CONTINUOUS-TIME META-LEARNING WITH FORWARD MODE DIFFERENTIATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Drawing inspiration from gradient-based meta-learning methods with infinitely small gradient steps, we introduce Continuous-Time Meta-Learning (COMLN), a meta-learning algorithm where adaptation follows the dynamics of a gradient vector field. Specifically, representations of the inputs are meta-learned such that a task-specific linear classifier is obtained as a solution of an ordinary differential equation (ODE). Treating the learning process as an ODE offers the notable advantage that the length of the trajectory is now continuous, as opposed to a fixed and discrete number of gradient steps. As a consequence, we can optimize the amount of adaptation necessary to solve a new task using stochastic gradient descent, in addition to learning the initial conditions as is standard practice in gradient-based meta-learning. Importantly, in order to compute the exact meta-gradients required for the outer-loop updates, we devise an efficient algorithm based on forward mode differentiation, whose memory requirements do not scale with the length of the learning trajectory, thus allowing longer adaptation in constant memory. We provide analytical guarantees for the stability of COMLN, we show empirically its efficiency in terms of runtime and memory usage, and we illustrate its effectiveness on a range of few-shot image classification problems.

1 INTRODUCTION

Among the existing meta-learning algorithms, gradient-based methods as popularized by Model-Agnostic Meta-Learning (MAML, Finn et al., 2017) have received a lot of attention over the past few years. They formulate the problem of learning a new task as an inner optimization problem, typically based on a few steps of gradient descent. An outer *meta*-optimization problem is then responsible for updating the meta-parameters of this learning process, such as the initialization of the gradient descent procedure. However since the updates at the outer level typically require backpropagating through the learning process, this class of methods has often been limited to only a few gradient steps of adaptation, due to memory constraints. Although solutions have been proposed to alleviate the memory requirements of these algorithms, including checkpointing (Baranchuk, 2019), using implicit differentiation (Rajeswaran et al., 2019), or reformulating the meta-learning objective (Flennerhag et al., 2018), they are generally either more computationally demanding, or only approximate the gradients of the meta-learning objective (Nichol et al., 2018; Flennerhag et al., 2020).

In this work, we propose a continuous-time formulation of gradient-based meta-learning, called *Continuous-Time Meta-Learning* (COMLN), where the adaptation is the solution of a differential equation (see Figure 1). Moving to continuous time allows us to devise a novel algorithm, based on forward mode differentiation, to efficiently compute the exact gradients for meta-optimization, no matter how long the adaptation to a new task might be. We show that using forward mode differentiation leads to a stable algorithm, unlike the counterpart of backpropagation in continuous time called the adjoint method (frequently used in the Neural ODE literature) which tends to be unstable with gradient vector fields. Moreover as the length of the adaptation trajectory is a continuous quantity, as opposed to a discrete number of gradient steps fixed ahead of time, we can treat the amount of adaptation in COMLN as a meta-parameter—on par with the initialization—which we can meta-optimize using stochastic gradient descent. We verify empirically that our method is both computationally and memory efficient, and we show that COMLN outperforms other standard meta-learning algorithms on few-shot image classification datasets.

(a) Gradient-based Meta-Learning

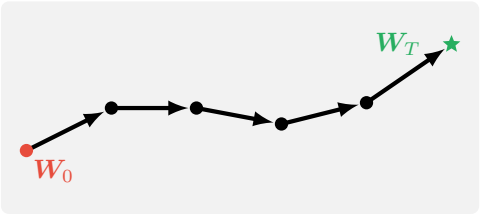

$$\boldsymbol{W}_{t+1} = \boldsymbol{W}_t - \alpha \nabla \mathcal{L}(\boldsymbol{W}_t)$$

(b) COMLN

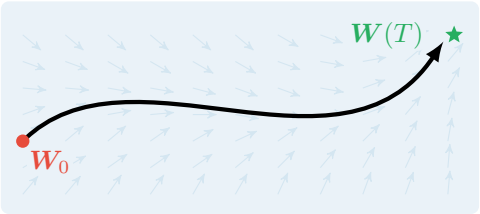

$$\frac{d\boldsymbol{W}}{dt} = -\nabla \mathcal{L}\big(\boldsymbol{W}(t)\big)$$

Figure 1: Illustration of the adaptation process in (a) a gradient-based meta-learning algorithm, such as ANIL (Raghu et al., 2019), where the adapted parameters $\boldsymbol{W}_T$ are given after T steps of gradient descent, and in (b) Continuous-Time Meta-Learning (COMLN), where the adapted parameters $\boldsymbol{W}(T)$ are the result of following the dynamics of a gradient vector field up to time T.

2 Background

In this work, we consider the problem of few-shot classification, that is the problem of learning a classification model with only a small number of training examples. More precisely for a classification task τ, we assume that we have access to a (small) training dataset $\mathcal{D}_\tau^{\text{train}} = \{(\boldsymbol{x}_m, \boldsymbol{y}_m)\}_{m=1}^M$ to fit a model on task τ, and a distinct test dataset $\mathcal{D}_\tau^{\text{test}}$ to evaluate how well this adapted model generalizes on that task. In the few-shot learning literature, it is standard to consider the problem of k-shot N-way classification, meaning that the model has to classify among N possible classes, and there are only k examples of each class in $\mathcal{D}_\tau^{\text{train}}$, so that overall the number of training examples is $M = kN$. We use the convention that the target labels $\boldsymbol{y}_m \in \{0,1\}^N$ are one-hot vectors.

2.1 Gradient-based meta-learning

Gradient-based meta-learning methods aim to learn an initialization such that the model is able to adapt to a new task via gradient descent. Such methods are often cast as a bi-level optimization process: adapting the task-specific parameters $\boldsymbol{\theta}$ in the inner loop, and training the (task-agnostic) meta-parameters Φ and initialization $\boldsymbol{\theta}_0$ in the outer loop. The meta-learning objective is:

$$\min_{\boldsymbol{\theta}_0, \Phi} \; \mathbb{E}_\tau \big[\mathcal{L}(\boldsymbol{\theta}_T^\tau, \Phi; \mathcal{D}_\tau^{\text{test}}) \big] \tag{1}$$

$$\text{s.t.} \;\; \boldsymbol{\theta}_{t+1}^\tau = \boldsymbol{\theta}_t^\tau - \alpha \nabla_{\boldsymbol{\theta}} \mathcal{L}(\boldsymbol{\theta}_t^\tau, \Phi; \mathcal{D}_\tau^{\text{train}}) \qquad \boldsymbol{\theta}_0^\tau = \boldsymbol{\theta}_0 \qquad \forall \tau \sim p(\tau), \tag{2}$$

where T is the number of inner loop updates. For example, in the case of MAML (Finn et al., 2017), there is no additional meta-parameter other than the initialization ($\Phi \equiv \emptyset$); in ANIL (Raghu et al., 2019), $\boldsymbol{\theta}$ are the parameters of the last layer, and Φ are the parameters of the shared embedding network; in CAVIA (Zintgraf et al., 2019), $\boldsymbol{\theta}$ are referred to as context parameters.

During meta-training, the model is trained over many tasks τ. The task-specific parameters $\boldsymbol{\theta}$ are learned via gradient descent on $\mathcal{D}_\tau^{\text{train}}$. The meta-parameters are then updated by evaluating the error of the trained model on the test dataset $\mathcal{D}_\tau^{\text{test}}$. At meta-testing time, the meta-trained model is adapted on $\mathcal{D}_\tau^{\text{train}}$—i.e. applying Equation 2 with the learned meta-parameters $\boldsymbol{\theta}_0$ and Φ.

2.2 Local sensitivity analysis of Ordinary Differential Equations

Consider the following (autonomous) Ordinary Differential Equation (ODE):

$$\frac{d\boldsymbol{z}}{dt} = g\big(\boldsymbol{z}(t); \boldsymbol{\theta}\big) \qquad \boldsymbol{z}(0) = \boldsymbol{z}_0(\boldsymbol{\theta}), \tag{3}$$

where the dynamics g and initial value $\boldsymbol{z}_0$ may depend on some external parameters $\boldsymbol{\theta}$, and integration is carried out from 0 to some time T. *Local sensitivity analysis* is the study of how the solution of this dynamical system responds to local changes in $\boldsymbol{\theta}$; this effectively corresponds to calculating the derivative $d\boldsymbol{z}(t)/d\boldsymbol{\theta}$. We present here two methods to compute this derivative, with a special focus on their memory efficiency.

Adjoint sensitivity method Based on the adjoint state (Pontryagin, 2018), and taking its root in control theory (Lions & Magenes, 2012), the *adjoint sensitivity method* (Bryson & Ho, 1969; Chavent et al., 1974) provides an efficient approach for evaluating derivatives of $\mathcal{L}\big(\boldsymbol{z}(T);\boldsymbol{\theta}\big)$, a function of $\boldsymbol{z}(T)$ the solution of the ODE in Equation 3. This method, popularized lately by the literature on Neural ODEs (Chen et al., 2018), requires the integration of the adjoint equation

$$\frac{d\boldsymbol{a}}{dt} = -\boldsymbol{a}(t)\frac{\partial g\big(\boldsymbol{z}(t);\boldsymbol{\theta}\big)}{\partial \boldsymbol{z}(t)} \qquad\qquad \boldsymbol{a}(T) = \frac{d\mathcal{L}\big(\boldsymbol{z}(T);\boldsymbol{\theta}\big)}{d\boldsymbol{z}(T)}, \tag{4}$$

backward in time. The adjoint sensitivity method can be viewed as a continuous-time counterpart to backpropagation, where the forward pass would correspond to integrating Equation 3 forward in time from 0 to T, and the backward pass to integrating Equation 4 backward in time from T to 0.

One possible implementation, reminiscent of backpropagation through time (BPTT), is to store the intermediate values of $\boldsymbol{z}(t)$ during the forward pass, and reuse them to compute the adjoint state during the backward pass. While several sophisticated checkpointing schemes have been proposed (Serban & Hindmarsh, 2003; Gholami et al., 2019), with different compute/memory trade-offs, the memory requirements of this approach typically grow with T; this is similar to the memory limitations standard gradient-based meta-learning methods suffer from as the number of gradient steps increases. An alternative is to augment the adjoint state $\boldsymbol{a}(t)$ with the original state $\boldsymbol{z}(t)$, and to solve this augmented dynamical system backward in time (Chen et al., 2018). This has the notable advantage that the memory requirements are now independent of T, since $\boldsymbol{z}(t)$ are no longer stored during the forward pass, but they are recomputed on the fly during the backward pass.

Forward sensitivity method While the adjoint method is related to reverse-mode automatic differentiation (backpropagation), the *forward sensitivity method* (Feehery et al., 1997; Leis & Kramer, 1988; Maly & Petzold, 1996; Caracotsios & Stewart, 1985), on the other hand, can be viewed as the continuous-time counterpart to forward (tangent-linear) mode differentiation (Griewank & Walther, 2008). This method is based on the fact that the derivative $\mathcal{S}(t) \triangleq d\boldsymbol{z}(t)/d\boldsymbol{\theta}$ is the solution of the so-called *forward sensitivity equation*

$$\frac{d\mathcal{S}}{dt} = \frac{\partial g\big(\boldsymbol{z}(t);\boldsymbol{\theta}\big)}{\partial \boldsymbol{z}(t)}\mathcal{S}(t) + \frac{\partial g\big(\boldsymbol{z}(t);\boldsymbol{\theta}\big)}{\partial \boldsymbol{\theta}} \qquad\qquad \mathcal{S}(0) = \frac{\partial \boldsymbol{z}_0}{\partial \boldsymbol{\theta}}. \tag{5}$$

This equation can be found throughout the literature in optimal control and system identification (Betts, 2010; Biegler, 2010). Unlike the adjoint method, which requires an explicit forward *and* backward pass, the forward sensitivity method only requires the integration forward in time of the original ODE in Equation 3, augmented by the sensitivity state $\mathcal{S}(t)$ with the dynamics above. The memory requirements of the forward sensitivity method do not scale with T either, but it now requires storing $\mathcal{S}(t)$, which may be very large; we will come back to this problem in Section 3.2. We will simply note here that in discrete-time, this is the same issue afflicting forward-mode training of RNNs with real-time recurrent learning (RTRL; Williams & Zipser, 1989), or other meta-learning algorithms (Sutton, 1992; Franceschi et al., 2017; Xu et al., 2018).

3 Continuous-time adaptation

In the limit of infinitely small steps, some optimization algorithms can be viewed as the solution trajectory of a differential equation. This point of view has often been taken to analyze their behavior (Platt & Barr, 1988; Wilson et al., 2016; Su et al., 2014; Orvieto & Lucchi, 2019). In fact, some optimization algorithms such as gradient descent with momentum have even been introduced initially from the perspective of dynamical systems (Polyak, 1964). As the simplest example, gradient descent with a constant step size $\alpha \to 0^+$ (i.e. α tends to 0 by positive values) corresponds to following the dynamics of an autonomous ODE called a *gradient vector field*

$$\boldsymbol{z}_{t+1} = \boldsymbol{z}_t - \alpha\nabla f(\boldsymbol{z}_t) \qquad \underset{\alpha\to 0^+}{\longrightarrow} \qquad \frac{d\boldsymbol{z}}{dt} = -\nabla f\big(\boldsymbol{z}(t)\big), \tag{6}$$

where the iterate $\boldsymbol{z}(t)$ is now a continuous function of time t. The solution of this dynamical system is uniquely defined by the choice of the initial condition $\boldsymbol{z}(0) = \boldsymbol{z}_0$.

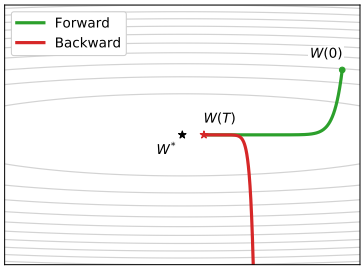

Figure 2: Numerical instability of the adjoint method applied to the gradient vector field of a quadratic loss function. The trajectory in green starting at $\boldsymbol{W}(0)$ corresponds to the integration of the dynamical system in Equation 8 forward in time up to T, and the trajectory in red starting at $\boldsymbol{W}(T)$ corresponds to its integration backward in time. Note that T was chosen so that $\boldsymbol{W}(T)$ does not reach the equilibrium/minimum of the loss $\boldsymbol{W}^\star$.

3.1 Continuous-time meta-learning

In gradient-based meta-learning, the task-specific adaptation with gradient descent may also be replaced by a gradient vector field in the limit of infinitely small steps. Inspired by prior work in meta-learning (Raghu et al., 2019; Javed & White, 2019), we assume that an embedding network f_Φ with meta-parameters Φ is shared across tasks, and only the parameters $\boldsymbol{W}$ of a task-specific linear classifier are adapted, starting at some initialization $\boldsymbol{W}_0$. Instead of being the result of a few steps of gradient descent though, the final parameters $\boldsymbol{W}(T)$ now correspond to integrating an ODE similar to Equation 6 up to a certain horizon T, with the initial conditions $\boldsymbol{W}(0) = \boldsymbol{W}_0$. We call this new meta-learning algorithm Continuous-Time Meta-Learning[1] (COMLN).

Treating the learning algorithm as a continuous-time process has the notable advantage that the adapted parameter $\boldsymbol{W}(T)$ is now differentiable wrt. the time horizon T (Wiggins, 2003, Chap. 7), in addition to being differentiable wrt. the initial conditions $\boldsymbol{W}_0$—which plays a central role in gradient-based meta-learning, as described in Section 2.1. This allows us to view T as a meta-parameter on par with Φ and $\boldsymbol{W}_0$, and to effectively optimize the amount of adaptation using stochastic gradient descent (SGD). The meta-learning objective of COMLN can be written as

$$\min_{\Phi, \boldsymbol{W}_0, T} \; \mathbb{E}_\tau \big[\mathcal{L}\big(\boldsymbol{W}_\tau(T); f_\Phi(\mathcal{D}_\tau^{\text{test}})\big) \big] \tag{7}$$

$$\text{s.t.} \;\; \frac{d\boldsymbol{W}_\tau}{dt} = -\nabla \mathcal{L}\big(\boldsymbol{W}_\tau(t); f_\Phi(\mathcal{D}_\tau^{\text{train}})\big) \qquad \boldsymbol{W}_\tau(0) = \boldsymbol{W}_0 \qquad \forall \tau \sim p(\tau), \tag{8}$$

where $f_\Phi(\mathcal{D}_\tau^{\text{train}}) = \{(f_\Phi(\boldsymbol{x}_m), \boldsymbol{y}_m) \mid (\boldsymbol{x}_m, \boldsymbol{y}_m) \in \mathcal{D}_\tau^{\text{train}}\}$ is the embedded training dataset, and $f_\Phi(\mathcal{D}_\tau^{\text{test}})$ is defined similarly for $\mathcal{D}_\tau^{\text{test}}$. In practice, adaptation is implemented using a numerical integration scheme based on an iterative discretization of the problem, such as Runge-Kutta methods. Although a complete discussion of numerical solvers is outside of the scope of this paper, we recommend (Butcher, 2008) for a comprehensive overview of numerical methods for solving ODEs.

3.2 The challenges of optimizing the meta-learning objective

In order to minimize the meta-learning objective of COMLN, it is common practice to use (stochastic) gradient methods; that requires computing its derivatives wrt. the meta-parameters, which we call *meta-gradients*. Our primary goal is to devise an algorithm whose memory requirements do not scale with the amount of adaptation T; this would contrast with standard gradient-based meta-learning methods that backpropagate through a sequence of gradient steps (similar to BPTT), where the intermediate parameters are stored during adaptation (i.e. $\boldsymbol{\theta}_t^\tau$ for all t in Equation 2). Since this objective involves the solution $\boldsymbol{W}(T)$ of an ODE, we can use either the adjoint method, or the forward sensitivity method, in order to compute the derivatives wrt. Φ and $\boldsymbol{W}_0$ (see Section 2.2).

Although the adjoint method has proven to be an effective strategy for learning Neural ODEs, in practice computing the state $\boldsymbol{W}(t)$ backward in time is numerically unstable when applied to a gradient vector field like the one in Equation 8, even for convex loss functions. Figure 2 shows an example where the trajectory of $\boldsymbol{W}(t)$ recomputed backward in time (in red) diverges significantly from the original trajectory (in green) on a quadratic loss function, even though the two should match exactly in theory since they follow the same dynamics. Intuitively, recomputing $\boldsymbol{W}(t)$ backward in time for a gradient vector field requires doing gradient *ascent* on the loss function, which is prone to compounding numerical errors; this is closely related to the loss of entropy observed by Maclaurin

[1] COMLN is pronounced *chameleon*

et al. (2015). This divergence makes the backward trajectory of $\boldsymbol{W}(t)$ unreliable to find the adjoint state, ruling out the adjoint sensitivity method for computing the meta-gradients in COMLN.

The forward sensitivity method addresses this shortcoming by avoiding the backward pass altogether. However, it can also be particularly expensive here in terms of memory requirements, since the sensitivity state $\mathcal{S}(t)$ in Section 2.2 now corresponds to Jacobian matrices, such as $d\boldsymbol{W}(t)/d\boldsymbol{W}_0$. As the size d of the feature vectors returned by f_Φ may be very large, this $Nd \times Nd$ Jacobian matrix would be almost impossible to store in practice; for example in our experiments, it can be as large as $d = 16{,}000$ for a ResNet-12 backbone. In Section 4.1, we will show how to apply forward sensitivity in a memory-efficient way, by leveraging the structure of the loss function. This is achieved by carefully decomposing the Jacobian matrices into smaller pieces that follow specific dynamics. We show in Appendix D that unlike the adjoint method, this process is stable.

3.3 Connection with Almost No Inner-Loop (ANIL)

Similarly to ANIL (Raghu et al., 2019), COMLN only adapts the parameters $\boldsymbol{W}$ of the last linear layer of the neural network. There is a deeper connection between both algorithms though: while our description of the adaptation in COMLN (Eq. 8) was independent of the choice of the ODE solver used to find the solution $\boldsymbol{W}(T)$ in practice, if we choose an explicit Euler scheme (Euler, 1913, roughly speaking, discretizing Equation 6 from right to left), then the adaptation of COMLN becomes functionally equivalent to ANIL. However, this equivalence can greatly benefit from the memory-efficient algorithm to compute the meta-gradients described in Section 4, based on the forward sensitivity method. This means that using the methods devised here for COMLN, we can effectively compute the meta-gradients of ANIL with a constant memory cost wrt. the number of gradient steps of adaptation, instead of relying on backpropagation (see also Section 4.2).

4 Memory-efficient meta-gradients

For some fixed task τ and $(\boldsymbol{x}_m, \boldsymbol{y}_m) \in \mathcal{D}_\tau^{\text{train}}$, let $\boldsymbol{\phi}_m = f_\Phi(\boldsymbol{x}_m) \in \mathbb{R}^d$ be the embedding of input $\boldsymbol{x}_m$ through the feature extractor f_Φ. Since we are confronted with a classification problem, the loss function of choice $\mathcal{L}(\boldsymbol{W})$ is typically the cross-entropy loss. Böhning (1992) showed that the gradient of the cross-entropy loss wrt. $\boldsymbol{W}$ can be written as

$$\nabla \mathcal{L}\big(\boldsymbol{W}; f_\Phi(\mathcal{D}_\tau^{\text{train}})\big) = \frac{1}{M} \sum_{m=1}^{M} (\boldsymbol{p}_m - \boldsymbol{y}_m)\boldsymbol{\phi}_m^\top, \tag{9}$$

where $\boldsymbol{p}_m = \mathrm{softmax}(\boldsymbol{W}\boldsymbol{\phi}_m)$ is the vector of probabilities returned by the neural network. The key observation here is that the gradient can be decomposed as a sum of M rank-one matrices, where the feature vectors $\boldsymbol{\phi}_m$ are independent of $\boldsymbol{W}$. Therefore we can fully characterize the gradient of the cross-entropy loss with M vectors $\boldsymbol{p}_m - \boldsymbol{y}_m \in \mathbb{R}^N$, as opposed to the full $N \times d$ matrix. This is particularly useful in the context of few-shot classification, where the number of training examples M is small, and typically significantly smaller than the embedding size d.

4.1 Decomposition of the meta-gradients

We saw in Section 3.2 that the forward sensitivity method was the only stable option to compute the meta-gradients of COMLN. However, naively applying the forward sensitivity equation would involve quantities that typically scale with d^2, which can be too expensive in practice. Using the structure of Equation 9, the Jacobian matrices appearing in the computation of the meta-gradients for COMLN can be decomposed in such a way that only small quantities will depend on time.

Meta-gradients wrt. $\boldsymbol{W}_0$ By the chain rule of derivatives, it is sufficient to compute the Jacobian matrix $d\boldsymbol{W}(T)/d\boldsymbol{W}_0$ in order to obtain the meta-gradient wrt. $\boldsymbol{W}_0$. We show in Appendix B.2 that the sensitivity state $d\boldsymbol{W}(t)/d\boldsymbol{W}_0$ can be decomposed as:

$$\frac{d\boldsymbol{W}(t)}{d\boldsymbol{W}_0} = \boldsymbol{I} - \sum_{i=1}^{M}\sum_{j=1}^{M} \boldsymbol{B}_t[i,j] \otimes \boldsymbol{\phi}_i \boldsymbol{\phi}_j^\top, \tag{10}$$

Table 1: Memory required to compute meta-gradients for different algorithms. Exact: the method returns the exact meta-gradients. Full net.: the whole network is adapted, with a number of meta-parameters $|\boldsymbol{\theta}|$. The requirements for checkpointing are taken from (Shaban et al., 2019). Note that typically $M \ll d$ in few-shot learning.

Model	**Exact**	**Full net.**	**Memory**
MAML (Finn et al., 2017)	✓	✓	$\mathcal{O}(\lvert\boldsymbol{\theta}\rvert \cdot T)$
ANIL (Raghu et al., 2019)	✓	✗	$\mathcal{O}(Nd \cdot T)$
Checkpointing (every $\sqrt{T}$ steps)	✓	✓	$\mathcal{O}(\lvert\boldsymbol{\theta}\rvert \cdot \sqrt{T})$
iMAML (Rajeswaran et al., 2019)	✗	✓	$\mathcal{O}(\lvert\boldsymbol{\theta}\rvert)$
Forward sensitivity (naive)	✓	✗	$\mathcal{O}(N^2d^2 + MNd^2)$
COMLN	✓	✗	$\mathcal{O}(M^2N^2 + M^3N)$

where $\otimes$ is the Kronecker product, and each $\boldsymbol{B}_t[i,j]$ is an $N \times N$ matrix, solution of the following system of ODEs[2]

$$\frac{d\boldsymbol{B}_t[i,j]}{dt} = \mathbb{1}(i=j)\boldsymbol{A}_i(t) - \boldsymbol{A}_i(t)\sum_{m=1}^{M}(\boldsymbol{\phi}_i^\top \boldsymbol{\phi}_m)\boldsymbol{B}_t[m,j] \qquad \boldsymbol{B}_0[i,j] = \mathbf{0}, \tag{11}$$

and $\boldsymbol{A}_i(t)$, defined in Appendix B.2, is also an $N \times N$ matrix that only depends on $\boldsymbol{W}(t)$ and $\boldsymbol{\phi}_i$. The main consequence of this decomposition is that we can simply integrate the augmented ODE in $\big\{\boldsymbol{W}(t), \boldsymbol{B}_t[i,j]\big\}$ up to T to obtain the desired Jacobian matrix, along with the adapted parameters $\boldsymbol{W}(T)$. Furthermore, in contrast to naively applying the forward sensitivity method (see Section 3.2), the M^2 matrices $\boldsymbol{B}_t[i,j]$ are significantly smaller than the full Jacobian matrix. In fact, we show in Appendix C that we can compute vector-Jacobian products—required for the chain rule—using only these smaller matrices, and without ever having to explicitly construct the full $Nd \times Nd$ Jacobian matrix $d\boldsymbol{W}(t)/d\boldsymbol{W}_0$ with Equation 10.

Meta-gradients wrt. Φ To backpropagate the error through the embedding network f_Φ, we need to first compute the gradients of the outer-loss wrt. the feature vectors $\boldsymbol{\phi}_m$. Again, by the chain rule, we can get these gradients with the Jacobian matrices $d\boldsymbol{W}(T)/d\boldsymbol{\phi}_m$. Similar to Equation 10, we can show that these Jacobian matrices can be decomposed as:

$$\frac{d\boldsymbol{W}(t)}{d\boldsymbol{\phi}_m} = -\Bigg[\boldsymbol{s}_m(t) \otimes \boldsymbol{I} + \sum_{i=1}^{M}\boldsymbol{B}_t[i,m]\boldsymbol{W}_0 \otimes \boldsymbol{\phi}_i + \sum_{i=1}^{M}\sum_{j=1}^{M}\boldsymbol{z}_t[i,j,m]\boldsymbol{\phi}_j^\top \otimes \boldsymbol{\phi}_i\Bigg], \tag{12}$$

where $\boldsymbol{s}_m(t)$ and $\boldsymbol{z}_t[i,j,m]$ are vectors of size N, that follow some dynamics; the exact form of this system of ODEs, as well as the proof of this decomposition, are given in Appendix B.3. Crucially, the only quantities that depend on time are small objects independent of the embedding size d. Following the same strategy as above, we can incorporate these vectors in the augmented ODE, and integrate it to get the necessary Jacobians. Once all the $d\boldsymbol{W}(t)/d\boldsymbol{\phi}_m$ are known, for all the training datapoints, we can apply standard backpropagation through f_Φ to obtain the meta-gradients wrt. Φ.

Meta-gradient wrt. T One of the major novelties of COMLN is the capacity to meta-learn the amount of adaptation using stochastic gradient descent. To compute the meta-gradient wrt. the time horizon T, we can directly borrow the results derived by Chen et al. (2018) in the context of Neural ODEs, and apply it to our gradient vector field in Equation 8 responsible for adaptation:

$$\frac{d\mathcal{L}\big(\boldsymbol{W}(T); f_\Phi(\mathcal{D}_\tau^{\text{test}})\big)}{dT} = -\Bigg[\frac{\partial\mathcal{L}\big(\boldsymbol{W}(T); f_\Phi(\mathcal{D}_\tau^{\text{test}})\big)}{\partial\boldsymbol{W}(T)}\Bigg]^\top \frac{\partial\mathcal{L}\big(\boldsymbol{W}(T); f_\Phi(\mathcal{D}_\tau^{\text{train}})\big)}{\partial\boldsymbol{W}(T)}. \tag{13}$$

The proof is available in Appendix B.4. Interestingly, we find that this involves the alignment between the vectors of partial derivatives of the inner-loss and the outer-loss at $\boldsymbol{W}(T)$, which appeared in different contexts in the meta-learning and the multi-task learning literature (Li et al., 2018; Rothfuss et al., 2019; Yu et al., 2020).

[2]Here we used the notation $\boldsymbol{B}_t[i,j]$ to make the dependence on t explicit, without overloading the notation. A more precise notation would be $\boldsymbol{B}[i,j](t)$.

Table 2: Few-shot classification on *mini*ImageNet & *tiered*ImageNet. The average accuracy (%) on 1,000 held-out meta-test tasks is reported with 95% confidence interval. ✓ denotes gradient-based meta-learning algorithms. $\star$ denotes baseline results we executed using the official implementations.

Model		**Backbone**	***mini*ImageNet 5-way**		***tiered*ImageNet 5-way**	
			1-shot	**5-shot**	**1-shot**	**5-shot**
MAML (Finn et al., 2017)	✓	Conv-4	48.70 ± 1.84	63.11 ± 0.92	51.67 ± 1.81	70.30 ± 1.75
ANIL (Raghu et al., 2019)	✓	Conv-4	46.30 ± 0.40	61.00 ± 0.60	49.35 ± 0.26	65.82 ± 0.12
Meta-SGD (Li et al., 2017)	✓	Conv-4	50.47 ± 1.87	64.03 ± 0.94	52.80 ± 0.44	62.35 ± 0.26
CAVIA (Zintgraf et al., 2019)	✓	Conv-4	51.82 ± 0.65	65.85 ± 0.55	$52.41 \pm 2.64^\star$	$67.55 \pm 2.05^\star$
iMAML (Rajeswaran et al., 2019)	✓	Conv-4	49.30 ± 1.88	$59.77 \pm 0.73^\star$	$38.54 \pm 1.37^\star$	$60.24 \pm 0.76^\star$
MetaOptNet-RR (Lee et al., 2019)		Conv-4	$\mathbf{53.23 \pm 0.59}$	69.51 ± 0.48	54.63 ± 0.67	$\mathbf{72.11 \pm 0.59}$
MetaOptNet-SVM (Lee et al., 2019)		Conv-4	52.87 ± 0.57	68.76 ± 0.48	$\mathbf{54.71 \pm 0.67}$	71.79 ± 0.59
COMLN (Ours)	✓	Conv-4	53.01 ± 0.62	$\mathbf{70.54 \pm 0.54}$	54.30 ± 0.69	71.35 ± 0.57
MAML (Finn et al., 2017)	✓	ResNet-12	49.92 ± 0.65	63.93 ± 0.59	55.37 ± 0.74	72.93 ± 0.60
ANIL (Raghu et al., 2019)	✓	ResNet-12	49.65 ± 0.65	59.51 ± 0.56	54.77 ± 0.76	69.28 ± 0.67
MetaOptNet-RR (Lee et al., 2019)		ResNet-12	61.41 ± 0.61	77.88 ± 0.46	65.36 ± 0.71	81.34 ± 0.52
MetaOptNet-SVM (Lee et al., 2019)		ResNet-12	$\mathbf{62.64 \pm 0.61}$	$\mathbf{78.63 \pm 0.46}$	$\mathbf{65.99 \pm 0.72}$	$\mathbf{81.56 \pm 0.53}$
COMLN (Ours)	✓	ResNet-12	59.26 ± 0.65	77.26 ± 0.49	62.93 ± 0.71	81.13 ± 0.53

4.2 Memory efficiency

Although naively applying the forward sensitivity method would be memory intensive, we have shown in Section 4.1 that the Jacobians can be carefully decomposed into smaller pieces. It turns out that even the parameters $\boldsymbol{W}(t)$ can be expressed using the vectors $\boldsymbol{s}_m(t)$ from the decomposition in Equation 12; see Appendix B.1 for details. As a consequence, to compute the adapted parameters $\boldsymbol{W}(T)$ as well as all the necessary meta-gradients, it is sufficient to integrate a dynamical system in $\big\{\boldsymbol{B}_t[i,j], \boldsymbol{s}_m(t), \boldsymbol{z}_t[i,j,m]\big\}$ (see Algorithms 1 & 2 in App. A.1), involving exclusively quantities that are independent of the embedding size d. Instead, the size of that system scales with M the total number of training examples, which is typically much smaller than d for few-shot classification.

Table 1 shows a comparison of the memory cost for different algorithms. It is important to note that contrary to other standard gradient-based meta-learning methods, the memory requirements of COMLN do not scale with the amount of adaptation T (i.e. the number of gradient steps in MAML & ANIL), while still returning the exact meta-gradients—unlike iMAML (Rajeswaran et al., 2019), which only returns an approximation of the meta-gradients. We verified empirically this efficiency, both in terms of memory and computation costs, in Section 5.2.

5 Experiments

For our embedding network f_Φ, we consider two commonly used architectures in meta-learning: Conv-4, a convolutional neural network with 4 convolutional blocks, and ResNet-12, a 12-layer residual network (He et al., 2016). Note that following Lee et al. (2019), ResNet-12 does not include a global pooling layer at the end of the network, leading to feature vectors with embedding dimension $d = 16,000$. Additional details about these architectures are given in Appendix E. To compute the adapted parameters and the meta-gradients in COMLN, we integrate the dynamical system described in Section 4.2 with a 4th order Runge-Kutta method with a Dormand Prince adaptive step size (Runge, 1895; Dormand & Prince, 1980); we will come back to the choice of this numerical solver in Section 5.2. Furthermore to ensure that $T > 0$, we parametrized it with an exponential activation.

5.1 Few-shot image classification

We evaluate COMLN on two standard few-shot image classification benchmarks: the *mini*ImageNet (Vinyals et al., 2016) and the *tiered*ImageNet datasets (Ren et al., 2018), both datasets being derived from ILSVRC-2012 (Russakovsky et al., 2015). The process for creating tasks follows the standard procedure from the few-shot classification literature (Santoro et al., 2016), with distinct classes

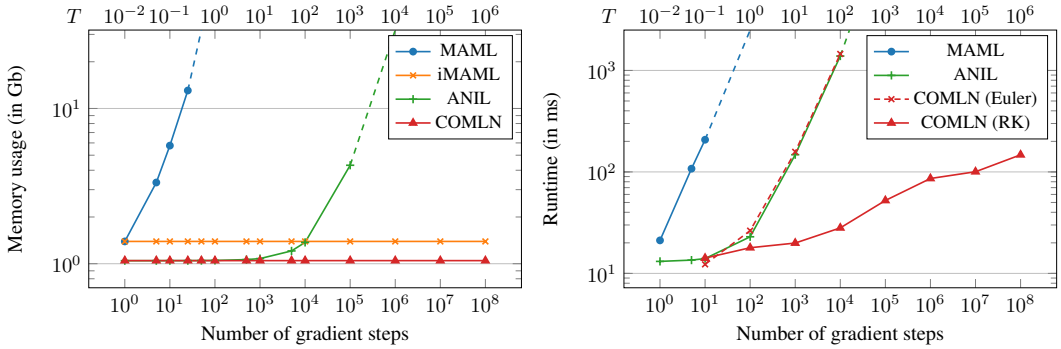

Figure 3: Empirical efficiency of COMLN on a single 5-shot 5-way task, with a Conv-4 backbone. (Left) Memory usage for computing the meta-gradients as a function of the number of inner-gradient steps. The extrapolated dashed lines correspond to the method reaching the memory capacity of a Tesla V100 GPU with 32Gb of memory. (Right) Average time taken (in ms) to compute the exact meta-gradients. The extrapolated dashed lines correspond to the method taking over 2 seconds.

between the different splits. *mini*Imagenet consists of 100 classes, split into 64 training classes, 16 validation classes, and 20 test classes. *tiered*ImageNet consists of 608 classes grouped into 34 high-level categories from ILSVRC-2012, split into 20 training, 6 validation, and 8 testing categories—corresponding to 351/97/160 classes respectively; Ren et al. (2018) argue that separating data according to high-level categories results in a more difficult and more realistic regime.

Table 2 shows the average accuracies of COMLN compared to various meta-learning methods, be it gradient-based or not. For both backbones, COMLN decisively outperforms all other gradient-based meta-learning methods. Compared to methods that explicitly backpropagate through the learning process, such as MAML or ANIL, the performance gain shown by COMLN could be credited to the longer adaptation T it learns, as opposed to a small number of gradient steps—usually about 10 steps; this was fully enabled by our memory-efficient method to compute meta-gradients, which does not scale with the length of adaptation anymore (see Section 4.2). We analyse the evolution of T during meta-training for these different settings in Appendix E.3. In almost all settings, COMLN is even closing the gap with a strong non-gradient-based method like MetaOptNet; the remainder may be explained in part by the training choices made by Lee et al. (2019) (see Appendix E for details).

5.2 Empirical efficiency of COMLN

In Section 4.2, we showed that our algorithm to compute the meta-gradients, based on forward differentiation, had a memory cost independent of the length of adaptation T. We verify this empirically in Figure 3, where we compare the memory required by COMLN and other methods to compute the meta-gradients on a single task, with a Conv-4 backbone (Figure 4 in Appendix E.2 shows similar results for ResNet-12). To ensure an aligned comparison between discrete and continuous time, we use a conversion corresponding to a learning rate $\alpha = 0.01$ in Equation 2; see Appendix E.2 for a justification. As expected, the memory cost increases for both MAML and ANIL as the number of gradient steps increases, while it remains constant for iMAML and COMLN. Interestingly, we observe that the cost of COMLN is equivalent to the cost of running ANIL for a small number of steps, showing that the additional cost of integrating the augmented ODE in Section 4.2 to compute the meta-gradients is minimal.

Increasing the length of adaptation also has an impact on the time it takes to compute the adapted parameters, and the meta-gradients. Figure 3 (right) shows how the runtime increases with the amount of adaptation for different algorithms. We see that the efficiency of COMLN depends on the numerical solver used. When we use a simple explicit-Euler scheme, the time taken to compute the meta-gradients matches the one of ANIL; this behavior empirically confirms our observation in Section 3.3. When we use an adaptive numerical solver, such as Runge-Kutta (RK) with a Dormand Prince step size, this computation can be significantly accelerated, thanks to the smaller number of function evaluations. In practice, we show in Appendix E.1 that the choice of the ODE solver has a very minimal impact on the accuracy.

6 RELATED WORK

We are interested in meta-learning (Bengio et al., 1991; Schmidhuber, 1987; Thrun & Pratt, 2012), and in particular we focus on gradient-based meta-learning methods (Finn, 2018), where the learning rule is based on gradient descent. While in MAML (Finn et al., 2017) the whole network was updated during this process, follow-up works have shown that it is generally sufficient to share most parts of the neural network, and to only adapt a few layers (Raghu et al., 2019; Chen et al., 2020b; Tian et al., 2020). Even though this hypothesis has been challenged recently (Arnold & Sha, 2021), COMLN also updates only the last layer of a neural network, and therefore can be viewed as a continuous-time extension of ANIL (Raghu et al., 2019); see also Section 3.3. With its shared embedding network across tasks, COMLN is also connected to metric-based meta-learning methods (Vinyals et al., 2016; Snell et al., 2017; Sung et al., 2018; Bertinetto et al., 2018; Lee et al., 2019).

Closely related to our work, Zhang et al. (2021) also introduce a formulation where the adaptation of prototypes follows a gradient vector field, but they finally opt for modeling it as a Neural ODE (Chen et al., 2018), possibly due to the challenges of applying the adjoint method we identified in Section 3.2. Zhou et al. (2021) also uses a gradient vector field to motivate a novel method with a closed-form adaptation, based on the NTK theory; COMLN still explicitly updates the parameters following the gradient vector field, since there is no closed-form solution of Eq. 8. As mentioned in Section 3, treating optimization as a continuous-time process has been used to analyze the convergence of different optimization algorithms, including the meta-optimization of MAML (Xu et al., 2021), or to introduce new meta-optimizers based on different integration schemes (Im et al., 2019). Guo et al. (2021) also uses meta-learning to learn new integration schemes for ODEs. Although this is a growing literature at the intersection of meta-learning and dynamical systems, our work is the first algorithm that uses a gradient vector field for adaptation in meta-learning.

Beyond the memory efficiency of our method, one of the main benefits of the continuous-time perspective is that COMLN is capable of learning when to stop the adaptation, as opposed to taking a number of gradient steps fixed ahead of time. However unlike Chen et al. (2020a), where the number of gradient steps are optimized (up to a maximal number) with variational methods, we incorporate the amount of adaptation as a (continuous) meta-parameter that can be learned using SGD. To compute the meta-gradients, which is known to be challenging for long sequences in gradient-based meta-learning, we use forward-mode differentiation as an alternative to backpropagation through the learning process, similar to prior work in meta-learning (Franceschi et al., 2017; Jiwoong Im et al., 2021) and hyperparameter optimization over long horizons (Micaelli & Storkey, 2021). This yields the exact meta-gradients in constant memory, without any assumption on the optimality of the inner optimization problem, which is necessary when using the normal equations (Bertinetto et al., 2018), or to apply implicit differentiation (Rajeswaran et al., 2019).

7 CONCLUSION AND FUTURE WORK

In this paper, we have introduced *Continuous-Time Meta-Learning* (COMLN), a novel algorithm that treats the adaptation in meta-learning as a continuous-time process, by following the dynamics of a gradient vector field up to a certain time horizon T. One of the major novelties of treating adaptation in continuous time is that the amount of adaptation T is now a continuous quantity, that can be viewed as a meta-parameter and can be learned using SGD, alongside the initial conditions and the parameters of the embedding network. In order to learn these meta-parameters, we have also introduced a novel practical algorithm based on forward mode automatic differentiation, capable of efficiently computing the exact meta-gradients using an augmented dynamical system. We have verified empirically that this algorithm was able to compute the meta-gradients in constant memory, making it the first gradient-based meta-learning approach capable of computing the exact meta-gradients with long sequences of adaptation using gradient methods. In practice, we have shown that COMLN significantly outperforms other standard gradient-based meta-learning algorithms.

In addition to having a single meta-parameter T that drives the adaptation of all possible tasks, the fact that the time horizon can be learned with SGD opens up new possibilities for gradient-based methods. For example, we could imagine treating T not as a shared meta-parameters, but as a task-specific parameter. This would allow the learning process to be more *adaptive*, possibly with different behaviors depending on the difficulty of the task. This is left as a future direction of research.

REPRODUCIBILITY STATEMENT

We provide in Appendix A.1 a full description in pseudo-code of the meta-training procedure (Algorithm 1), along with the exact dynamics of the ODE (Algorithm 2) and the projection operations (Algorithms 3 & 4) to avoid explicitly building the Jacobian matrices to compute Jacobian-vector products (see Section 4.1).

We also provide in Appendix A.2 a snippet of code in JAX (Bradbury et al., 2018) to compute the adapted parameters $\boldsymbol{W}(T)$, as well as all the necessary objects $\left\{\boldsymbol{B}_t[i,j], \boldsymbol{s}_m(t), \boldsymbol{z}_t[i,j,m]\right\}$ to compute all the meta-gradients (see Section 4.2). We also give in Code Snippet 2 the code to compute the meta-gradients wrt. the initialization $\boldsymbol{W}_0$ and the integration time T. Computing the meta-gradients wrt. Φ involves non-minimal dependencies on Haiku (Hennigan et al., 2020), and therefore is omitted here. The full code is available in the *Supplementary Materials*.

Data generation & hyperparameters We used the *mini*ImageNet and *tiered*ImageNet datasets provided by Lee et al. (2019) in order to create the 1-shot 5-way and 5-shot 5-way tasks for both datasets. During evaluation, for each setting, a fixed set of $1,000$ tasks were sampled; this means that both architectures for COMLN have been evaluated using exactly the same data, to ensure direct comparison across backbones. A full description of all the hyperparameters used in COMLN is given in Appendix E.

Reproducibility of baseline results To the best of our ability, we have tried to report baseline results from existing work, to limit as much as possible the bias induced by running our own baseline experiments. The references of those works are given in Table 3. We still had to run CAVIA and iMAML on the remaining settings, since these results have not been reported in the literature. For both methods, we used the data generation described above.

- *CAVIA*: We used the official implementation[3]. We used the hyperparameters reported in (Zintgraf et al., 2019) for *mini*ImageNet, and an architecture with 64 filters.
- *iMAML*: We used the official implementation[4]. We used the hyperparameters reported in (Rajeswaran et al., 2019) for *mini*ImageNet 1-shot 5-way.

Table 3: References for the results provided in Table 2: (Liu et al., 2019), (Oh et al., 2021), (Aimen et al., 2021), (Arnold et al., 2021), and are reported in their respective references (under *Model*). Recall that $\star$ denotes baseline results we executed using the official implementations.

Model		**Backbone**	***mini*ImageNet 5-way**		***tiered*ImageNet 5-way**	
			1-shot	**5-shot**	**1-shot**	**5-shot**
MAML (Finn et al., 2017)	✓	Conv-4	48.70 ± 1.84	63.11 ± 0.92	51.67 ± 1.81	70.30 ± 1.75
ANIL (Raghu et al., 2019)	✓	Conv-4	46.30 ± 0.40	61.00 ± 0.60	49.35 ± 0.26	65.82 ± 0.12
Meta-SGD (Li et al., 2017)	✓	Conv-4	50.47 ± 1.87	64.03 ± 0.94	52.80 ± 0.44	62.35 ± 0.26
CAVIA (Zintgraf et al., 2019)	✓	Conv-4	51.82 ± 0.65	65.85 ± 0.55	$52.41 \pm 2.64^{\star}$	$67.55 \pm 2.05^{\star}$
iMAML (Rajeswaran et al., 2019)	✓	Conv-4	49.30 ± 1.88	$59.77 \pm 0.73^{\star}$	$38.54 \pm 1.37^{\star}$	$60.24 \pm 0.76^{\star}$
MetaOptNet-RR (Lee et al., 2019)		Conv-4	$\mathbf{53.23 \pm 0.59}$	69.51 ± 0.48	54.63 ± 0.67	$\mathbf{72.11 \pm 0.59}$
MetaOptNet-SVM (Lee et al., 2019)		Conv-4	52.87 ± 0.57	68.76 ± 0.48	$\mathbf{54.71 \pm 0.67}$	71.79 ± 0.59
COMLN (Ours)	✓	Conv-4	53.01 ± 0.62	$\mathbf{70.54 \pm 0.54}$	54.30 ± 0.69	71.35 ± 0.57
MAML (Finn et al., 2017)	✓	ResNet-12	49.92 ± 0.65	63.93 ± 0.59	55.37 ± 0.74	72.93 ± 0.60
ANIL (Raghu et al., 2019)	✓	ResNet-12	49.65 ± 0.65	59.51 ± 0.56	54.77 ± 0.76	69.28 ± 0.67
MetaOptNet-RR (Lee et al., 2019)		ResNet-12	61.41 ± 0.61	77.88 ± 0.46	65.36 ± 0.71	81.34 ± 0.52
MetaOptNet-SVM (Lee et al., 2019)		ResNet-12	$\mathbf{62.64 \pm 0.61}$	$\mathbf{78.63 \pm 0.46}$	$\mathbf{65.99 \pm 0.72}$	$\mathbf{81.56 \pm 0.53}$
COMLN (Ours)	✓	ResNet-12	59.26 ± 0.65	77.26 ± 0.49	62.93 ± 0.71	81.13 ± 0.53

[3] `https://github.com/lmzintgraf/cavia/`
[4] `https://github.com/aravindr93/imaml_dev`

REFERENCES

Aroof Aimen, Sahil Sidheekh, and Narayanan C Krishnan. Task Attended Meta-Learning for Few-Shot Learning. *arXiv preprint*, 2021.

Sébastien MR Arnold and Fei Sha. Embedding Adaptation is Still Needed for Few-Shot Learning. *arXiv preprint*, 2021.

Sébastien MR Arnold, Guneet S Dhillon, Avinash Ravichandran, and Stefano Soatto. Uniform Sampling over Episode Difficulty. *arXiv preprint*, 2021.

Dmitry Baranchuk. Memory Efficient MAML, 2019. URL `https://github.com/dbaranchuk/memory-efficient-maml`.

Yoshua Bengio, Samy Bengio, Jocelyn Cloutier, and Jan Gecsei. Learning a Synaptic Learning Rule. *International Joint Conference on Neural Networks*, 1991.

Luca Bertinetto, Joao F Henriques, Philip HS Torr, and Andrea Vedaldi. Meta-learning with Differentiable Closed-Form Solvers. *arXiv preprint*, 2018.

John T Betts. *Practical Methods for Optimal Control and Estimation Using Nonlinear Programming*. SIAM, 2010.

Lorenz T. Biegler. *Nonlinear Programming*. Society for Industrial and Applied Mathematics, January 2010.

Dankmar Böhning. Multinomial Logistic Regression Algorithm. *Annals of the Institute of Statistical Mathematics*, 1992.

James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, and Skye Wanderman-Milne. JAX: composable transformations of Python+NumPy programs, 2018. URL `http://github.com/google/jax`.

A. E. Bryson and Y. C. Ho. *Applied Optimal Control*. Blaisdell, New York, 1969.

John Charles Butcher. *Numerical Methods for Ordinary Differential Equations*. Wiley, 2008.

Makis Caracotsios and Warren E Stewart. Sensitivity analysis of initial value problems with mixed odes and algebraic equations. *Computers & Chemical Engineering*, 9(4):359–365, 1985.

G Chavent, RE Goodson, and M Polis. Identification of parameter distributed systems. *Identification of function parameters in partial differential equations*, pp. 31–48, 1974.

Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David Duvenaud. Neural Ordinary Differential Equations. *Advances in Neural Information Processing Systems*, 2018.

Xinshi Chen, Hanjun Dai, Yu Li, Xin Gao, and Le Song. Learning To Stop While Learning To Predict. In *International Conference on Machine Learning*, 2020a.

Yutian Chen, Abram L Friesen, Feryal Behbahani, Arnaud Doucet, David Budden, Matthew W Hoffman, and Nando de Freitas. Modular Meta-Learning with Shrinkage. *Neural Information Processing Systems*, 2020b.

John R Dormand and Peter J Prince. A family of embedded Runge-Kutta formulae. *Journal of computational and applied mathematics*, 1980.

Leonhard Euler. De integratione aequationum differentialium per approximationem. *Opera Omnia*, 1913.

William F Feehery, John E Tolsma, and Paul I Barton. Efficient sensitivity analysis of large-scale differential-algebraic systems. *Applied Numerical Mathematics*, 25(1):41–54, 1997.

Chelsea Finn. *Learning to Learn with Gradients*. PhD thesis, UC Berkeley, 2018.

Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks. *International Conference on Machine Learning (ICML)*, 2017.

Sebastian Flennerhag, Pablo G Moreno, Neil D Lawrence, and Andreas Damianou. Transferring knowledge across learning processes. *arXiv preprint*, 2018.

Sebastian Flennerhag, Andrei A Rusu, Razvan Pascanu, Francesco Visin, Hujun Yin, and Raia Hadsell. Meta-Learning with Warped Gradient Descent. *International Conference on Learning Representations*, 2020.

Luca Franceschi, Michele Donini, Paolo Frasconi, and Massimiliano Pontil. Forward and reverse gradient-based hyperparameter optimization. In *International Conference on Machine Learning*, 2017.

Golnaz Ghiasi, Tsung-Yi Lin, and Quoc V Le. Dropblock: A regularization method for convolutional networks. In *Neural Information Processing Systems*, 2018.

Amir Gholami, Kurt Keutzer, and George Biros. ANODE: unconditionally accurate memory-efficient gradients for neural odes. *CoRR*, abs/1902.10298, 2019. URL `http://arxiv.org/abs/1902.10298`.

Andreas Griewank and Andrea Walther. *Evaluating Derivatives*. Society for Industrial and Applied Mathematics, January 2008.

Yue Guo, Felix Dietrich, Tom Bertalan, Danimir T Doncevic, Manuel Dahmen, Ioannis G Kevrekidis, and Qianxiao Li. Personalized Algorithm Generation: A Case Study in Meta-Learning ODE Integrators. *arXiv preprint*, 2021.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016.

Tom Hennigan, Trevor Cai, Tamara Norman, and Igor Babuschkin. Haiku: Sonnet for JAX, 2020. URL `http://github.com/deepmind/dm-haiku`.

Daniel Jiwoong Im, Yibo Jiang, and Nakul Verma. Model-Agnostic Meta-Learning using Runge-Kutta Methods. *arXiv preprint*, 2019.

Khurram Javed and Martha White. Meta-Learning Representations for Continual Learning. In *Advances in Neural Information Processing Systems*, 2019.

Daniel Jiwoong Im, Cristina Savin, and Kyunghyun Cho. Online hyperparameter optimization by Real-Time Recurrent Learning. *arXiv preprint*, 2021.

Kwonjoon Lee, Subhransu Maji, Avinash Ravichandran, and Stefano Soatto. Meta-learning with Differentiable Convex Optimization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019.

Jorge R Leis and Mark A Kramer. The simultaneous solution and sensitivity analysis of systems described by ordinary differential equations. *ACM Transactions on Mathematical Software (TOMS)*, 14(1):45–60, 1988.

Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy Hospedales. Learning to Generalize: Meta-Learning for Domain Generalization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2018.

Zhenguo Li, Fengwei Zhou, Fei Chen, and Hang Li. Meta-SGD: Learning to learn quickly for few-shot learning. *arXiv preprint*, 2017.

Jacques Louis Lions and Enrico Magenes. *Non-homogeneous boundary value problems and applications: Vol. 1*, volume 181. Springer Science & Business Media, 2012.

Yanbin Liu, Juho Lee, Minseop Park, Saehoon Kim, Eunho Yang, Sung Ju Hwang, and Yi Yang. Learning to Propagate Labels: Transductive Propagation Network for Few-Shot Learning. *International Conference on Learning Representations*, 2019.

Dougal Maclaurin, David Duvenaud, and Ryan Adams. Gradient-based hyperparameter optimization through reversible learning. In *International conference on machine learning*, pp. 2113–2122. PMLR, 2015.

Timothy Maly and Linda R Petzold. Numerical methods and software for sensitivity analysis of differential-algebraic systems. *Applied Numerical Mathematics*, 20(1-2):57–79, 1996.

Paul Micaelli and Amos Storkey. Gradient-based Hyperparameter Optimization Over Long Horizons. In *Neural Information Processing Systems*, 2021.

Alex Nichol, Joshua Achiam, and John Schulman. On First-Order Meta-Learning Algorithms. *arXiv preprint*, 2018.

Jaehoon Oh, Hyungjun Yoo, ChangHwan Kim, and Se-Young Yun. BOIL: Towards Representation Change for Few-Shot Learning. *International Conference on Learning Representations*, 2021.

Antonio Orvieto and Aurelien Lucchi. Shadowing Properties of Optimization Algorithms. In *Advances in Neural Information Processing Systems*, 2019.

John Platt and Alan Barr. Constrained Differential Optimization. In *Neural Information Processing Systems*, 1988.

Boris T Polyak. Some methods of speeding up the convergence of iteration methods. *USSR computational mathematics and mathematical physics*, 1964.

Lev Semenovich Pontryagin. *Mathematical theory of optimal processes*. Routledge, 2018.

Aniruddh Raghu, Maithra Raghu, Samy Bengio, and Oriol Vinyals. Rapid learning or feature reuse? towards understanding the effectiveness of maml. *arXiv preprint*, 2019.

Aravind Rajeswaran, Chelsea Finn, Sham M Kakade, and Sergey Levine. Meta-Learning with Implicit Gradients. In *Advances in Neural Information Processing Systems*, 2019.

Mengye Ren, Eleni Triantafillou, Sachin Ravi, Jake Snell, Kevin Swersky, Joshua B Tenenbaum, Hugo Larochelle, and Richard S Zemel. Meta-learning for semi-supervised few-shot classification. In *International Conference on Learning Representations*, 2018.

Jonas Rothfuss, Dennis Lee, Ignasi Clavera, Tamim Asfour, and Pieter Abbeel. ProMP: Proximal Meta-Policy Search. *International Conference on Learning Representations*, 2019.

Carl Runge. Über die numerische auflösung von differentialgleichungen. *Mathematische Annalen*, 1895.

Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015.

Andrei A Rusu, Dushyant Rao, Jakub Sygnowski, Oriol Vinyals, Razvan Pascanu, Simon Osindero, and Raia Hadsell. Meta-learning with Latent Embedding Optimization. *arXiv preprint*, 2018.

Adam Santoro, Sergey Bartunov, Matthew Botvinick, Daan Wierstra, and Timothy Lillicrap. One-shot learning with memory-augmented neural networks. *arXiv preprint*, 2016.

Jürgen Schmidhuber. *Evolutionary principles in self-referential learning, or on learning how to learn: the meta-meta-... hook*. PhD thesis, Technische Universität München, 1987.

Radu Serban and Alan C Hindmarsh. Cvodes: An ode solver with sensitivity analysis capabilities. Technical report, Technical Report UCRL-JP-200039, Lawrence Livermore National Laboratory, 2003.

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.

Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical Networks for Few-shot Learning. In *Advances in Neural Information Processing Systems*, 2017.

Weijie Su, Stephen Boyd, and Emmanuel Candes. A Differential Equation for Modeling Nesterov's Accelerated Gradient Method: Theory and Insights. *Advances in Neural Information Processing Systems*, 2014.

Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip HS Torr, and Timothy M Hospedales. Learning to compare: Relation network for few-shot learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1199–1208, 2018.

Richard S. Sutton. Adapting bias by gradient descent: An incremental version of delta-bar-delta. In William R. Swartout (ed.), *Proceedings of the 10th National Conference on Artificial Intelligence, San Jose, CA, USA, July 12-16, 1992*, pp. 171–176. AAAI Press / The MIT Press, 1992.

Sebastian Thrun and Lorien Pratt. *Learning to learn*. Springer Science & Business Media, 2012.

Yonglong Tian, Yue Wang, Dilip Krishnan, Joshua B Tenenbaum, and Phillip Isola. Rethinking Few-Shot Image Classification: a Good Embedding Is All You Need? 2020.

Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. *Neural Information Processing Systems*, 29:3630–3638, 2016.

Stephen Wiggins. *Introduction to Applied Nonlinear Dynamical Systems and Chaos*, volume 2. Springer, 2003.

Ronald J. Williams and David Zipser. A learning algorithm for continually running fully recurrent neural networks. *Neural Computation*, 1(2):270–280, June 1989.

Ashia C Wilson, Benjamin Recht, and Michael I Jordan. A Lyapunov Analysis of Momentum Methods in Optimization. *arXiv preprint*, 2016.

Ruitu Xu, Lin Chen, and Amin Karbasi. Meta Learning in the Continuous Time Limit . In *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*. PMLR, 2021.

Zhongwen Xu, Hado van Hasselt, and David Silver. Meta-gradient reinforcement learning. In Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pp. 2402–2413, 2018.

Tianhe Yu, Saurabh Kumar, Abhishek Gupta, Sergey Levine, Karol Hausman, and Chelsea Finn. Gradient Surgery for Multi-Task Learning. *Neural Information Processing Systems*, 2020.

Baoquan Zhang, Xutao Li, Yunming Ye, Shanshan Feng, and Rui Ye. MetaNODE: Prototype Optimization as a Neural ODE for Few-Shot Learning. *arXiv preprint*, 2021.

Yufan Zhou, Zhenyi Wang, Jiayi Xian, Changyou Chen, and Jinhui Xu. Meta-Learning with Neural Tangent Kernels. *International Conference on Learning Representations*, 2021.

Luisa Zintgraf, Kyriacos Shiarli, Vitaly Kurin, Katja Hofmann, and Shimon Whiteson. Fast context adaptation via meta-learning. In *International Conference on Machine Learning*, pp. 7693–7702. PMLR, 2019.