiation and, for the first time applied to this setting, the equilibrium propagation theorem. We present the mathematical foundations behind such methods, introduce the gradient estimation algorithms in detail, and compare the competitive advantages of the different approaches.**

## 1 Introduction ---

Recent years have witnessed an explosion of breakthroughs fueled by deep learning in many scientific fields such as computer vision (Krizhevsky, Sutskever, & Hinton, 2012), natural language processing (Brown et al., 2020), game playing (Mnih et al., 2015), and biology (Jumper et al., 2021). There are many lessons to learn from these advances. Among them is the surprising effectiveness of gradient descent: updating the millions, or even billions, of parameters of a deep learning model through greedy gradient-following updates turns out to be extremely powerful and cheap, thanks to the backpropagation of errors algorithm (Linnainmaa, 1976; Werbos, 1982; Rumelhart, Hinton, & Williams, 1986).

In its standard form, backpropagation provides an efficient way of computing gradients in neural networks, but its applicability is limited to acyclic directed computational graphs whose nodes are explicitly defined.

Feedforward neural networks or unfolded-in-time recurrent neural networks are prime examples of such graphs. However, there exists a wide range of computations that are easier to describe through what they achieve rather than by the exact sequence of calculations they perform and that thus do not fulfill the requirements of backpropagation. This includes outputs of algorithmic solvers that provably minimize some cost function (Djolonga & Krause, 2017; Wang, Donti, Wilder, & Kolter, 2019; Vlastelica, Paulus, Musil, Martius, & Rolínek, 2020), of learning processes that do loss minimization (MacKay, 1992; Bengio, 2000), and even of physical systems, such as biological neural networks (Hopfield, 1984) or electrical circuits (Wyatt & Standley, 1989; Kendall, Pantone, Manickavasagam, Bengio, & Scellier, 2020; Scellier, 2021), reaching a steady state. Gradient descent based on naive backpropagation cannot improve those computations, as the algorithm is not directly applicable.

In this review, we frame learning such implicitly defined systems as a bilevel optimization problem and review how to compute the corresponding gradients through implicit differentiation methods. We then present a less explored alternative approach, which relies on the equilibrium propagation theorem, recently discovered by Scellier and Bengio (2017).

Our review is organized as follows:

- In section 2, we formalize bilevel optimization and discuss some examples in which it appears in machine learning from a historical perspective. We then focus on hyperparameter optimization and metalearning to highlight the challenges behind solving bilevel optimization problems.
- Behind the tools of interest for this review are the concept of implicit function and the so-called implicit function theorem. In section 3, we provide insight into why this notion is so fundamental and present the implicit function theorem in detail.
- Section 4 is the core of the review: we introduce two gradient-based approaches to solve bilevel optimization problems. The first class of methods leverages the differentiation formula provided by the implicit function theorem, while the second one builds on another theorem, the equilibrium propagation theorem. We present the mathematical foundations of the two approaches and show how to transform them into efficient gradient estimation algorithms. We theoretically analyze the quality of the gradient those algorithms produce as a function of the different sources of approximation they introduce.
- In section 5, we compare the different algorithms we presented in the previous section with each other, exhibiting their qualities and limitations. We then discuss when these methods shine by contrasting them with backpropagation through the optimization process.

This review has two levels of reading: one for readers interested in learning new gradient estimation methods and one for readers who wish to know the mathematical foundations behind them. For this reason, we mark all theory-oriented sections with the symbol †. They can be skipped without hindering the understanding of the rest.

## 2 Bilevel Optimization in Machine Learning ---

**2.1 Bilevel Optimization.** The high-level description of bilevel optimization that we have briefly sketched contains two elements: an inner optimization process that describes what the system does and an outer loss function that ultimately measures how good the result of this process is. We now make this formulation more precise.

We denote by  $\phi$  the parameters that are optimized by the inner process to minimize the inner loss function  $L^{\text{in}}$ . The system we consider has some parameters  $\theta$  that we want to learn. We assume that they modify the behavior of the system through  $L^{\text{in}}$ . The computation performed by the system is then

$$\phi_{\theta}^* \in \arg \min_{\phi} L^{\text{in}}(\phi, \theta).$$

We use the subscript  $\theta$  to underline that  $\phi_{\theta}^*$  is an implicit function of  $\theta$  (as  $L^{\text{in}}$  depends on  $\theta$ ), as it can be any local minimizer of the inner loss  $L^{\text{in}}$ . Note that we do not make any assumption on how to obtain  $\phi_{\theta}^*$ , as we only assume that it minimizes a loss function.

The outer loss  $L^{\text{out}}$  measures the quality of the output  $\phi_{\theta}^*$  of the system and plays the usual role of a loss function in machine learning. We can then frame learning of the parameters  $\theta$  as the minimization of the outer loss, which leads to the bilevel optimization problem that we study in this review:

$$\begin{aligned} & \min_{\theta} L^{\text{out}}(\phi_{\theta}^*, \theta) \\ & \text{s.t. } \phi_{\theta}^* \in \arg \min_{\phi} L^{\text{in}}(\phi, \theta). \end{aligned} \tag{2.1}$$

**2.2 Historical Perspective.** Bilevel optimization was introduced in the 1930s by von Stackelberg (1934) in the context of two-player games with a leader and a follower, and later extensively studied in the field of optimization as a way to model optimization problems that contain different objectives (Bard, 1998). Closer to the learning formulation of interest to this review is bilevel optimization as introduced for the training of recurrent neural networks in the late 1980s. Instead of describing neural dynamics by their dynamics and then backpropagating through them, the neural activity

is assumed to converge to a steady state. This view led to the introduction of the recurrent backpropagation algorithm (Almeida, 1990; Pineda, 1987), one of the algorithms we review in section 4. Often, converging dynamics are described as minimizing an energy function (Hopfield, 1984; Cohen & Grossberg, 1983; Scellier & Bengio, 2017; Whittington & Bogacz, 2017). This offers stability guarantees and allows connecting to physical systems such as resistive (Millar, 1951; Hutchinson, Koch, Luo, & Mead, 1988; Wyatt & Standley, 1989; Kendall et al., 2020) or flow (Stern, Hexner, Rocks, & Liu, 2021) networks.

This implicit characterization of entire neural networks, or layers of them, introduced in the early days of deep learning has regained considerable interest recently (Amos & Kolter, 2017; Djolonga & Krause, 2017; Agrawal et al., 2019; Gould, Hartley, & Campbell, 2021). Notably, a class of such implicit networks called deep equilibrium models<sup>1</sup> (Bai, Kolter, & Koltun, 2019; Bai, Koltun, & Kolter, 2020) have achieved state-of-the-art performance in many problem domains. These results demonstrate that the performance of large deep feedforward neural networks can be matched by neural networks with far fewer parameters when the computations they perform are iterated repeatedly until equilibrium. As we will see in more detail, this results in large memory savings not only during inference but also during learning. Bilevel optimization also appears in many other forms in modern machine learning, going from hyperparameter optimization and metalearning, to generative adversarial networks (Goodfellow et al., 2014; Metz, Poole, Pfau, & Sohl-Dickstein, 2017) and reinforcement learning (Pfau & Vinyals, 2016; Rajeswaran, Mordatch, & Kumar, 2020; Nikishin, Abachi, Agarwal, & Bacon, 2022). We zoom in on hyperparameter optimization and metalearning in the next section as this is one of the problems for which bilevel optimization is mostly used now. We refer curious readers to appendix A for a more extensive presentation of some existing formulations.

## 2.3 Hyperparameter Optimization and Meta-Learning.

*2.3.1 Hyperparameter Optimization.* Let us consider the following problem: we want to find the parameters  $\theta$  of a learning algorithm, its hyperparameters, that generate model parameters  $\phi$ , which generalize well. We measure generalization performance by testing the learned model on held-out data. Furthermore, as is conventionally done, we assume that model parameters are obtained by maximum a posteriori estimation (MacKay, 1992; Foo, Do, & Ng, 2007) or, alternatively, by regularized empirical risk minimization (Bengio, 2000; Goutte & Larsen, 1998). This leads to the following

---

<sup>1</sup>See Kolter, Duvenaud, and Johnson (2021) for a tutorial on the topic.

bilevel optimization problem:

$$\begin{aligned} & \min_{\theta} L(\phi_{\theta}^*, \mathcal{D}^{\text{val}}) \\ \text{s.t. } & \phi_{\theta}^* \in \arg \min_{\phi} L(\phi, \mathcal{D}^{\text{train}}) + R(\phi, \theta), \end{aligned} \quad (2.2)$$

where  $L$  is the negative log likelihood that measures the discrepancy between the predictions of a neural network parameterized by  $\phi$  and the target outputs on a data set  $\mathcal{D}$ ,  $\mathcal{D}^{\text{train}}$  is the training set,  $\mathcal{D}^{\text{val}}$  is a held-out data set and  $R(\phi, \theta)$  is the negative log prior (in the Bayesian view) or a regularizing term on  $\phi$  (in the frequentist view). For instance, a very common choice is to take  $R(\phi, \theta) = \lambda \|\phi\|^2$ ; in this case, the hyperparameters are  $\theta = \{\lambda\}$ . A zoo of different interactions between  $\phi$  and  $\theta$  can be considered, and we mention a few of them in section A.2.

When  $\theta$  is low-dimensional, black-box optimization methods such as grid or random search (Bergstra & Bengio, 2012) can search for the best hyperparameters. However, this becomes intractable for high-dimensional hyperparameters. Alternatively, one could backpropagate through the training trajectory, but this does not scale well with the number of updates, as the entire history of parameters must be stored during training and then revisited in reverse-time order. The implicit methods we present in section 4 do not suffer from these limitations. They can scale to a large number of hyperparameters and long training procedures.

**2.3.2 Metalearning.** The previous formulation can be extended to metalearning (Thrun & Pratt, 1998; Schmidhuber, 1987; Bengio, Bengio, & Cloutier, 1990; Finn, Abbeel, & Levine, 2017; Bertinetto, Henriques, Torr, & Vedaldi, 2019) by considering several tasks. The goal is now to learn meta-parameters  $\theta$  that yield a learning algorithm that generalizes well on a family of tasks: ideally, the algorithm will achieve low loss on unseen tasks, which are assumed to be similar to those encountered during metalearning. The corresponding optimization problem is then:

$$\begin{aligned} & \min_{\theta} \mathbb{E}_{\tau} \left[ L^{\text{out}}(\phi_{\tau, \theta}^*, \theta, \mathcal{D}_{\tau}^{\text{val}}) \right] \\ \text{s.t. } & \phi_{\tau, \theta}^* \in \arg \min_{\phi} L^{\text{in}}(\phi, \theta, \mathcal{D}_{\tau}^{\text{train}}), \end{aligned} \quad (2.3)$$

where  $L^{\text{in}}$  and  $L^{\text{out}}$  are the same kind of loss used for hyperparameter optimization with the difference that the data on which they are defined are now dependent on the task  $\tau$ . In practice, this is solved by stochastic gradient descent on the expected outer loss over the task distribution: one task (or more) is sampled, and the gradient corresponding to that task is

approximated in the same way it would be for hyperparameter optimization. Thus, black-box optimization methods and backpropagation through training suffer from the same problems we already highlighted.

### 3 The Implicit Function Theorem

---

Studying implicit functions is about understanding the relationship between two variables  $y$  and  $x$  when they are linked together through an equation  $f(x, y) = 0$ . The first apparitions of implicit functions can be traced back to Descartes (1637) and Newton (1670), who considered the behavior of some specific curves (Krantz & Parks, 2003). Cauchy (1831) laid down the theoretical foundations behind the implicit function theorem, and the extended modern multivariate version of the theorem was introduced by Ulysse Dini in lecture notes<sup>2</sup> supporting his teaching on infinitesimal analysis at the University of Pisa during the academic year 1877–78 (Scarpello & Ritelli, 2002). We end this historical note with a citation from Euler (1748, as translated by John D. Blanton) that perfectly captures why implicit functions are relevant in mathematics in general and particularly relates to the philosophy behind bilevel optimization:

Indeed frequently algebraic functions cannot be expressed explicitly. For example, consider the function  $Z$  of  $z$  defined by the equation,  $z^5 = az^2Z^3 - bz^4Z^2 + cz^3Z - 1$ . Even if this equation cannot be solved, still it remains true that  $Z$  is equal to some expression composed of the variable  $z$  and constants, and for this reason  $Z$  shall be a function of  $z$ .

Implicit functions are inherent to bilevel optimization as the function  $\phi_\theta^*$  used in equation 2.1 satisfies  $\partial_\phi L^{\text{in}}(\phi_\theta^*, \theta) = 0$ . Understanding how implicit functions behave is therefore crucial; this is what the implicit function theorem brings. More precisely, it contains two statements: it ensures  $\theta \mapsto \phi_\theta^*$  exists locally and yields an analytical formula for the outer gradient  $\nabla_\theta$  associated with our problem:

$$\begin{aligned} \nabla_\theta^\top &:= \frac{d}{d\theta} L^{\text{out}}(\phi_\theta^*, \theta) \\ &= \frac{\partial L^{\text{out}}}{\partial \theta}(\phi_\theta^*, \theta) - \frac{\partial L^{\text{out}}}{\partial \phi}(\phi_\theta^*, \theta) \left( \frac{\partial^2 L^{\text{in}}}{\partial \phi^2}(\phi_\theta^*, \theta) \right)^{-1} \frac{\partial^2 L^{\text{in}}}{\partial \theta \partial \phi}(\phi_\theta^*, \theta). \quad (3.1) \end{aligned}$$

We derive this formula in section 3.2.

The attentive reader will have noticed that we are using two different notations for derivatives in the outer gradient formula. Let us clarify the convention we follow. We use  $\partial_x$  to denote partial derivatives with respect

---

<sup>2</sup>There is, however, no trace of the implicit function theorem in the 69 original papers Dini published.

to  $x$  and  $d_x$  for total derivatives. There is no difference between partial and total derivatives when the function depends on only one variable. In the multivariate case, this is different. We here use the  $\partial_x$  notation when the derivative is straightforward to calculate, as for the gradient of a loss function, and  $d_x$  when the function has some hidden dependency on  $x$ , as it occurs for implicit functions. We consider both partial and total derivatives of scalar functions to be row vectors, so that  $\partial_\theta L^{\text{out}}$  is a row vector of size  $|\theta|$ ,  $\partial_\phi L^{\text{out}}$  a row vector of size  $|\phi|$ ,  $\partial_\phi^2 L^{\text{in}}$  a squared matrix of size  $|\phi| \times |\phi|$ , and  $\partial_\theta \partial_\phi L^{\text{in}}$  a matrix of size  $|\phi| \times |\theta|$ .

Using the outer gradient  $\nabla_\theta$  for gradient descent would in principle yield an efficient algorithm to solve our bilevel optimization problem with the nice property that it requires knowing only  $\phi_\theta^*$ . Unlike backpropagation-through-time, storing the sequence of intermediate parameter values generated by the learning algorithm is no longer needed. However, computing the outer gradient requires computing the Hessian  $\partial_\phi^2 L^{\text{in}}(\phi_\theta^*, \theta)$ , a second-order derivative, and inverting it. Those two operations are costly and often intractable in large-scale machine learning problems. We therefore need to approximate the outer gradient  $\nabla_\theta$  if we want to use it for practical purposes. This is what the methods we present in section 4 do.

The rest of the section is dedicated to explaining in further detail the statements and consequences of the implicit function theorem for our bilevel optimization problem. It can be skipped on a first reading without impairing the understanding of the rest of the review.

The usual formulation of the implicit function theorem (Dontchev & Rockafellar, 2009) encompasses both the existence statement and the differentiation formula. We present and discuss next the two parts separately for clarity.

**3.1 Existence of Implicit Functions †.** In the bilevel optimization formulation, equation 2.1, we used the implicit function  $\phi_\theta^*$  without ensuring that it is correctly defined. The first part of the implicit function theorem ensures that such a function exists.

**Theorem 1** (Existence of implicit functions; Dontchev & Rockafellar, 2009). *Let  $f$  be continuously differentiable and  $(\bar{\phi}, \bar{\theta})$  be such that  $f(\bar{\phi}, \bar{\theta}) = 0$ . If the Jacobian matrix  $\partial_\phi f(\bar{\phi}, \bar{\theta})$  is invertible, there exists a unique continuous implicit function  $\theta \mapsto \phi_\theta^*$  defined in a neighborhood of  $\bar{\theta}$  such that  $\phi_{\bar{\theta}}^* = \bar{\phi}$  and which verifies  $f(\phi_\theta^*, \theta) = 0$  for all  $\theta$  in that neighborhood.*

Once applied to the constraint  $f = \partial_\phi L^{\text{in}}$  that follows from the local minimality constraint in equation 2.1, the invertibility condition becomes an invertibility condition on the Hessian  $\partial_\phi^2 L^{\text{in}}(\phi_\theta^*, \theta)$  and the implicit function verifies  $\partial_\phi L^{\text{in}}(\phi_\theta^*, \theta) = 0$  on the neighborhood on which it is defined. Note that if  $\bar{\phi}$  is a minimizer of  $L^{\text{in}}(\cdot, \bar{\theta})$ , then  $\phi_{\bar{\theta}}^*$  will also be as long as

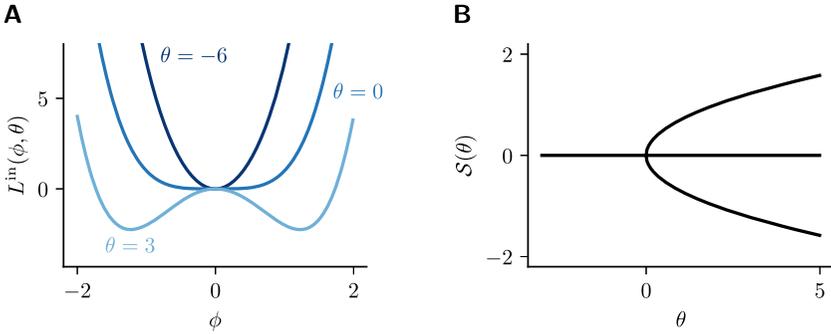


Figure 1: (A) Visualization of the function  $L^{\text{in}} : (\phi, \theta) \mapsto \phi^4 - \theta \phi^2$  for several  $\theta$  values. (B) When  $\theta$  equals 0, the Hessian at  $\phi = 0$  is noninvertible, which implies that there is no implicit function defined around  $\phi = 0$ , as shown on the graph of the solution mapping  $S(\theta) := \{\phi \mid \partial_\phi L^{\text{in}}(\phi, \theta) = 0\}$  associated with the equilibrium condition  $\partial_\phi L^{\text{in}}(\phi, \theta) = 0$ .

$L^{\text{in}}$  is twice continuously differentiable.<sup>3</sup> The implicit function theorem is purely local in the sense that several implicit functions can cohabit for a given  $\theta$  but in different regions of the  $\phi$  space, as shown on Figure 1B for  $\theta > 0$ . This is why we use the notation  $\phi_\theta^* \in \arg \min_\phi L^{\text{in}}(\phi, \theta)$  in (2.1): the problem is still well defined even if there exist several local minima for the same  $\theta$ .

The main assumption of theorem 1 applied to bilevel optimization is the invertibility of the Hessian at  $(\bar{\phi}, \bar{\theta})$ . Without this assumption, the graph associated with the minimizers can split, as illustrated in the following example. Let  $L^{\text{in}}(\phi, \theta) := \phi^4 - \theta \phi^2$  for  $\phi$  and  $\theta$  real variables. We plot the graph of this function for several  $\theta$  values on Figure 1A. The Hessian, here a second-order derivative, is null when  $(\phi, \theta) = (0, 0)$  (hence not invertible). A branching behavior occurs at this point since there exists a unique minimizer (which is also the only stationary point) of the function at  $\phi = 0$  when  $\theta$  is negative and 3 otherwise (see Figure 1B). The graph associated with the implicit functions therefore splits into three branches at  $\theta = 0$ , making it impossible to properly define an implicit function in this neighborhood.

**3.2 Analytical Formula for the Outer Gradient †.** Once we know that an implicit function exists, we would like to know how it locally reacts to

<sup>3</sup>This can be obtained by remarking that the smallest eigenvalue of an invertible Hessian is strictly positive and the smallest eigenvalue of  $\partial_\phi^2 L^{\text{in}}(\phi_\theta^*, \theta)$  is a continuous function of  $\theta$ . This implies that for  $\theta$  in the neighborhood of  $\bar{\theta}$  considered in Theorem 1, the smallest eigenvalue of  $\partial_\phi^2 L^{\text{in}}(\phi_\theta^*, \theta)$  is strictly positive and hence that  $\phi_\theta^*$  is a local minimizer of  $L^{\text{in}}$  for every  $\theta$  in this neighborhood.

changes in  $\theta$ , that is, if it is differentiable and, if so, what its derivative is. This is what the second part of the implicit function theorem brings.

**Theorem 2** (Differentiating implicit functions; Dontchev & Rockafellar, 2009). *Under the assumptions of theorem 1, the implicit function  $\phi_\theta^*$  defined in that theorem is differentiable and*

$$\frac{d\phi_\theta^*}{d\theta} = - \left( \frac{\partial f}{\partial \phi}(\phi_\theta^*, \theta) \right)^{-1} \frac{\partial f}{\partial \theta}(\phi_\theta^*, \theta).$$

**Proof.** The derivation of the previous formula is relatively straightforward once we know the differentiable implicit function exists as it only requires differentiating through the constraint using the chain rule: as  $f(\phi_\theta^*, \theta) = 0$  for all  $\theta$  on which  $\phi_\theta^*$  is defined, we have

$$\begin{aligned} 0 &= \frac{d}{d\theta} f(\phi_\theta^*, \theta) \\ &= \frac{\partial f}{\partial \theta}(\phi_\theta^*, \theta) + \frac{\partial f}{\partial \phi}(\phi_\theta^*, \theta) \frac{d\phi_\theta^*}{d\theta}, \end{aligned}$$

which yields the desired formula after rearranging the different terms.  $\square$

We can then use this formula to obtain  $d_\theta \phi_\theta^*$  for our bilevel optimization problem:

$$\frac{d\phi_\theta^*}{d\theta} = - \left( \frac{\partial^2 L^{\text{in}}}{\partial \phi^2}(\phi_\theta^*, \theta) \right)^{-1} \frac{\partial^2 L^{\text{in}}}{\partial \theta \partial \phi}(\phi_\theta^*, \theta). \tag{3.2}$$

Together with the chain rule, this is just what we need to obtain an expression for the outer gradient:

$$\begin{aligned} \nabla_\theta^\top &= \frac{d}{d\theta} L^{\text{out}}(\phi_\theta^*, \theta) \\ &= \frac{\partial L^{\text{out}}}{\partial \theta}(\phi_\theta^*, \theta) + \frac{\partial L^{\text{out}}}{\partial \phi}(\phi_\theta^*, \theta) \frac{d\phi_\theta^*}{d\theta} \\ &= \frac{\partial L^{\text{out}}}{\partial \theta}(\phi_\theta^*, \theta) - \frac{\partial L^{\text{out}}}{\partial \phi}(\phi_\theta^*, \theta) \left( \frac{\partial^2 L^{\text{in}}}{\partial \phi^2}(\phi_\theta^*, \theta) \right)^{-1} \frac{\partial^2 L^{\text{in}}}{\partial \theta \partial \phi}(\phi_\theta^*, \theta). \end{aligned} \tag{3.3}$$

Note that the implicit function theorem and all methods we present here only require the stationarity condition,  $\partial_\phi L^{\text{in}}(\phi_\theta^*, \theta) = 0$ , to be satisfied, and not the more restrictive minimality assumption,  $\phi_\theta^* \in \arg \min_\phi L^{\text{in}}(\phi, \theta)$ . The

methods we introduce in the next section can therefore be easily be extended to solve any optimization problem of the form

$$\begin{aligned} \min_{\theta} L^{\text{out}}(\phi_{\theta}^*, \theta) \\ \text{s.t. } \partial_{\phi} L^{\text{in}}(\phi_{\theta}^*, \theta) = 0. \end{aligned} \quad (3.4)$$

#### 4 Approximations of the Outer Gradient

---

As we mentioned in the last section, computing the outer gradient using its analytical formula,

$$\nabla_{\theta}^{\top} = \frac{\partial L^{\text{out}}}{\partial \theta}(\phi_{\theta}^*, \theta) - \frac{\partial L^{\text{out}}}{\partial \phi}(\phi_{\theta}^*, \theta) \left( \frac{\partial^2 L^{\text{in}}}{\partial \phi^2}(\phi_{\theta}^*, \theta) \right)^{-1} \frac{\partial^2 L^{\text{in}}}{\partial \theta \partial \phi}(\phi_{\theta}^*, \theta), \quad (4.1)$$

is not feasible in most practical applications of bilevel optimization: we need approximations. Different methods exist to do so. We classify them into two different categories: implicit differentiation methods that approximate the outer gradient by directly using the analytical formula equation 4.1, obtained with the implicit function theorem and equilibrium propagation, which leverages an alternative formulation for the outer gradient that we will later present. Note that we have here written the derivative with respect to outer parameters  $\theta$ , but everything can be transposed to derivatives with respect to inputs, thus allowing us to backpropagate through implicitly defined layers in deep architectures (Amos & Kolter, 2017; Gould et al., 2021).

In the following, we provide the intuition behind the different methods, exhibit their fundamental similarities and differences, and compare their theoretical guarantees.

##### 4.1 Implicit Differentiation.

*4.1.1 Gradient Computation as Minimization of a Quadratic Form.* Computing and inverting Hessians are costly operations (respectively, quadratic and cubic in the size of the differentiated parameter) so the inverse Hessian term in equation 4.1 must often be approximated in practice. The first key insight needed for those methods is to iteratively approximate the row vector,

$$\pi^* := \frac{\partial L^{\text{out}}}{\partial \phi}(\phi_{\theta}^*, \theta) \left( \frac{\partial^2 L^{\text{in}}}{\partial \phi^2}(\phi_{\theta}^*, \theta) \right)^{-1}, \quad (4.2)$$

by minimizing the quadratic form,

$$\pi \mapsto \frac{1}{2} \pi \frac{\partial^2 L^{\text{in}}}{\partial \phi^2}(\phi_\theta^*, \theta) \pi^\top - \pi \frac{\partial L^{\text{out}}}{\partial \phi}(\phi_\theta^*, \theta)^\top. \tag{4.3}$$

As we are using row vectors, the quantity  $\pi \frac{\partial L^{\text{out}}}{\partial \phi}(\phi_\theta^*, \theta)^\top$  corresponds to a dot product. If  $\phi_\theta^*$  is a nonflat local minimizer of  $L^{\text{in}}$ , then the invertible Hessian condition needed in theorem 1 is satisfied and the quadratic form, equation 4.3, is positive definite so it has a unique minimizer, which is  $\pi^*$ .

*4.1.2 Choice of the Optimizer.* Naively minimizing the quadratic form (4.3) does not yet lead to a practical algorithm. Let us take the example of gradient descent. An update would take the form

$$\pi \leftarrow \pi - \alpha \left( \pi \frac{\partial^2 L^{\text{in}}}{\partial \phi^2}(\phi_\theta^*, \theta) - \frac{\partial L^{\text{out}}}{\partial \phi}(\phi_\theta^*, \theta) \right) \tag{4.4}$$

with  $\alpha$  the learning rate. Evaluating equation 4.4 appears to require computing the Hessian  $\frac{\partial^2 L^{\text{in}}}{\partial \phi^2}$  and then multiplying it with the vector  $\pi$ , an operation with quadratic complexity that would render the method impractical. However, there is a way of obtaining the update above without ever having to explicitly calculate the Hessian: a cleverer implementation exploits the fact that all we need is a Hessian-vector product. Remarkably, computing such products has the same complexity as computing gradients (Pearlmutter, 1994). Gradient descent, and in fact many other optimization procedures, can therefore be executed efficiently. We call this process the second phase, whereas the first phase consists of computing  $\phi_\theta^*$ .

Implicit differentiation methods take different forms depending on the choice of the optimizer. When gradient descent is chosen as in equation 4.4, this leads to recurrent backpropagation, also known as the Almeida-Pineda algorithm (Almeida, 1990; Pineda, 1987).<sup>4</sup> The very same update of recurrent backpropagation can be obtained from different perspectives. For example, it can be derived starting from the Neumann series formulation of the inverse of a matrix (Liao et al., 2018; Lorraine, Vicol, & Duvenaud,

---

<sup>4</sup>The usual way of deriving recurrent backpropagation is by using iterative updates to find the solution of the linear system  $\pi \frac{\partial^2 L^{\text{in}}}{\partial \phi^2}(\hat{\phi}, \theta) = \frac{\partial L^{\text{out}}}{\partial \phi}(\hat{\phi}, \theta)$ . Although this is equivalent to gradient descent on the quadratic form when applied to bilevel optimization, this view allows considering the more general case in which equilibrium states are not necessarily minimizers of a loss function. Here, we use the quadratic form minimization view as it makes the comparison to other methods easier.

2020). We have

$$\begin{aligned} \left( \frac{\partial^2 L^{\text{in}}}{\partial \phi^2}(\phi_\theta^*, \theta) \right)^{-1} &= \alpha \left( \alpha \frac{\partial^2 L^{\text{in}}}{\partial \phi^2}(\phi_\theta^*, \theta) \right)^{-1} \\ &= \alpha \sum_{i=0}^{\infty} \left( \text{Id} - \alpha \frac{\partial^2 L^{\text{in}}}{\partial \phi^2}(\phi_\theta^*, \theta) \right)^i \end{aligned} \quad (4.5)$$

whenever the absolute eigenvalues of  $(\text{Id} - \alpha \frac{\partial^2 L^{\text{in}}}{\partial \phi^2}(\phi_\theta^*, \theta))$  are strictly smaller than one (which requires  $\alpha$  small enough). We cannot use this formula alone as it still requires computing the Hessian, but we can use it to iteratively approximate  $\pi^*$  using

$$\pi \leftarrow \pi \left( \text{Id} - \alpha \frac{\partial^2 L^{\text{in}}}{\partial \phi^2}(\phi_\theta^*, \theta) \right) + \alpha \frac{\partial L^{\text{out}}}{\partial \phi}(\phi_\theta^*, \theta), \quad (4.6)$$

which is exactly the same update as equation 4.4. This is why we used the same notation for the learning rate and the rescaling parameter even though we introduced those two parameters from different contexts. Alternatively, this update also appears in truncated backpropagation (Williams & Peng, 1990) when gradient descent on  $L^{\text{in}}$  has reached a minimum for several steps (Shaban, Cheng, Hatch, & Boots, 2019). Backpropagating through the last iteration takes exactly the same form as equation 4.4, but it requires storing the intermediate states in memory as opposed to recurrent backpropagation.

Gradient descent is a very general algorithm. Since we want to minimize a specific class of function, one may ask whether more tailored optimization procedures might be more efficient. This is what the conjugate gradient method provides (we refer to Shewchuk, 1994, for more details on the algorithm), while still only requiring Hessian-vector products.

Note that we can obtain first-order approximations of the outer gradient by limiting the number of steps in the second phase. If we skip the second minimization and approximate the result by  $\hat{\pi} = 0$ , the corresponding approximate outer gradient will be equal to the direct derivative  $\partial_\theta L^{\text{out}}(\phi_\theta^*, \theta)$ . If we perform only one step and take  $\alpha = 1$ , we approximate the Hessian with the identity (Luketina, Berglund, Greff, & Raiko, 2016). The amount of computing attributed to the second phase therefore progressively transforms a first-order approximation toward the true value of the gradient.

**4.1.3 Some Practical Considerations.** In practice, we rarely directly minimize equation 4.3 as we do not have access to an exact minimizer  $\phi_\theta^*$  of the inner loss, only to an estimate  $\hat{\phi}$ . Instead, we use the estimated version of the quadratic form (replacing  $\phi_\theta^*$  by  $\hat{\phi}$ ) as shown in algorithm 1.

---

**Algorithm 1:** Side-by-Side Comparison of Implicit Differentiation (Section 4.1) and Equilibrium Propagation (Section 4.2) Methods.

---

**Result:** Approximate solution  $\theta$  of the bilevel optimization problem 2.1

**for**  $i = 1, \dots, n$  **do**

Minimize  $L^{\text{in}}(\phi, \theta)$  with respect to  $\phi$  and note  $\hat{\phi}$  the approximate result;

**Implicit differentiation**

Minimize the quadratic form

$$\pi \mapsto \frac{1}{2} \pi \frac{\partial^2 L^{\text{in}}}{\partial \phi^2}(\hat{\phi}, \theta) \pi^\top - \pi \frac{\partial L^{\text{out}}}{\partial \phi}(\hat{\phi}, \theta)^\top$$

through, e.g., gradient descent or conjugate gradient, and note  $\hat{\pi}$  the approximate result

Estimate the outer gradient with

$$\hat{\nabla}_\theta^\top := \frac{\partial L^{\text{out}}}{\partial \theta}(\hat{\phi}, \theta) - \hat{\pi} \frac{\partial^2 L^{\text{in}}}{\partial \theta \partial \phi}(\hat{\phi}, \theta)$$

**Equilibrium propagation**

Minimize  $\mathcal{L}(\phi, \theta, \beta) = L^{\text{in}}(\phi, \theta) + \beta L^{\text{out}}(\phi, \theta)$  with respect to  $\phi$  for some small nonzero  $\beta$  value (potentially for more  $\beta$  values if needed), starting from  $\hat{\phi}$ , and note  $\hat{\phi}_\beta$  the approximate result

Estimate the outer gradient with

$$\hat{\nabla}_\theta^\top := \frac{1}{\beta} \left( \frac{\partial \mathcal{L}}{\partial \theta}(\hat{\phi}_\beta, \theta, \beta) - \frac{\partial \mathcal{L}}{\partial \theta}(\hat{\phi}, \theta, 0) \right)$$

or with an estimator that uses more points

Update  $\theta$  using  $\hat{\nabla}_\theta$ ;

---

In many applications,  $L^{\text{in}}$  is an empirical risk, that is, the average of some loss evaluated on many different data samples. In this case, it might not be possible to compute Hessian-vector products for all the data at once. To work around this issue, we can resort to stochastic updates on the quadratic (taking a random subset of the data for each step), as noted in the lecture notes of Grosse (2021).

*4.1.4 Robustness to Nonoptimality †.* The local minimizer  $\phi_\theta^*$  is almost always approximated in practice. A natural question to ask is whether the methods introduced above are robust to this approximation. In other words, we may ask how good  $\hat{\nabla}_\theta$  is compared to  $\nabla_\theta$ , with

$$\hat{\nabla}_\theta^\top = \frac{\partial L^{\text{out}}}{\partial \theta}(\hat{\phi}, \theta) - \hat{\pi} \frac{\partial^2 L^{\text{in}}}{\partial \theta \partial \phi}(\hat{\phi}, \theta), \tag{4.7}$$

as in algorithm 1. In equation 4.7,  $\hat{\pi}$  is obtained by iteratively minimizing the estimated version of the quadratic form, equation 4.3, that consists

of replacing  $\phi_\theta^*$  by  $\hat{\phi}$ . Its estimation will therefore be the other source of approximation.

We are doing approximate gradient descent at the outer level, which will result in approximate solutions to the bilevel optimization problem. d'Aspremont (2008) and Friedlander & Schmidt (2012) have shown that the error made in solving a convex optimization problem with inexact gradients can be linked to the gradient approximation error  $\|\widehat{\nabla}_\theta - \nabla_\theta\|$ . Motivated by those results, we present a theoretical bound on the error made in estimating the outer gradient  $\nabla_\theta$  with  $\widehat{\nabla}_\theta$  depending on the quality of  $\hat{\phi}$  and  $\hat{\pi}$ .

**Assumption 1.** Suppose that there exist positive real numbers  $(\mu, \rho, B, L, M)$  such that:

- i.  $L^{\text{in}}$  is twice continuously differentiable, and  $L^{\text{out}}$  is continuously differentiable.
- ii.  $L^{\text{in}}$  is  $\mu$ -strongly convex as a function of  $\phi$ .
- iii. The second-order derivatives (Hessian and cross-derivatives) of  $L^{\text{in}}$  are  $\rho$ -Lipschitz as functions of  $\phi$ .
- iv. As functions of  $\phi$ ,  $L^{\text{out}}$  is  $B$ -Lipschitz,  $L$ -smooth and  $\partial_\theta L^{\text{out}}$  is  $M$ -Lipschitz.

**Theorem 3** (Error bound for implicit differentiation methods; Pedregosa, 2016). *Let  $\phi_\theta^*$  be a minimizer of  $L^{\text{in}}$  and  $\hat{\phi}$  be its approximated value. Let  $\delta \in ]0, \frac{\mu}{2\rho}[$  be an upper bound on the corresponding approximation error:*

$$\|\phi_\theta^* - \hat{\phi}\| \leq \delta.$$

*Let  $\hat{\pi}$  be an approximation of  $\pi^*$  computed by one of the implicit differentiation methods and  $\delta' > 0$  be an upper bound of its approximation error:*

$$\|\hat{\pi} - \partial_\phi L^{\text{out}}(\hat{\phi}, \theta) \partial_\phi^2 L^{\text{in}}(\hat{\phi}, \theta)^{-1}\| \leq \delta'.$$

*Then, under assumption 1, there exists a constant  $C$  such that*

$$\|\nabla_\theta - \widehat{\nabla}_\theta\| \leq C(\delta + \delta').$$

The quantities  $\delta$  and  $\delta'$  measure the error made in the two phases of the algorithms, where the first phase consists in finding a minimum of  $L^{\text{in}}$  and the second one in minimizing the local quadratic form. Theorem 3 shows that the approximation error in the outer gradient grows linearly with those two errors. Assumption 1 ensures that the problem we are considering and its derivatives are well defined (i and ii) and that  $L^{\text{in}}$  and  $L^{\text{out}}$  are regular enough (iii and iv).

We present a proof of theorem 3, along with a discussion on how to transform the global convexity assumptions into local ones in section B.2.

**4.2 Equilibrium Propagation.** Instead of differentiating through the implicit functions, we can resort to the mathematical result known as equilibrium propagation (Scellier & Bengio, 2017), which reformulates the outer gradient in a way that is easier to estimate numerically. While equilibrium propagation was originally presented in the context of energy-based recurrent neural network learning, the result is far more general. As we discuss next, equilibrium propagation can be applied to solve general bilevel optimization problems.

*4.2.1 Equilibrium Propagation Theorem.* The first step in equilibrium propagation consists of breaking up the hierarchy of losses and mixing  $L^{\text{in}}$  and  $L^{\text{out}}$  in an augmented loss:

$$\mathcal{L}(\phi, \theta, \beta) := L^{\text{in}}(\phi, \theta) + \beta L^{\text{out}}(\phi, \theta). \tag{4.8}$$

The nudging strength  $\beta$  is a scalar that controls the strength of the mix; when it is equal to 0, we retrieve the inner learning problem. We denote by  $\phi_{\theta, \beta}^* \in \arg \min_{\phi} \mathcal{L}(\phi, \theta, \beta)$  the different minimizers of  $\mathcal{L}$ . We can now introduce the equilibrium propagation result.

**Theorem 4** (Equilibrium propagation; Scellier & Bengio, 2017; Scellier, 2021). *Let  $L^{\text{in}}$  and  $L^{\text{out}}$  be two twice continuously differentiable functions. Let  $\bar{\phi}$  be a stationary point of  $\mathcal{L}(\cdot, \bar{\theta}, \bar{\beta})$ ,*

$$\frac{\partial \mathcal{L}}{\partial \phi}(\bar{\phi}, \bar{\theta}, \bar{\beta}) = 0,$$

*such that  $\partial_{\phi}^2 \mathcal{L}(\bar{\phi}, \bar{\theta}, \bar{\beta})$  is invertible. Then there exists a neighborhood of  $(\bar{\theta}, \bar{\beta})$  and a continuously differentiable function  $(\theta, \beta) \mapsto \phi_{\theta, \beta}^*$  such that  $\phi_{\bar{\theta}, \bar{\beta}}^* = \bar{\phi}$ , and for every  $(\theta, \beta)$  in this neighborhood, we have*

$$\frac{\partial \mathcal{L}}{\partial \phi}(\phi_{\theta, \beta}^*, \theta, \beta) = 0$$

and

$$\frac{d}{d\theta} \frac{\partial \mathcal{L}}{\partial \beta}(\phi_{\theta, \beta}^*, \theta, \beta) = \frac{d}{d\beta} \frac{\partial \mathcal{L}}{\partial \theta}(\phi_{\theta, \beta}^*, \theta, \beta)^{\top}.$$

**Proof.** The existence part in the equilibrium propagation theorem directly follows from theorem 1 using  $f = \partial_{\phi} \mathcal{L}$ . Obtaining the differentiation formula is not as complicated as it may appear at first glance. The first step consists of applying the symmetry of second-order derivative result, also known as Schwarz’s theorem:

$$\frac{d}{d\theta} \frac{d}{d\beta} \mathcal{L}(\phi_{\theta, \beta}^*, \theta, \beta) = \frac{d}{d\beta} \frac{d}{d\theta} \mathcal{L}(\phi_{\theta, \beta}^*, \theta, \beta)^\top.$$

We then apply the chain rule on both sides of the previous equation and use the equilibrium condition  $\partial_\phi \mathcal{L}(\phi_{\theta, \beta}^*, \theta, \beta) = 0$  to simplify the derivatives. For the left-hand side of the previous equation, it yields

$$\begin{aligned} \frac{d}{d\theta} \frac{d}{d\beta} \mathcal{L}(\phi_{\theta, \beta}^*, \theta, \beta) &= \frac{d}{d\theta} \left[ \frac{\partial \mathcal{L}}{\partial \beta}(\phi_{\theta, \beta}^*, \theta, \beta) + \frac{\partial \mathcal{L}}{\partial \phi}(\phi_{\theta, \beta}^*, \theta, \beta) \frac{d\phi_{\theta, \beta}^*}{d\beta} \right] \\ &= \frac{d}{d\theta} \frac{\partial \mathcal{L}}{\partial \beta}(\phi_{\theta, \beta}^*, \theta, \beta). \end{aligned}$$

The right-hand side can be simplified in the same way, which gives the desired formula.  $\square$

The equilibrium propagation result can be used to reformulate the outer gradient  $\nabla_\theta$  by remarking that  $\partial_\beta \mathcal{L} = L^{\text{out}}$  and  $\phi_{\theta, \beta}^*|_{\beta=0} = \phi_\theta^*$ . We then have

$$\nabla_\theta = \frac{d}{d\beta} \frac{\partial \mathcal{L}}{\partial \theta}(\phi_{\theta, \beta}^*, \theta, \beta)|_{\beta=0}. \quad (4.9)$$

Theorem 4 uses a stationary condition on the vector  $\phi$ , but more general versions of equilibrium propagation exist for stationary distributions or trajectories (see Scellier, 2021, for more details). A very similar gradient estimate has also been derived when  $\phi$  is a discrete quantity, and the inner and outer losses are expectations measured over a continuous distribution (Hazan, Keshet, & McAllester, 2010; Song, Schwing, Zemel, & Urtasun, 2016).

There is a deep connection between the equilibrium propagation and implicit differentiation approaches: the quantity  $\pi^*$  that we defined in equation 4.2 is actually indirectly computed in equilibrium propagation, since  $\pi^* = d_\beta \phi_{\theta, \beta}^*|_{\beta=0}$ . The trajectories in the second phases of equilibrium propagation and implicit differentiation methods can also be shown to be closely related, when gradient descent is used in the second phase for the two methods (Scellier & Bengio, 2019).

*4.2.2 Numerical Estimation of  $\nabla_\theta$ .* The formula provided by the equilibrium propagation theorem might not appear useful at first. Closer inspection, however, reveals that it offers a new way of numerically estimating the outer gradient. The outer gradient is a derivative of a scalar function with respect to a vector  $\theta$  and is thus hard to estimate numerically, in particular when  $\theta$  is high-dimensional. By contrast, the right-hand side of equation 4.9 is the derivative of a vector-valued function with respect to a scalar, which

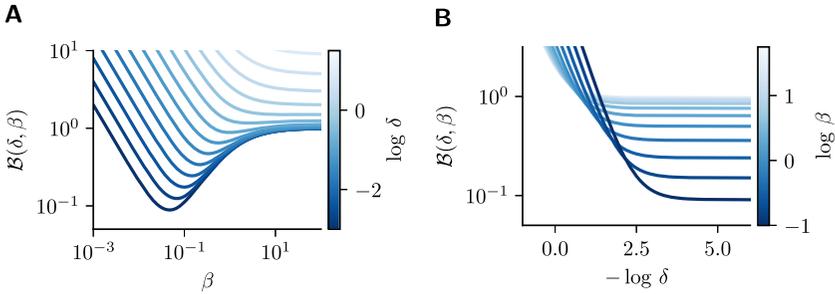


Figure 2: Visualization of the bound  $\mathcal{B}$  on the gradient estimation error obtained in theorem 5, as a function of  $\beta$  (A) and as a function of  $\delta = \delta'$  (B) (adapted from Zuchet et al., 2021).

can be readily estimated with finite difference techniques. The simplest finite difference estimator is

$$\widehat{\nabla}_\theta^\top := \frac{1}{\beta} \left( \frac{\partial \mathcal{L}}{\partial \theta}(\hat{\phi}_\beta, \theta, \beta) - \frac{\partial \mathcal{L}}{\partial \theta}(\hat{\phi}_0, \theta, 0) \right), \tag{4.10}$$

where  $\hat{\phi}_0$  and  $\hat{\phi}_\beta$  are the approximated values of  $\phi_{\theta,0}^*$  and  $\phi_{\theta,\beta}^*$ . This formula yields a two-phase algorithm that is detailed in algorithm 1. The approximation of the outer gradient can be refined by adding more points to the estimator, for instance, by resorting to the central or forward finite difference estimators. The idea is to collect the value of  $\partial_\theta \mathcal{L}$  at different values, for example,  $-\beta$  and  $\beta$  for the central one with three points (as in Laborieux et al., 2021) or  $0, \beta, 2\beta, \dots$  for the forward ones. We provide more details in appendix C.

**4.2.3 Robustness to Nonoptimality †.** As for implicit differentiation methods, it is possible to bound the error made by the two-point equilibrium propagation estimator, equation 4.10. There are two sources of error: the approximation of the two solutions and the one rooted in the finite difference scheme. When  $\beta$  gets smaller, the finite difference error gets smaller. On the other side, decreasing  $\beta$  increases the sensitivity of the estimation to noise or inaccurate minimizations. Theorem 5 quantifies it; we visualize the result on Figure 2.

**Assumption 2.** Assume that  $L^{\text{in}}$  and  $L^{\text{out}}$  are three-times continuously differentiable. Additionally, suppose that there exist positive real numbers  $(B^{\text{in}}, B^{\text{out}}, L, \mu, \rho, \sigma)$  such that  $L^{\text{in}}$  and  $L^{\text{out}}$ , as functions of  $\phi$ , verify the following properties:

- i.  $\partial_\theta L^{\text{in}}$  is  $B^{\text{in}}$ -Lipschitz and  $\partial_\theta L^{\text{out}}$  is  $B^{\text{out}}$ -Lipschitz.
- ii.  $L^{\text{in}}$  and  $L^{\text{out}}$  are  $L$ -smooth and  $\mu$ -strongly convex.

- iii. Their Hessians are  $\rho$ -Lipschitz.
- iv.  $\partial_{\phi} \partial_{\theta} L^{\text{in}}$  and  $\partial_{\phi} \partial_{\theta} L^{\text{out}}$  are  $\sigma$ -Lipschitz.

**Theorem 5** (Error bound for equilibrium propagation; Zucchet, Schug, von Oswald, Zhao, & Sacramento, 2021). *Let  $\beta > 0$  and  $(\delta, \delta')$  be such that*

$$\|\hat{\phi}_0 - \phi_{\theta,0}^*\| \leq \delta$$

and

$$\|\hat{\phi}_{\beta} - \phi_{\theta,\beta}^*\| \leq \delta'.$$

Then, under assumption 2, there exists a  $\theta$ -dependent constant  $C$  such that

$$\|\widehat{\nabla}_{\theta} - \nabla_{\theta}\| \leq \frac{B^{\text{in}}(\delta + \delta')}{\beta} + B^{\text{out}}\delta' + C \frac{\beta}{1 + \beta} =: \mathcal{B}(\delta, \delta', \beta).$$

As for implicit differentiation methods, we can obtain a more local version of theorem 5 by replacing the strong convexity assumption of  $L^{\text{in}}$  by a non-flat minimum assumption. The behavior of the estimator when the Hessian of  $L^{\text{in}}$  at  $\hat{\phi}$  is not positive definite is, however, quite different from the other kind of methods. Assuming that  $L^{\text{in}}$  is bounded from below, the second phase ends up in a nearby basin of attraction in the worst case. The gradient estimator will then converge to some finite value, as opposed to implicit differentiation methods that will diverge.

**4.2.4 Comparison with Implicit Differentiation Methods †.** It is not yet possible to compare the bounds from theorems 3 and 5 as the bound for equilibrium propagation is still  $\beta$ -dependent. We can remove this dependency through the following corollary.

**Corollary 1** (Corollary of theorem 5, Zucchet et al., 2021). *Under assumption 2, if we additionally suppose that for every  $\beta > 0$ , we approximate the two solutions with precision  $\delta$  and  $\delta'$  and if  $(\delta + \delta') < C/B^{\text{in}}$ , the best achievable bound in theorem 5 is smaller than*

$$B^{\text{out}}\delta' + 2\sqrt{CB^{\text{in}}(\delta + \delta')}.$$

The error made in the two-point equilibrium propagation estimator is therefore  $O(\sqrt{\delta + \delta'})$ , which implies that implicit differentiation methods are theoretically less sensitive to approximations in the two phases than equilibrium propagation. We compare in more detail the two approaches in the next section.

## 5 Comparison of the Different Approaches

---

Now that we have introduced implicit methods for bilevel optimization, the questions that come next are in which conditions they are useful and which one to pick. The objective of this section is to provide insight into where the different methods shine but not to give a definitive answer to such questions.

**5.1 Alternative Methods.** Backpropagation through time (Werbos, 1990) can be used to compute gradients when the process used to estimate  $\phi_\theta^*$  is a sequence of differentiable operations. In most settings, it is impossible to store the entire sequence of intermediate parameters produced by the algorithm in memory. The standard workaround to this problem is to run (truncate) the backward pass for a limited number of steps (Jaeger, 2002; Shaban et al., 2019) or to use a checkpointing strategy (Gruslys, Munos, Danihelka, Lanctot, & Graves, 2016). Whenever backpropagation or its truncated version is applicable, it is often a strong alternative to the implicit methods studied here; it is difficult to rule out a priori one class of methods over the other without experimenting with both.

There are, however, a number of clearly identifiable scenarios in which the methods discussed in this review may be preferable. Perhaps most important, it is not always possible to write the underlying optimization algorithm as a differentiable program. For example, an algorithmic solver can provably minimize a smooth loss function, but its inner process is not necessarily differentiable. In such cases, automatic differentiation is not an option, and implicit methods are in general the only gradient-based methods available. Furthermore, even when the learning algorithm is technically differentiable, it may generate chaotic sequences of parameters, which render gradients extremely noisy (Metz et al., 2019). In such situations, the methods studied here may lead to an implicit form of regularization of the learning process by selecting outer parameters that are less prone to inducing chaos. More work is needed to investigate this hypothesis.

There is growing interest in physically plausible learning algorithms, where optimization is performed by a physical system evolving in time (Millar, 1951; Kendall et al., 2020; Stern et al., 2021; Scellier, Mishra, Bengio, & Ollivier, 2022). It is generally impossible to implement backpropagation in such systems, as this would entail going back in (physical) time; even for reversible processes, it is difficult to conceive backpropagation through time, since the computations performed in the forward and backward phases of this algorithm are not the same. Provided that the process that governs the time evolution of the parameters is differentiable, forward differentiation (also known as real-time recurrent learning; see Williams & Zipser, 1989) is the classic alternative to backpropagation that avoids going backward in time. However, in its original form, this algorithm is typically infeasible to implement as well. First, its memory requirements scale

Table 1: Summary of the Comparison between Methods Rooted in Implicit Differentiation.

	Efficiency	Robustness	Simplicity
First-order approximation	+	+++	+++
Recurrent backpropagation	++	++	+
Conjugate gradients	+++	+	+
Equilibrium propagation	++	++	++

with the dimension of  $\theta$  multiplied by the dimension of  $\phi$ , which results in a huge memory cost. This is in fact a concern for most standard computer implementations as well. Second, the algorithm requires computing Jacobian-vector products, which may or may not be difficult to calculate in physical systems. Most of the concerns outlined above apply equally when looking at backpropagation or forward differentiation as biological learning algorithms. Some approximations have been developed to circumvent these limitations (Sutton, 1992; Tallec & Ollivier, 2018; Bellec et al., 2020; Marschall, Cho, & Savin, 2020; Menick et al., 2021).

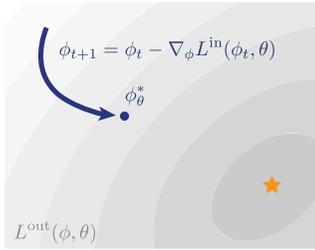
**5.2 Comparison of the Different Implicit Methods.** Finally, we compare the methods presented in the previous sections. More concretely, we consider methods that use first-order (FO) approximations of the outer gradient,<sup>5</sup> equilibrium propagation (EP), recurrent backpropagation (RBP), and the conjugate gradient (CG) method. We determine use cases for the different methods based on three criteria: efficiency when all the theoretical assumptions are met, robustness to violation of the assumptions, and simplicity of the methods in terms of the computational elements involved. The result of the comparison is summarized in Table 1, and a visual comparison of the algorithm is provided in Figure 3.

*5.2.1 Efficiency under Met Assumptions.* We build our comparison on the theoretical analysis presented in the previous section (theorem 3 and corollary 1) and want to figure out which method produces the best estimate. It assumes that  $\hat{\phi}$  is sufficiently close to  $\phi_0^*$  so that all the methods are properly defined.

First-order methods here suffer as the estimation they provide cannot be refined to get closer to the outer gradient. For the remaining methods, we have to consider two things: how sensitive the gradient estimate is to the approximations made in the two phases and which optimizers are used. Regarding the first point, implicit differentiation methods (CG and RBP)

<sup>5</sup>Note that as we mentioned in section 4.1, those methods can be seen as implicit differentiation methods with extremely short second phases.

**A** Implicit differentiation



$$\pi \mapsto \frac{1}{2} \pi H \pi^\top + \pi g$$

**B** Equilibrium propagation

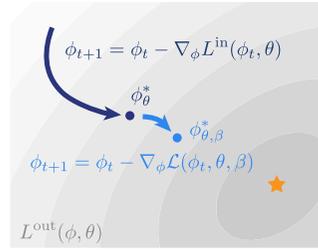


Figure 3: Visual comparison of implicit differentiation (A) and equilibrium propagation (B). The goal of bilevel optimization is to search for outer parameters  $\theta$  such that the corresponding  $\phi_\theta^*$  minimizes the outer loss  $L^{\text{out}}$  (loss magnitude is plotted in gray, the darker the smaller; the star denotes a minimum of  $L^{\text{out}}$ ). The dark blue curve represents the inner optimization dynamics, here chosen to be gradient descent, which is the same for the two approaches. The gradient calculation however differs. (A) An auxiliary quadratic form minimization problem is solved ( $H$  here denotes the Hessian matrix  $\partial_\phi^2 L^{\text{in}}(\phi_\theta^*, \theta)$  and  $g$  the gradient  $\partial_\phi L^{\text{out}}(\phi_\theta^*, \theta)^\top$ ) and the output  $\pi^*$  is then used to estimate the outer gradient. (B) The inner objective is nudged toward the outer objective through the augmented loss  $\mathcal{L}$ , which is then minimized (light blue curve). The outer gradient is then estimated by contrasting partial derivatives evaluated at  $\phi_\theta^*$  and  $\phi_{\theta,\beta}^*$ .

outclass equilibrium propagation. All methods minimize the same objective in the first phase, so it is reasonable to consider that they use the same optimizer. For the second phase, RBP uses gradient descent, CG conjugate gradients, and EP whatever optimizer is best suited to the augmented objective. This implies that the algorithm with the best guarantees is CG as the optimizer it uses in the second phase is extremely efficient (hence more efficient than the one EP would use). The comparison between EP and RBP depends on the problem considered and requires empirical evidence.

*5.2.2 Robustness to Violated Assumptions.* In the previous paragraph, we assumed that we are sufficiently close to minimizing the inner loss so that all implicit methods are properly justified. We now look at how they behave when those conditions are not met. First-order methods here shine as they just perform a crude approximation and do not rely on those assumptions. In principle, both CG and RBP would have a diverging second phase if the Hessian is not positive definite, but in practice, it seems that CG is much more unstable (Liao et al., 2018; Lorraine et al., 2020; Grosse, 2021). EP does not have diverging second phase as long as  $\beta \geq 0$  and the inner and outer losses are bounded from below.

*5.2.3 Simplicity of the Computational Elements.* The methods we compare here require different computational elements. While every method requires computing partial derivatives with respect to the outer parameters, approximate first-order methods stand out in their simplicity of implementation. In particular, these methods do not even require storing the result of the first phase. On the other hand, implicit differentiation methods are the most complex to implement as they involve calculating Hessian-vector products. In digital computers, automatic differentiation software offers efficient implementations of this operation (e.g., Abadi et al., 2016; Paszke et al., 2019; Bradbury et al., 2018). However, implementing Hessian-vector products can be challenging in large-scale distributed systems, neuromorphic hardware, or more exotic analog physical systems. Arguably, it is also hard to conceive such an operation as being biologically plausible. Remarkably, equilibrium propagation only requires contrasting partial derivatives and, therefore avoids computing such Hessian-vector products. Recent developments (Scellier et al., 2022) on equilibrium propagation have shown that the outer gradient can still be estimated when the inner loss function underlying the (bio)physical system dynamics and its partial derivatives are unknown, as long as the parameters  $\theta$  can be externally controlled. This considerably widens the scope of systems in which equilibrium propagation can be applied.

*5.2.4 Which Method to Choose?* First-order methods tend to work best off-the-shelf, without extensive tuning, so they are a good choice if performance is not the most important criterion. When performance is important and inner optimization is easy enough so that it is possible to closely approximate a local minimum of the inner loss function, the conjugate gradient method is the best one. If it reveals to be too unstable, recurrent backpropagation might solve those instability issues. Finally, if computing Hessian-vector products is not an option but performance is still important, equilibrium propagation is worth consideration.

## 6 Conclusion

---

We have presented bilevel optimization in a broad machine learning context and discussed gradient-based methods to solve such problems. Framing learning as bilevel optimization generalizes the traditional cost-minimization view of learning to computations that are not necessarily explicitly described and that therefore cannot be learned through gradient descent with backpropagated errors. The implicit methods we reviewed here, rooted in either implicit differentiation or equilibrium propagation, allow computing gradients for such problems using local information, and sometimes using only elementary operations. These properties may turn out to be of particular importance for the development of biological

theories of learning, as well as for the development of next-generation learning machines.

## Appendix A: Some Bilevel Optimization Learning Problems \_\_\_\_\_

We here review different learning problems that fit in the bilevel optimization framework.

### A.1 Energy-Based Neural Networks.

*A.1.1 Explicit Description of Neural Networks.* Neural networks are usually described through the computations that they perform to process an input signal  $x$ . For a feedforward neural network, it usually takes the following form

$$\phi^0 = x, \quad \phi^{l+1} = \rho(W^l \phi^l + b^l), \quad (\text{A.1})$$

where  $\phi^l$  corresponds to the activity of the neurons from the  $l$ th layer,  $\rho$  to a nonlinear activation function,  $W^l$  to the weights connecting layer  $l$  to layer  $l + 1$ , and  $b^l$  to the biases of layer  $l$ . In the supervised learning framework, the activity  $\phi^L$  at the very last layer is compared to a desired output  $y$  through a cost function  $C(\phi^L, y)$ . The backpropagation algorithm (Linnainmaa, 1976; Werbos, 1982; Rumelhart et al., 1986) propagates the error measured at the last layer toward the first layers of the network to efficiently compute gradients and then learn the weights of the network.

*A.1.2 Energy-Based Description.* An alternative description of neural networks is to consider that the activity is an equilibrium of some energy function  $E$ . Going from an explicit to an implicit description is easy for the feedforward neural network:<sup>6</sup> The energy function

$$E(\phi, \theta, x) := \frac{1}{2} \|\phi^0 - x\|^2 + \frac{1}{2} \sum_{l=0}^{L-1} \|\phi^{l+1} - \rho(W^l \phi^l + b^l)\|^2 \quad (\text{A.2})$$

has only one global minimizer, which is the neuronal activity  $\phi_\theta^*$ , as computed through the feedforward processing already described.

The energy-based formulation is more than a mathematical reformulation. For example, the energy, equation A.2, is derived from an approximate probabilistic approach in the predictive coding framework (Rao & Ballard, 1999; Whittington & Bogacz, 2017). Other types of energy functions also exist, such as the Hopfield energy (Hopfield, 1984; Scellier & Bengio, 2017),

---

<sup>6</sup>Similar manipulations can be done in general to obtain an implicit description of a system from an explicit one

and encompasses computations that cannot be formulated explicitly. Although the term *energy* has a physical meaning, physical networks can minimize other quantities than the physical energy, such as the co-content for electrical circuits (Millar, 1951; Kendall et al., 2020).

Under this paradigm, learning under supervision can be formulated as the following bilevel optimization:

$$\begin{aligned} \min_{\theta} \mathbb{E}_{(x,y)} [C(\phi_{\theta}^*, y)] \\ \text{s.t. } \phi_{\theta}^* \in \arg \min_{\phi} E(\phi, \theta, x). \end{aligned} \quad (\text{A.3})$$

Backpropagation is not generally applicable to compute gradients associated with this optimization problem. This is why we introduce implicit methods in this review.

**A.2 Hyperparameter Optimization and Metalearning.** We have briefly introduced bilevel optimization for hyperparameter optimization and metalearning in section 2.3. Recall that in this context, bilevel optimization generally takes the form

$$\begin{aligned} \min_{\theta} \mathbb{E}_{\tau} [L^{\text{out}}(\phi_{\tau,\theta}^*, \theta, \mathcal{D}_{\tau}^{\text{val}})] \\ \text{s.t. } \phi_{\tau,\theta}^* \in \arg \min_{\phi} L^{\text{in}}(\phi, \theta, \mathcal{D}_{\tau}^{\text{train}}), \end{aligned} \quad (\text{A.4})$$

where expectation is taken over multiple tasks  $\tau$  for metalearning and over a single one for hyperparameter optimization (the expectation then disappears).

The purpose of this section is to underline the diversity of interactions between inner and outer parameters, which are also referred to as base and metaparameters in metalearning. We mentioned in section 2.3 that the outer parameters can be the parameters of a quadratic regularization in the context of hyperparameter optimization (Goutte & Larsen, 1998; Bengio, 2000), but the very same regularizer can be used in metalearning (Rajeswaran et al., 2020; Zucchet et al., 2021). In metalearning, the center of the regularization is also metalearned, providing a rough idea of which base parameter configuration works well on the task distribution. Another example is when the metaparameters are the weights of a hypernetwork (Ha, Dai, & Le, 2017) that take the inner parameters as input to produce the weights of the network that processes incoming data (Lorraine & Duvenaud, 2018; MacKay et al., 2019; Rusu et al., 2019; Zhao, Kobayashi, Sacramento, & von Oswald, 2020). The task-specific modification can also be done at the neuron level (Zintgraf, Shiarlis, Kurin, Hofmann, & Whiteson, 2019; Mudrakarta, Sandler, Zhmoginov, & Howard, 2019; Zucchet et al., 2021), while keeping the weights of the neural network shared across tasks. Alternatively, the

task-shared outer parameters can be the weights of a neural network that acts as a feature extractor that will help a task-specific classifier or regressor parameterized by the base parameters to solve the task at hand (Raghu, Raghu, Bengio, & Vinyals, 2020; Lee, Maji, Ravichandran, & Soatto, 2019; Bertinetto et al., 2019).

**A.3 Generative Adversarial Networks.** Generative adversarial networks (Goodfellow et al., 2014; Metz, Poole, Pfau, & Sohl-Dickstein, 2017) consist of a generative and a discriminative network that are learned in an adversarial fashion. The discriminator, parameterized by  $\phi$  has to distinguish between samples generated by the generator and samples coming from the true data distribution  $p(x)$ . On the other side, the objective of the generative model  $g_\theta$  is to generate samples that fool a perfect discriminator  $D_{\phi_\theta^*}$ . The corresponding optimization bilevel optimization problem is

$$\begin{aligned} \min_{\theta} & -\mathbb{E}_{z \sim \mathcal{N}(0,1)} [\log D_{\phi_\theta^*}(g_\theta(z))] \\ \text{s.t. } \phi_\theta^* & \in \arg \min_{\phi} \mathbb{E}_{x \sim p(x)} [\log D_{\phi}(x)] + \mathbb{E}_{z \sim \mathcal{N}(0,1)} [\log (1 - D_{\phi}(g_\theta(z)))] . \end{aligned} \quad (\text{A.5})$$

**A.4 Actor-Critic.** Some reinforcement learning problems can be formulated as bilevel optimization, such as actor-critics (Konda & Tsitsiklis, 2000). The objective here is to learn an actor, which is a policy that tries to maximize the expected reward received while interacting with an environment. It receives help from a critic, an action-value function, which gives better feedback to the actor than the reward only. Following (Pfau & Vinyals, 2016; Yang, Chen, Hong, & Wang, 2019; Zhou, Li, Yang, Wang, & Hospedales, 2020; Hong, Wai, Wang, & Yang, 2020), training an actor-critic can be formulated as

$$\begin{aligned} \max_{\theta} & \mathbb{E}_{s \sim \rho, a \sim \pi_\theta(\cdot|s)} [Q_{\phi_\theta^*}(s, a)] \\ \text{s.t. } \phi_\theta^* & \in \arg \min_{\phi} \mathbb{E}_{s \sim \rho, a \sim \pi_\theta(\cdot|s)} [(Q_\phi(s, a) - Q^{\pi_\theta}(s, a))^2], \end{aligned} \quad (\text{A.6})$$

where  $\rho$  is the initial state distribution,  $\pi_\theta(\cdot|s)$  is the policy distribution parameterized by  $\theta$  (the actor),  $Q^{\pi_\theta}$  its corresponding Q-function, and  $Q_\phi$  the approximate Q network (the critic). A similar formulation exists in model-based reinforcement learning where the critic is replaced by a model that tries to predict the feature (Rajeswaran, Mordatch, & Kumar, 2020).

**A.5 Stackelberg Games.** Interestingly, the last two examples can be given a game-theoretic interpretation through the notion of Stackelberg games. Stackelberg games (von Stackelberg, 1934) are a class of games where two players, a leader and a follower, play with hierarchical order. The leader  $\theta$  has a strategic advantage: it plays first and knows what will

be the perfect answer of the follower  $\phi$ . In the bilevel optimization framework, the follower's best response minimizes the inner loss, and the leader optimizes the outer loss knowing the perfect answer of the follower  $\phi_\theta^*$ .

For generative adversarial networks, the generator is the leader and the discriminator the follower in the Stackelberg game terminology (Fiez, Chasnov, & Ratliff, 2020). For actor-critic methods, the actor is the leader, as we ultimately want to get a good working policy, and the critic is the follower (Zheng, Fiez, Alumbaugh, Chasnov, & Ratliff, 2021).

## Appendix B: Theoretical Analysis for Implicit Differentiation Methods

---

**B.1 Proof of Theorem 3.** This proof is inspired from Pedregosa (2016), which proves a very similar result under local assumptions, and Rajeswaran et al. (2019), which uses global assumptions and studies the regularized inner loss we studied in section 2.3. The proof uses the general formulation of the former with the stronger assumptions of the latter, in the goal of making the proof as insightful as possible.

We first rewrite the assumptions and the statement of the theorem.

**Assumption 1.** Suppose that there exist positive real numbers  $(\mu, \rho, B, L, M)$  such that:

- i.  $L^{\text{in}}$  is twice continuously differentiable, and  $L^{\text{out}}$  is continuously differentiable.
- ii.  $L^{\text{in}}$  is  $\mu$ -strongly convex as a function of  $\phi$ .
- iii. The second-order derivatives (Hessian and cross derivatives) of  $L^{\text{in}}$  are  $\rho$ -Lipschitz as functions of  $\phi$ .
- iv. As functions of  $\phi$ ,  $L^{\text{out}}$  is  $B$ -Lipschitz,  $L$ -smooth and  $\partial_\theta L^{\text{out}}$  is  $M$ -Lipschitz,

**Theorem 3.** Let  $\phi_\theta^*$  be a minimizer of  $L^{\text{in}}$  and  $\hat{\phi}$  be its approximated value. Let  $\hat{\pi}$  be an approximation of  $\pi^*$  computed by one of the implicit differentiation methods. Let  $\delta \in ]0, \frac{\mu}{2\rho}[$  be such that

$$\|\phi_\theta^* - \hat{\phi}\| \leq \delta$$

and  $\delta' > 0$  such that

$$\|\hat{\pi} - \partial_\phi L^{\text{out}}(\hat{\phi}, \theta) \partial_\phi^2 L^{\text{in}}(\hat{\phi}, \theta)^{-1}\| \leq \delta'.$$

Then, under assumption 1, there exists a constant  $C$  such that

$$\|\nabla_\theta - \hat{\nabla}_\theta\| \leq C(\delta + \delta').$$

The main idea of the proof is to show that the outer gradient estimation error introduced by the implicit differentiation algorithm comes from two sources: the fixed-point approximation error and the finite number of steps in the estimation of  $\hat{\pi}$ . Bounding the impact of the first source will be straightforward, but the second one requires more work. This stems from the fact that implicit differentiation methods do not directly approximate  $\pi^*$  in their second phase but only the proxy  $\partial_\phi L^{\text{out}}(\hat{\phi}, \theta) \partial_\phi^2 L^{\text{in}}(\hat{\phi}, \theta)^{-1}$ . We therefore need to quantify how far the proxy is from  $\pi^*$ . This can be done by remarking that the two are solutions of two similar linear systems. Lemma 1 is a result from perturbed linear systems theory that will allow us to upper-bound the distance between the two.

**Lemma 1** (Theorem 7.2; Higham, 2002). *Let  $Ax = b$  and  $A'x' = b'$  two linear systems with  $\|A - A'\| \leq \varepsilon_A$  and  $\|b - b'\| \leq \varepsilon_b$ . If  $\varepsilon_A \|A^{-1}\| < 1$ ; then*

$$\|x - x'\| \leq \frac{\|A^{-1}\|}{1 - \varepsilon_A \|A^{-1}\|} (\varepsilon_b + \|A^{-1}b\| \varepsilon_A).$$

**Proof.** Consider the quantity  $A(x' - x)$ . It is equal to

$$\begin{aligned} A(x' - x) &= A'x' + (A - A')x' - Ax \\ &= b' + (A - A')x' - b \\ &= b' - b + (A - A')x + (A - A')(x' - x). \end{aligned}$$

Then,

$$x' - x = A^{-1} (b' - b + (A - A')A^{-1}b + (A - A')(x' - x))$$

so

$$\|x' - x\| \leq \|A^{-1}\| (\varepsilon_b + \varepsilon_A \|A^{-1}b\| + \varepsilon_A \|x' - x\|),$$

which yields the required result after subtracting  $\varepsilon_A \|A^{-1}\| \|x' - x\|$  to both sides. □

With this result, we can now prove theorem 3.

**Proof of Theorem 3.** Recall that

$$\nabla_\theta^\top = \frac{\partial L^{\text{out}}}{\partial \theta}(\phi_\theta^*, \theta) - \pi^* \frac{\partial^2 L^{\text{in}}}{\partial \theta \partial \phi}(\phi_\theta^*, \theta)$$

for

$$\pi^* = \frac{\partial L^{\text{out}}}{\partial \phi}(\phi_\theta^*, \theta) \left( \frac{\partial^2 L^{\text{in}}}{\partial \phi^2}(\phi_\theta^*, \theta) \right)^{-1}$$

is estimated with

$$\widehat{\nabla}_\theta^\top = \frac{\partial L^{\text{out}}}{\partial \theta}(\hat{\phi}, \theta) - \hat{\pi} \frac{\partial^2 L^{\text{in}}}{\partial \theta \partial \phi}(\hat{\phi}, \theta).$$

We introduce the shorthand  $\chi^* := \partial_\theta \partial_\phi L^{\text{in}}(\phi_\theta^*, \theta)$  and  $\hat{\chi}$  its estimated counterpart. We then have

$$\begin{aligned} \|\nabla_\theta - \widehat{\nabla}_\theta\| &\leq \|\partial_\theta L^{\text{out}}(\phi_\theta^*, \theta) - \partial_\theta L^{\text{out}}(\hat{\phi}, \theta)\| + \|\pi^* \chi^* - \hat{\pi} \hat{\chi}\| \\ &\leq \underbrace{\|\partial_\theta L^{\text{out}}(\phi_\theta^*, \theta) - \partial_\theta L^{\text{out}}(\hat{\phi}, \theta)\|}_{a)} + \underbrace{\|\pi^* (\chi^* - \hat{\chi})\|}_{b)} + \underbrace{\|(\pi^* - \hat{\pi}) \hat{\chi}\|}_{c)}. \end{aligned}$$

We bound each term:

1. From the Lipschitz continuity of  $\partial_\theta L^{\text{out}}$  comes

$$\|\partial_\theta L^{\text{out}}(\phi_\theta^*, \theta) - \partial_\theta L^{\text{out}}(\hat{\phi}, \theta)\| \leq M \|\phi_\theta^* - \hat{\phi}\| \leq M \delta.$$

2. The  $\mu$ -strong convexity of  $L^{\text{in}}$  and the  $B$ -Lipschitz continuity of  $L^{\text{out}}$  imply that

$$\|\pi^*\| \leq \|\partial_\phi L^{\text{out}}(\phi_\theta^*, \theta)\| \|\partial_\phi^2 L^{\text{in}}(\phi_\theta^*, \theta)^{-1}\| \leq B \frac{1}{\mu}.$$

Using the Lipschitz continuity of the cross-derivatives of  $L^{\text{in}}$ , we have

$$\|\hat{\chi} - \chi^*\| \leq \rho \delta$$

so

$$\|(\chi^* - \hat{\chi})\pi^*\| \leq \rho \delta \|\pi^*\| \leq \rho \delta \frac{B}{\mu}.$$

3. Lipschitz continuity of  $\partial_\theta L^{\text{in}}$  yields  $\|\partial_\phi \partial_\theta L^{\text{in}}(\hat{\phi}, \theta)\| \leq M$ . With the symmetry of the cross-derivatives and the fact that the norm of a matrix equals the norm of its transpose, we have  $\|\hat{\chi}\| \leq M$  and

$$\|\hat{\chi}(\pi^* - \hat{\pi})\| \leq M \|\pi^* - \hat{\pi}\|.$$

In the term  $\|\hat{\pi} - \pi^*\|$ , we still take in account the error made in the fixed point approximation. We can separate it with

$$\begin{aligned} \|\hat{\pi} - \pi^*\| &\leq \|\hat{\pi} - \partial_\phi L^{\text{out}}(\hat{\phi}, \theta) \partial_\phi^2 L^{\text{in}}(\hat{\phi}, \theta)^{-1}\| \\ &\quad + \|\partial_\phi L^{\text{out}}(\hat{\phi}, \theta) \partial_\phi^2 L^{\text{in}}(\hat{\phi}, \theta)^{-1} - \pi^*\| \\ &\leq \delta' + \|\partial_\phi L^{\text{out}}(\hat{\phi}, \theta) \partial_\phi^2 L^{\text{in}}(\hat{\phi}, \theta)^{-1} - \pi^*\|. \end{aligned}$$

The second term now only depends on the fixed-point approximation error and can be bounded using lemma 1. Due to the  $\rho$ -Hessian Lipschitz property of  $L^{\text{in}}$ ,

$$\varepsilon_A := \|\partial_\phi^2 L^{\text{in}}(\phi_\theta^*, \theta) - \partial_\phi^2 L^{\text{in}}(\hat{\phi}, \theta)\| \leq \rho\delta.$$

The use of the lemma is then justified by the upper-bound assumption on  $\delta$ :

$$\varepsilon_A \|\partial_\phi^2 L^{\text{in}}(\phi_\theta^*, \theta)^{-1}\| \leq \rho\delta/\mu \leq 1/2 < 1.$$

The smoothness of  $L^{\text{in}}$  implies

$$\varepsilon_b := \|\partial_\phi L^{\text{out}}(\phi_\theta^*, \theta) - \partial_\phi L^{\text{in}}(\hat{\phi}, \theta)\| \leq L\delta.$$

We can now apply the lemma, which yields

$$\begin{aligned} \|\partial_\phi L^{\text{out}}(\hat{\phi}, \theta) \partial_\phi^2 L^{\text{in}}(\hat{\phi}, \theta)^{-1} - \pi^*\| &\leq \frac{\mu^{-1}}{1 - 1/2} (L\delta + \|\pi^*\|\rho\delta) \\ &\leq \frac{2\delta}{\mu} \left( L + \frac{B\rho}{\mu} \right). \end{aligned}$$

We have therefore proved

$$\|\hat{\chi}(\hat{\pi} - \pi^*)\| \leq M \left( \delta' + \frac{2\delta}{\mu} \left( L + \frac{B\rho}{\mu} \right) \right).$$

Gathering the three bounds gives

$$\|\nabla_\theta - \widehat{\nabla}_\theta\| \leq M\delta + \frac{B\rho}{\mu}\delta + M \left( \delta' + \frac{2\delta}{\mu} \left( L + \frac{B\rho}{\mu} \right) \right).$$

Choosing

$$C := M + \frac{B\rho + 2ML}{\mu} + \frac{2MB\rho}{\mu^2}$$

finishes the proof. □

**B.2 Extension to Local Assumptions.** In theorem 3, we assumed the strong convexity of  $L^{\text{in}}$  to get a bound on the outer gradient estimation error. We can get a more local version of it if we only assume that the Hessian at a minimum  $\phi^*$  of  $L^{\text{in}}$  is positive-definite, that is, the minimum is not flat. The idea is to show that when the Hessian is continuous and it is positive-definite at  $\phi_\theta^*$ ,  $L^{\text{in}}$  will be strongly convex in a neighborhood of  $\phi_\theta^*$ , which allows going back to the assumptions of theorem 3. This is formalized in fact 1:

Let  $\phi^*$  be a local minimum of  $L^{\text{in}}$  such that  $\partial_{\phi}^2 L^{\text{in}}(\phi^*, \theta)$  is positive definite. Note  $\mu$  its smallest (strictly positive) eigenvalue. If  $L^{\text{in}}$  is  $\rho$ -Lipschitz Hessian, then  $L^{\text{in}}$  is  $\mu/2$ -strongly convex on the ball of radius  $\mu/2\rho$  centered on  $\phi^*$ .

Interestingly, the Hessian of  $L^{\text{in}}$  at  $\hat{\phi}$  is not necessarily positive semidefinite when  $\hat{\phi}$  outside the ball is centered in  $\hat{\phi}$  with radius  $\mu/\rho$  (with the notations of assumption 1 and fact 1). In this case, the quadratic form, equation 4.3, is not bounded from below and procedures that try to minimize it will diverge.

### Appendix C: Equilibrium Propagation Estimators with Multiple Points

---

In section 4.2, we have presented a way to estimate the outer gradient formula given by the equilibrium propagation theorem using two points. Recall that the equilibrium propagation allows reformulating the outer gradient  $\nabla_{\theta}$  as

$$\nabla_{\theta} = \frac{d}{d\beta} \frac{\partial \mathcal{L}}{\partial \theta}(\phi_{\theta, \beta}^*, \theta, \beta) \Big|_{\beta=0}.$$

The simplest finite different estimator is the two-point estimator that we presented in section 4.2:

$$\hat{\nabla}_{\theta}^{\top} = \frac{1}{\beta} \left( \frac{\partial \mathcal{L}}{\partial \theta}(\phi_{\theta, \beta}^*, \theta, \beta) - \frac{\partial \mathcal{L}}{\partial \theta}(\phi_{\theta, 0}^*, \theta, 0) \right).$$

We now derive an estimator that uses several points to make a more accurate estimation of the derivative.

**C.1 Forward Finite Differences.** The objective of this paragraph is to derive the  $p$ -forward finite difference learning rule that uses  $p$  points to get a finer approximation of the outer gradient. Consider the values of the function  $f : t \mapsto \partial_{\theta} \mathcal{L}(\phi_{\theta, t}^*, \theta, t)$  for  $t \in \{0, \beta, \dots, (p-1)\beta\}$ . We seek to find a linear combination of those measurements that approximates  $\nabla_{\theta} = d_{\beta} \partial_{\theta} \mathcal{L}(\phi_{\theta, \beta}^*, \theta, \beta)|_{\beta=0} = f'(0)$ , that is, find a vector  $\alpha \in \mathbb{R}^p$  such that

$$\sum_{i=0}^{p-1} \alpha_i f(i\beta) = \beta f'(0) + O(\beta^p). \tag{C.1}$$

Taylor series approximation (around  $\beta = 0$ ) and an inversion of the summation indices yield

$$\begin{aligned}
 \sum_{i=0}^{p-1} \alpha_i f(i\beta) &= \sum_{i=0}^{p-1} \alpha_i \left( \sum_{k=0}^{p-1} f^{(k)}(0) \frac{(i\beta)^k}{k!} + O(\beta^p) \right) \\
 &= \sum_{k=0}^{p-1} \sum_{i=0}^{p-1} \alpha_i f^{(k)}(0) \frac{(i\beta)^k}{k!} + O(\beta^p) \\
 &= \sum_{k=0}^{p-1} f^{(k)}(0) \frac{\beta^k}{k!} \sum_{i=0}^{p-1} \alpha_i i^k + O(\beta^p).
 \end{aligned}
 \tag{C.2}$$

In equations C.1 and C.2, we have two polynomials in  $\beta$  that we want to be equal so all their coefficients have to be the same. We hence need to solve

$$\binom{i^k}{i,k} \alpha = \begin{pmatrix} 0 \\ 1 \\ 0 \\ \vdots \end{pmatrix},
 \tag{C.3}$$

where  $\binom{i^k}{i,k}$  is a  $p \times p$  invertible Vandermonde matrix. The resolution of such a system can easily be done numerically. The values of  $\alpha$  for small  $p$  are:

$p$	$\alpha_0$	$\alpha_1$	$\alpha_2$	$\alpha_3$	$\alpha_4$
2	-1	1	0	0	0
3	-3/2	2	-1/2	0	0
4	-11/6	3	-3/2	1/3	0
5	-25/12	4	-3	4/3	-1/4

Note that the 2-forward finite difference learning rule is the same as the finite difference one. In the following, we assume that  $\alpha$  satisfies equation C.3. The resulting algorithm is presented in algorithm 2.

**C.2 Why Forward Finite Differences?** There exist different kinds of finite difference estimators that use multiple points.<sup>7</sup> We chose to present the forward difference ones above as they are the ones that only use estimates for positive  $\beta$  values. We illustrate why this may be important on an example.

When using  $p = 3$  points, the bias reduction obtained with a forward estimate is similar to that obtained with the symmetric or central estimate,

---

<sup>7</sup>See [https://en.wikipedia.org/wiki/Finite\\_difference](https://en.wikipedia.org/wiki/Finite_difference).

---

**Algorithm 2:**  $p$ -Forward Finite Difference Learning Rule.
 

---

**Result:** Estimation of  $\nabla_\theta$ 

For every  $i \in \{0, \dots, (p-1)\beta\}$ , minimize  $\phi \mapsto \mathcal{L}(\phi, \theta, i\beta)$ , starting from the solution of last step, and note  $\hat{\phi}_{i\beta}$  the result;

Estimate  $\nabla_\theta$ , the derivative of  $\theta \mapsto L^{\text{out}}(\phi_\theta^*, \theta)$ , using

$$(\widehat{\nabla}_\theta^p)^\top = \frac{1}{\beta} \sum_{i=0}^{p-1} \alpha_i \frac{\partial \mathcal{L}}{\partial \theta}(\hat{\phi}_{i\beta}, \theta, i\beta).$$

Return  $\widehat{\nabla}_\theta^p$ ;

---

$$\frac{1}{2\beta} \left( \frac{\partial \mathcal{L}}{\partial \theta}(\phi_{\theta, \beta}^*, \theta, \beta) - \frac{\partial \mathcal{L}}{\partial \theta}(\phi_{\theta, -\beta}^*, \theta, -\beta) \right), \quad (\text{C.4})$$

that is used in Laborieux et al. (2021). However, negative  $\beta$  values can prove to be problematic. To illustrate that, consider  $L^{\text{in}}$  and  $L^{\text{out}}$ , two  $\mu$ -strongly convex and  $L$ -smooth functions (e.g.,  $L^{\text{in}}(\phi) = L\|\phi\|^2/2$  and  $L^{\text{out}}(\phi) = \mu\|\phi\|^2/2$ ). Then  $\mathcal{L}$  (here, equal to  $(L + \beta\mu)\|\phi\|^2/2$ ) is not bounded from below when  $\beta < -L/\mu$  so  $\phi_{\theta, -\beta}^*$  does not exist anymore and the estimate diverges. When using negative  $\beta$  values, one therefore has to be careful that all the phases converge.

**C.3 When Are Multiple Points Estimators Worth It?** One can think that adding more points will always lead to a more precise estimation of the outer gradient. If the approximations  $\hat{\phi}_{i\beta}$  of the minimizers  $\phi_{i\beta, \theta}^*$  are perfect, it will always be the case as the bias is a  $O(\beta^p)$  with  $p$  the number of points. Using the terminology used in the proof of theorem 5, this means that the finite difference error decreases when the number of points increases. However, when we can only obtain approximate minimizers, adding more points aggregates the fixed-point approximation errors made in each phase, thus potentially making the estimation error bigger. Whether more points would be useful is therefore a practical matter.

---

### Acknowledgments

This research was supported by an Ambizione grant (PZ00P3 186027) from the Swiss National Science Foundation and an ETH research grant (ETH-23 21-1) awarded to J.S. We thank Benjamin Scellier, Johannes von Oswald, and Simon Schug for their detailed comments on the manuscript for this review.

## References

---

- Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., . . . Zheng, X. (2016). TensorFlow: A system for large-scale machine learning. In *Proceedings of the 12th USENIX Symposium on Operating Systems Design and Implementation*.
- Agrawal, A., Amos, B., Barratt, S., Boyd, S., Diamond, S., & Kolter, J. Z. (2019). Differentiable convex optimization layers. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, & R. Garnett (Eds.), *Advances in neural information processing systems*, 32. Red Hook, NY: Curran.
- Almeida, L. B. (1990). A learning rule for asynchronous perceptrons with feedback in a combinatorial environment. In K. Suzuki (Ed.), *Artificial neural networks: Concept learning*. Rijeka, Croatia: InTech.
- Amos, B., & Kolter, J. Z. (2017). Optnet: Differentiable optimization as a layer in neural networks. In *Proceedings of the International Conference on Machine Learning*.
- Bai, S., Kolter, J. Z., & Koltun, V. (2019). Deep equilibrium models. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, & R. Garnett (Eds.), *Advances in neural information processing systems*, 32. Red Hook, NY: Curran.
- Bai, S., Koltun, V., & Kolter, J. Z. (2020). Multiscale deep equilibrium models. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, & H. Lin (Eds.), *Advances in neural information processing systems*, 33. Red Hook, NY: Curran.
- Bard, J. F. (1998). *Practical bilevel optimization*. Berlin: Springer.
- Bellec, G., Scherr, F., Subramoney, A., Hajek, E., Salaj, D., Legenstein, R., & Maass, W. (2020). A solution to the learning dilemma for recurrent networks of spiking neurons. *Nature Communications*, 11(1), 33. 10.1038/s41467-020-17236-y
- Bengio, Y. (2000). Gradient-based optimization of hyperparameters. *Neural Computation*, 12(8). 10.1162/089976600300015187
- Bengio, Y., Bengio, S., & Cloutier, J. (1990). *Learning a synaptic learning rule*. Université de Montréal, Département d'Informatique et de Recherche opérationnelle technical report.
- Bergstra, J., & Bengio, Y. (2012). Random search for hyper-parameter optimization. *Journal of Machine Learning Research*, 13, 281–305.
- Bertinetto, L., Henriques, J. F., Torr, P. H. S., & Vedaldi, A. (2019). Metalearning with differentiable closed-form solvers. In *Proceedings of the International Conference on Learning Representations*.
- Bradbury, J., Frostig, R., Hawkins, P., Johnson, M. J., Leary, C., Maclaurin, D., & Wanderman-Milne, S. (2018). *JAX: Composable transformations of Python+ NumPy programs*. <http://github.com/google/jax>
- Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., . . . Amodei, D. (2020). Language models are few-shot learners. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, & H. Lin (Eds.), *Advances in neural information processing systems*, 33. Red Hook, NY: Curran.
- Cauchy, A. L. (1831). *Turin memoir*.
- Cohen, M. A., & Grossberg, S. (1983). Absolute stability of global pattern formation and parallel memory storage by competitive neural networks. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-13(5). 10.1109/TSMC.1983.6313075

- d'Aspremont, A. (2008). Smooth optimization with approximate gradient. *SIAM Journal on Optimization*, 19(3).
- Descartes, R. (1637). *Géométrie*.
- Djulonga, J., & Krause, A. (2017). I. Guyon, Y. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, & R. Garnett (Eds.), Differentiable learning of submodular models. In *Advances in neural information processing systems*, 30. Red Hook, NY: Curran.
- Dontchev, A. L. & Rockafellar, R. T. (2009). *Implicit functions and solution mappings*. Berlin: Springer.
- Euler, L. (1748). *Introductio in analysi n infinitorum*.
- Fiez, T., Chasnov, B., & Ratliff, L. (2020). Implicit learning dynamics in stackelberg games: Equilibria characterization, convergence analysis, and empirical study. In *Proceedings of the International Conference on Machine Learning*.
- Finn, C., Abbeel, P., & Levine, S. (2017). Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of the International Conference on Machine Learning*.
- Foo, C.-S., Do, C. B., & Ng, A. Y. (2007). Efficient multiple hyperparameter learning for log-linear models. In J. Platt, D. Koller, Y. Singer, & S. Roweis (Eds.) *Advances in neural information processing systems*, 20. Cambridge, MA: MIT Press.
- Friedlander, M. P., & Schmidt, M. (2012). Hybrid deterministic-stochastic methods for data fitting. *SIAM Journal on Scientific Computing*, 34(3). 10.1137/110830629
- Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., . . . Bengio, Y. (2014). Generative adversarial networks. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, & K. Q. Weinberger (Eds.), *Advances in neural information processing systems*, 27. Red Hook, NY: Curran.
- Gould, S., Hartley, R., & Campbell, D. J. (2021). Deep declarative networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44, 3988–4004.
- Goutte, C., & Larsen, J. (1998). Adaptive regularization of neural networks using conjugate gradient. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*. Piscataway, NJ: IEEE.
- Grosse, R. (2021). *Lecture 11: Bilevel optimization, CSC 2541: Neural net training dynamics*. University of Toronto.
- Gruslys, A., Munos, R., Danihelka, I., Lanctot, M., & Graves, A. (2016). Memory-efficient backpropagation through time. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, & R. Garnett (Eds.), *Advances in neural information processing systems*, 29. Red Hook, NY: Curran.
- Ha, D., Dai, A., & Le, Q. V. (2017). HyperNetworks. In *Proceedings of the International Conference on Learning Representations*.
- Hazan, T., Keshet, J., & McAllester, D. A. (2010). Direct loss minimization for structured prediction. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, & R. Garnett (Eds.), *Advances in neural information processing systems*, 29. Red Hook, NY: Curran.
- Higham, N. J. (2002). *Accuracy and stability of numerical algorithms* (2nd ed.). Philadelphia: Society for Industrial and Applied Mathematics.
- Hong, M., Wai, H.-T., Wang, Z., & Yang, Z. (2020). *A two-timescale framework for bilevel optimization: Complexity analysis and application to Actor-Critic*. arXiv:2007.05170.

- Hopfield, J. J. (1984). Neurons with graded response have collective computational properties like those of two-state neurons. In *Proceedings of the National Academy of Sciences*, 81(10). 10.1073/pnas.81.10.3088
- Hutchinson, J., Koch, C., Luo, J., & Mead, C. (1988). Computing motion using analog and binary resistive networks. *Computer*, 21(3). 10.1109/2.31
- Jaeger, H. (2002). *A tutorial on training recurrent neural networks, covering BPPT, RTRL, EKF and the "echo state network" approach* (GMD Report 159). St. Augustin: German National Research Center for Information Technology.
- Jumper, J., Evans, R., Pritzel, A., Green, T., Figurnov, M., Ronneberger, O., Tunyasuvunakool, K., . . . Hassabis, D. (2021). Highly accurate protein structure prediction with AlphaFold. *Nature*, 596(7873). 10.1038/s41586-021-03819-2
- Kendall, J., Pantone, R., Manickavasagam, K., Bengio, Y., & Scellier, B. (2020). *Training end-to-end analog neural networks with equilibrium propagation*. arXiv:2006.01981.
- Kolter, Z., Duvenaud, D., & Johnson, M. (2021). *Deep implicit layers: Neural ODEs, deep equilibrium models, and beyond*. <http://implicit-layers-tutorial.org>
- Konda, V. R., & Tsitsiklis, J. N. (2000). Actor-critic algorithms. In T. Leen, T. Dietterich, & V. Tresp (Eds.), *Advances in neural information processing systems*, 13. Red Hook, NY: Curran.
- Krantz, S. G., & Parks, H. R. (2003). *The implicit function theorem: History, theory and applications*. Berlin: Springer.
- Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. In F. Pereira, C. J. C. Burges, L. Bottou, & K. Q. Weinberger (Eds.), *Advances in neural information processing systems*, 25. Red Hook, NY: Curran.
- Laborieux, A., Ernout, M., Scellier, B., Bengio, Y., Grollier, J., & Querlioz, D. (2021). Scaling equilibrium propagation to deep ConvNets by drastically reducing its gradient estimator bias. *Frontiers in Neuroscience*, 15. 10.3389/fnins.2021.633674
- Lee, K., Maji, S., Ravichandran, A., & Soatto, S. (2019). Meta-learning with differentiable convex optimization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. Piscataway, NJ: IEEE.
- Liao, R., Xiong, Y., Fetaya, E., Zhang, L., Yoon, K., Pitkow, X., . . . Zemel, R. (2018). Reviving and improving recurrent back-propagation. In *Proceedings of the International Conference on Machine Learning*.
- Linnainmaa, S. (1976). Taylor expansion of the accumulated rounding error. *BIT Numerical Mathematics*, 16. 10.1007/BF01931367
- Lorraine, J., & Duvenaud, D. (2018). *Stochastic hyperparameter optimization through hypernetworks*. arXiv:1802.09419.
- Lorraine, J., Vicol, P., & Duvenaud, D. (2020). Optimizing millions of hyperparameters by implicit differentiation. In *Proceedings of the International Conference on Artificial Intelligence and Statistics*.
- Luketina, J., Berglund, M., Greff, K., & Raiko, T. (2016). Scalable gradient-based tuning of continuous regularization hyperparameters. In *Proceedings of the International Conference on Machine Learning*.
- MacKay, D. J. C. (1992). A practical Bayesian framework for backpropagation networks. *Neural Computation*, 4(3).
- MacKay, M., Vicol, P., Lorraine, J., Duvenaud, D., & Grosse, R. (2019). *Self-tuning networks: Bilevel optimization of hyperparameters using structured best-response functions*. arXiv:1903.03088.

- Marschall, O., Cho, K., & Savin, C. (2020). A unified framework of online learning algorithms for training recurrent neural networks. *Journal of Machine Learning Research*, 21(135).
- Menick, J., Elsen, E., Evcı, U., Osindero, S., Simonyan, K., & Graves, A. (2021). Practical real time recurrent learning with a sparse approximation. In *Proceedings of the International Conference on Learning Representations*.
- Metz, L., Maheswaranathan, N., Nixon, J., Freeman, C. D., & Sohl-Dickstein, J. (2019). Understanding and correcting pathologies in the training of learned optimizers. In *Proceedings of the International Conference on Machine Learning*.
- Metz, L., Poole, B., Pfau, D., & Sohl-Dickstein, J. (2017). Unrolled generative adversarial networks. In *Proceedings of the International Conference on Artificial Intelligence and Statistics*.
- Millar, W. (1951). Some general theorems for non-linear systems possessing resistance. *London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 42. 10.1080/14786445108561361
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., . . . Hassabis, D. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540). 10.1038/nature14236
- Mudrakarta, P. K., Sandler, M., Zhmoginov, A., & Howard, A. (2019). K for the price of 1: Parameter-efficient multi-task and transfer learning. In *Proceedings of the International Conference on Learning Representations*.
- Newton, I. (1670). *De methodis serierum et fluxionu*.
- Nikishin, E., Abachi, R., Agarwal, R., & Bacon, P.-L. (2022). Control-oriented model-based reinforcement learning with implicit differentiation. In *Proceedings of the AAAI Conference on Artificial Intelligence*. Palo Alto, CA: AAAI Press.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., . . . Chintala, S. (2019). PyTorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, & R. Garnett (Eds.), *Advances in neural information processing systems*, 32. Red Hook, NY: Curran.
- Pearlmutter, B. A. (1994). Fast exact multiplication by the Hessian. *Neural Computation*, 6(1). 10.1162/neco.1994.6.1.147
- Pedregosa, F. (2016). Hyperparameter optimization with approximate gradient. In *Proceedings of the International Conference on Machine Learning*.
- Pfau, D., & Vinyals, O. (2016). *Connecting generative adversarial networks and actor-critic methods*. arXiv:1610.01945.
- Pineda, F. J. (1987). Generalization of back-propagation to recurrent neural networks. *Physical Review Letters*, 59(19). 10.1103/PhysRevLett.59.2229
- Raghu, A., Raghu, M., Bengio, S., & Vinyals, O. (2020). Rapid learning or feature reuse? Towards understanding the effectiveness of MAML. In *Proceedings of the International Conference on Learning Representations*.
- Rajeswaran, A., Finn, C., Kakade, S., & Levine, S. (2019). Meta-learning with implicit gradients. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, & R. Garnett (Eds.), *Advances in neural information processing systems*, 32. Red Hook, NY: Curran.
- Rajeswaran, A., Mordatch, I., & Kumar, V. (2020). A game theoretic framework for model based reinforcement learning. In *Proceedings of the International Conference on Machine Learning*.

- Rao, R. P. N., & Ballard, D. H. (1999). Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, 2(1). 10.1038/4580
- Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning representations by back-propagating errors. *Nature*, 323(6088). 10.1038/323533a0
- Rusu, A. A., Rao, D., Sygnowski, J., Vinyals, O., Pascanu, R., Osindero, S., & Hadsell, R. (2019). Meta-learning with latent embedding optimization. In *Proceedings of the International Conference on Learning Representations*.
- Scarpello, G. M., & Ritelli, D. (2002). A historical outline of the theorem of implicit functions. *Divulgaciones Matemáticas*, 10(2).
- Scellier, B. (2021). *A deep learning theory for neural networks grounded in physics*. PhD diss., Université de Montréal.
- Scellier, B., & Bengio, Y. (2017). Equilibrium propagation: Bridging the gap between energy-based models and backpropagation. *Frontiers in Computational Neuroscience*, 11. 10.3389/fncom.2017.00024
- Scellier, B., & Bengio, Y. (2019). Equivalence of equilibrium propagation and recurrent backpropagation. *Neural Computation*, 31(2). 10.1162/neco\_a\_01160
- Scellier, B., Mishra, S., Bengio, Y., & Ollivier, Y. (2022). *Agnostic physics-driven deep learning*. arXiv:2205.15021.
- Schmidhuber, J. (1987). *Evolutionary principles in self-referential learning, or on learning how to learn: the meta-meta...-hook*. Diploma thesis, Institut für Informatik, Technische Universität München.
- Shaban, A., Cheng, C.-A., Hatch, N., & Boots, B. (2019). Truncated backpropagation for bilevel optimization. In *Proceedings of the International Conference on Artificial Intelligence and Statistics*.
- Shewchuk, J. R. (1994). *An introduction to the conjugate gradient method without the agonizing pain*. Pittsburgh, PA: School of Computer Science, Carnegie Mellon University.
- Song, Y., Schwing, A. G., Zemel, R. S., & Urtasun, R. (2016). Training deep neural networks via direct loss minimization. In *Proceedings of the International Conference on Machine Learning*.
- Stern, M., Hexner, D., Rocks, J. W., & Liu, A. J. (2021). Supervised learning in physical networks: From machine learning to learning machines. *Physical Review X*, 11(2). 10.1103/PhysRevX.11.021045
- Sutton, R. S. (1992). Adapting bias by gradient descent: An incremental version of delta-bar-delta. In *Proceedings of the National Conference on Artificial Intelligence*. Palo Alto, CA: AAAI Press.
- Tallic, C., & Ollivier, Y. (2018). Unbiased online recurrent optimization. In *Proceedings of the International Conference on Learning Representations*.
- Thrun, S., & Pratt, L. (1998). *Learning to learn*. Berlin: Springer.
- Vlastelica, M., Paulus, A., Musil, V., Martius, G., & Rolínek, M. (2020). Differentiation of blackbox combinatorial solvers. In *Proceedings of the International Conference on Learning Representations*.
- von Stackelberg, H. (1934). *Market structure and equilibrium*. Berlin: Springer.
- Wang, P.-W., Donti, P. L., Wilder, B., & Kolter, Z. (2019). SATNet: Bridging deep learning and logical reasoning using a differentiable satisfiability solver. In *Proceedings of the International Conference on Machine Learning*.

- Werbos, P. J. (1982). Applications of advances in nonlinear sensitivity analysis. In A. V. Balakrishnan & M. Thoma (Eds.), *System modeling and optimization*. Berlin: Springer-Verlag.
- Werbos, P. J. (1990). Backpropagation through time: What it does and how to do it. *Proceedings of the IEEE*, 78(10). 10.1109/5.58337
- Whittington, J. C. R., & Bogacz, R. (2017). An approximation of the error backpropagation algorithm in a predictive coding network with local Hebbian synaptic plasticity. *Neural Computation*, 29(5). 10.1162/NECO\_a\_00949
- Williams, R. J., & Peng, J. (1990). An efficient gradient-based algorithm for on-line training of recurrent network trajectories. *Neural Computation*, 2(4). 10.1162/neco.1990.2.4.490
- Williams, R. J., & Zipser, D. (1989). A learning algorithm for continually running fully recurrent neural networks. *Neural Computation*, 1(2). 10.1162/neco.1989.1.2.270
- Wyatt, J. L., & Standley, D. L. (1989). Criteria for robust stability in a class of lateral inhibition networks coupled through resistive grids. *Neural Computation*, 1(1). 10.1162/neco.1989.1.1.58
- Yang, Z., Chen, Y., Hong, M., & Wang, Z. (2019). On the global convergence of actor-critic: A case for linear quadratic regulator with ergodic cost. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, & R. Garnett (Eds.), *Advances in neural information processing systems*, 32. Red Hook, NY: Curran.
- Zhao, D., Kobayashi, S., Sacramento, J., & von Oswald, J. (2020). Meta-learning via hypernetworks. In *Workshop on Meta-Learning at NeurIPS*.
- Zheng, L., Fiez, T., Alumbaugh, Z., Chasnov, B., & Ratliff, L. J. (2021). Stackelberg actor-critic: Game-theoretic reinforcement learning algorithms. In *Proceedings of the AAAI Conference on Artificial Intelligence*. Palo Alto, CA: AAAI Press.
- Zhou, W., Li, Y., Yang, Y., Wang, H., & Hospedales, T. M. (2020). Online meta-critic learning for off-policy actor-critic methods. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, & H. Lin (Eds.) *Advances in neural information processing systems*, 33. Red Hook, NY: Curran.
- Zintgraf, L., Shiarlis, K., Kurin, V., Hofmann, K., & Whiteson, S. (2019). Fast context adaptation via meta-learning. In *Proceedings of the International Conference on Learning Representations*.
- Zuchet, N., Schug, S., von Oswald, J., Zhao, D., & Sacramento, J. (2021). *A contrastive rule for meta-learning*. arXiv:2104.01677.