arXiv:2104.01677v3 [cs.LG] 3 Oct 2022

A contrastive rule for meta-learning

Nicolas Zucchet[*]
Department of Computer Science
ETH Zürich
nzucchet@inf.ethz.ch

Simon Schug[*]
Institute of Neuroinformatics
University of Zürich & ETH Zürich
sschug@ethz.ch

Johannes von Oswald[*]
Department of Computer Science
ETH Zürich
voswaldj@ethz.ch

Dominic Zhao
Institute of Neuroinformatics
University of Zürich & ETH Zürich
dozhao@ethz.ch

João Sacramento
Institute of Neuroinformatics
University of Zürich & ETH Zürich
rjoao@ethz.ch

Abstract

Humans and other animals are capable of improving their learning performance as they solve related tasks from a given problem domain, to the point of being able to learn from extremely limited data. While synaptic plasticity is generically thought to underlie learning in the brain, the precise neural and synaptic mechanisms by which learning processes improve through experience are not well understood. Here, we present a general-purpose, biologically-plausible meta-learning rule which estimates gradients with respect to the parameters of an underlying learning algorithm by simply running it twice. Our rule may be understood as a generalization of contrastive Hebbian learning to meta-learning and notably, it neither requires computing second derivatives nor going backwards in time, two characteristic features of previous gradient-based methods that are hard to conceive in physical neural circuits. We demonstrate the generality of our rule by applying it to two distinct models: a complex synapse with internal states which consolidate task-shared information, and a dual-system architecture in which a primary network is rapidly modulated by another one to learn the specifics of each task. For both models, our meta-learning rule matches or outperforms reference algorithms on a wide range of benchmark problems, while only using information presumed to be locally available at neurons and synapses. We corroborate these findings with a theoretical analysis of the gradient estimation error incurred by our rule.[1]

1 Introduction

The seminal study of Harlow [1] established that humans and non-human primates can become better at learning when presented with a series of learning tasks which share a certain common structure. To achieve this, the brain must extract and encode whichever aspects are common within a problem domain, in such a way that future learning performance is improved. This capacity, which we refer to as meta-learning, confers great evolutionary advantage to an organism over another that must

[*]Equal contribution; arbitrary ordering.
[1]Code available at https://github.com/smonsays/contrastive-meta-learning

36th Conference on Neural Information Processing Systems (NeurIPS 2022).

face new tasks starting from *tabula rasa*. The neural and synaptic basis of this higher-order form of learning is largely unknown and theories are notably scarce [2]. The present work focuses on developing one such theory.

Formally, we define learning as the optimization of a data-dependent objective function with respect to learnable parameters, following the prevalent view in machine learning [3]. Meta-learning can be straightforwardly accommodated for in this framework by first specifying a learning algorithm through a set of meta-parameters, and then measuring post-learning performance through a meta-objective function [4–8]. Formulated as such, meta-learning corresponds to a hierarchical optimization problem, where lower-level parameters are optimized to learn the specifics of each task, and meta-parameters are adapted over tasks to improve overall learning performance.

An essential question in this framework is how to optimize meta-parameters. In current deep learning practice, meta-parameters are almost always learned by backpropagation-through-learning, an instance of backpropagation-through-time [9]. While a number of biologically-plausible designs [3, 10–13] have been developed for the standard error backpropagation algorithm for feedforward neural networks [14, 15], backpropagation-through-learning suffers from a number of issues which appear to be fundamentally difficult to overcome in biological circuits. For example, when learning involves optimizing synaptic connection weights – as it is presumed to be the case in the brain – implementing backpropagation-through-learning would entail backtracking through a sequence of synaptic changes in reverse-time order, while carrying out operations which would require knowledge of all synaptic weights to be available at a single synapse. This is clearly at odds with what is currently known about synaptic plasticity. Thus, calculating meta-parameter gradients by backpropagation is both computationally expensive, and hard to reconcile with biological constraints.

Here we present a meta-learning rule for adapting meta-parameters which does not exhibit such issues. Instead of backpropagating through a learning process, our rule estimates meta-parameter gradients by running the underlying learning algorithm twice: learning a task is followed by a second run to solve an augmented learning problem which includes the meta-objective. Our rule has a number of appealing properties: (1) it runs forward in time, making the learning rule causal; (2) implementing it only requires temporarily buffering one intermediate state; (3) it does not evaluate second derivatives, thus avoiding accessing information that is non-local to a parameter; and (4) it approximates meta-gradients as accurately as needed. Furthermore, our rule is generically applicable and it can be used to learn any meta-parameter which influences the meta-objective function.

The local and causal nature of our rule allows us to develop a theory of meta-plastic synapses, which slowly consolidate information over tasks in their internal hidden states or in their synaptic weights. We show through experiments that, when governed by our meta-learning rule, such slow adaptation processes result in improved learning performance in a variety of benchmark problems and network architectures, from deep convolutional to recurrent spiking neural networks, on both supervised and reinforcement learning paradigms. Moreover, we find that our meta-learning rule performs as well or better than reference methods, including backpropagation-through-learning, and we provide a theoretical bound for its meta-gradient estimation error which is confirmed by our experimental findings. Thus, our results demonstrate that gradient-based meta-learning is possible with local learning rules, and suggest ways by which slower synaptic processes in the brain optimize the performance of faster learning processes.

2 Background and problem setup

The goal of meta-learning is to improve the performance of a learning algorithm through experience. We begin by formalizing this goal as a mathematical optimization problem and outlining its solution with standard gradient-based methods. The approach we present below underlies a large body of work studying meta-learning in neural networks [e.g., 7, 16–18]. We also discuss why these standard methods may be deemed unsatisfactory as models of meta-learning in the brain.

Problem setup. Formally, we wish to optimize the meta-parameters θ of an algorithm which learns to solve a given task τ by changing the parameters ϕ of a model. Each task is drawn from a distribution $p(\tau)$ representing the problem domain and comes with an associated loss function $L_\tau^{\text{learn}}(\phi, \theta)$, which depends on some data D_τ^{learn}. The goal of learning is to minimize this loss while keeping the meta-parameters θ fixed; we denote the outcome of learning task τ by $\phi^*_{\theta,\tau}$. The subscript

θ in $\phi^*_{\theta,\tau}$ is here to emphasize that the solution of a task implicitly depends on the meta-parameters θ used during learning. Learning performance is then evaluated by measuring again a loss function $L^{\text{eval}}_\tau(\phi^*_{\theta,\tau}, \theta)$, defined on new evaluation data D^{eval}_τ from the same task. The meta-objective is this evaluation loss, averaged over tasks. Hence, we formalize meta-learning as a bilevel optimization problem, which can be compactly written as follows:

$$\min_\theta \ \mathbb{E}_{\tau\sim p(\tau)}\left[L^{\text{eval}}_\tau(\phi^*_{\theta,\tau}, \theta)\right] \qquad \text{s.t.} \ \ \phi^*_{\theta,\tau} \in \arg\min_\phi L^{\text{learn}}_\tau(\phi, \theta). \tag{1}$$

In this paper, we approach problem (1) with stochastic gradient descent, which uses meta-gradient information to update meta-parameters after learning a task (or a minibatch of tasks) presented by the environment. For a given task τ we thus need to compute the meta-gradient

$$\nabla_{\theta,\tau} := \left(\frac{\mathrm{d}}{\mathrm{d}\theta} L^{\text{eval}}_\tau(\phi^*_{\theta,\tau}, \theta)\right)^\top. \tag{2}$$

The implicit dependence of $\phi^*_{\theta,\tau}$ on the meta-parameters θ complicates the computation of the meta-gradient; differentiating through the learning algorithm efficiently is a central question in gradient-based meta-learning. We next review two major known ways of doing so.

Review of backpropagation-through-learning. A common strategy followed in previous work [cf. 19] is to replace the solution $\phi^*_{\theta,\tau}$ to a learning task by the result $\phi_{\theta,\tau,T}$ obtained after applying a differentiable learning algorithm for T time steps, not necessarily until convergence. One advantage of this formulation is that the computational graph for $\phi_{\theta,\tau,T}$ is explicitly available. Thus, backpropagation can be invoked to compute the meta-gradient $\nabla_{\theta,\tau}$, yielding what we refer to as backpropagation-through-learning. This approach is hardly biologically-plausible, as it requires storing and revisiting the parameter trajectory $\{\phi_t\}_{t=1}^T$ backwards in time, from $t = T$ to $t = 0$. Moreover, when the learning algorithm which produces $\phi_{\theta,\tau,T}$ is itself gradient-based, as it typically is in deep learning, differentiating through learning gives rise to second derivatives. These second-order terms involve cross-parameter dependencies that are difficult to resolve with local processes.

Review of implicit differentiation. An alternative line of methods [20–23] approaches problem (1) through the implicit function theorem [24]. This theorem provides conditions under which the meta-gradient $\nabla_{\theta,\tau}$ is well-defined, while also providing a formula for it. Over backpropagation-through-learning, this approach has the advantages that it does not require storing parameter trajectories $\{\phi_t\}_{t=1}^T$, and that it is agnostic to which algorithm is used to learn a task. However, the meta-gradient formula provided by the implicit function theorem is difficult to evaluate directly for neural network models, as it includes the inverse learning loss Hessian. This makes it hard to design biologically-plausible meta-learning algorithms based directly on the implicit meta-gradient expression. We refer to Section S2 for more details and an expanded discussion on this class of meta-learning methods.

3 Contrastive meta-learning

Here we present a new meta-learning rule which is generically applicable to meta-learning problems of the form (1). Our rule is gradient-following, and therefore scalable to neural network problems involving high-dimensional meta-parameters, while being simpler to conceive in biological neural circuits than the standard gradient-based methods reviewed in the previous section.

To derive our meta-learning rule we first introduce an auxiliary objective function which mixes the two levels of the bilevel optimization problem (1):

$$\mathcal{L}_\tau(\phi, \theta, \beta) = L^{\text{learn}}_\tau(\phi, \theta) + \beta L^{\text{eval}}_\tau(\phi, \theta). \tag{3}$$

We refer to $\mathcal{L}_\tau(\phi, \theta, \beta)$ as the augmented loss function. This auxiliary loss depends on a new scalar parameter $\beta \in \mathbb{R}$, which we call the nudging strength. Positive values of β nudge learning towards the meta-objective associated with task τ. Thus, we can define a family of auxiliary learning problems through the augmented loss $\mathcal{L}_\tau$ by varying the nudging strength β away from zero. We denote the solutions to these auxiliary learning problems by

$$\phi^*_{\theta,\beta,\tau} \in \arg\min_\phi \mathcal{L}_\tau(\phi, \theta, \beta), \tag{4}$$

and we use $\hat{\phi}_{\theta,\beta,\tau}$ to distinguish approximate model parameters found in practice with some learning algorithm from the true minimizers $\phi^*_{\theta,\beta,\tau}$. Note that for the special case of $\beta = 0$, we recover a solution $\phi^*_{\theta,0,\tau}$ of the original learning task defined by $L^{\text{learn}}_\tau(\phi, \theta)$.

Our contrastive meta-learning rule prescribes the following change to the meta-parameters θ after encountering learning task τ:

$$\Delta_{\theta,\tau} := -\frac{1}{\beta}\left(\frac{\partial \mathcal{L}_\tau}{\partial \theta}(\hat{\phi}_{\theta,\beta,\tau}, \theta, \beta) - \frac{\partial \mathcal{L}_\tau}{\partial \theta}(\hat{\phi}_{\theta,0,\tau}, \theta, 0)\right)^\top. \tag{5}$$

This rule contrasts information over two model parameter settings, $\hat{\phi}_{\theta,0,\tau}$ and $\hat{\phi}_{\theta,\beta,\tau}$; it may be understood as a generalization to meta-learning of a classical recurrent neural network learning algorithm known as contrastive Hebbian learning [25–29]. Intuitively, as we compute the solution to the augmented learning problem with $\beta > 0$, we nudge our learning algorithm towards a parameter setting $\hat{\phi}_{\theta,\beta,\tau}$ that would have been better in terms of the meta-objective — that we wish our algorithm had actually reached, without needing the meta-objective to influence the learning process.

Our rule implements meta-learning by gradient descent when the learning solutions $\hat{\phi}_{\theta,0,\tau}$ and $\hat{\phi}_{\theta,\beta,\tau}$ are exact and as $\beta \to 0$. This important property can be shown by invoking the equilibrium propagation theorem [29, 30] discovered and proved by Scellier and Bengio; we restate this result and present the technical conditions for applying it to meta-learning in Section S1. Critically, $\Delta_{\theta,\tau}$ estimates the meta-gradient $\nabla_{\theta,\tau}$ using only partial derivative information and without ever directly calculating the total derivative in (2). Depending on the model, partial derivatives of the augmented loss $\mathcal{L}_\tau$ may be easy to calculate analytically and implement, or they may require dedicated neural circuits for their evaluation; we return to this point in the next section.

We recall that the two points $\hat{\phi}_{\theta,0,\tau}$ and $\hat{\phi}_{\theta,\beta,\tau}$ which appear in (5) respectively correspond to approximate solutions of the original and the augmented learning problems. Thus, the information required to implement our rule can be collected causally by invoking the learning algorithm for a second time, after the actual task has been learned, while buffering information across the two runs. In contrast to backpropagation-through-learning, this process runs forward in time, it only requires keeping a single intermediate state in short-term memory, and it is entirely agnostic to which underlying learning algorithm is used. Moreover, as we will show in the theoretical results, its precision can be varied; the same rule can produce both coarse- and fine-grained meta-gradient estimates as needed, by varying the amount of resources spent in learning and by controlling the nudging strength β.

4 Models

In the previous section, our contrastive meta-learning rule was presented in its general form. We now describe two concrete neural models that provide complementary views on how meta-learning could be conceived in the brain. We study the specific meta-learning rules arising from the application of the update (5) to each case and discuss their implementation with biological neural circuitry.

4.1 Synaptic consolidation as meta-learning

We first use our general contrastive meta-learning rule (5) to derive meta-plasticity rules for a complex synapse model which has been featured in prior meta-learning [22, 31] and continual learning [32, 33] work. Biological synapses are complex devices which comprise components that adapt at multiple time scales. Beyond changes induced by standard long-term potentiation and depression protocols lasting minutes to several hours, synapses exhibit activity-dependent plasticity at much longer time scales [34–36]. While previous work has focused on characterizing memory retention in more realistic synapse models, here we study how such slow synaptic consolidation processes may support fast future learning through our contrastive meta-learning rule.

In the model we consider, besides a synaptic weight ϕ which influences postsynaptic activity, each synapse has an internal consolidated state ω towards which the weight is attracted whenever the synapse changes. We further allow the attraction strength λ to vary over synapses; its reciprocal λ^{-1} plays a role similar to a learning rate. For this model the meta-parameters are therefore $\theta = \{\lambda, \omega\}$.

We model the interaction between these three components through a quadratic function, which is added to the task-specific learning loss $l_\tau^{\text{learn}}(\phi)$:

$$L_\tau^{\text{learn}}(\phi, \theta) = l_\tau^{\text{learn}}(\phi) + \frac{1}{2} \sum_{i=1}^{|\phi|} \lambda_i (\omega_i - \phi_i)^2. \tag{6}$$

In machine learning terms, we regularize the learning loss with a quadratic regularizer. On the other hand, the evaluation loss function $L_\tau^{\text{eval}}(\phi)$ depends only on the synaptic weights ϕ such that the meta-parameters θ only influence learning, not prediction.

The partial derivatives which appear in our contrastive meta-learning rule (5) can be analytically obtained for this synaptic model. A calculation yields the meta-plasticity rules

$$\Delta_{\omega,\tau} = \frac{\lambda}{\beta} \left(\hat{\phi}_{\theta,\beta,\tau} - \hat{\phi}_{\theta,0,\tau} \right) \quad \text{and} \quad \Delta_{\lambda,\tau} = \frac{1}{2\beta} \left[(\hat{\phi}_{\theta,0,\tau} - \omega)^2 - (\hat{\phi}_{\theta,\beta,\tau} - \omega)^2 \right], \tag{7}$$

where all operations are carried out elementwise. Contrastive meta-learning thus offers a principled way to slowly (over learning tasks) consolidate information in the internal states of complex synapses to improve future learning performance. Critically, it leads to meta-plasticity rules that are entirely local to a synapse and are independent of the method used to learn. Our meta-plasticity rules can thus be flexibly applied to improve the performance of any learning algorithm, including a host of biologically-plausible learning rules, from precise neuron-specific error backpropagation circuits [37, 38] to stochastic perturbation reinforcement rules [39]. The only requirement our theory makes is that learning corresponds to the optimization of an objective.

4.2 Learning by top-down modulation

The second model that we consider is inspired by the modulatory role that is attributed to top-down inputs from higher- to lower-order brain areas. Such modulatory inputs often feature in neural theories of attention and contextual processing [40–42]. Here, we explore the possibility that they subserve fast learning of new tasks. We incorporate this insight into a simple meta-learning model, where learning a task τ corresponds to finding the right pattern of task-specific modulation $\phi^*_{\theta,\tau}$, and meta-learning corresponds to changing synaptic weights θ. Unlike in the complex synapse model presented in the previous section, here we interpret the task-specific parameters ϕ_τ as patterns of neural activity, not synaptic weights. This implies that, if meta-learning succeeds, it becomes possible to learn new tasks on the fast neural time scale without evoking synaptic plasticity.

More concretely, we take as modulatory inputs a multiplicative gain g and an adaptive threshold b per neuron, as done in previous work [43, 44]. Rapid (input-dependent) multiplicative and additive modulation of the sensitivity of the neural input-output response curve $\sigma(x)$ is typically observed in cortical neurons [45]. There exist a number of biophysical mechanisms which allow top-down inputs to modulate $\sigma(x)$ [e.g., 46]. Assuming a simple linear-threshold neuron model with weights θ, this yields the response $\sigma(x) = g(\theta \cdot x - b)_+$ to some input x, where $(\cdot)_+$ denotes the positive-part operation. In this model, there are only few learnable parameters $\phi = \{g, b\}$, as they scale with the number of neurons and not with the number of synaptic connections.

We apply contrastive meta-learning to this model by changing synaptic weights θ according to our rule (5). For this model, partial derivatives of the augmented loss function correspond to the usual derivatives with respect to model parameters that are routinely evaluated to learn deep neural networks; our rule simply asks to compute them twice. We therefore build upon existing theories of learning by backpropagation-of-error in the brain and assume that some mechanism for neuron-specific spatial error backpropagation is available, for example via prediction error neural subpopulations [37] or dendritic error representations [11, 38, 47], or by invoking equilibrium propagation again [29].

5 Theoretical and experimental analyses

In the following, we theoretically analyze the approximation error incurred by our contrastive meta-learning rule before empirically testing it on a suite of meta-learning problems. The objective of our experiments is twofold. First, we aim to confirm our theoretical results and demonstrate the performance of contrastive meta-learning on standard machine learning benchmarks. Second, we want to illustrate the generality of our approach by applying it to various supervised and reinforcement meta-learning problems as well as to a more biologically realistic neuron and plasticity model.

5.1 Theoretical analysis of the meta-gradient approximation error

The contrastive meta-learning rule (5) only provides an approximation to the meta-gradient. This approximation can be improved by refining the two learning solutions $\hat{\phi}_{\theta,0,\tau}$ and $\hat{\phi}_{\theta,\beta,\tau}$ through additional computation or by using a better learning algorithm, and by decreasing the nudging strength β, as prescribed by the equilibrium propagation theorem. In Theorem 1, we theoretically analyze how the meta-gradient estimate (5) benefits from such improvements (see Fig. 1A for a visualization of the result, Section S3 for a proof and empirical verification of our theoretical results). We find that the refinement of the learning solutions must be coupled to a decrease in β: too small β greatly detracts from the quality of the meta-gradient estimate when the solutions are not improved, while better approximations are inefficient if β is not decreased accordingly.

Theorem 1 (Informal)**.** *Let $\beta > 0$ and δ be such that $\|\hat{\phi}_{\theta,0,\tau} - \phi^*_{\theta,0,\tau}\| \leq \delta$ and $\|\hat{\phi}_{\theta,\beta,\tau} - \phi^*_{\theta,\beta,\tau}\| \leq \delta$. Then, under regularity and convexity assumptions, there exists a constant C such that*

$$\|-\Delta_{\theta,\tau} - \nabla_{\theta,\tau}\| \leq C \left(\frac{1+\beta}{\beta}\,\delta + \frac{\beta}{1+\beta} \right) =: \mathcal{B}(\delta, \beta).$$

5.2 Contrastive meta-learning is a high-performance meta-optimization algorithm

As a first set of experiments, we study a supervised meta-optimization problem based on the entire CIFAR-10 image dataset [48]. In these experiments the goal is to meta-learn a set of hyperparameters (meta-parameters) such that generalization performance improves. This problem is a common testbed for assessing the ability of a meta-learning algorithm to optimize a given meta-objective [23]; it can be thought of as a limiting case of full meta-learning, as there are learnable meta-parameters, but only one task. As the meta-objective we take the cross-entropy loss l evaluated on a held-out dataset D^{eval}: $L^{\text{eval}}(\phi) = \frac{1}{|D^{\text{eval}}|} \sum_{(x,y)\in D^{\text{eval}}} l(x, y, \phi)$, where x is an image input and y its label. We equip a convolutional deep neural network with our synaptic model (6), meta-learning only the per-synapse regularization strength λ, keeping ω fixed at zero: $L^{\text{learn}}(\phi, \lambda) = \frac{1}{|D^{\text{learn}}|} \sum_{(x,y)\in D^{\text{learn}}} l(x, y, \phi) + \frac{1}{2} \sum_{i=1}^{|\phi|} \lambda_i \phi_i^2$. We learn the weights ϕ by stochastic gradient descent paired with backpropagation. Additional details and analyses may be found in Section S4.1.

We benchmark our meta-plasticity rule (7) against implicit gradient-based meta-learning methods, which are considered state-of-the-art for this type of problem [23] (see Section S2 for a review). More concretely, recurrent backpropagation (RBP [49, 50]; also known as the Neumann series approximation [23, 51]) and the conjugate gradient method (CG) [21, 52] correspond to two different numerical schemes for calculating the meta-gradient; T1-T2 [53] is an approximate method which neglects complicated terms, thus introducing a non-reducible bias in the meta-gradient estimate. Critically, unlike our contrastive meta-learning rule (CML), this method offers no control over the meta-gradient error.

Table 1: Meta-learning a per-synapse regularization strength meta-parameter (cf. Section 4.1) on CIFAR-10. Average accuracies (acc.) $\pm$ s.e.m. over 10 seeds.

Method	Evaluation acc. (%)	Test acc. (%)
T1-T2	$64.77^{\pm 0.40}$	$62.57^{\pm 0.31}$
CG	$57.65^{\pm 1.51}$	$57.51^{\pm 0.98}$
RBP	$64.92^{\pm 1.32}$	$62.14^{\pm 0.97}$
CML	$74.43^{\pm 0.53}$	$66.94^{\pm 0.25}$
No meta	$60.06^{\pm 0.37}$	$60.13^{\pm 0.38}$
TBPTL	$73.17^{\pm 0.27}$	$65.35^{\pm 0.36}$

We find that our meta-learning rule outperforms all three baseline implicit differentiation methods in terms of both evaluation-set and actual generalization (test-set) performance, cf. Tab. 1. As a side result, we confirm the instability of CG in deep learning reported in ref. [51, 54]. We note that the hyperparameters of all four methods were independently and carefully set (cf. Section S4.1). These strong results on a modern deep learning benchmark, involving stochastic approximate learning, demonstrate that contrastive meta-learning is a scalable, highly effective meta-optimization algorithm. Moreover, Theorem 1 is in excellent qualitative agreement with our experiments, cf. Fig. 1.

To further contextualize our findings, we provide results for training the same network without meta-learning, where we performed a conventional hyperparameter search over a scalar regularization strength hyperparameter shared by all synapses. This simple approach yields only a moderate evaluation and test accuracy.

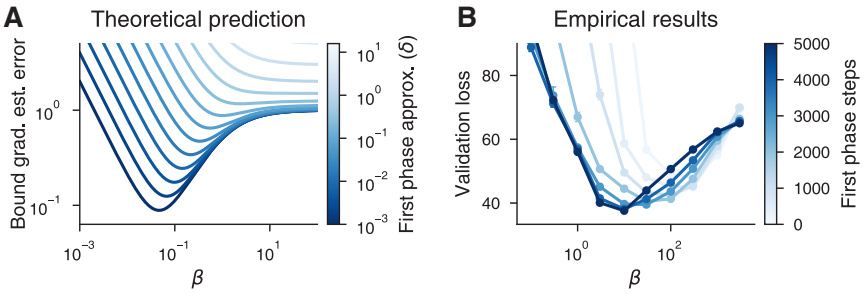

Figure 1: (A) Visualization of the theoretical bound $\mathcal{B}$ on the meta-gradient estimation error from Theorem 1 as a function of the nudging strength β. Better approximations of the solutions (smaller δ) improve the quality of the meta-gradient, as they enable using smaller values of β. (B) Confirmation of the qualitative findings of the theory on deep learning experiments. We show results for a hyperparameter meta-learning problem, where a per-synapse regularization strength is meta-learned (cf. Section 4.1) on CIFAR-10 with rule (7). The validation loss is a proxy for the quality of the gradient and the number of steps in the first phase is a proxy for $-\log \delta$.

As all methods incur numerical errors when computing the meta-gradient, a comparison to using the analytical solution for the meta-gradient would be desirable. Since this is intractable in this case and running full backpropagation-through-learning requires too much memory, we evaluate truncated backpropagation-through-learning (TBPTL) with the maximal truncation window we can fit on a single graphics processing unit (in our case 200 out of 5000 steps). The resulting evaluation accuracy and test accuracy outperform other implicit gradient-based meta-learning methods but are still surpassed by our method.

5.3 Contrastive meta-learning enables visual few-shot learning

The ability to learn new object classes based on only a few examples is a hallmark of human intelligence [55] and a prime application of meta-learning. We test whether our contrastive meta-learning rule is able to turn into a few-shot learner a standard visual system, a convolutional deep neural network learned by gradient descent and error backpropagation. Furthermore, we ask how our contrastive meta-learning rule fares against other gradient-based meta-learning algorithms which rely on backpropagation-through-learning and implicit differentiation to compute gradients. To that end, we focus on two widely-studied few-shot image classification problems based on miniImageNet [56] and the Omniglot [57] datasets. To further facilitate comparisons, we reproduce exactly the experimental setup of ref. [18], which has been adopted in a large number of studies.

Briefly, during meta-learning, N-way K-shot tasks are created on-the-fly by sampling N classes at random from a fixed pool of classes, and then splitting the data into task-specific learning D_τ^{learn} (with K examples per class for learning) and evaluation D_τ^{eval} sets, used to define the corresponding loss functions L_τ^{learn} and L_τ^{eval}. The meta-objective is then simply the task-averaged evaluation loss, measured after learning. The performance of the learning algorithm is tested on new tasks consisting of classes that were not seen during meta-learning. We provide all experimental details in Section S4.2.

As reference methods, we compare against the well-known model-agnostic meta-learning (MAML) algorithm [18], which relies on backpropagation-through-learning to meta-learn an initial set of weights, starting from which a few gradient steps should succeed; this is conceptually similar to meta-learning the consolidated state ω of our complex synapses. We also include results obtained with its first-order approximation FOMAML (as well as a closely related algorithm known as Reptile [58]), which, like the T1-T2 algorithm of the previous section, excludes all second-order terms from the meta-gradient estimate to simplify the update, at the expense of introducing a bias. Finally,

Table 2: One-shot miniImageNet learning. Averages over 5 seeds $\pm$ std.

Method	Test acc. (%)
MAML [18]	48.70 $^{\pm 1.84}$
FOMAML [18]	48.07 $^{\pm 1.75}$
Reptile [58]	49.97 $^{\pm 0.32}$
iMAML [22]	48.96 $^{\pm 1.84}$
CML (synaptic)	48.43 $^{\pm 0.43}$
CML (modulatory)	49.80 $^{\pm 0.40}$

we compare to the implicit MAML (iMAML) algorithm [22], which corresponds exactly to meta-learning our consolidated synaptic state ω, but with implicit differentiation methods.

When applied to the problem domain of miniImageNet one-shot learning tasks, the performance of all meta-learning algorithms we consider here is closely clustered together, cf. Tab. 2. In particular, meta-learning the consolidated states ω of our complex synapses with implicit differentiation (iMAML) or our local update (7) leads to comparable performance. Interestingly, we further find that miniImageNet one-shot learning performance is significantly improved when using the modulatory model described in Section 4.2, despite the low dimensionality of the task-specific variable ϕ. This is in line with other results suggesting that highly efficient visual learning of new categories may be possible without necessarily engaging synaptic plasticity [43]. On Omniglot (see Section S4.2 for additional variants), the situation is comparable, except that on its 20-way 1-shot variant, the performance gap between first- and second-order methods widens. In line with our theory, our contrastive meta-learning rule performs close to (second-order) implicit differentiation, showing that despite its simplicity and locality our rule is able to accurately estimate meta-gradients.

Table 3: Omniglot character few-shot learning. Test set classification accuracy (%) averaged over 5 seeds $\pm$ std.

Method	20-way 1-shot	20-way 5-shot
MAML [18]	$95.8^{\pm 0.3}$	$98.9^{\pm 0.2}$
FOMAML [18]	$89.4^{\pm 0.5}$	$97.9^{\pm 0.1}$
Reptile [58]	$89.43^{\pm 0.14}$	$97.12^{\pm 0.32}$
iMAML [22]	$94.46^{\pm 0.42}$	$98.69^{\pm 0.1}$
CML (synaptic)	$94.16^{\pm 0.12}$	$98.06^{\pm 0.26}$
CML (modulatory)	$94.24^{\pm 0.39}$	$98.60^{\pm 0.27}$

5.4 Contrastive meta-learning enables meta-plasticity in a recurrent spiking network

For the experiments described on the previous sections we used simple artificial neuron models and backpropagation-of-error to learn. We now move closer to a biological neuron and plasticity model and consider meta-learning in a recurrently-connected neural network of leaky integrate-and-fire neurons with plastic synapses. We study a simple few-shot regression problem [18], where the aim is to quickly learn to approximate sinusoidal functions which differ in their phase and amplitude (for additional details see Section S4.3). For each task, we measure the mean squared error on 10 samples for the learning loss and 10 samples for the evaluation loss. We implement synaptic plasticity using the local e-prop rule [59] and use a population of 100 Poisson neurons to encode inputs, see Fig. 2A. As our contrastive meta-learning rule (5) is agnostic to the specifics of the learning process, we can augment the model with our synaptic consolidation model and apply the meta-plasticity rules derived in (7). Fig. 2B illustrates how the learning process improves with increasing number of tasks encountered, eventually consolidating a sinusoidal prior that can be quickly adapted to the specifics of a task from few examples, cf. Fig. 2C.

We compare our method to a standard baseline where updates are computed by back-propagating through the synaptic plasticity process (backpropagation-through-learning; BPTL) using surrogate gradients to handle spiking nonlinearities [60] similar to previous work on spiking neuron meta-learning [61]. Since full BPTL requires reducing the number of learning steps compared to our method due to memory constraints, we also include TBPTL with the same number of 500 learning steps and a truncation window of 100 steps. In both cases, we find competitive performance for our method, see Tab. 4.

Table 4: Few-shot learning of sinusoidal functions with a recurrent spiking neural network. Avg. mean squared error (MSE) over 10 seeds $\pm$ s.e.m.

Method	Validation MSE	Test MSE
BPTL + BPTT	$0.17^{\pm 0.01}$	$0.41^{\pm 0.10}$
BPTL + e-prop	$0.52^{\pm 0.05}$	$0.72^{\pm 0.08}$
TBPTL + e-prop	$0.27^{\pm 0.07}$	$0.50^{\pm 0.11}$
CML + e-prop	$0.23^{\pm 0.04}$	$0.23^{\pm 0.04}$

5.5 Contrastive meta-learning improves reward-based learning

Finally, we demonstrate how contrastive meta-learning can be applied in the challenging setting of reward-based learning, second nature to most animals. Reward-based learning clearly demonstrates hallmarks of meta-learning as animals are capable of flexibly remapping reward representations when task contingencies change [62, 63]. Inspired by this, we aim to meta-learn a value function on a

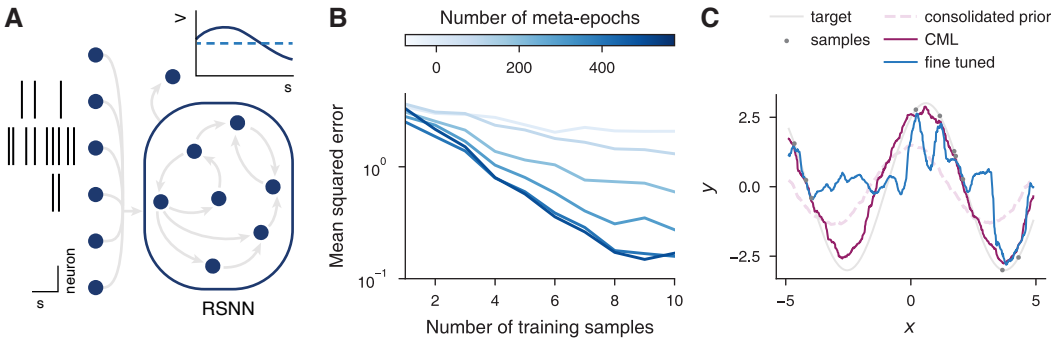

Figure 2: (A) A network of recurrently-connected leaky-integrate and fire neurons is tasked with learning sinusoids on an input encoding of Poisson spike trains. Its prediction is the voltage of the output neuron averaged over time. (B) Learning performance from few examples measured as the mean squared error on evaluation examples during a learning episode improves as more tasks are encountered over the course of meta-learning. (C) Meta-plasticity encodes information on the consolidated synaptic component (dashed) which results in improved learning performance (purple), compared to a naive network learning from scratch (blue).

family of reward-based learning tasks that can be quickly adapted to predict the expected reward of the actions available to the agent in a particular task.

Specifically, we consider the wheel bandit problem introduced by [64] with the meta-learning setup previously studied in refs. [65, 66]. On each task, an agent is presented with a sequence of context coordinates randomly drawn from a unit circle for each of which it has to choose among 5 actions to receive a stochastic reward. Hidden to the agent, a task-specific radius δ tiles the context space into a low- and a high-reward region depending on which the optimal action to take changes (see Section S4.4).

Table 5: Cumulative regret on the wheel bandit problem for different δ. Values normalized by the cumulative regret of a uniformly random agent. Avgs. over 50 seeds $\pm$ s.e.m.

δ	0.5	0.9	0.99
NeuralLinear [64]	$0.95^{\pm 0.02}$	$4.65^{\pm 0.18}$	$49.63^{\pm 2.41}$
MAML	$0.45^{\pm 0.01}$	$1.02^{\pm 0.76}$	$15.21^{\pm 1.69}$
CML (synaptic)	$0.40^{\pm 0.02}$	$0.82^{\pm 0.02}$	$12.27^{\pm 1.02}$
CML (modulatory)	$0.42^{\pm 0.01}$	$1.83^{\pm 0.11}$	$16.46^{\pm 1.80}$

The goal of meta-learning is to discern the general structure of the low- and high-reward region across tasks whereas the goal of learning becomes to identify the task-specific radius δ of the current task. During meta-learning, we randomly sample tasks $\delta \sim \mathcal{U}(0, 1)$ and generate a dataset by choosing actions randomly. Data from each task is split into training and evaluation data, effectively creating a sparse regression problem where only the outcome of a randomly chosen action can be observed for a particular context. After meta-learning, we evaluate the cumulative regret obtained by an agent that chooses his actions greedily with respect to its predicted rewards and adapts its fast parameters on the observed context, action, reward triplets stored in a replay buffer.

We use both our synaptic consolidation and modulatory network models to meta-learn the value function using our contrastive rule. We compare our two models to MAML and the non-meta-learned baseline, NeuralLinear, from ref. [64], which performed among the best in their large-scale comparison. Tab. 5 shows the cumulative regret obtained on different task parametrizations δ in the online evaluation after meta-learning (extended table in Section S4.4). Meta-learning clearly improves upon the non-meta-learned baseline with both our models performing comparably to MAML. This improvement is more pronounced for tasks with larger δ within which it is more difficult to discover the high-reward region.

6 Discussion

We have presented a general-purpose meta-learning rule which allows estimating meta-gradients from local information only, and we have demonstrated its versatility studying two neural models on a range of meta-learning problems. The competitive performance we observed suggests that contrastive meta-

learning is a worthy contender to biologically-implausible machine learning algorithms – especially for problems involving long learning trajectories, as demonstrated by the strong results on supervised meta-optimization. At its core, our method relies on contrasting the outcome of two different learning episodes. Despite its conceptual simplicity this requires complex synaptic machinery which is able to buffer these outcomes in a way accessible to synaptic consolidation.

According to our top-down modulation model the goal of synaptic plasticity in primary brain areas is *not* to learn a specific task, in contrast to more traditional theories of learning. Instead, we postulate that the goal of synaptic plasticity is to make it possible to learn any given task by modulating the sensitivity of primary-area neurons in a task-dependent manner. This view is consistent with the experimental findings of Fritz et al. [67], who observed the rapid formation of task-dependent receptive fields in the primary auditory cortex of ferrets, as the animals learned several tasks, presumably due to changes in top-down signals originating in frontal cortex. Together with the strong results of the modulatory model in the challenging setting of visual one-shot learning and recent studies in continual learning problems [68–71] this shows the practical effectiveness of task-dependent modulation. Complementary to the interaction of the frontal cortex with primary cortical areas, the prefrontal cortex might similarly modulate the striatum during reward-based learning. Whereas classical dopamine-based learning posits that reward prediction errors are used subcortically to learn the reward structure of a task, recent work has demonstrated that reward can similarly affect prefrontal representations to quickly infer the current task identity and switch the context provided to the striatum [72]. More broadly, viewing synaptic plasticity as meta-learning is also consistent with recent modeling work casting the prefrontal cortex as a meta-reinforcement learning system [73].

Reflecting on how our meta-learning rule can be implemented in the brain, we conjecture that the hippocampal formation plays a central role in coordinating the two phases as well as creating the augmented learning problem. First, some mechanism must signal that a switch from learning problem to augmented learning problem has occurred, corresponding to the sign switch in our rule (5). We argue that the hippocampus is well positioned for signaling such a switch to cortical synapses. A recent experimental study shows that the hippocampus is at least able to control cortical synaptic consolidation [74] but further evidence would be needed to support our hypothesis.

Second, we conjecture that the creation of the augmented learning problem at the heart of our meta-gradient estimation algorithm might itself critically rely on the hippocampus. In all our experiments, this second learning problem consisted simply of new data, presented to the learning algorithm to evaluate how well learning went. Transferring additional data into cortical networks, putatively during sleep and wakeful rest, fits well with the role that is classically attributed to the hippocampus in systems consolidation and complementary learning systems theories [75, 76]. We thus speculate that the hippocampus 'prescribes' additional learning problems to the cortex, which serve the purpose of testing its generalization performance. By showing that a second 'sleep' learning phase enables meta-learning with simple plasticity rules, our results lend further credit to complementary learning systems theory, as well as to the hypothesis that dreams have evolved to assist generalization [77].

Lastly, this view of the cortex as a contrastive meta-learning system aided by the hippocampus may also help elucidate how the brain learns from an endless, non-stationary stream of data. Current artificial neural networks notoriously struggle to strike a balance between learning new knowledge and retaining old one in such continual learning problems, in particular when the data are not independent and identically distributed nor structured into clearly delineated tasks [78]. Interestingly, recent investigations have shown that meta-learning can greatly improve continual learning performance [79–83]. While details vary, the essence of these methods is to blend in past (replay) data with new data in a meta-objective function. This amounts to a different instantiation of our bilevel optimization problem (1), resulting in an augmented learning problem in which past and present data are intermixed, for which the hippocampus would again appear to be ideally positioned.

Acknowledgments and Disclosure of Funding

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 João Sacramento. Johannes von Oswald is funded by the Swiss Data Science Center (J.v.O. P18-03). We thank Angelika Steger, Benjamin Scellier, Greg Wayne, Abhishek Banerjee, Blake A. Richards, Nicol Harper, Thomas Akam, Mohamady El-Gaby, Rafal Bogacz, Giacomo Indiveri, Jean-Pascal Pfister, Mark van Rossum, Maciej Wołczyk, Seijin Kobayashi and Alexander Meulemans for discussions and feedback, and Charlotte Frenkel for assistance in our implementation of e-prop.

References

[1] Harry F. Harlow. The formation of learning sets. *Psychological Review*, 56(1):51, 1949.

[2] Johanni Brea and Wulfram Gerstner. Does computational neuroscience need new synaptic learning paradigms? *Current Opinion in Behavioral Sciences*, 11:61–66, 2016.

[3] Blake A. Richards, Timothy P. Lillicrap, Philippe Beaudoin, Yoshua Bengio, Rafal Bogacz, Amelia Christensen, Claudia Clopath, Rui Ponte Costa, Archy de Berker, Surya Ganguli, Colleen J. Gillon, Danijar Hafner, Adam Kepecs, Nikolaus Kriegeskorte, Peter Latham, Grace W. Lindsay, Kenneth D. Miller, Richard Naud, Christopher C. Pack, Panayiota Poirazi, Pieter Roelfsema, João Sacramento, Andrew Saxe, Benjamin Scellier, Anna C. Schapiro, Walter Senn, Greg Wayne, Daniel Yamins, Friedemann Zenke, Joel Zylberberg, Denis Therien, and Konrad P. Kording. A deep learning framework for neuroscience. *Nature Neuroscience*, 22(11): 1761–1770, 2019.

[4] Jürgen Schmidhuber. *Evolutionary principles in self-referential learning, or on learning how to learn: the meta-meta-... hook*. Diploma thesis, Institut für Informatik, Technische Universität München, 1987.

[5] Yoshua Bengio, Samy Bengio, and Jocelyn Cloutier. Learning a synaptic learning rule. Technical report, Université de Montréal, Département d'Informatique et de Recherche opérationnelle, 1990.

[6] David J. Chalmers. The evolution of learning: an experiment in genetic connectionism. In David S. Touretzky, Jeffrey L. Elman, Terrence J. Sejnowski, and Geoffrey E. Hinton, editors, *Connectionist Models*, pages 81–90. Morgan Kaufmann, 1991.

[7] Sebastian Thrun and Lorien Pratt. *Learning to learn*. Springer US, 1998.

[8] Sepp Hochreiter, A. Steven Younger, and Peter R. Conwell. Learning to learn using gradient descent. In *International Conference on Artificial Neural Networks*, Lecture Notes in Computer Science. Springer, 2001.

[9] Paul J. Werbos. Backpropagation through time: what it does and how to do it. *Proceedings of the IEEE*, 78(10):1550–1560, 1990.

[10] James C. R. Whittington and Rafal Bogacz. Theories of error back-propagation in the brain. *Trends in Cognitive Sciences*, 23(3):235–250, 2019.

[11] Blake A. Richards and Timothy P. Lillicrap. Dendritic solutions to the credit assignment problem. *Current Opinion in Neurobiology*, 54:28–36, 2019.

[12] Pieter R. Roelfsema and Anthony Holtmaat. Control of synaptic plasticity in deep cortical networks. *Nature Reviews Neuroscience*, 19(3):166–180, 2018.

[13] Timothy P. Lillicrap, Adam Santoro, Luke Marris, Colin J. Akerman, and Geoffrey Hinton. Backpropagation and the brain. *Nature Reviews Neuroscience*, 21(6):335–346, 2020.

[14] Paul J. Werbos. *Beyond regression: new tools for prediction and analysis in the behavioral sciences*. Ph.D. thesis, Harvard University, 1974.

[15] David E. Rumelhart, Geoffrey E. Hinton, and Ronald J. Williams. Learning representations by back-propagating errors. *Nature*, 323(6088):533–536, 1986.

[16] Richard S. Sutton. Adapting bias by gradient descent: An incremental version of delta-bar-delta. In *National Conference on Artificial Intelligence*, 1992.

[17] Marcin Andrychowicz, Misha Denil, Sergio Gomez, Matthew W. Hoffman, David Pfau, Tom Schaul, Brendan Shillingford, and Nando de Freitas. Learning to learn by gradient descent by gradient descent. In *Advances in Neural Information Processing Systems*, 2016.

[18] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *International Conference on Machine Learning*, 2017.

[19] Timothy Hospedales, Antreas Antoniou, Paul Micaelli, and Amos Storkey. Meta-learning in neural networks: a survey. *arXiv preprint arXiv:2004.05439*, 2020.

[20] Yoshua Bengio. Gradient-based optimization of hyperparameters. *Neural Computation*, 12(8): 1889–1900, 2000.

[21] Fabian Pedregosa. Hyperparameter optimization with approximate gradient. In *International Conference on Machine Learning*, 2016.

[22] Aravind Rajeswaran, Chelsea Finn, Sham Kakade, and Sergey Levine. Meta-learning with implicit gradients. In *Advances in Neural Information Processing Systems*, 2019.

[23] Jonathan Lorraine, Paul Vicol, and David Duvenaud. Optimizing millions of hyperparameters by implicit differentiation. In *International Conference on Artificial Intelligence and Statistics*, 2020.

[24] Asen L. Dontchev and R. Tyrrell Rockafellar. *Implicit Functions and Solution Mappings*. Springer, NY, 2009.

[25] Carsten Peterson and James R. Anderson. A mean field theory learning algorithm for neural networks. *Complex Systems*, 1:995–1019, 1987.

[26] Javier R. Movellan. Contrastive Hebbian learning in the continuous Hopfield model. In *Connectionist Models*, pages 10–17. Elsevier, 1991.

[27] Pierre Baldi and Fernando Pineda. Contrastive learning and neural oscillations. *Neural Computation*, 3(4):526–545, 1991.

[28] Randall C. O'Reilly. Biologically plausible error-driven learning using local activation differences: The generalized recirculation algorithm. *Neural Computation*, 8(5):895–938, 1996.

[29] Benjamin Scellier and Yoshua Bengio. Equilibrium propagation: bridging the gap between energy-based models and backpropagation. *Frontiers in Computational Neuroscience*, 11, 2017.

[30] Benjamin Scellier. *A deep learning theory for neural networks grounded in physics*. PhD Thesis, Université de Montréal, 2021.

[31] Yutian Chen, Abram L. Friesen, Feryal Behbahani, Arnaud Doucet, David Budden, Matthew W. Hoffman, and Nando de Freitas. Modular meta-learning with shrinkage. In *Advances in Neural Information Processing Systems*, 2020.

[32] Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. In *International Conference on Machine Learning*, 2017.

[33] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A. Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, Demis Hassabis, Claudia Clopath, Dharshan Kumaran, and Raia Hadsell. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences of the United States of America*, 114(13):3521–3526, 2017.

[34] Wickliffe C. Abraham. Metaplasticity: tuning synapses and networks for plasticity. *Nature Reviews Neuroscience*, 9(5):387–387, 2008.

[35] Stefano Fusi, Patrick J. Drew, and Larry F. Abbott. Cascade models of synaptically stored memories. *Neuron*, 45(4):599–611, 2005.

[36] Lorric Ziegler, Friedemann Zenke, David B. Kastner, and Wulfram Gerstner. Synaptic consolidation: from synapses to behavioral modeling. *Journal of Neuroscience*, 35(3):1319–1334, 2015.

[37] James C. R. Whittington and Rafal Bogacz. An approximation of the error backpropagation algorithm in a predictive coding network with local Hebbian synaptic plasticity. *Neural Computation*, 29(5):1229–1262, 2017.

[38] Alexandre Payeur, Jordan Guerguiev, Friedemann Zenke, Blake A. Richards, and Richard Naud. Burst-dependent synaptic plasticity can coordinate learning in hierarchical circuits. *Nature Neuroscience*, 24(7):1010–1019, 2021.

[39] Xiaohui Xie and H. Sebastian Seung. Learning in neural networks by reinforcement of irregular spiking. *Physical Review E*, 69(4), 2004.

[40] Earl K. Miller and Jonathan D. Cohen. An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience*, 24(1):167–202, 2001.

[41] Rajeev V. Rikhye, Ralf D. Wimmer, and Michael M. Halassa. Toward an integrative theory of thalamic function. *Annual Review of Neuroscience*, 41(1):163–183, 2018.

[42] Heather K. Titley, Nicolas Brunel, and Christian Hansel. Toward a neurocentric view of learning. *Neuron*, 95(1):19–32, 2017.

[43] Luisa Zintgraf, Kyriacos Shiarli, Vitaly Kurin, Katja Hofmann, and Shimon Whiteson. Fast context adaptation via meta-learning. In *International Conference on Machine Learning*, 2019.

[44] Ethan Perez, Florian Strub, Harm de Vries, Vincent Dumoulin, and Aaron C. Courville. Film: visual reasoning with a general conditioning layer. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.

[45] Katie A. Ferguson and Jessica A. Cardin. Mechanisms underlying gain modulation in the cortex. *Nature Reviews Neuroscience*, 21(2):80–92, 2020.

[46] Matthew E. Larkum, Walter Senn, and Hans-R. Lüscher. Top-down dendritic input increases the gain of layer 5 pyramidal neurons. *Cerebral Cortex*, 14(10):1059–1070, 2004.

[47] João Sacramento, Rui P. Costa, Yoshua Bengio, and Walter Senn. Dendritic cortical microcircuits approximate the backpropagation algorithm. In *Advances in Neural Information Processing Systems*, 2018.

[48] Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, 2009.

[49] Luís B. Almeida. Backpropagation in perceptrons with feedback. In Rolf Eckmiller and Christoph v.d. Malsburg, editors, *Neural Computers*, pages 199–208. Springer Berlin Heidelberg, 1989.

[50] Fernando J. Pineda. Recurrent backpropagation and the dynamical approach to adaptive neural computation. *Neural Computation*, 1(2):161–172, 1989.

[51] Renjie Liao, Yuwen Xiong, Ethan Fetaya, Lisa Zhang, KiJung Yoon, Xaq Pitkow, Raquel Urtasun, and Richard Zemel. Reviving and improving recurrent back-propagation. In *International Conference on Machine Learning*, 2018.

[52] Chuan-sheng Foo, Chuong B. Do, and Andrew Y. Ng. Efficient multiple hyperparameter learning for log-linear models. In *Advances in Neural Information Processing Systems*, 2007.

[53] Jelena Luketina, Mathias Berglund, Klaus Greff, and Tapani Raiko. Scalable gradient-based tuning of continuous regularization hyperparameters. In *International Conference on Machine Learning*, 2016.

[54] Amirreza Shaban, Ching-An Cheng, Nathan Hatch, and Byron Boots. Truncated back-propagation for bilevel optimization. In *International Conference on Artificial Intelligence and Statistics*, 2019.

[55] Brenden M. Lake, Ruslan Salakhutdinov, and Joshua B. Tenenbaum. Human-level concept learning through probabilistic program induction. *Science*, 350(6266):1332–1338, 2015.

[56] Sachin Ravi and Hugo Larochelle. Optimization as a model for few-shot learning. In *International Conference on Learning Representations*, 2016.

[57] Brenden M. Lake, Ruslan Salakhutdinov, Jason Gross, and Joshua B. Tenenbaum. One shot learning of simple visual concepts. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, 2011.

[58] Alex Nichol, Joshua Achiam, and John Schulman. On first-order meta-learning algorithms. *arXiv preprint arXiv:1803.02999*, 2018.

[59] Guillaume Bellec, Franz Scherr, Anand Subramoney, Elias Hajek, Darjan Salaj, Robert Legenstein, and Wolfgang Maass. A solution to the learning dilemma for recurrent networks of spiking neurons. *Nature Communications*, 11(1):3625, 2020.

[60] Emre O. Neftci, Hesham Mostafa, and Friedemann Zenke. Surrogate gradient learning in spiking neural networks: bringing the power of gradient-based optimization to spiking neural networks. *IEEE Signal Processing Magazine*, 36(6):51–63, 2019.

[61] Guillaume Bellec, Darjan Salaj, Anand Subramoney, Robert Legenstein, and Wolfgang Maass. Long short-term memory and learning-to-learn in networks of spiking neurons. *Advances in Neural Information Processing Systems*, 2018.

[62] Abhishek Banerjee, Giuseppe Parente, Jasper Teutsch, Christopher Lewis, Fabian F. Voigt, and Fritjof Helmchen. Value-guided remapping of sensory cortex by lateral orbitofrontal cortex. *Nature*, 585(7824):245–250, 2020.

[63] Veronika Samborska, James Butler, Mark Walton, Timothy E.J. Behrens, and Thomas Akam. Complementary task representations in hippocampus and prefrontal cortex for generalising the structure of problems. *bioRxiv*, 2021.

[64] Carlos Riquelme, George Tucker, and Jasper Snoek. Deep Bayesian bandits showdown: an empirical comparison of Bayesian deep networks for Thompson sampling. In *International Conference on Learning Representations*, 2018.

[65] Marta Garnelo, Jonathan Schwarz, Dan Rosenbaum, Fabio Viola, Danilo J. Rezende, S. M. Ali Eslami, and Yee Whye Teh. Neural processes. *arXiv preprint arXiv:1807.01622*, 2018.

[66] Sachin Ravi and Alex Beatson. Amortized Bayesian meta-learning. In *International Conference on Learning Representations*, 2019.

[67] Jonathan B. Fritz, Stephen V. David, Susanne Radtke-Schuller, Pingbo Yin, and Shihab A. Shamma. Adaptive, behaviorally gated, persistent encoding of task-relevant auditory information in ferret frontal cortex. *Nature Neuroscience*, 13(8):1011–1019, 2010.

[68] Nicolas Y. Masse, Gregory D. Grant, and David J. Freedman. Alleviating catastrophic forgetting using context-dependent gating and synaptic stabilization. *Proceedings of the National Academy of Sciences*, 115(44):E10467–E10475, October 2018.

[69] Yeming Wen, Dustin Tran, and Jimmy Ba. BatchEnsemble: an alternative approach to efficient ensemble and lifelong learning. In *International Conference on Learning Representations*, 2020.

[70] Johannes von Oswald, Christian Henning, Benjamin F. Grewe, and João Sacramento. Continual learning with hypernetworks. In *International Conference on Learning Representations*, 2020.

[71] Ben Tsuda, Kay M. Tye, Hava T. Siegelmann, and Terrence J. Sejnowski. A modeling framework for adaptive lifelong learning with transfer and savings through gating in the prefrontal cortex. *Proceedings of the National Academy of Sciences*, 117(47):29872–29882, 2020.

[72] Marta Blanco-Pozo, Thomas Akam, and Mark Walton. Dopamine reports reward prediction errors, but does not update policy, during inference-guided choice. preprint, Neuroscience, June 2021. URL `http://biorxiv.org/lookup/doi/10.1101/2021.06.25.449995`.

[73] Jane X. Wang, Zeb Kurth-Nelson, Dharshan Kumaran, Dhruva Tirumala, Hubert Soyer, Joel Z. Leibo, Demis Hassabis, and Matthew Botvinick. Prefrontal cortex as a meta-reinforcement learning system. *Nature Neuroscience*, 21(6):860–868, 2018.

[74] Guy Doron, Jiyun N. Shin, Naoya Takahashi, Moritz Drüke, Christina Bocklisch, Salina Skenderi, Lisa de Mont, Maria Toumazou, Julia Ledderose, Michael Brecht, Richard Naud, and Matthew E. Larkum. Perirhinal input to neocortical layer 1 controls learning. *Science*, 370 (6523):eaaz3136, 2020.

[75] James L. McClelland, Bruce L. McNaughton, and Randall C. O'Reilly. Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, 102(3): 419–457, 1995.

[76] Dharshan Kumaran, Demis Hassabis, and James L. McClelland. What learning systems do intelligent agents need? Complementary learning systems theory updated. *Trends in Cognitive Sciences*, 20(7):512–534, 2016.

[77] Erik Hoel. The overfitted brain: Dreams evolved to assist generalization. *Patterns*, 2(5):100244, 2021.

[78] Raia Hadsell, Dushyant Rao, Andrei A. Rusu, and Razvan Pascanu. Embracing change: continual learning in deep neural networks. *Trends in Cognitive Sciences*, 24(12):1028–1040, 2020.

[79] Khurram Javed and Martha White. Meta-learning representations for continual learning. In *Advances in Neural Information Processing Systems*, 2019.

[80] Matthew Riemer, Ignacio Cases, Robert Ajemian, Miao Liu, Irina Rish, Yuhai Tu, and Gerald Tesauro. Learning to learn without forgetting by maximizing transfer and minimizing interference. In *International Conference on Learning Representations*, 2019.

[81] Gunshi Gupta, Karmesh Yadav, and Liam Paull. Look-ahead meta learning for continual learning. In *Advances in Neural Information Processing Systems*, 2020.

[82] Shawn Beaulieu, Lapo Frati, Thomas Miconi, Joel Lehman, Kenneth O. Stanley, Jeff Clune, and Nick Cheney. Learning to continually learn. *arXiv preprint arXiv:2002.09571*, 2020.

[83] Johannes von Oswald, Dominic Zhao, Seijin Kobayashi, Simon Schug, Massimo Caccia, Nicolas Zucchet, and João Sacramento. Learning where to learn: Gradient sparsity in meta and continual learning. In *Advances in Neural Information Processing Systems*, 2021.

Supplementary Materials

Nicolas Zucchet*, Simon Schug*, Johannes von Oswald*, Dominic Zhao, João Sacramento

Table of Contents

S1 Derivation of the contrastive meta-learning rule **2**
S1.1 Equilibrium propagation theorem . . . 2
S1.2 The contrastive meta-learning rule . . . 2
S1.3 Application to the complex synapse model . . . 3
S1.4 Application to the top-down modulation model . . . 3

S2 Review of implicit gradient methods for meta-learning **5**

S3 Theoretical results **7**
S3.1 Meta-gradient estimation error bound . . . 7
S3.2 Proof of Theorem 1 . . . 7
S3.3 Proof of technical lemmas . . . 10
S3.4 A corollary of Theorem 1 . . . 13
S3.5 Verification of the theoretical results on an analytical problem . . . 14
S3.6 Verification the theoretical results on a simple hyperparameter optimization task . 16

S4 Experimental details **19**
S4.1 Supervised meta-optimization . . . 19
S4.2 Few-shot image classification . . . 20
S4.3 Few-shot regression in recurrent spiking network . . . 21
S4.4 Meta-reinforcement learning . . . 22

S5 Additional details **29**
S5.1 Compute resources . . . 29
S5.2 Software and libraries . . . 29
S5.3 Datasets . . . 29

S1 Derivation of the contrastive meta-learning rule

Our contrastive meta-learning rule relies on the equilibrium propagation theorem [29, 30]. We review this result and how we use it to derive the different instances of our rule.

S1.1 Equilibrium propagation theorem

First, we restate the equilibrium propagation theorem as presented in Scellier [30]. Recall the definition of the augmented loss

$$\mathcal{L}(\phi, \theta, \beta) = L^{\text{learn}}(\phi, \theta) + \beta L^{\text{eval}}(\phi, \theta). \tag{8}$$

Note that compared to the main text, we omit the subscript τ for conciseness. Given the augmented loss, the equilibrium propagation theorem states the following:

Theorem S1 (Equilibrium propagation)**.** *Let L^{learn} and L^{eval} be two twice continuously differentiable functions. Let ϕ^* be a fixed point of $\mathcal{L}(\,\cdot\,, \bar{\theta}, \bar{\beta})$, i.e.*

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

such that $\partial^2_\phi \mathcal{L}(\phi^*, \bar{\theta}, \bar{\beta})$ 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 $\phi^*_{\bar{\theta},\bar{\beta}} = \phi^*$ and for every (θ, β) in this neighborhood

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

Furthermore,

$$\frac{\mathrm{d}}{\mathrm{d}\theta}\frac{\partial \mathcal{L}}{\partial \beta}\left(\phi^*_{\theta,\beta}, \theta, \beta\right) = \frac{\mathrm{d}}{\mathrm{d}\beta}\frac{\partial \mathcal{L}}{\partial \theta}\left(\phi^*_{\theta,\beta}, \theta, \beta\right)^\top.$$

Proof. The first point follows from the implicit function theorem [24]. Let (θ, β) be in a neighborhood of $(\bar{\theta}, \bar{\beta})$ in which $\phi^*_{\theta,\beta}$ is differentiable.

The symmetry of second order derivatives of a scalar function implies that

$$\frac{\mathrm{d}}{\mathrm{d}\theta}\frac{\mathrm{d}}{\mathrm{d}\beta}\mathcal{L}\left(\phi^*_{\theta,\beta}, \theta, \beta\right) = \frac{\mathrm{d}}{\mathrm{d}\beta}\frac{\mathrm{d}}{\mathrm{d}\theta}\mathcal{L}\left(\phi^*_{\theta,\beta}, \theta, \beta\right)^\top. \tag{9}$$

We then simplify the two sides of the equation. First, we look at the left-hand side and simplify $\mathrm{d}_\beta \mathcal{L}(\phi^*_{\theta,\beta}, \theta, \beta)$ using the chain rule and the fixed point condition

$$\begin{aligned}\frac{\mathrm{d}}{\mathrm{d}\beta}\mathcal{L}(\phi^*_{\theta,\beta}, \theta, \beta) &= \frac{\partial \mathcal{L}}{\partial \beta}(\phi^*_{\theta,\beta}, \theta, \beta) + \frac{\partial \mathcal{L}}{\partial \phi}(\phi^*_{\theta,\beta}, \theta, \beta)\frac{\mathrm{d}\phi^*_{\theta,\beta}}{\mathrm{d}\beta} \\ &= \frac{\partial \mathcal{L}}{\partial \beta}(\phi^*_{\theta,\beta}, \theta, \beta).\end{aligned} \tag{10}$$

Similarly, the $\mathrm{d}_\theta \mathcal{L}(\phi^*_{\theta,\beta}, \theta, \beta)$ term on the right-hand side is equal to $\partial_\theta \mathcal{L}(\phi^*_{\theta,\beta}, \theta, \beta)$ and we obtain the required result[2]. □

S1.2 The contrastive meta-learning rule

Equilibrium propagation can be used to compute the gradient associated with the bilevel optimization problem studied in this paper

$$\min_\theta \; L^{\text{eval}}(\phi^*_\theta) \quad \text{s.t.} \;\; \phi^*_\theta \in \arg\min_\phi \; L^{\text{learn}}(\phi, \theta). \tag{11}$$

To do so, we first characterize ϕ^*_θ through the stationarity condition

$$\frac{\partial L^{\text{learn}}}{\partial \phi}(\phi^*_\theta, \theta) = 0. \tag{12}$$

[2]Note that we use ∂ to denote partial derivatives and d to denote total derivatives.

As $L^{\text{learn}}(\phi, \theta) = \mathcal{L}(\phi, \theta, 0)$ we can define an implicit function $\phi^*_{\theta,\beta}$ if $\partial^2_\phi L^{\text{learn}}(\phi^*_\theta, \theta)$ is invertible for which $\phi^*_{\theta,0} = \phi^*_\theta$, and that satisfies, for β close to 0,

$$\frac{\partial \mathcal{L}}{\partial \phi}(\phi^*_{\theta,\beta}, \theta, \beta) = 0. \tag{13}$$

The gradient associated with the bilevel optimization problem (11) is then equal to

$$\nabla_\theta := \left(\frac{\mathrm{d}}{\mathrm{d}\theta} L^{\text{eval}}(\phi^*_\theta, \theta)\right)^\top = \frac{\mathrm{d}}{\mathrm{d}\theta}\frac{\partial \mathcal{L}}{\partial \beta}(\phi^*_{\theta,\beta}, \theta, \beta)\Big|^\top_{\beta=0} = \frac{\mathrm{d}}{\mathrm{d}\beta}\frac{\partial \mathcal{L}}{\partial \theta}(\phi^*_{\theta,\beta}, \theta, \beta)\Big|_{\beta=0}. \tag{14}$$

Since β is a scalar, we can use finite difference methods to efficiently estimate

$$\Delta\theta = -\widehat{\nabla}_\theta = -\frac{1}{\beta}\left(\frac{\partial \mathcal{L}}{\partial \theta}(\hat{\phi}_\beta, \theta, \beta) - \frac{\partial \mathcal{L}}{\partial \theta}(\hat{\phi}_0, \theta, 0)\right)^\top, \tag{15}$$

in which $\hat{\phi}_0$ and $\hat{\phi}_\beta$ denote the estimates of $\phi^*_{\theta,0}$ and $\phi^*_{\theta,\beta}$. If those estimates are exact, we are guaranteed that the update converges to the true gradient. In some of our experiments, we use a more sophisticated center difference approximation similar to [87], that is

$$\Delta\theta^{\text{sym}} = -\widehat{\nabla}^{\text{sym}}_\theta = -\frac{1}{2\beta}\left(\frac{\partial \mathcal{L}}{\partial \theta}(\hat{\phi}_\beta, \theta, \beta) - \frac{\partial \mathcal{L}}{\partial \theta}(\hat{\phi}_{-\beta}, \theta, -\beta)\right)^\top. \tag{16}$$

We refer to (16) as the symmetric variant of our contrastive rule. When the estimates for the fixed points are exact, it reduces the meta-gradient estimation bias from $O(\beta)$ for the forward difference above to $O(\beta^2)$ at the expense of having to run a third phase.

S1.3 Application to the complex synapse model

We can now derive the meta-learning rules for the complex synapse model of Section 4.1. Recall that

$$L^{\text{learn}}(\phi, \theta) = l^{\text{learn}}(\phi) + \frac{1}{2}\sum_{i=1}^{|\phi|} \lambda_i(\omega_i - \phi_i)^2 \tag{17}$$

and

$$L^{\text{eval}}(\phi) = l^{\text{eval}}(\phi), \tag{18}$$

where $l^{\text{eval}}(\phi)$ and $l^{\text{learn}}(\phi)$ are two data-dependent loss functions.

For the complex synapse model, only the learning loss depends on the meta-parameters, hence $\partial_\theta \mathcal{L} = \partial_\theta L^{\text{learn}}$ and

$$\begin{aligned}\frac{\partial \mathcal{L}}{\partial \omega}(\phi, \theta, \beta) &= \lambda(\omega - \phi)\\ \frac{\partial \mathcal{L}}{\partial \lambda}(\phi, \theta, \beta) &= \frac{1}{2}(\omega - \phi)^2,\end{aligned} \tag{19}$$

where all the operations are carried out elementwise. Plugging the last equation in the contrastive update (15) yields

$$\begin{aligned}\Delta\omega &= -\frac{\lambda}{\beta}\left((\omega - \hat{\phi}_\beta) - (\omega - \hat{\phi}_0)\right) = \frac{\lambda}{\beta}\left(\hat{\phi}_\beta - \hat{\phi}_0\right)\\ \Delta\lambda &= -\frac{1}{2\beta}\left((\omega - \hat{\phi}_\beta)^2 - (\omega - \hat{\phi}_0)^2\right) = \frac{1}{2\beta}\left((\omega - \hat{\phi}_0)^2 - (\omega - \hat{\phi}_\beta)^2\right).\end{aligned} \tag{20}$$

S1.4 Application to the top-down modulation model

The structure of the learning and evaluation losses for the top-down modulation model is the following:

$$\begin{aligned}L^{\text{learn}}(\phi, \theta) &= l^{\text{learn}}(h(\phi, \theta))\\ L^{\text{eval}}(\phi, \theta) &= l^{\text{eval}}(h(\phi, \theta)),\end{aligned} \tag{21}$$

where $l^{\text{learn}}(\psi)$ and $l^{\text{eval}}(\psi)$ are two data-driven losses that use learning and evaluation datasets to evaluate the performance of a network parametrized by ψ, and $h(\phi, \theta)$ produces the parameters ψ by modulating a base network parametrized by θ. Specifically, we modulate the rectified linear unit (ReLU) activation function for each neuron i with a gain g_i and shift b_i, $\sigma_\phi(x_i) = g_i((\theta \cdot x)_i - b_i)_+$, with the gain and shift parameters of all neurons defining the fast parameters $\phi = \{g, b\}$.

Applying our contrastive update (15) to this model we obtain the following equations:

$$\begin{aligned}\Delta\theta &= -\frac{1}{\beta}\left(\frac{\partial \mathcal{L}}{\partial \theta}(\hat{\phi}_\beta, \theta, \beta) - \frac{\partial \mathcal{L}}{\partial \theta}(\hat{\phi}_0, \theta, 0)\right)^\top \\ &= -\frac{1}{\beta}\left(\frac{\partial [l^{\text{learn}} + \beta l^{\text{eval}}]}{\partial \psi}(h(\hat{\phi}_\beta, \theta))\frac{\partial h}{\partial \theta}(\hat{\phi}_\beta, \theta) - \frac{\partial l^{\text{learn}}}{\partial \psi}(h(\hat{\phi}_0, \theta))\frac{\partial h}{\partial \theta}(\hat{\phi}_0, \theta)\right)^\top .\end{aligned} \tag{22}$$

Let us now decompose what this update means. The losses $l^{\text{learn}}(\psi)$ and $[l^{\text{learn}} + \beta l^{\text{eval}}](\psi)$ measure the performance of a network parametrized by ψ on the learning data, and on a weighted mix of learning and evaluation data. The derivatives $\partial_\psi l^{\text{learn}}$ and $\partial_\psi [l^{\text{learn}} + \beta l^{\text{eval}}]$ can therefore be computed using the backpropagation-of-error algorithm, or any biologically plausible alternative to it. Those derivatives are then multiplied by $\partial_\theta h$, which is a diagonal matrix as the modulation does not combine weights together, but only individually changes them. As a result, the update (22) contrasts two elementwise modulated gradients with respect to the weights.

S2 Review of implicit gradient methods for meta-learning

The gradient associated with the bilevel optimization problem of Eq. 11 can be calculated analytically using the implicit function theorem [24]. This insight forms the basis for implicit gradient methods for meta-learning which we shortly review in the following. We additionally provide a comparison of the computational and memory complexity of different meta-learning methods in Table S1.

As for the derivation of the contrastive meta-learning rule, we start by characterizing the implicit function ϕ_θ^* of θ through its corresponding first-order stationarity condition

$$\frac{\partial L^{\text{learn}}}{\partial \phi}(\phi_\theta^*, \theta) = 0. \tag{23}$$

Then, when the Hessian $\partial_\phi^2 L^{\text{learn}}(\phi_\theta^*, \theta)$ is invertible, we have

$$\begin{aligned} \frac{\mathrm{d}}{\mathrm{d}\theta} L^{\text{eval}}(\phi_\theta^*, \theta) &= \frac{\partial L^{\text{eval}}}{\partial \theta}(\phi_\theta^*, \theta) + \frac{\partial L^{\text{eval}}}{\partial \phi} \frac{\mathrm{d}\phi_\theta^*}{\mathrm{d}\theta} \\ &= \frac{\partial L^{\text{eval}}}{\partial \theta}(\phi_\theta^*, \theta) - \frac{\partial L^{\text{eval}}}{\partial \phi} \left(\frac{\partial^2 L^{\text{learn}}}{\partial \phi^2}(\phi_\theta^*, \theta) \right)^{-1} \frac{\partial^2 L^{\text{learn}}}{\partial \phi \partial \theta}(\phi_\theta^*, \theta), \end{aligned} \tag{24}$$

where in the first line we used the chain rule and in the second line the differentiation formula provided by the implicit function theorem [24].

In most practical applications, ϕ is high dimensional rendering the computation and inversion of the Hessian $\partial_\phi^2 L^{\text{learn}}(\phi_\theta^*, \theta)$ intractable. In order to obtain a practical algorithm, implicit gradient methods numerically approximate the row vector

$$\mu := -\frac{\partial L^{\text{eval}}}{\partial \phi}(\phi_\theta^*, \theta) \left(\frac{\partial^2 L^{\text{learn}}}{\partial \phi^2}(\phi_\theta^*, \theta) \right)^{-1}. \tag{25}$$

The simplest algorithm, T1-T2 [53], replaces the inverse Hessian by the identity, i.e. $\mu \approx \partial_\phi L^{\text{eval}}(\phi_\theta^*, \theta)$, which yields an estimate relying only on first derivatives.

The recurrent backpropagation algorithm [RBP, 49–51], also known as Neumann series approximation [23, 51], builds on the insight that μ is the solution of the linear system

$$x \frac{\partial^2 L^{\text{learn}}}{\partial \phi^2}(\phi_\theta^*, \theta) = -\frac{\partial L^{\text{eval}}}{\partial \phi}(\phi_\theta^*, \theta). \tag{26}$$

which can be solved via fixed point iteration.

Finally, μ can be seen as the solution of the optimization problem

$$\min_x \; x \frac{\partial^2 L^{\text{learn}}}{\partial \phi^2}(\phi_\theta^*, \theta) x^\top + x \frac{\partial L^{\text{eval}}}{\partial \phi}(\phi_\theta^*, \theta)^\top \tag{27}$$

when the Hessian of L^{learn} is positive definite. This optimization problem can be efficiently solved via the conjugate gradient method [22, 91].

The three algorithms described above provide different estimates for μ but all follow the same basic procedure: (1) minimize the learning loss to approximate ϕ_θ^*; (2) estimate μ; and (3) update the meta-parameters using (24) with the estimated μ.

Compared to our contrastive meta-learning rule, these algorithms require a second phase that is completely different from the first one and which involves second derivatives (apart from the biased T1-T2). Additionally, as mentioned in Section 5.2, the conjugate gradient method is faster in theory, but was reported to be unstable by several studies [51, 54]. We confirm those findings in our experiments (cf. Section S3.6 and S4.1).

Table S1: Comparison of computational and memory complexity of meta-learning methods. T denotes the number of steps in the base learning process and K refers to steps taken in an algorithm-specific second phase. "HVP" abbreviates "Hessian-vector product" and "cross der. VP" denotes "cross derivative vector product". An algorithm is "exact in the limit" if it computes the meta-gradient or can approximate it with arbitrary precision given enough compute. The algorithms compared in this table are contrastive meta-learning (CML), conjugate gradients (CG; used in iMAML), recurrent backpropagation (RBP), T1-T2, backpropagation-through-learning (BPTL; used in MAML), its truncated version (TBPTL) and its first-order version where all Hessians are replaced by the identity (FOBPTL; also known as FOMAML) and Reptile. The first four algorithms assume that the base learning process reaches an equilibrium, whereas the last four require no such assumption. * Reptile is not a general-purpose meta-learning method as it is restricted to meta-learn the initialization of the learning process.

Method	# gradients w.r.t.		# 2nd-order terms		Memory	Exact in the limit
	ϕ	θ	HVP	cross der. VP		
CML (ours)	$T+K$	2	0	0	$\mathcal{O}(\lvert\phi\rvert+\lvert\theta\rvert)$	✓
CG [21, 22, 52]	$T+1$	1	K	1	$\mathcal{O}(\lvert\phi\rvert+\lvert\theta\rvert)$	✓
RBP [23, 49–51]	$T+1$	1	K	1	$\mathcal{O}(\lvert\phi\rvert+\lvert\theta\rvert)$	✓
T1-T2 [53]	$T+1$	1	0	1	$\mathcal{O}(\lvert\phi\rvert+\lvert\theta\rvert)$	✗
BPTL [18]	$T+1$	$T+1$	T	0	$\mathcal{O}(T\lvert\phi\rvert+\lvert\theta\rvert)$	✓
TBPTL [54]	$T+1$	$K+1$	K	0	$\mathcal{O}(K\lvert\phi\rvert+\lvert\theta\rvert)$	✗
FOBPTL [18]	$T+1$	$T+1$	0	0	$\mathcal{O}(T\lvert\phi\rvert+\lvert\theta\rvert)$	✗
Reptile* [58]	T	0	0	0	$\mathcal{O}(\lvert\phi\rvert)$	✗

S3 Theoretical results

The contrastive meta-learning rule (15) only provides an approximation $\widehat{\nabla}_\theta$ to the meta-gradient ∇_θ due to the limited precision of the fixed points and the finite difference estimator. We can in principle arbitrarily improve the approximation by spending more compute to refine the quality of the solutions $\hat{\phi}_0$ and $\hat{\phi}_\beta$ and decreasing the nudging strength β. The purpose of this section is to theoretically analyze the impact of such a refinement on the quality of the meta-gradient estimate. We state Theorem 1 formally, present a corollary of this result, and verify that it holds experimentally.

S3.1 Meta-gradient estimation error bound

We start by upper bounding the meta-gradient estimation error $\|\widehat{\nabla}_\theta - \nabla_\theta\|$, given the value of β and the error made in the approximation of the solutions of the lower-level learning process. Two conflicting phenomena impact the estimation error. First, our meta-learning rule uses potentially inexact solutions. Second, the finite difference approximation of the β-derivative yields the so-called finite difference error. To study those two errors in more detail, we introduce

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

the contrastive estimate of the meta-gradient ∇_θ, but evaluated at the exact solutions $\phi^*_{\theta,0}$ and $\phi^*_{\theta,\beta}$ (recall that $\widehat{\nabla}_\theta$ has the same structure, but it is evaluated on the approximate solutions $\hat{\phi}_0$ and $\hat{\phi}_\beta$). Equipped with $\widehat{\nabla}^*_\theta$, we now have a way to quantify the two errors described above: $\|\nabla_\theta - \widehat{\nabla}^*_\theta\|$ measures the finite difference error and $\|\widehat{\nabla}^*_\theta - \widehat{\nabla}_\theta\|$ measures the solution approximation induced error, that is the consequence of the imperfect solutions.

Informally, higher β values will reduce the sensitivity to crude approximations to the lower-level solutions while increasing the finite difference error. Theorem 1 theoretically justifies this intuition under the idealized regime of strong convexity and smoothness defined in Assumption 1. This result holds for every rule induced by equilibrium propagation.

Assumption 1. *Assume that L^{learn} and L^{eval} are three-times continuously differentiable and that they, as functions of ϕ, verify the following properties.*

i. $\partial_\theta L^{\text{learn}}$ is B^{learn}-Lipschitz and $\partial_\theta L^{\text{eval}}$ is B^{eval}-Lipschitz.

ii. L^{learn} and L^{eval} are L-smooth and μ-strongly convex.

iii. their Hessians are ρ-Lipschitz.

iv. $\partial_\phi \partial_\theta L^{\text{learn}}$ and $\partial_\phi \partial_\theta L^{\text{eval}}$ are σ-Lipschitz.

Theorem 1 (Formal)**.** *Let $\beta > 0$ and (δ, δ') be such that*

$$\|\phi^*_{\theta,0} - \hat{\phi}_0\| \leq \delta, \quad \text{and} \quad \|\phi^*_{\theta,\beta} - \hat{\phi}_\beta\| \leq \delta'.$$

Under Assumption 1, there exists a θ-dependent constant C such that

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

If we additionally assume that θ lies in a compact set, we can choose C to be independent of θ.

We visualize our bound in Fig. S1, as a function of β and of the solution approximation errors δ and δ'. When δ and δ' are fixed, the estimation error quickly increases when β deviates from its optimal value and it saturates for large β values (cf. Fig. S1A and B). A better solution approximation naturally improves the quality of the meta-gradient estimate for β held constant (cf. Fig. S1C). However, the benefits saturate above some β-dependent value: investing extra compute in the approximation of the fixed point does not pay off if β is not decreased accordingly.

S3.2 Proof of Theorem 1

As mentioned above, Theorem 1 can be proved by individually bounding the two kind of errors that compose the meta-gradient estimation error, that are the finite difference error and the solution approximation induced error.

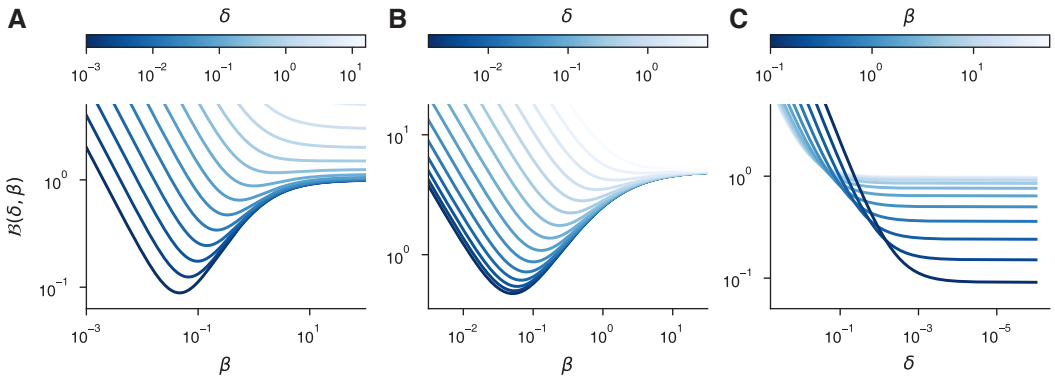

Figure S1: Theorem 1 ($C = 1$), as a function of β (A, B) and as a function of $\delta = \delta'$ (C). (A) We take $B^{\text{learn}} = B^{\text{eval}} = 1$ and $\delta = \delta'$. (B) Bound for the setting in which δ' is fixed to 0.01 and L^{eval} is independent of θ (as for the complex synapse model). (C) We use the same setting as for (A).

The $B^{\text{learn}}(\delta + \delta')/\beta + B^{\text{eval}}\delta'$ part of the bound stems from the solution approximation error, and can be obtained by using the assumption that the partial derivatives of L^{learn} and L^{eval} are Lipschitz continuous.

Bounding the finite difference error requires more work. We use Taylor's theorem to show that $\widehat{\nabla}^*_\theta - \nabla_\theta$ is equal to some integral remainder. It then remains to bound what is inside the integral remainder, which is the second order derivative $\mathrm{d}^2_\beta \partial_\theta \mathcal{L}(\phi^*_\beta, \beta)$. This is done in the Lemmas presented in this section: Lemma 1 allows us to get uniform bounds, Lemmas 2 and 3 bound the first and second order derivatives of $\beta \mapsto \phi^*_\beta$ and Lemma 4 bounds $\mathrm{d}^2_\beta \partial_\theta \mathcal{L}(\phi^*_\beta, \beta)$ with the norm of the two derivatives we have just bounded. We present the proofs for those four lemmas in Section S3.3.

Lemma 1. *Under Assumption 1.ii, if θ lies in a compact set $\mathcal{D}$ the function $(\theta, \beta) \mapsto \phi^*_{\theta,\beta}$ is uniformly bounded.*

Lemma 2. *Under Assumption 1.ii, there exists a θ-dependent constant R s.t., for every positive β,*

$$\left\| \frac{\mathrm{d}\phi^*_\beta}{\mathrm{d}\beta} \right\| \leq \frac{LR}{(1+\beta)^2 \mu}.$$

If we additionally assume that θ lies in a compact set, we can choose R to be independent of θ.

Remark 1. *A side product of the proof of Lemma 2 is a bound on the distance between the minimizer of $\mathcal{L}$ and the minimizers of L^{learn} and L^{eval}. We have*

$$\|\phi^*_\beta - \phi^*_\infty\| \leq \frac{1}{1+\beta}$$

and

$$\|\phi^*_\beta - \phi^*_0\| \leq \frac{\beta}{1+\beta}$$

up to some constant factors.

Lemma 3. *Under Assumptions 1.ii and 1.iii,*

$$\left\| \frac{\mathrm{d}^2\phi^*_\beta}{\mathrm{d}\beta^2} \right\| \leq \frac{\rho}{\mu} \left\| \frac{\mathrm{d}\phi^*_\beta}{\mathrm{d}\beta} \right\|^2 + \frac{2L}{(1+\beta)\mu} \left\| \frac{\mathrm{d}\phi^*_\beta}{\mathrm{d}\beta} \right\|.$$

When Lemma 3 is combined with Lemma 2,

$$\left\| \frac{\mathrm{d}^2\phi^*_\beta}{\mathrm{d}\beta^2} \right\| \leq \frac{1}{(1+\beta)^3}. \tag{28}$$

up to some constant factor.

Lemma 4. *Under Assumptions 1.ii, 1.iii and 1.iv, there exists a constant M such that*

$$\left\|\frac{\mathrm{d}^2}{\mathrm{d}\beta^2}\frac{\partial\mathcal{L}}{\partial\theta}(\phi^*_\beta,\beta)\right\| \leq M\left(\left\|\frac{\mathrm{d}\phi^*_\beta}{\mathrm{d}\beta}\right\| + (1+\beta)\left(\left\|\frac{\mathrm{d}\phi^*_\beta}{\mathrm{d}\beta}\right\|^2 + \left\|\frac{\mathrm{d}^2\phi^*_\beta}{\mathrm{d}\beta^2}\right\|\right)\right).$$

We can now prove Theorem 1 using the four lemmas that we have just presented. Note that we omit the θ-dependency whenever θ is fixed, for the sake of conciseness.

Proof of Theorem 1. We separate the sources of error within the meta-gradient estimation error using the triangle inequality:

$$\|\widehat{\nabla}_\theta - \nabla_\theta\| \leq \underbrace{\|\widehat{\nabla}_\theta - \widehat{\nabla}^*_\theta\|}_{a)} + \underbrace{\|\widehat{\nabla}^*_\theta - \nabla_\theta\|}_{b)}, \tag{29}$$

and bound the two terms separately:

a) Recall that

$$\widehat{\nabla}_\theta = \frac{1}{\beta}\left(\frac{\partial\mathcal{L}}{\partial\theta}(\hat{\phi}_\beta,\beta) - \frac{\partial\mathcal{L}}{\partial\theta}(\hat{\phi}_0,0)\right) \tag{30}$$

and that a similar formula holds for $\widehat{\nabla}^*_\theta$ (evaluated at the true solutions instead of the approximations). It follows

$$\|\widehat{\nabla}_\theta - \widehat{\nabla}^*_\theta\| \leq \frac{1}{\beta}\left(\left\|\frac{\partial\mathcal{L}}{\partial\theta}(\hat{\phi}_\beta,\beta) - \frac{\partial\mathcal{L}}{\partial\theta}\left(\phi^*_\beta,\beta\right)\right\| + \left\|\frac{\partial\mathcal{L}}{\partial\theta}(\hat{\phi}_0,0) - \frac{\partial\mathcal{L}}{\partial\theta}(\phi^*_0,0)\right\|\right). \tag{31}$$

Since $\phi \mapsto \partial_\theta\mathcal{L}(\phi,\beta)$ is a $(B^{\mathrm{learn}} + \beta B^{\mathrm{eval}})$-Lipschitz function as a sum of $\partial_\theta L^{\mathrm{learn}}$ and $\partial_\theta L^{\mathrm{eval}}$, two Lipschitz continuous functions with constants B^{learn} and B^{eval},

$$\|\widehat{\nabla}_\theta - \widehat{\nabla}^*_\theta\| \leq \frac{B^{\mathrm{learn}} + \beta B^{\mathrm{eval}}}{\beta}\|\hat{\phi}_\beta - \phi^*_\beta\| + \frac{B^{\mathrm{learn}}}{\beta}\|\hat{\phi}_0 - \phi^*_0\| \tag{32}$$

$$\leq \frac{B^{\mathrm{learn}} + \beta B^{\mathrm{eval}}}{\beta}\delta' + \frac{B^{\mathrm{learn}}}{\beta}\delta. \tag{33}$$

b) Taylor's theorem applied to $\beta \mapsto \partial_\theta\mathcal{L}(\phi^*_\beta,\beta)$ up to the first order of differentiation yields

$$\frac{\partial\mathcal{L}}{\partial\theta}\left(\phi^*_\beta,\beta\right) = \frac{\partial\mathcal{L}}{\partial\theta}\left(\phi^*_0,0\right) + \beta\frac{\mathrm{d}}{\mathrm{d}\beta}\frac{\partial\mathcal{L}}{\partial\theta}\left(\phi^*_0,0\right) + \int_0^\beta(\beta-t)\frac{\mathrm{d}^2}{\mathrm{d}\beta^2}\frac{\partial\mathcal{L}}{\partial\theta}\left(\phi^*_t,t\right)\mathrm{d}t. \tag{34}$$

The equilibrium propagation theorem (Theorem S1), which is applicable thanks to Assumption 1.ii, gives

$$\nabla_\theta = \frac{\mathrm{d}}{\mathrm{d}\beta}\frac{\partial\mathcal{L}}{\partial\theta}(\phi^*_0,0), \tag{35}$$

hence

$$\|\widehat{\nabla}^*_\theta - \nabla_\theta\| = \left\|\int_0^\beta(\beta-t)\frac{\mathrm{d}^2}{\mathrm{d}\beta^2}\frac{\partial\mathcal{L}}{\partial\theta}\left(\phi^*_t,t\right)\mathrm{d}t\right\|. \tag{36}$$

Using the integral version of the Cauchy-Schwartz inequality, we have

$$\|\widehat{\nabla}^*_\theta - \nabla_\theta\| \leq \int_0^\beta(\beta-t)\left\|\frac{\mathrm{d}^2}{\mathrm{d}\beta^2}\frac{\partial\mathcal{L}}{\partial\theta}(\phi^*_t,t)\right\|\mathrm{d}t. \tag{37}$$

We now use Lemma 4 combined with Lemmas 2 and 3 to bound $\mathrm{d}^2_\beta\partial_\theta\mathcal{L}(\phi^*_t,t)$. We focus on the β dependencies and omitting constant factors:

$$\begin{aligned}\left\|\frac{\mathrm{d}^2}{\mathrm{d}\beta^2}\frac{\partial\mathcal{L}}{\partial\theta}(\phi^*_t,t)\right\| &\leq \left\|\frac{\mathrm{d}\phi^*_t}{\mathrm{d}\beta}\right\| + (1+t)\left(\left\|\frac{\mathrm{d}\phi^*_t}{\mathrm{d}\beta}\right\|^2 + \left\|\frac{\mathrm{d}^2\phi^*_t}{\mathrm{d}\beta^2}\right\|\right)\\ &\leq \frac{1}{(1+t)^2} + (1+t)\left(\frac{1}{(1+t)^3} + \frac{1}{(1+t)^4}\right)\\ &\leq (1+t)^{-2}.\end{aligned}$$

It follows that

$$\begin{aligned}\|\widehat{\nabla}^*_\theta - \nabla_\theta\| &\leq \int_0^\beta \frac{(\beta - t)}{(1+t)^2}\mathrm{d}t \\ &= (1+\beta)\int_0^\beta \frac{1}{(1+t)^2}\mathrm{d}t - \int_0^\beta \frac{1}{(1+t)}\mathrm{d}t \\ &= (1+\beta)\frac{\beta}{1+\beta} - \ln(1+\beta) \\ &\leq \beta - \frac{\beta}{1+\beta} \\ &= \frac{\beta^2}{1+\beta}.\end{aligned} \tag{38}$$

where the inequality comes from the well-known $\ln(x) \geq 1 - \frac{1}{x}$ inequality for positive x (applied to $x = 1 + \beta$). There hence exists a constant C such that

$$\|\widehat{\nabla}^*_\theta - \nabla_\theta\| \leq C\frac{\beta}{1+\beta}. \tag{39}$$

If θ lies in a compact set, the bound in Lemma 2 is uniform over θ. This is the only constant factor that depends on θ, so the bound is uniform. □

S3.3 Proof of technical lemmas

In this section, we prove the four technical lemmas that we need for Theorem 1.

Proof of Lemma 1

Lemma 1. *Under Assumption 1.ii, if θ lies in a compact set $\mathcal{D}$ the function $(\theta, \beta) \mapsto \phi^*_{\theta,\beta}$ is uniformly bounded.*

Proof. Let $\alpha \in [0, 1]$. Define

$$\mathcal{L}'(\phi, \theta, \alpha) := (1-\alpha)L^{\mathrm{learn}}(\phi, \theta) + \alpha L^{\mathrm{eval}}(\phi, \theta). \tag{40}$$

As L^{learn} and L^{eval} are strongly-convex, there exists a unique minimizer $\phi^{*\prime}_{\theta,\alpha}$ of $\phi \mapsto \mathcal{L}'(\phi, \theta, \alpha)$. The implicit function theorem ensures that the function $(\theta, \alpha) \mapsto \phi^{*\prime}_{\theta,\alpha}$, defined on $\mathcal{D} \times [0, 1]$, is continuous. As $\mathcal{D} \times [0, 1]$ is a compact set, $\phi^{*\prime}_{\theta,\alpha}$ is then uniformly bounded. Now, remark that

$$\mathcal{L}(\phi, \theta, \beta) = (1+\beta)\mathcal{L}'\left(\phi, \theta, \frac{\beta}{1+\beta}\right) \tag{41}$$

and thus $\phi^*_{\theta,\beta} = \phi^{*\prime}_{\theta,\beta/(1+\beta)}$. It follows that $\phi^*_{\theta,\beta}$ is uniformly bounded. □

Proof of Lemma 2

Lemma 2. *Under Assumption 1.ii, there exists a θ-dependent constant R s.t., for every positive β,*

$$\left\|\frac{\mathrm{d}\phi^*_\beta}{\mathrm{d}\beta}\right\| \leq \frac{LR}{(1+\beta)^2\mu}.$$

If we additionally assume that θ lies in a compact set, we can choose R to be independent of θ.

Proof. The function $\phi \mapsto \mathcal{L}(\phi, \beta)$ is $(1+\beta)\mu$-strongly convex so its Hessian $\partial^2_\phi \mathcal{L}$ is invertible and its inverse has a spectral norm upper bounded by $1/((1+\beta)\mu)$. The use of the implicit function theorem follows and gives

$$\begin{aligned}\left\|\mathrm{d}_\beta \phi^*_\beta\right\| &= \left\| - \left(\partial^2_\phi \mathcal{L}(\phi^*_\beta, \beta)\right)^{-1} \partial_\beta \partial_\phi \mathcal{L}(\phi^*_\beta)\right\| \\ &= \left\| - \left(\partial^2_\phi \mathcal{L}(\phi^*_\beta, \beta)\right)^{-1} \partial_\phi L^{\mathrm{eval}}(\phi^*_\beta)\right\| \\ &\leq \frac{1}{(1+\beta)\mu}\|\partial_\phi L^{\mathrm{eval}}(\phi^*_\beta)\|.\end{aligned} \tag{42}$$

It remains to bound the gradient of L^{eval}. Since $\beta \mapsto \phi^*_\beta$ is continuous and has finite limits in 0 and ∞ (namely the minimizers of L^{learn} and L^{eval}), it evolves in a bounded set. There hence exists a positive constant R such that, for all positive β,

$$\max\left(\left\|\phi^*_\beta - \phi^*_0\right\|, \left\|\phi^*_\beta - \phi^*_\infty\right\|\right) \leq \frac{R}{2}. \tag{43}$$

If θ lies in a compact set, Lemma 1 guarantees that there exists such a constant that doesn't depend on the choice of θ. We then bound the gradient of L^{eval} using the smoothness properties of L^{learn} and L^{eval}, either directly

$$\|\partial_\phi L^{\text{eval}}(\phi^*_\beta)\| \leq L\|\phi^*_\beta - \phi^*_\infty\| \leq \frac{LR}{2} \tag{44}$$

or indirectly, using the fixed point condition $\partial_\phi \mathcal{L}(\phi^*_\beta, \beta) = 0$,

$$\|\partial_\phi L^{\text{eval}}(\phi^*_\beta)\| = \frac{1}{\beta}\|-\partial_\phi L^{\text{learn}}(\phi^*_\beta)\| \leq \frac{L\|\phi^*_\beta - \phi^*_0\|}{\beta} \leq \frac{LR}{2\beta}. \tag{45}$$

The required result is finally obtained by remarking

$$\|\partial_\phi L^{\text{eval}}(\phi^*_\beta)\| \leq \min\left(1, \frac{1}{\beta}\right)\frac{LR}{2} \leq \frac{LR}{1+\beta}. \tag{46}$$

□

Proof of Remark 1 We now prove Remark 1, which directly follows from the previous proof. Recall that we have just proved

$$\|\partial_\phi L^{\text{eval}}(\phi^*_\beta)\| \leq \frac{LR}{1+\beta}. \tag{47}$$

With the strong convexity of L^{eval}, the gradient is also lower bounded

$$\|\partial_\phi L^{\text{eval}}(\phi^*_\beta)\| \geq \mu\|\phi^*_\beta - \phi^*_\infty\|, \tag{48}$$

meaning that

$$\|\phi^*_\beta - \phi^*_\infty\| \leq \frac{LR}{\mu(1+\beta)}. \tag{49}$$

Similarly, one can show that

$$\|\phi^*_0 - \phi^*_\beta\| \leq \frac{\beta}{1+\beta} \tag{50}$$

up to some constant factor. This can be proved with

$$\|\phi^*_0 - \phi^*_\beta\| \leq \frac{\|\partial_\phi L^{\text{learn}}(\phi^*_\beta)\|}{\mu} = \frac{\beta\|\partial_\phi L^{\text{eval}}(\phi^*_\beta)\|}{\mu} \leq \frac{\beta LR}{(1+\beta)\mu}. \tag{51}$$

Proof of Lemma 3

Lemma 3. *Under Assumptions 1.ii and 1.iii,*

$$\left\|\frac{\mathrm{d}^2\phi^*_\beta}{\mathrm{d}\beta^2}\right\| \leq \frac{\rho}{\mu}\left\|\frac{\mathrm{d}\phi^*_\beta}{\mathrm{d}\beta}\right\|^2 + \frac{2L}{(1+\beta)\mu}\left\|\frac{\mathrm{d}\phi^*_\beta}{\mathrm{d}\beta}\right\|.$$

Proof. The starting point of the proof is the implicit function theorem, that we differentiate with respect to β as a product of functions

$$\begin{aligned}\frac{\mathrm{d}^2\phi^*_\beta}{\mathrm{d}\beta^2} &= \frac{\mathrm{d}}{\mathrm{d}\beta}\left(-\left(\frac{\partial^2\mathcal{L}}{\partial\phi^2}(\phi^*_\beta, \beta)\right)^{-1}\frac{\partial L^{\text{eval}}}{\partial\phi}(\phi^*_\beta)\right) \\ &= -\underbrace{\left(\frac{\mathrm{d}}{\mathrm{d}\beta}\frac{\partial^2\mathcal{L}}{\partial\phi^2}(\phi^*_\beta, \beta)^{-1}\right)\frac{\partial L^{\text{eval}}}{\partial\phi}(\phi^*_\beta)}_{a)} - \underbrace{\frac{\partial^2\mathcal{L}}{\partial\phi^2}(\phi^*_\beta, \beta)^{-1}\left(\frac{\mathrm{d}}{\mathrm{d}\beta}\frac{\partial L^{\text{eval}}}{\partial\phi}(\phi^*_\beta)\right)}_{b)}.\end{aligned} \tag{52}$$

We now individually calculate and bound each term.

a) The differentiation of the inverse of a matrix gives

$$a) = -\partial_\phi^2 \mathcal{L}(\phi_\beta^*, \beta)^{-1} \left(\mathrm{d}_\beta \partial_\phi^2 \mathcal{L}(\phi_\beta^*, \beta)\right) \partial_\phi^2 \mathcal{L}(\phi_\beta^*, \beta)^{-1} \partial_\phi L^{\mathrm{eval}}(\phi_\beta^*), \tag{53}$$

which we can rewrite as

$$a) = \partial_\phi^2 \mathcal{L}(\phi_\beta^*, \beta)^{-1} \left(\mathrm{d}_\beta \partial_\phi^2 \mathcal{L}(\phi_\beta^*, \beta)\right) \mathrm{d}_\beta \phi_\beta^*. \tag{54}$$

The derivative term in the middle of the right hand side is equal to

$$\begin{aligned} \mathrm{d}_\beta \partial_\phi^2 \mathcal{L}(\phi_\beta^*, \beta) &= \mathrm{d}_\beta \left[\partial_\phi^2 L^{\mathrm{learn}}(\phi_\beta^*) + \beta \partial_\phi^2 L^{\mathrm{eval}}(\phi_\beta^*)\right] \\ &= \mathrm{d}_\beta \partial_\phi^2 L^{\mathrm{learn}}(\phi_\beta^*) + \beta \mathrm{d}_\beta \partial_\phi^2 L^{\mathrm{eval}}(\phi_\beta^*) + \partial_\phi^2 L^{\mathrm{eval}}(\phi_\beta^*). \end{aligned} \tag{55}$$

Using the Lipschitz continuity of the Hessians,

$$\left\|\mathrm{d}_\beta \partial_\phi^2 L^{\mathrm{learn}}(\phi_\beta^*) + \beta \mathrm{d}_\beta \partial_\phi^2 L^{\mathrm{eval}}(\phi_\beta^*)\right\| \leq (1+\beta)\rho \left\|\mathrm{d}_\beta \phi_\beta^*\right\|. \tag{56}$$

We can upper bound the norm of the Hessian of L^{eval} by L as L^{eval} is L-smooth. The last two equations hence give

$$\left\|\mathrm{d}_\beta \partial_\phi^2 \mathcal{L}(\phi_\beta^*, \beta)\right\| \leq (1+\beta)\rho \left\|\mathrm{d}_\beta \phi_\beta^*\right\| + L. \tag{57}$$

We finally have

$$\begin{aligned} \|a)\| &\leq \frac{1}{\mu(1+\beta)} \left((1+\beta)\rho \left\|\mathrm{d}_\beta \phi_\beta^*\right\| + L\right) \left\|\mathrm{d}_\beta \phi_\beta^*\right\| \\ &\leq \frac{\rho}{\mu} \left\|\mathrm{d}_\beta \phi_\beta^*\right\|^2 + \frac{L}{(1+\beta)\mu} \left\|\mathrm{d}_\beta \phi_\beta^*\right\|. \end{aligned} \tag{58}$$

b) With the chain rule,

$$\mathrm{d}_\beta \partial_\phi L^{\mathrm{eval}}(\phi_\beta^*) = \partial_\phi^2 L^{\mathrm{eval}}(\phi_\beta^*) \mathrm{d}_\beta \phi_\beta^* \tag{59}$$

so

$$\begin{aligned} \|b)\| &\leq \left\|\partial_\phi^2 \mathcal{L}(\phi_\beta^*, \beta)^{-1}\right\| \left\|\partial_\phi^2 L^{\mathrm{eval}}(\phi_\beta^*)\right\| \left\|\mathrm{d}_\beta \phi_\beta^*\right\| \\ &\leq \frac{L}{(1+\beta)\mu} \left\|\mathrm{d}_\beta \phi_\beta^*\right\|. \end{aligned} \tag{60}$$

□

Proof of Lemma 4

Lemma 4. *Under Assumptions 1.ii, 1.iii and 1.iv, there exists a constant M such that*

$$\left\|\frac{\mathrm{d}^2}{\mathrm{d}\beta^2} \frac{\partial \mathcal{L}}{\partial \theta}(\phi_\beta^*, \beta)\right\| \leq M \left(\left\|\frac{\mathrm{d}\phi_\beta^*}{\mathrm{d}\beta}\right\| + (1+\beta)\left(\left\|\frac{\mathrm{d}\phi_\beta^*}{\mathrm{d}\beta}\right\|^2 + \left\|\frac{\mathrm{d}^2 \phi_\beta^*}{\mathrm{d}\beta^2}\right\|\right)\right).$$

Proof. We want to bound the norm of $\mathrm{d}_\beta^2 \partial_\theta \mathcal{L}(\phi_\beta^*, \beta)$. The first order derivative can be calculated with the chain rule of differentiation

$$\mathrm{d}_\beta \partial_\theta \mathcal{L}(\phi_\beta^*, \beta) = \partial_\beta \partial_\theta \mathcal{L}(\phi_\beta^*, \beta) + \partial_\phi \partial_\theta \mathcal{L}(\phi_\beta^*, \beta) \mathrm{d}_\beta \phi_\beta^*. \tag{61}$$

We then once again differentiate this equation with respect to β. The $\partial_\beta \partial_\theta \mathcal{L}(\phi_\beta^*, \beta)$ term has in fact, due to the nature of $\mathcal{L}$, no direct dependence on β and is equal to $\partial_\theta L^{\mathrm{eval}}(\phi_\beta^*)$. Hence

$$\mathrm{d}_\beta \partial_\beta \partial_\theta \mathcal{L}(\phi_\beta^*, \beta) = \partial_\phi \partial_\theta L^{\mathrm{eval}}(\phi_\beta^*) \mathrm{d}_\beta \phi_\beta^*. \tag{62}$$

Differentiating the other term yields

$$\begin{aligned} \mathrm{d}_\beta \left[\partial_\phi \partial_\theta \mathcal{L}(\phi_\beta^*, \beta) \mathrm{d}_\beta \phi_\beta^*\right] = \left[\partial_\beta \partial_\phi \partial_\theta \mathcal{L}(\phi_\beta^*, \beta) + \partial_\phi^2 \partial_\theta \mathcal{L}(\phi_\beta^*, \beta) \otimes \mathrm{d}_\beta \phi_\beta^*\right] \mathrm{d}_\beta \phi_\beta^* + \\ \partial_\phi \partial_\theta \mathcal{L}(\phi_\beta^*, \beta) \mathrm{d}_\beta^2 \phi_\beta^*. \end{aligned} \tag{63}$$

Therefore,

$$\begin{aligned} \mathrm{d}_\beta^2 \partial_\theta \mathcal{L}(\phi_\beta^*, \beta) = 2\partial_\phi \partial_\theta L^{\mathrm{eval}}(\phi_\beta^*) \mathrm{d}_\beta \phi_\beta^* + \partial_\phi^2 \partial_\theta \mathcal{L}(\phi_\beta^*, \beta) \otimes \mathrm{d}_\beta \phi_\beta^* \otimes \mathrm{d}_\beta \phi_\beta^* + \\ \partial_\phi \partial_\theta \mathcal{L}(\phi_\beta^*, \beta) \mathrm{d}_\beta^2 \phi_\beta^*. \end{aligned} \tag{64}$$

We now individually bound each term:

– due to Assumption 1.i, $\phi \mapsto \partial_\theta L^{\text{eval}}(\phi)$ is B^{eval}-Lipschitz continuous, so $\|\partial_\phi \partial_\theta L^{\text{eval}}\| \leq B^{\text{eval}}$ and

$$\left\|2\partial_\phi \partial_\theta L^{\text{eval}}(\phi^*_\beta) \mathrm{d}_\beta \phi^*_\beta\right\| \leq 2B^{\text{eval}} \left\|\mathrm{d}_\beta \phi^*_\beta\right\|. \tag{65}$$

– similarly to the previous point,

$$\left\|\partial_\phi \partial_\theta \mathcal{L}(\phi^*_\beta) \mathrm{d}^2_\beta \phi^*_\beta\right\| \leq (B^{\text{learn}} + \beta B^{\text{eval}}) \left\|\mathrm{d}^2_\beta \phi^*_\beta\right\|. \tag{66}$$

– Assumption 1.iv ensures that $\phi \mapsto \partial_\phi \partial_\theta \mathcal{L}(\phi, \beta)$ is $(1+\beta)\sigma$-Lipschitz continous and

$$\left\|\partial^2_\phi \partial_\theta \mathcal{L}(\phi^*_\beta, \beta) \otimes \mathrm{d}_\beta \phi^*_\beta \otimes \mathrm{d}_\beta \phi^*_\beta\right\| \leq (1+\beta)\sigma \left\|\mathrm{d}_\beta \phi^*_\beta\right\|^2. \tag{67}$$

Take $M := \max(2B^{\text{eval}}, B^{\text{learn}}, \sigma)$: we now have the desired result. □

S3.4 A corollary of Theorem 1

Theorem 1 highlights the importance of considering β as a hyperparameter of the learning rule that needs to be adjusted to yield the best possible meta-gradient estimate. Corollary 1 removes the dependence in β and considers the best achievable bound under given fixed point approximation errors.

Corollary 1. *Under Assumption 1, if we suppose that for every strictly positive β we approximate the two fixed points with precision δ and δ' and if $(\delta + \delta') < C/B^{\text{learn}}$, the best achievable bound in Theorem 1 is smaller than*

$$B^{\text{eval}}\delta' + 2\sqrt{CB^{\text{learn}}(\delta + \delta')}$$

and is attained for β equal to

$$\beta^*(\delta, \delta') = \frac{\sqrt{B^{\text{learn}}(\delta + \delta')}}{\sqrt{C} - \sqrt{B^{\text{learn}}(\delta + \delta')}}.$$

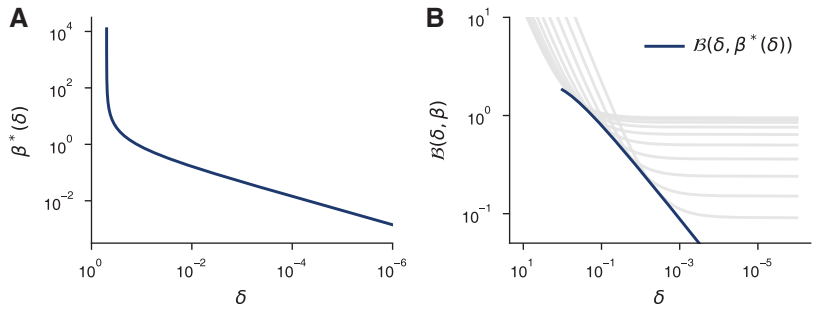

Figure S2: Visualization of Corollary 1. (A) β value that minimizes the bound, as a function of $\delta = \delta'$. (B) Best achievable bound as a function of $\delta = \delta'$ in blue (more precisely the one before the last upper bound in the proof). The grey lines are the bounds from Theorem 1 we laid out on Fig. S1C.

The most limiting part of the bound depends on the sum $\delta + \delta'$ and not on the individual quantities, suggesting that the two errors should be of the same magnitude to avoid unnecessary computations.

Proof. The β derivative of the bound $\mathcal{B}$ obtained in Theorem 1 is

$$\frac{\partial \mathcal{B}}{\partial \beta}(\delta, \delta', \beta) = -\frac{B^{\text{learn}}(\delta + \delta')}{\beta^2} + \frac{C}{(1+\beta)^2} \tag{68}$$

and vanishes for β verifying

$$\beta\left(\sqrt{C} - \sqrt{B^{\text{learn}}(\delta + \delta')}\right) = \sqrt{B^{\text{learn}}(\delta + \delta')}. \tag{69}$$

As $(\delta + \delta') < C/B^{\text{learn}}$, the previous criterion is met when β is equal to the positive

$$\beta^* := \frac{\sqrt{B^{\text{learn}}(\delta + \delta')}}{\sqrt{C} - \sqrt{B^{\text{learn}}(\delta + \delta')}}. \tag{70}$$

The optimal bound is then

$$\begin{aligned}\mathcal{B}(\delta, \delta', \beta^*) &= B^{\text{eval}}\delta' + \sqrt{B^{\text{learn}}(\delta + \delta')}\left(\sqrt{C} - \sqrt{B^{\text{learn}}(\delta + \delta')}\right) + \sqrt{CB^{\text{learn}}(\delta + \delta')} \\ &\leq B^{\text{eval}}\delta' + 2\sqrt{CB^{\text{learn}}(\delta + \delta')}.\end{aligned}$$

□

S3.5 Verification of the theoretical results on an analytical problem

We investigate a quadratic approximation of the complex synapse model, in which everything can be calculated in closed form and where the assumptions needed for the theory hold. Define L^{learn} and L^{eval} as follows[3]:

$$\begin{aligned}L^{\text{learn}}(\phi, \omega) &= \frac{1}{2}(\phi - \phi^l)^\top H(\phi - \phi^l) + \frac{\lambda}{2}\|\phi - \omega\|^2 \\ L^{\text{eval}}(\phi) &= \frac{1}{2}(\phi - \phi^e)^\top H(\phi - \phi^e)\end{aligned}$$

where λ is a scalar that controls the strength of the regularization that we consider fixed, ϕ^l and ϕ^e two vectors and H a positive definite diagonal matrix. The rationale behind this approximation is the following: the data-driven learning and evaluation losses share the same curvature but have different minimizers, respectively ϕ^l and ϕ^e. The matrix H then models the Hessian and we consider it diagonal for simplicity. Thanks to the quadratic approximation, many quantities involved in our contrastive meta-learning rule can be calculated in closed form.

Calculation of the finite difference error. A formula for the minimizer of $\mathcal{L} = L^{\text{learn}} + \beta L^{\text{eval}}$ can be derived analytically. The derivative of $\mathcal{L}$ vanishes if and only if

$$\left((1 + \beta)H + \lambda \text{Id}\right)\phi - H\phi^l - \beta H\phi^e - \lambda\omega = 0,$$

hence

$$\phi^*_{\theta,\beta} = \left((1 + \beta)\text{Id} + \lambda H^{-1}\right)^{-1}\left(\phi^l + \beta\phi^e + \lambda H^{-1}\omega\right).$$

λH^{-1} is an interesting quantity in this example. It acts as the effective per-coordinate regularization strength: regularization will be stronger on flat directions.

The meta-gradient calculation follows. As

$$\partial_\omega \mathcal{L}(\phi, \omega, \beta) = -\lambda(\phi - \omega),$$

the use of the equilibrium propagation theorem (Theorem S1) gives

$$\begin{aligned}\nabla_\omega &= \frac{\text{d}}{\text{d}\beta}\frac{\partial \mathcal{L}}{\partial \omega}(\phi^*_{\theta,\beta}, \omega, \beta)\bigg|_{\beta=0} \\ &= \frac{\partial^2 L^{\text{learn}}}{\partial\phi\partial\omega}(\phi^*_{\theta,0}, \omega)\ \frac{\text{d}\phi^*_{\theta,\beta}}{\text{d}\beta}\bigg|_{\beta=0} + 0 \\ &= -\lambda\ \frac{\text{d}\phi^*_{\theta,\beta}}{\text{d}\beta}\bigg|_{\beta=0}.\end{aligned}$$

[3] In our experiments, we take the dimension of the parameter space N to be equal to 50. The Hessian is taken to be $\text{diag}(1, ..., 1/N)$. ω is randomly generated according to $\omega \sim \mathcal{N}(0, \sigma_\omega)$ with $\sigma_\omega = 2$. ϕ^l and ϕ^e are drawn around $\phi^\tau \sim \mathcal{N}(0, \sigma_\tau)$ (with $\sigma_\tau = 1$).

It now remains to calculate the derivative of $\phi^*_{\theta,\beta}$ with respect to β using the formula of $\phi^*_{\theta,\beta}$:

$$\begin{aligned}\frac{\mathrm{d}\phi^*_{\theta,\beta}}{\mathrm{d}\beta} &= \left((1+\beta)\mathrm{Id} + \lambda H^{-1}\right)^{-1}\phi^e \\ &\quad - \left((1+\beta)\mathrm{Id} + \lambda H^{-1}\right)^{-2}\left(\phi^l + \beta\phi^e + \lambda H^{-1}\omega\right) \\ &= \left((1+\beta)\mathrm{Id} + \lambda H^{-1}\right)^{-2}\left((1+\beta)\phi^e + \lambda H^{-1}\phi^e - \phi^l - \beta\phi^e - \lambda H^{-1}\omega\right) \\ &= \left((1+\beta)\mathrm{Id} + \lambda H^{-1}\right)^{-2}\left((\phi^e - \phi^l) + \lambda H^{-1}(\phi^e - \omega)\right)\end{aligned}$$

Define $\psi := (\phi^e - \phi^l) + \lambda H^{-1}(\phi^e - \omega)$; the meta-gradient finally is

$$\nabla_\omega = -\lambda(\mathrm{Id} + \lambda H^{-1})^{-2}\psi.$$

We can now calculate the finite difference error. Recall the equilibrium propagation estimate at fixed points

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

In this formulation, it is equal to

$$\begin{aligned}\widehat{\nabla}^*_\omega &= -\frac{\lambda}{\beta}(\phi^*_{\theta,\beta} - \phi^*_{\theta,0}) \\ &= -\lambda\left((\mathrm{Id} + \lambda H^{-1})((1+\beta)\mathrm{Id} + \lambda H^{-1})\right)^{-1}\psi \\ &= (\mathrm{Id} + \lambda H^{-1})\left((1+\beta)\mathrm{Id} + \lambda H^{-1})\right)^{-1}\nabla_\omega.\end{aligned}$$

The finite difference can now be lower and upper bounded. First,

$$\nabla_\omega - \widehat{\nabla}^*_\omega = \beta\left((1+\beta)\mathrm{Id} + \lambda H^{-1}\right)^{-1}\nabla_\omega.$$

Introduce μ the smallest eigenvalue of H and L its largest one. We then have

$$\frac{\mu\beta}{(1+\beta)\mu + \lambda}\|\nabla_\omega\| \leq \|\nabla_\omega - \widehat{\nabla}^*_\omega\| \leq \frac{L\beta}{(1+\beta)L + \lambda}\|\nabla_\omega\|. \tag{71}$$

This shows that the finite difference error part of Theorem 1 is tight and, in this case, accurately describes the behavior of the finite difference error as a function of β.

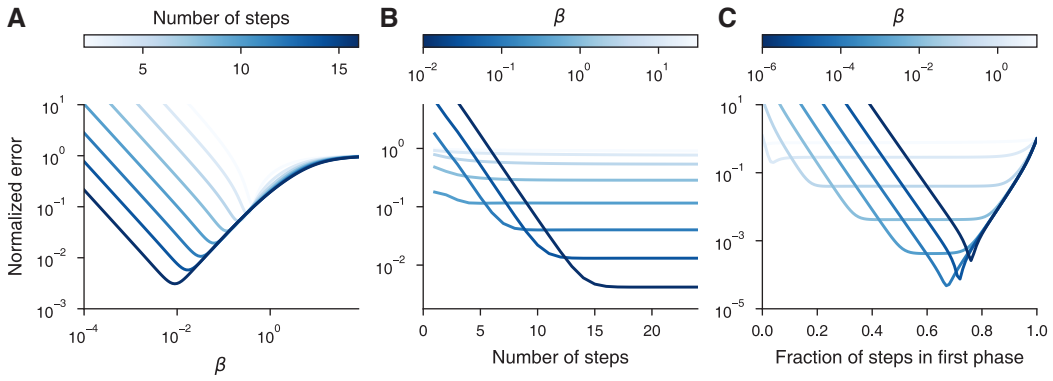

Figure S3: Empirical verification of the theoretical results on an analytical quadratic approximation of the synaptic model. We plot the normalized error between the meta-gradient estimate $\widehat{\nabla}_\theta$ and the true one ∇_θ, as a function of β (A), of the number of steps in the two phases (which is a proxy for $-\log\delta$ and $-\log\delta'$ used in the theory) (B), and as a function of the allocation of the computational resources between the two phases, the total number of steps being fixed to 100 (C).

Empirical results. The solution approximation induced error part of the bound cannot be treated analytically as it depends on δ and δ', which are in essence empirical quantities. We cannot directly

control them either. Instead, we use the number of gradient descent steps to minimize $\mathcal{L}$ as a proxy, that is closely related to $-\log\delta$ when gradient descent has a linear convergence rate. We choose the number of steps to be the same in the two phases, for the sake of simplicity, even though it may not be optimal. We plot the evolution of the normalized error

$$\frac{\|\nabla_\omega - \widehat{\nabla}_\omega\|}{\|\nabla_\omega\|}$$

between the meta-gradient and the contrastive estimate (15) in Fig. S3. The qualitative behavior of this error, as a function of β (Fig. S3A) and of number of steps (Fig. S3B), is accurately captured by Theorem 1 (compare with Fig. S1A and C).

We finish the study of this quadratic model by probing the (δ, δ') space in a different way, by fixing the total number of steps and then modifying the allocation across the two phases (Fig. S3C). The best achievable error, as a function of β, decreases before some β^* value and then increases, following the predictions of Theorem 1: too small β values turn out to hurt performance when the solutions cannot be approximated arbitrarily well. Interestingly, the error plateaus for large β values and the size of the plateau decreases with β until reaching a critical value where it disappears. A conservative choice in practice is therefore to overestimate β, as it reduces the meta-gradient estimation sensitivity to a sub-optimal allocation, with only a minor degradation in the best achievable quality.

S3.6 Verification the theoretical results on a simple hyperparameter optimization task

We now move to a more complicated setting that is closer to problems of practical interest, and in which we are not guaranteed that the assumptions of the theory hold. Still, it is simple enough such that we can calculate the exact value of the meta-gradient ∇_θ using the analytical formula (24). This problem is a single-task regularization-strength learning problem [20, 21, 23, 52, 94] on the Boston housing dataset [98] (70% learning and 30% evaluation split). We study a nonlinear neural network model f_ϕ with a small hidden layer (20 neurons, hyperbolic tangent transfer function). The bilevel optimization problem we are solving here is the one we consider in Section 5.2, that is:

$$\begin{aligned} &\min_\lambda \frac{1}{|D^{\text{eval}}|} \sum_{(x,y)\,\in\, D^{\text{eval}}} l(f_{\phi^*_\lambda}(x), y) \\ &\text{s.t. } \phi^*_\lambda \in \arg\min_\phi \frac{1}{|D^{\text{learn}}|} \sum_{(x,y)\,\in\, D^{\text{learn}}} l(f_\phi(x), y) + \frac{1}{2}\sum_{i=1}^{|\phi|} \lambda_i \phi_i^2. \end{aligned} \tag{72}$$

Meta-gradient estimation error. We plot the normalized error between the meta-gradient estimate and its true value on Fig S4. The qualitative behavior closely matches the one we obtained for the quadratic analytical model in the last section, as well as the ones predicted by our theory.

Comparison with other implicit gradient methods. We also use this problem to directly compare the meta-gradient approximation error made by our contrastive meta-learning rule (CML) to other implicit gradient methods, namely recurrent backpropagation (RBP) and conjugate gradient (CG). To make the comparison fair, we pick the hyperparameters that yield the smallest error for each method (β and the parameters of the optimizer minimizing the second phase for CML, a scaling parameter for RBP, none for CG). Fig S5 characterizes the meta-gradient estimation errors by the different methods.

We first perfectly solve the first phase so that $\hat{\phi}_\lambda = \phi^*_\lambda$ and compare how efficient the second phase of those algorithms is. The Hessian of the learning loss $\partial^2_\phi L^{\text{learn}}$ is positive definite as shown in Fig. S6A. The conjugate gradient method is therefore much more efficient than the other methods as its assumptions are met, and quickly reaches the numerical accuracy limit. Our rule compares favorably to recurrent backpropagation, even though the theoretical bound is weaker ($\sqrt{\delta'}$ for our rule compared to δ' for implicit methods [21]). A possible explanation comes from the fact that we are using Nesterov accelerated gradient descent for the second phase of our contrastive update, whereas the fixed point iteration of RBP is a form of gradient descent.

We repeat our analysis in the more realistic setting in which $\hat{\phi}_\lambda$ is not equal to ϕ^*_λ. We use the same number of steps in the two phases (and the same estimate for $\hat{\phi}_\lambda$ and find that recurrent

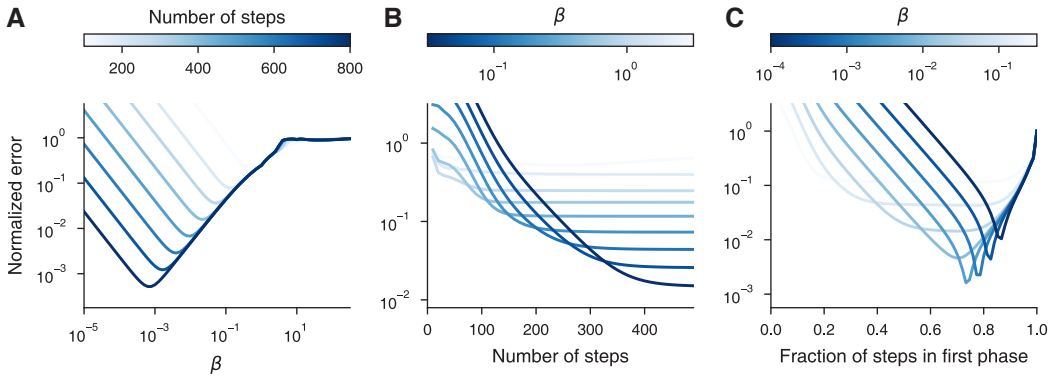

Figure S4: Empirical verification of the theoretical results on a regularization-strength learning problem on the Boston dataset. We plot the normalized error between the meta-gradient estimate $\widehat{\nabla}_\lambda$ and the true one ∇_λ, as a function of β (A), of the number of steps in the two phases (which is a proxy for $-\log\delta$ and $-\log\delta'$ used in the theory) (B), and as a function of the allocation of the computational resources between the two phases, the total number of steps being fixed to 750 (C).

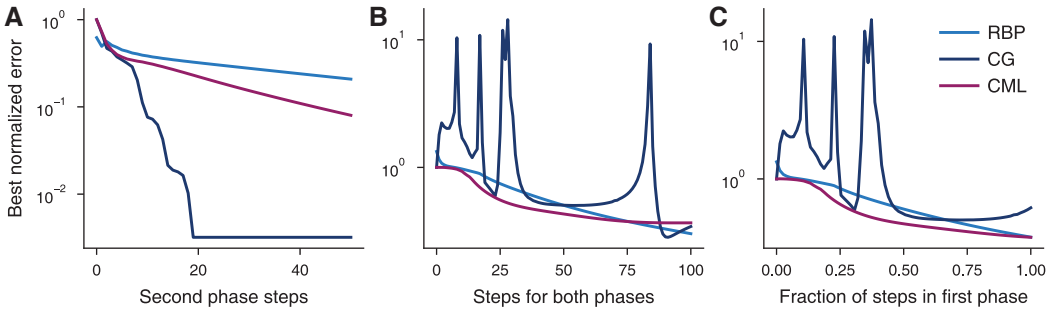

Figure S5: Comparison of the meta-gradient estimation errors provided our contrastive meta-learning rule (CML), recurrent backpropagation (RBP) and conjugate gradients (CG), on a regularization-strength learning problem on the Boston dataset. The hyperparameters for each value of the x axis, such that the normalized error is minimized. (A) Error as a function of the number of steps in the second phase, the first phase being perfectly solved. (B) Error as a function of the number of steps performed in the two phase, which is fixed. (C) Error as a function of the fraction of steps in the first phase, the total number of steps for the two phases being fixed to 75.

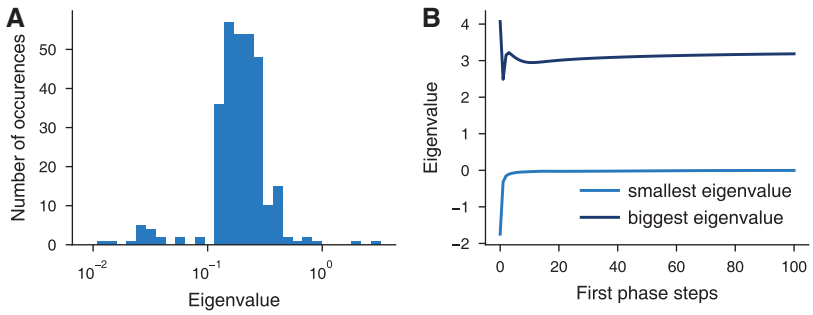

Figure S6: Eigenvalues of the Hessian of the learning loss $\partial^2_\phi L^{\text{learn}}(\hat{\phi}_\lambda, \lambda)$ on the regularization-strength learning problem on the Boston dataset. (A) Spectrum of the Hessian when the first phase is perfectly solved, i.e., $\hat{\phi}_\lambda = \phi^*_\lambda$. (B) Smallest and biggest eigenvalue of the Hessian, as a function of the number of steps in the first phase. The higher the number of steps, the close $\hat{\phi}_\lambda$ is to ϕ^*_λ.

backpropagation and our contrastive meta-learning rule improve their estimate of the meta-gradient as the number of steps for both phases increases, cf. Fig. S5.B. In contrast, conjugate gradient is unstable when the number of steps is low. In Fig. S6B, we check whether the needed assumptions are satisfied by plotting the smallest eigenvalue of the Hessian of the learning loss $\partial_\phi^2 L^{\text{learn}}(\hat{\phi}_\lambda, \lambda)$, as a function of the number of steps. We find that this eigenvalue is negative on the range of the number of steps we consider, the Hessian is therefore not positive definite so the conjugate gradient method cannot approximate μ well. We obtain the same qualitative behavior when we fix the total number of steps, and vary the faction of steps in the first phase, cf. Fig. S5C.

S4 Experimental details

S4.1 Supervised meta-optimization

Task details. For the supervised meta-optimization experiments we meta-learn parameter-wise l2-regularization strengths ($\omega = 0$) on the CIFAR-10 image classification task [48] starting each learning phase from a fixed neural network initialization. The dataset comprises 60000 32x32 RGB images divided into 10 classes, with 6000 images per class. We split the 50000 training images randomly in half to obtain a training set and a validation set for meta-learning and use the remaining 10000 test images for testing. In Tab. S2 we report additional results on the simpler MNIST image classification task [100] for which we use the same data splitting strategy.

Table S2: Meta-learning a per-synapse regularization strength meta-parameter (cf. Section 4.1) on MNIST. Average accuracies (acc.) $\pm$ s.e.m. over 10 seeds.

Method	Validation acc. (%)	Test acc. (%)
T1-T2	$98.70^{\pm 00.08}$	$97.63^{\pm 00.03}$
CG	$97.02^{\pm 00.28}$	$96.96^{\pm 00.15}$
RBP	$99.53^{\pm 00.01}$	$97.31^{\pm 00.02}$
CML	$99.45^{\pm 00.16}$	$97.92^{\pm 00.11}$

Additional results. We perform an additional experiment investigating how the number of lower-level parameter updates affects the meta-learning performance of our method and comparison methods. We consider a simplified data regime for this experiment, using a random subset of 1000 examples of CIFAR-10 split into 50 samples for the learning loss and 950 samples for the evaluation loss which allows us to fit all samples into a single batch during learning and meta-learning. Results shown in Fig. S7 demonstrate that our contrastive meta-learning rule is able to fit the meta-parameters to the validation set across different number of lower-level parameter updates while competing methods require more updates to obtain similar performance. We found the conjugate gradient method (CG) to be unstable in this setting. To obtain these results, we tuned the hyperparameters for each method for each number of lower-level parameter updates.

Architecture details. For CIFAR-10 experiments we use a modified version of the classic LeNet-5 model [101] where we insert batch normalization layers [102] before each nonlinearity and replace the hyperbolic tangent nonlinearities with rectified linear units. For MNIST experiments we use a feedforward neural network with 5 hidden layers of size 256 and hyperbolic tangent nonlinearity.

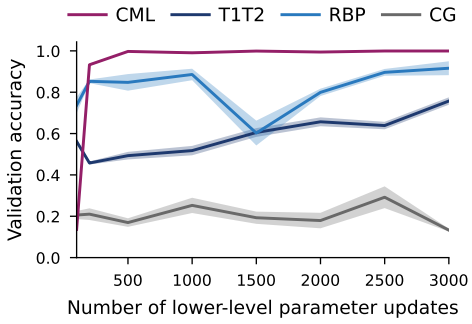

Figure S7: Dependence of the final validation accuracy on the number of lower-level parameter updates obtained after meta-optimizing per-synapse regularization strength meta-parameters on a subset of CIFAR-10. A random subset of 1000 examples from CIFAR-10 are split into 50 samples for the learning loss and 950 samples for the evaluation loss. Mean over 10 random seeds with error bars indicating ± 1 s.e.m.

Hyperparameters. We perform a comprehensive random hyperparameter search for each method with the search space for CIFAR-10 experiments specified in Tab. S5 and the search space for MNIST experiments specified in Tab. S6.

In Fig. 1B, we furthermore investigate the interaction of β and the number of first phase steps on the validation loss, keeping all other hyperparameters fixed. We compare it to the corresponding theoretical prediction visualized in Fig. S1B, for the case where L^{eval} is independent of θ and δ' is fixed.

S4.2 Few-shot image classification

Task details. We follow the standard experimental setup [18, 22, 103, 104] for our Omniglot [57] and miniImageNet [56] experiments.

Additional results. The results of related work reported in the main text (Tabs. 2 and 3) are taken from the original papers, except for Omniglot first-order MAML which is reported in ref. [58]. In Tab. S3 we provide results for additional 5-way Omniglot variants that are easier than the 20-way ones studied in the main text.

Architecture details. For Omniglot, we use max-pooling instead of stride in the convolutional layers, as we found the latter led to optimization instabilities, as previously reported [109]. We evaluate the statistics of batch normalization units [102] on the test set as in ref. [18], which yields a transductive classifier. More complex architectures whereby a second modulatory neural network which generates task-specific parameters is explicitly modeled [a hypernetwork; 110–112] can be easily accommodated into our framework, but here for simplicity we implement our top-down modulation model by taking advantage of existing batch normalization layers in our neural networks and consider the gain and shift parameters of these units as well as the synaptic weights and biases of the output layer as our task-specific parameters ϕ.

Optimization details. We used the symmetric version of our contrastive rule for meta-learning and the Kaiming scheme for parameter initialization [113]. The task-specific learning and evaluation losses are both taken to be the cross-entropy with dataset splits into learning and evaluation data following the setup considered by Finn et al. [18]. In order to stabilize results, we used Polyak averaging [114] for the meta-parameters to compute final performance. Specifically, we started averaging meta-parameters after a certain number of meta-parameter updates (5 for Omniglot, 50 for miniImageNet). Note that the performance of the non-averaged meta-parameters performs only slightly differently averaged over iterations but is considerably more noisy.

Hyperparameters. We perform a comprehensive grid-search over hyperparameters with search ranges and optimal hyperparameters found reported in Tab. S7.

Table S3: Few-shot learning of Omniglot characters. We report results obtained with contrastive meta-learning for the synaptic and modulatory models. We present test set classification accuracy (%) averaged over 5 seeds $\pm$ std.

Method	5-way 1-shot	5-way 5-shot	20-way 1-shot	20-way 5-shot
MAML [18]	$98.7^{\pm 0.4}$	$99.9^{\pm 0.1}$	$95.8^{\pm 0.3}$	$98.9^{\pm 0.2}$
First-order MAML [18]	$98.3^{\pm 0.5}$	$99.2^{\pm 0.2}$	$89.4^{\pm 0.5}$	$97.9^{\pm 0.1}$
Reptile [58]	$97.68^{\pm 0.04}$	$99.48^{\pm 0.06}$	$89.43^{\pm 0.14}$	$97.12^{\pm 0.32}$
iMAML [22]	$99.16^{\pm 0.35}$	$99.67^{\pm 0.12}$	$94.46^{\pm 0.42}$	$98.69^{\pm 0.1}$
CML (synaptic)	$98.11^{\pm 0.34}$	$99.49^{\pm 0.16}$	$94.16^{\pm 0.12}$	$98.06^{\pm 0.26}$
CML (modulatory)	$98.05^{\pm 0.06}$	$99.45^{\pm 0.04}$	$94.24^{\pm 0.39}$	$98.60^{\pm 0.27}$

S4.3 Few-shot regression in recurrent spiking network

Task details. We consider a standard sinusoidal 10-shot regression problem. For each task a sinusoid with random amplitude sampled uniformly from $[0.1, 5.0]$ and random phase sampled uniformly from $[0, \pi]$ is generated. 10 data points are drawn uniformly from the range $[-5, 5]$ both for the learning loss and for the evaluation loss.

Architecture details. We encode the input with a population of 100 neurons similar to Bellec et al. [61]. Each neuron i has a Gaussian response field with the mean values μ_i evenly distributed in the range $[0, 1]$ across neurons and a fixed variance $\sigma^2 = 0.0002$. The firing probability of each neuron at a single time step is given by $p_i = \exp(\frac{-(\mu_i - z)^2}{2\sigma^2})$ where z is the input value standardized to the range $[0, 1]$. We generate 20 time steps for each data point by sampling spikes from a Bernoulli distribution given the firing probabilities p_i for each neuron.

We use a singe-layer recurrent spiking neural network with leaky integrate and fire neurons that follow the time-discretized dynamics with step size $\Delta t = 1.0$ (notation taken from Bellec et al. [59]):

$$h_j^{t+1} = \alpha h_j^t + \sum_{i \neq j} W_{ji}^{\text{rec}} z_i^t + \sum_i W_{ji}^{\text{in}} x_i^{t+1} - z_j^t v_{\text{th}} \tag{73}$$

$$z_j^t = \Theta(h_j^t - v_{\text{th}}) \tag{74}$$

$$y_k^{t+1} = \kappa y_k^t + \sum_j W_{kj}^{\text{out}} z_j^t \tag{75}$$

where $W^{\text{in}}, W^{\text{rec}}, W^{\text{out}}$ are the synaptic input, recurrent and output weights, $\alpha = \exp(-\frac{\Delta t}{\tau_{\text{hidden}}})$ and $\kappa = \exp(-\frac{\Delta t}{\tau_{\text{out}}})$ are decay factors with $\tau_{\text{hidden}} = \tau_{\text{out}} = 30.0$, $v_{\text{th}} = 0.1$ is the threshold potential, and $\Theta(\cdot)$ denotes the Heaviside step function. The weights are initialized using the Kaiming normal scheme [113] and scaled down by a factor of $0.1, 0.01, 0.1$ for $W^{\text{in}}, W^{\text{rec}}, W^{\text{out}}$ respectively.

Optimization details. The weights are updated according to e-prop [59]:

$$\Delta W_{kj}^{\text{out}} \propto \sum_t (y_k^{*,t} - y_k^t) \sum_{t' \leq t} (\kappa^{t-t'} z_j^{t'}) \tag{76}$$

$$\Delta W_{ji}^{\text{rec}} \propto \sum_t (\sum_k W_{kj}^{\text{out}} (y_k^{*,t} - y_k^t)) \sum_{t' \leq t} (\kappa^{t-t'} h_j^{t'} \sum_{t'' \leq t'} (\alpha^{t'-t''} z_i^{t''})) \tag{77}$$

$$\Delta W_{ji}^{\text{in}} \propto \sum_t (\sum_k W_{kj}^{\text{out}} (y_k^{*,t} - y_k^t)) \sum_{t' \leq t} (\kappa^{t-t'} h_j^{t'} \sum_{t'' \leq t'} (\alpha^{t'-t''} x_i^{t''})) \tag{78}$$

The loss is computed as the mean-squared error between the target and the prediction given by the average output over time. Note that the output of the network is non-spiking. We add a regularization term to the loss that is computed as the mean squared difference between the average neuron firing rate and a target rate and decrease the learning rate for updating W^{out} with e-prop by a factor of 0.1. We use the symmetric version of our contrastive rule to obtain meta-updates.

Comparison methods. We compare our method to a standard baseline where both fast parameter updates and slow meta-parameter updates are computed by backpropagating through the synaptic plasticity process (BPTT+BPTT) using surrogate gradients to handle spiking nonlinearities [60]. As this biologically-implausible process is computationally expensive, we restrict the number of update steps on the learning loss to 10 changes as done by prior work [18]. For a second comparison method, we compute the fast parameter updates using the e-prop update stated above and use backpropagation through 10 e-prop updates for meta-parameter updates (BPTT+e-prop).

Hyperparameters. For each method we employ an extensive random hyperparameter search over the search space defined in Tab. S8 using a meta-validation set to select the optimal set of hyperparameters.

S4.4 Meta-reinforcement learning

Task details. The contextual wheel bandit was introduced by Riquelme et al. [64] to parametrize the task difficulty of a contextual bandit task in terms of its exploration-exploitation trade-off. Each task consists of a sequence of context coordinates X randomly drawn from the unit circle and a scalar radius $\delta \in [0, 1]$. The radius δ tiles the unit circle into a low-reward region and a high-reward region, see Fig. S8. If the current context lies within the low-reward region, $\|X\| \leq \delta$, all actions $a \in \{1, 2, 3, 4\}$ return reward $r \sim \mathcal{N}(1.0, 0.01^2)$ except for the last action $a = 5$ which returns $r \sim \mathcal{N}(1.2, 0.01^2)$ and is thus optimal. If the current context lies within the high-reward region, one of the first four actions is optimal returning a high reward $r \sim \mathcal{N}(50.0, 0.01^2)$, the last arm still returns $r \sim \mathcal{N}(1.2, 0.01^2)$ and the remaining arms return $r \sim \mathcal{N}(1.0, 0.01^2)$. Which of the first four actions returns the high reward depends on the quadrant of the high-reward region in which the current context lies. Action 1 is optimal in the upper right quadrant, action 2 in the lower right quadrant, action 3 in the upper left quadrant and action 4 in the lower left quadrant.

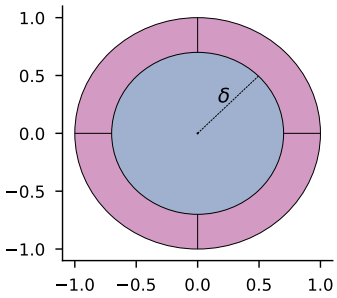

Figure S8: The wheel bandit task tiles the context space into an inner low-reward region (blue) and a high-reward outer rim (purple). Across tasks the radius δ of the inner low-reward region is varied. The high-reward region is divided into 4 quadrants, depending on which the optimal action changes.

Additional results. Following previous work [65, 66], we treat the contextual wheel bandit as a meta-learning problem. During meta-learning, we sample $M = 64$ different wheel tasks $\{\delta_i\}_{i=1}^{M}$, $\delta_i \sim \mathcal{U}(0, 1)$ for each of which we sample a sequence of $N = 562$ random contexts, random actions and corresponding rewards $\{(X_j, a_j, r_j)\}_{j=1}^{N}$. We use 512 observations for the learning loss and 50 observations for the evaluation loss. Both the learning l^{learn} and the evaluation loss l^{eval} are measured as the mean-squared error between observed reward and the predicted value for the corresponding action. For each meta-learning step, we randomly sample from the M tasks such that specific tasks may be encountered multiple times. After meta-learning, we evaluate the agent online on a long episode with 80000 contexts and track its cumulative reward relative to the cumulative reward obtained by an agent that chooses its actions at random. As done by Riquelme et al. [64], each action is initially explored twice before choosing actions according to the agent's policy. The extended results over more settings for δ can be seen in Tab. S4. Results for the NeuralLinear baseline reported here and in the main text are taken from the original paper [64].

Architecture results. For all methods including the synaptic consolidation and modulatory network model, we consider as a base a multilayer perceptron with ReLU nonlinearites and two hidden layers with 100 units. For the modulatory network model, each hidden unit is multiplied by a gain and shifted by a bias prior to applying the nonlinearity. The non-meta-learned baseline NeuralLinear we report from Riquelme et al. [64] additionally uses a Bayesian regression head for each action and applies Thompson sampling to choose actions.

Optimization and evaluation details. During online evaluation, we take the greedy action with respect to the predicted expected rewards on each context and store each observation (X_j, a_j, r_j) in

Table S4: Cumulative regret on the wheel bandit problem for different values of δ. Values are normalized by the cumulative regret of a uniformly random agent. Averages over 50 seeds $\pm$ s.e.m.

δ	0.5	0.7	0.9	0.95	0.99
NeuralLinear [64]	$0.95^{\pm 0.02}$	$1.60^{\pm 0.03}$	$4.65^{\pm 0.18}$	$9.56^{\pm 0.36}$	$49.63^{\pm 2.41}$
MAML	$0.45^{\pm 0.01}$	$0.62^{\pm 0.03}$	$1.02^{\pm 0.76}$	$1.56^{\pm 0.62}$	$15.21^{\pm 1.69}$
CML (synaptic)	$0.40^{\pm 0.02}$	$0.45^{\pm 0.01}$	$0.82^{\pm 0.02}$	$1.42^{\pm 0.07}$	$12.27^{\pm 1.02}$
CML (modulatory)	$0.42^{\pm 0.01}$	$0.65^{\pm 0.03}$	$1.83^{\pm 0.11}$	$3.68^{\pm 0.59}$	$16.46^{\pm 1.80}$

a replay buffer. This data is used to train the fast parameters every t_f contexts for t_s steps, where t_f, t_s are hyperparameters tuned for every method.

Hyperparameters. We perform a comprehensive random hyperparameter search for each method with the search space specified in Tab. S9. Optimal parameters are selected on 5 validation tasks with $\delta = 0.95$.

Table S5: Hyperparameter search space for the supervised meta-optimization experiment on CIFAR-10. For all methods 500 samples were randomly drawn from the search space and Asynchronous HyperBand from ray tune [120] was used for scheduling with a grace period of 10. Best found parameters are marked in bold.

Hyperparameter	CML	CG
batch_size	500	500
β	$\{0.01, 0.03, 0.1, 0.3, 1.0, 3.0, \mathbf{10.0}\}$	-
λ	$\{10^{-5}, 10^{-4}, 10^{-3}, 10^{-2}, \mathbf{10^{-1}}\}$	$\{\mathbf{10^{-5}}, 10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}\}$
lr_inner	$\{0.0001, \mathbf{0.0003}, 0.001, 0.003, 0.01, 0.03, 0.1\}$	$\{0.0001, 0.0003, \mathbf{0.001}, 0.003, 0.01, 0.03, 0.1\}$
lr_nudged	$\{\mathbf{0.0001}, 0.0003, 0.001, 0.003, 0.01, 0.03, 0.1\}$	-
lr_outer	$\{0.0001, 0.0003, 0.001, 0.003, 0.01, \mathbf{0.03}, 0.1\}$	$\{0.0001, 0.0003, \mathbf{0.001}, 0.003, 0.01, 0.03, 0.1\}$
optimizer_inner	{adam, **sgd_nesterov_0.9**}	{adam, **sgd_nesterov_0.9**}
optimizer_outer	adam	adam
steps_cg	-	$\{\mathbf{100}, 500, 1000, 2000\}$
steps_inner	$\{2000, 3000, \mathbf{5000}\}$	$\{\mathbf{2000}, 3000, 5000\}$
steps_nudged	$\{100, \mathbf{200}, 500\}$	-
steps_outer	100	100

Hyperparameter	NSA	T1T2
batch_size	500	500
λ	$\{\mathbf{10^{-5}}, 10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}\}$	$\{10^{-5}, \mathbf{10^{-4}}, 10^{-3}, 10^{-2}, 10^{-1}\}$
lr_inner	$\{0.0001, 0.0003, \mathbf{0.001}, 0.003, 0.01, 0.03, 0.1\}$	$\{0.0001, 0.0003, \mathbf{0.001}, 0.003, 0.01, 0.03, 0.1\}$
lr_outer	$\{0.0001, 0.0003, 0.001, \mathbf{0.003}, 0.01, 0.03, 0.1\}$	$\{0.0001, 0.0003, \mathbf{0.001}, 0.003, 0.01, 0.03, 0.1\}$
nsa_alpha	$\{0.000001, 0.000003, 0.00001, 0.00003, 0.0001, \mathbf{0.0003}\}$	-
optimizer_inner	{adam, **sgd_nesterov_0.9**}	{adam, **sgd_nesterov_0.9**}
optimizer_outer	adam	adam
steps_inner	$\{\mathbf{2000}, 3000, 5000\}$	$\{2000, \mathbf{3000}, 5000\}$
steps_nsa	500	-
steps_outer	100	100

Hyperparameter	TBPTL	no-meta
batch_size	500	500
λ	$\{10^{-5}, \mathbf{10^{-4}}, 10^{-3}, 10^{-2}, 10^{-1}\}$	$\{10^{-5}, 10^{-4}, 10^{-3}, \mathbf{10^{-2}}, 10^{-1}\}$
lr_inner	$\{0.0001, 0.0003, \mathbf{0.001}, 0.003, 0.01, 0.03, 0.1\}$	$\{0.0001, 0.0003, \mathbf{0.001}, 0.003, 0.01, 0.03, 0.1\}$
lr_outer	$\{0.0001, 0.0003, \mathbf{0.001}, 0.003, 0.01, 0.03, 0.1\}$	-
optimizer_inner	{adam, **sgd_nesterov_0.9**}	{adam, **sgd_nesterov_0.9**}
optimizer_outer	adam	-
steps_inner	{2000, **3000**, 5000}	5000
steps_outer	100	0

Table S6: Hyperparameter search space for the supervised meta-optimization experiment on MNIST. For all methods 500 samples were randomly drawn from the search space and Asynchronous HyperBand from ray tune [120] was used for scheduling with a grace period of 10. Best found parameters are marked in bold.

Hyperparameter	CML	CG
`batch_size`	500	500
β	$\{0.01, 0.03, 0.1, 0.3, 1.0, \mathbf{3.0}, 10.0\}$	-
λ	$\{\mathbf{0.00001}, 0.0001, 0.001, 0.01, 0.1\}$	$\{0.00001, \mathbf{0.0001}, 0.001, 0.01, 0.1\}$
`lr_inner`	$\{0.0001, \mathbf{0.0003}, 0.001, 0.003, 0.01, 0.03, 0.1\}$	$\{\mathbf{0.0001}, 0.0003, 0.001, 0.003, 0.01, 0.03, 0.1\}$
`lr_nudged`	$\{0.0001, \mathbf{0.0003}, 0.001, 0.003, 0.01, 0.03, 0.1\}$	-
`lr_outer`	$\{\mathbf{0.0001}, 0.0003, 0.001, 0.003, 0.01, 0.03, 0.1\}$	$\{0.0001, \mathbf{0.0003}, 0.001, 0.003, 0.01, 0.03, 0.1\}$
`optimizer_inner`	{**adam**, sgd_nesterov_0.9}	{**adam**, sgd_nesterov_0.9}
`optimizer_outer`	adam	adam
`steps_cg`	-	$\{\mathbf{100}, 250, 500, 1000, 2000\}$
`steps_inner`	2000	2000
`steps_nudged`	$\{100, \mathbf{200}, 500\}$	-
`steps_outer`	100	100

Hyperparameter	NSA	T1T2
`batch_size`	500	500
λ	$\{0.00001, \mathbf{0.0001}, 0.001, 0.01, 0.1\}$	$\{\mathbf{0.00001}, 0.0001, 0.001, 0.01, 0.1\}$
`lr_inner`	$\{0.0001, 0.0003, \mathbf{0.001}, 0.003, 0.01, 0.03, 0.1\}$	$\{0.0001, 0.0003, 0.001, 0.003, \mathbf{0.01}, 0.03, 0.1\}$
`lr_outer`	$\{0.0001, 0.0003, 0.001, 0.003, \mathbf{0.01}, 0.03, 0.1\}$	$\{\mathbf{0.0001}, 0.0003, 0.001, 0.003, 0.01, 0.03, 0.1\}$
`nsa_alpha`	$\{0.000001, 0.000003, 0.00001, 0.00003, 0.0001, 0.0003, \mathbf{0.001}, 0.003\}$	-
`optimizer_inner`	{adam, **sgd_nesterov_0.9**}	{adam, **sgd_nesterov_0.9**}
`optimizer_outer`	adam	adam
`steps_inner`	2000	2000
`steps_nsa`	200	-
`steps_outer`	100	100

Table S7: Hyperparameter search space for the few-shot image classification experiments on Omniglot and miniImageNet. Best found parameters are marked in bold.

Hyperparameter	Omni-5W-1s	Omni-5W-5s	Omni-20W-1s	Omni-20W-5s	miniImageNet
`batch_size`	32	32	16	16	4
β	{0.01, 0.03,**0.1**, 0.3, 1.}	{0.01, 0.03,**0.1**, 0.3, 1.}	{0.01, **0.03**, 0.1, 0.3, 1.}	{0.01, 0.03, **0.1**, 0.3, 1.}	{**0.01**, 0.03, 0.1, 0.3, 1.}
λ	{0.1, 0.25,**0.5**}	{**0.1**, 0.25, 0.5}	{0.1, 0.25, **0.5**}	{0.1, **0.25**, 0.5}	{0.1, 0.25, **0.5**}
`lr_inner`	0.01	0.01	0.01	0.01	0.01
`lr_outer`	{**0.01**, 0.001 }	{**0.01**, 0.001 }	{**0.01**, 0.001 }	{**0.01**, 0.001 }	0.001
`optimizer_inner`	gd_nesterov_0.9	gd_nesterov_0.9	gd_nesterov_0.9	gd_nesterov_0.9	gd_nesterov_0.9
`optimizer_outer`	adam	adam	adam	adam	adam
`steps_inner`	{50, 100, 150, **200**}	{50, 100, 150, **200**}	{50, 100, 150, **200**}	{50, 100, 150, **200**}	{50, **100**, 150, 200}
`steps_nudged`	{50, **100**, 150, 200}	{50, **100**, 150, 200}	{50, **100**, 150, 200}	{50, 100, 150, **200**}	{25, 50, **75**, 100}
`steps_outer`	3750	3750	3750	3750	25000

Table S8: Hyperparameter search space for the sinusoidal fewshot regression experiment. For all methods 500 samples were randomly drawn from the search space and the Asynchronous HyperBand scheduler from ray tune was used with a grace period of 10 [120]. Best found parameters are marked in bold.

Hyperparameter	CML + e-prop
`activity_reg_strength`	$\{10^{-1}, 10^{-2}, 10^{-3}, 10^{-4}, \mathbf{10^{-5}}, 10^{-6}\}$
`activity_reg_target`	$\{0.05, 0.1, \mathbf{0.2}\}$
`batch_size`	$\{1, 5, \mathbf{10}\}$
β	$\{0.01, 0.03, 0.1, 0.3, 1.0, \mathbf{3.0}, 10.0\}$
λ	$\{10^{0}, 10^{-1}, \mathbf{10^{-2}}, 10^{-3}, 10^{-4}, 10^{-5}, 10^{-6}\}$
`lr_inner`	$\{0.0001, 0.0003, \mathbf{0.001}, 0.003, 0.01, 0.03, 0.1\}$
`lr_nudged`	$\{0.0001, 0.0003, 0.001, \mathbf{0.003}, 0.01, 0.03, 0.1\}$
`lr_outer`	$\{0.0001, 0.0003, 0.001, \mathbf{0.003}, 0.01, 0.03, 0.1\}$
`meta_batch_size`	$\{1, 10, \mathbf{25}\}$
`optimizer_inner`	{**adam**, sgd_nesterov_0.9}
`optimizer_outer`	adam
`steps_inner`	500
`steps_nudged`	$\{50, \mathbf{100}, 200\}$
`steps_outer`	1000

Hyperparameter	BPTT + e-prop	BPTT + BPTT	TBPTL + e-prop
`activity_reg_strength`	$\{10^{-1}, \mathbf{10^{-2}}, 10^{-3}, 10^{-4}, 10^{-5}, 10^{-6}\}$	$\{10^{-1}, 10^{-2}, 10^{-3}, 10^{-4}, 10^{-5}, \mathbf{10^{-6}}\}$	$\{10^{-1}, 10^{-2}, 10^{-3}, 10^{-4}, 10^{-5}, \mathbf{10^{-6}}\}$
`activity_reg_target`	$\{\mathbf{0.05}, 0.1, 0.2\}$	$\{0.05, \mathbf{0.1}, 0.2\}$	$\{0.05, \mathbf{0.1}, 0.2\}$
`batch_size`	$\{\mathbf{1}, 5, 10\}$	$\{1, 5, \mathbf{10}\}$	$\{\mathbf{1}, 5, 10\}$
`lr_inner`	$\{\mathbf{0.0001}, 0.0003, 0.001, 0.003, 0.01, 0.03, 0.1\}$	$\{0.0001, \mathbf{0.0003}, 0.001, 0.003, 0.01, 0.03, 0.1\}$	$\{\mathbf{0.0001}, 0.0003, 0.001, 0.003, 0.01, 0.03, 0.1\}$
`lr_outer`	$\{0.0001, \mathbf{0.0003}, 0.001, 0.003, 0.01, 0.03, 0.1\}$	$\{0.0001, \mathbf{0.0003}, 0.001, 0.003, 0.01, 0.03, 0.1\}$	$\{0.0001, 0.0003, \mathbf{0.001}, 0.003, 0.01, 0.03, 0.1\}$
`meta_batch_size`	$\{1, \mathbf{10}, 25\}$	$\{1, \mathbf{10}, 25\}$	$\{\mathbf{1}, 10, 25\}$
`optimizer_inner`	{adam, **sgd_nesterov_0.9**}	{adam, **sgd_nesterov_0.9**}	sgd
`optimizer_outer`	adam	adam	adam
`steps_inner`	10	10	500
`steps_outer`	1000	1000	1000

Table S9: Hyperparameter search space for the wheel bandit experiment. For all methods 1000 samples were randomly drawn from the search space. Best found parameters are marked in bold.

Hyperparameter	CML (synaptic)	CML (modulatory)	MAML
`batch_size`	512	512	512
β	$\{0.01, 0.03, 0.1, \mathbf{0.3}, 1.0, 3.0, 10.0\}$	$\{0.01, 0.03, 0.1, 0.3, 1.0, 3.0, \mathbf{10.0}\}$	-
λ	$\{10^{-6}, 10^{-5}, \ldots, \mathbf{10^{3}}\}$	-	-
`lr_inner`	$\{\mathbf{0.0001}, 0.0003, 0.001, 0.003, 0.01, 0.03, 0.1\}$	$\{0.0001, \mathbf{0.0003}, 0.001, 0.003, 0.01, 0.03, 0.1\}$	$\{0.0001, 0.0003, 0.001, 0.003, \mathbf{0.01}, 0.03, 0.1\}$
`lr_nudged`	$\{0.0001, 0.0003, 0.001, 0.003, 0.01, \mathbf{0.03}, 0.1\}$	$\{\mathbf{0.0001}, 0.0003, 0.001, 0.003, 0.01, 0.03, 0.1\}$	-
`lr_online`	$\{\mathbf{0.0001}, 0.0003, 0.001, 0.003, 0.01, 0.03, 0.1\}$	$\{\mathbf{0.0001}, 0.0003, 0.001, 0.003, 0.01, 0.03, 0.1\}$	$\{0.0001, 0.0003, \mathbf{0.001}, 0.003, 0.01, 0.03, 0.1\}$
`lr_outer`	$\{0.0001, 0.0003, 0.001, 0.003, 0.01, \mathbf{0.03}, 0.1\}$	$\{0.0001, 0.0003, 0.001, 0.003, 0.01, \mathbf{0.03}, 0.1\}$	$\{0.0001, 0.0003, 0.001, 0.003, 0.01, 0.03, \mathbf{0.1}\}$
`meta_batch_size`	$\{\mathbf{8}, 16, 32\}$	$\{8, \mathbf{16}, 32\}$	$\{8, 16, \mathbf{32}\}$
`optimizer_inner`	adam	{adam, **sgd**, sgd_nesterov_0.9}	sgd
`optimizer_online`	adam	{adam, sgd, **sgd_nesterov_0.9**}	sgd
`optimizer_outer`	adam	{**adam**, adamw}	adam
`steps_inner`	$\{100, \mathbf{250}, 500, 1000\}$	$\{100, 250, 500, \mathbf{1000}\}$	$\{5, \mathbf{10}, 50, 100\}$
`steps_nudged`	$\{\mathbf{100}, 250, 500, 1000\}$	$\{\mathbf{100}, 250, 500, 1000\}$	-
`steps_outer`	6400	6400	6400
t_f	$\{20, \mathbf{50}, 100\}$	$\{20, 50, \mathbf{100}\}$	$\{20, 50, \mathbf{100}\}$
t_s	$\{50, 100, \mathbf{250}, 500, 1000\}$	$\{50, 100, 250, \mathbf{500}, 1000\}$	$\{5, 10, 50, \mathbf{100}\}$

S5 Additional details

S5.1 Compute resources

We used Linux workstations with 2 Nvidia RTX 3090 and 4 Nvidia RTX 3070 GPUs during development and conducted hyperparameter searches and larger experiments on up to 3 Linux servers with 8 Nvidia RTX 3090 GPUs with 24 GB memory each. Most of the experiments and corresponding hyperparameter scans presented take less than a few hours to complete on a single server. The more challenging recurrent spiking network and miniImageNet experiments require approximately 2-5 days to complete. During development we conducted many more hyperparameter scans over the course of several months.

S5.2 Software and libraries

For the results produced in this paper we relied on free and open-source software. We implemented our experiments in Python using PyTorch [121, BSD-style license], JAX [122, Apache License 2.0], Ray [120, Apache License 2.0] and NumPy [123, BSD-style license]. For the visual few-shot classification dataset splits we used the Torchmeta library [124, MIT license] and for the generation of plots we used matplotlib [125, BSD-style license].

S5.3 Datasets

We conducted our experiments with the public domain datasets Boston housing [98, MIT License], MNIST [100, GNU GPL v3.0], Omniglot [57] (MIT license), miniImageNet [56] (custom MIT/ImageNet license) and CIFAR-10 (MIT license) [48].

Supplementary References

[29] Benjamin Scellier and Yoshua Bengio. Equilibrium propagation: bridging the gap between energy-based models and backpropagation. *Frontiers in Computational Neuroscience*, 11, 2017.

[30] Benjamin Scellier. *A deep learning theory for neural networks grounded in physics*. PhD Thesis, Université de Montréal, 2021.

[24] Asen L. Dontchev and R. Tyrrell Rockafellar. *Implicit Functions and Solution Mappings*. Springer, NY, 2009.

[87] Axel Laborieux, Maxence Ernoult, Benjamin Scellier, Yoshua Bengio, Julie Grollier, and Damien Querlioz. Scaling equilibrium propagation to deep convnets by drastically reducing its gradient estimator bias. *Frontiers in Neuroscience*, 14, 2021.

[53] Jelena Luketina, Mathias Berglund, Klaus Greff, and Tapani Raiko. Scalable gradient-based tuning of continuous regularization hyperparameters. In *International Conference on Machine Learning*, 2016.

[51] Renjie Liao, Yuwen Xiong, Ethan Fetaya, Lisa Zhang, KiJung Yoon, Xaq Pitkow, Raquel Urtasun, and Richard Zemel. Reviving and improving recurrent back-propagation. In *International Conference on Machine Learning*, 2018.

[23] Jonathan Lorraine, Paul Vicol, and David Duvenaud. Optimizing millions of hyperparameters by implicit differentiation. In *International Conference on Artificial Intelligence and Statistics*, 2020.

[91] C. Goutte and J. Larsen. Adaptive regularization of neural networks using conjugate gradient. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*, 1998.

[22] Aravind Rajeswaran, Chelsea Finn, Sham Kakade, and Sergey Levine. Meta-learning with implicit gradients. In *Advances in Neural Information Processing Systems*, 2019.

[54] Amirreza Shaban, Ching-An Cheng, Nathan Hatch, and Byron Boots. Truncated back-propagation for bilevel optimization. In *International Conference on Artificial Intelligence and Statistics*, 2019.

[94] David J. C. MacKay. A practical Bayesian framework for backpropagation networks. *Neural Computation*, 4(3), 1992.

[20] Yoshua Bengio. Gradient-based optimization of hyperparameters. *Neural Computation*, 12(8): 1889–1900, 2000.

[52] Chuan-sheng Foo, Chuong B. Do, and Andrew Y. Ng. Efficient multiple hyperparameter learning for log-linear models. In *Advances in Neural Information Processing Systems*, 2007.

[21] Fabian Pedregosa. Hyperparameter optimization with approximate gradient. In *International Conference on Machine Learning*, 2016.

[98] David Harrison Jr. and Daniel L. Rubinfeld. Hedonic housing prices and the demand for clean air. *Journal of Environmental Economics and Management*, 5(1):81–102, 1978.

[48] Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, 2009.

[100] Yann LeCun. The MNIST database of handwritten digits. *Available at http://yann. lecun. com/exdb/mnist*, 1998.

[101] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.

[102] Sergey Ioffe and Christian Szegedy. Batch normalization: accelerating deep network training by reducing internal covariate shift. In *International Conference on Machine Learning*, 2015.

[103] Adam Santoro, Sergey Bartunov, Matthew Botvinick, Daan Wierstra, and Timothy Lillicrap. Meta-learning with memory-augmented neural networks. In *International Conference on Machine Learning*, 2016.

[104] Oriol Vinyals, Charles Blundell, Timothy P. Lillicrap, Koray Kavukcuoglu, and Daan Wierstra. Matching networks for one shot learning. In *Advances in Neural Information Processing Systems*, 2016.

[18] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *International Conference on Machine Learning*, 2017.

[57] Brenden M. Lake, Ruslan Salakhutdinov, Jason Gross, and Joshua B. Tenenbaum. One shot learning of simple visual concepts. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, 2011.

[56] Sachin Ravi and Hugo Larochelle. Optimization as a model for few-shot learning. In *International Conference on Learning Representations*, 2016.

[58] Alex Nichol, Joshua Achiam, and John Schulman. On first-order meta-learning algorithms. *arXiv preprint arXiv:1803.02999*, 2018.

[109] Antreas Antoniou, Harrison Edwards, and Amos Storkey. How to train your MAML. In *International Conference on Learning Representations*, 2019.

[110] David Ha, Andrew Dai, and Quoc V. Le. HyperNetworks. In *International Conference on Learning Representations*, 2017.

[111] Andrei A. Rusu, Dushyant Rao, Jakub Sygnowski, Oriol Vinyals, Razvan Pascanu, Simon Osindero, and Raia Hadsell. Meta-learning with latent embedding optimization. In *International Conference on Learning Representations*, 2019.

[112] Dominic Zhao, Seijin Kobayashi, João Sacramento, and Johannes von Oswald. Meta-learning via hypernetworks. In *Workshop on Meta-Learning at NeurIPS*, 2020.

[113] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: surpassing human-level performance on ImageNet classification. In *Proceedings of the IEEE International Conference on Computer Vision*, 2015.

[114] Boris T. Polyak and Anatoli B. Juditsky. Acceleration of stochastic approximation by averaging. *SIAM Journal on Control and Optimization*, 30(4):838–855, 1992.

[59] Guillaume Bellec, Franz Scherr, Anand Subramoney, Elias Hajek, Darjan Salaj, Robert Legenstein, and Wolfgang Maass. A solution to the learning dilemma for recurrent networks of spiking neurons. *Nature Communications*, 11(1):3625, 2020.

[60] Emre O. Neftci, Hesham Mostafa, and Friedemann Zenke. Surrogate gradient learning in spiking neural networks: bringing the power of gradient-based optimization to spiking neural networks. *IEEE Signal Processing Magazine*, 36(6):51–63, 2019.

[65] Marta Garnelo, Jonathan Schwarz, Dan Rosenbaum, Fabio Viola, Danilo J. Rezende, S. M. Ali Eslami, and Yee Whye Teh. Neural processes. *arXiv preprint arXiv:1807.01622*, 2018.

[66] Sachin Ravi and Alex Beatson. Amortized Bayesian meta-learning. In *International Conference on Learning Representations*, 2019.

[64] Carlos Riquelme, George Tucker, and Jasper Snoek. Deep Bayesian bandits showdown: an empirical comparison of Bayesian deep networks for Thompson sampling. In *International Conference on Learning Representations*, 2018.

[120] Richard Liaw, Eric Liang, Robert Nishihara, Philipp Moritz, Joseph E. Gonzalez, and Ion Stoica. Tune: A research platform for distributed model selection and training. *arXiv preprint arXiv:1807.05118*, 2018.

[121] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. PyTorch: an imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*, 2019.

[122] James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. JAX: composable transformations of Python+NumPy programs, 2018. URL `http://github.com/google/jax`.

[123] Charles R. Harris, K. Jarrod Millman, Stéfan J. van der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J. Smith, Robert Kern, Matti Picus, Stephan Hoyer, Marten H. van Kerkwijk, Matthew Brett, Allan Haldane, Jaime Fernández del Río, Mark Wiebe, Pearu Peterson, Pierre Gérard-Marchant, Kevin Sheppard, Tyler Reddy, Warren Weckesser, Hameer Abbasi, Christoph Gohlke, and Travis E. Oliphant. Array programming with NumPy. *Nature*, 585(7825):357–362, 2020.

[124] Tristan Deleu, Tobias Würfl, Mandana Samiei, Joseph Paul Cohen, and Yoshua Bengio. Torchmeta: A meta-learning library for PyTorch. *arXiv preprint arXiv:1909.06576*, 2019.

[125] J. D. Hunter. Matplotlib: A 2D graphics environment. *Computing in Science & Engineering*, 9(3):90–95, 2007.