
On Training Implicit Models

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 This paper focuses on the principles of training implicit models of infinite layers.
2 Specifically, previous works employ the implicit differentiation and solve the *exact*
3 gradient for the backward propagation. However, *is it necessary to compute such*
4 *an exact gradient (which is usually quite expensive) for training?* To this end,
5 we propose a novel gradient estimate for these implicit models, named *phantom*
6 *gradient*, that 1) forgoes the costly approximation of the exact gradient; and 2)
7 provides an update direction (empirically) preferable to the implicit model training.
8 We theoretically analyze the condition under which a descent direction of the loss
9 landscape could be found, and provide two specific instantiations of the phantom
10 gradient based on unrolling and the Neumann series. Experiments on large-scale
11 vision tasks demonstrate that these lightweight phantom gradients significantly
12 accelerate the backward passes in training implicit models (roughly $1.7\times$ speedup),
13 and even boost the performance over approaches based on the exact gradient.

14 1 Introduction

15 Conventional neural networks are typically constructed by explicitly stacking multiple linear and
16 non-linear operators in a feed-forward manner. Recently, the implicitly-defined models [1, 2, 3, 4, 5]
17 have attracted increasing attentions and are able to match the state-of-the-art level results by explicit
18 models on several vision [4, 5] and natural language processing [3] tasks. Specifically, these works
19 treat the evolution of the intermediate hidden states as a certain form of dynamical system, such as
20 fixed-point equations [3, 4, 5] or an ordinary differential equation (ODE) flow [1, 2], which represents
21 infinite latent states. The forward passes of implicit models are therefore formulated as solving
22 these underlying dynamics, by either black-box ODE solvers [1, 2] or root-finding algorithms [3, 4].
23 As for the backward passes, however, directly differentiating through the forward pass trajectories
24 could induce a heavy memory overhead [6, 7]. To this end, researchers have developed several
25 approaches based on the implicit differentiation, such as solving a Jacobian-based linear fixed-point
26 equation for the backward pass of deep equilibrium (DEQ) models [3], which eventually makes
27 the backpropagation trajectories independent of the forward pass ones, allowing one to train these
28 implicit models with essentially constant memory consumption (as we only need to store the final
29 output and the layer itself, without any intermediate states).

30 However, in order to estimate the exact gradient promised by the implicit differentiation, these
31 implicit models still have to rely on black-box solvers (*e.g.*, ODE solvers or root-solving algorithms),
32 whose iterative nature usually makes the gradient computation very costly in practice (*e.g.*, over 30
33 steps for large-scale DEQ models). In this work, we investigate the question of whether an accurate
34 gradient estimate is necessary for training implicit models. We found that *a first-order oracle that*
35 *produces good “gradient estimates” is enough to efficiently and effectively train the model*, without
36 the need to precisely (and laboriously) compute the exact gradient, as in prior works [3, 4, 8, 9].

37 As an application of our principle, we develop a framework in which a balanced trade-off is made
38 between the precision and conditioning of the gradient estimate. Specifically, we name our gradient

39 estimate as the *phantom gradient*, and provide the general condition under which the phantom
 40 gradient can provide a *descent direction* of the loss landscape. We further propose two instantiations
 41 of the phantom gradient in the context of DEQ models, which are based on the the simple fixed-point
 42 unrolling and the Neumann series analysis. Importantly, we show that our proposed instantiations
 43 satisfy the descent condition, and the stochastic gradient descent (SGD) algorithm based on the
 44 phantom gradient enjoys a sound convergence property as long as the relevant hyperparameters
 45 (e.g., the damping factor) are wisely selected. Note that our proposed method only directly affects,
 46 and thus accelerates, the backward formulation of these implicit models, with the forward pass
 47 formulation (i.e., the root-solving process) and inference-time behavior mostly intact.

48 We conduct an extensive set of synthetic, ablation, and large-scale experiments to both analyze the
 49 theoretical properties of the phantom gradient and validate its performance on large-scale tasks, such
 50 as CIFAR-10 [10] and 224×224 ImageNet [11] classification tasks. Overall, our results suggest that:
 51 1) the phantom gradient estimates a descent direction; 2) it is applicable to large-scale tasks and is
 52 capable of achieving a strong performance that is comparable with or better than when using exact
 53 gradients; and 3) it significantly shortens the training time needed for implicit models, by a factor
 54 of $1.4 \sim 1.7\times$. We believe these theoretical and empirical results provide strong evidence for the
 55 effectiveness of training implicit models with the inexact and lightweight phantom gradient.

56 2 Method

57 2.1 Inspection of Implicit Differentiation

58 In this work, we primarily focus on the formulation of root-solving-based implicit models, represented
 59 by the DEQ models [3]. Specifically, given a non-linear layer \mathcal{F} , the output of the implicit model is
 60 characterized as the solution \mathbf{h}^* to the following fixed-point equation:

$$\mathbf{h}^* = \mathcal{F}(\mathbf{h}^*, \mathbf{z}), \quad (1)$$

61 where $\mathbf{z} \in \mathbb{R}^{d_u+d_\theta}$ is the union of the module’s input $\mathbf{u} \in \mathbb{R}^{d_u}$ and parameters $\boldsymbol{\theta} \in \mathbb{R}^{d_\theta}$, i.e., $\mathbf{z}^\top =$
 62 $[\mathbf{u}^\top, \boldsymbol{\theta}^\top]$. Here, \mathbf{u} is usually the projection of the original data point $\mathbf{x} \in \mathbb{R}^{d_x}$, e.g., $\mathbf{u} = \mathcal{M}(\mathbf{x})$. In
 63 this section, we assume \mathcal{F} is a contractive mapping w.r.t. \mathbf{h} so that its Lipschitz constant L_h w.r.t. \mathbf{h}
 64 is less than one (i.e., $L_h < 1$), a setting that has been analyzed in numerous prior works [12, 13, 14].

65 To perform backpropagation through the module induced by Eq. (1), we need to calculate the Jacobian
 66 matrix of \mathbf{h}^* w.r.t. the projected input (as well as parameters) \mathbf{z} . By Implicit Function Theorem,

$$\frac{\partial \mathbf{h}^*}{\partial \mathbf{z}} = \frac{\partial \mathcal{F}}{\partial \mathbf{z}} \Big|_{\mathbf{h}^*} \left(\mathbf{I} - \frac{\partial \mathcal{F}}{\partial \mathbf{h}} \Big|_{\mathbf{h}^*} \right)^{-1} \quad (2)$$

67 The fixed point \mathbf{h}^* of Eq. (1) is then passed to a post-processing function \mathcal{G} to predict $\tilde{\mathbf{y}} = \mathcal{G}(\mathbf{h}^*)$. In
 68 the generic learning scenario, the training objective is to minimize the following expected loss:

$$\mathcal{R}(\boldsymbol{\theta}) = \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{P}} [\mathcal{L}(\tilde{\mathbf{y}}(\mathbf{x}, \boldsymbol{\theta}), \mathbf{y})], \quad (3)$$

69 where \mathbf{y} is the groundtruth corresponding to the training example \mathbf{x} , and \mathcal{P} is the data distribution.
 70 Here, we omit the parameters of \mathcal{G} , because given the output \mathbf{h}^* of the implicit module \mathcal{F} , training the
 71 post-processing part \mathcal{G} is the same as training explicit neural networks. The most crucial component
 72 is the gradient of the loss function \mathcal{L} w.r.t. the input vector $\mathbf{z}^\top = [\mathbf{u}^\top, \boldsymbol{\theta}^\top]$, which is used to train
 73 both the implicit module \mathcal{F} and the input projection module \mathcal{M} . Using Eq. (2) with the condition
 74 $\mathbf{h} = \mathbf{h}^*$, we have

$$\frac{\partial \mathcal{L}}{\partial \mathbf{u}} = \frac{\partial \mathcal{F}}{\partial \mathbf{u}} \left(\mathbf{I} - \frac{\partial \mathcal{F}}{\partial \mathbf{h}} \right)^{-1} \frac{\partial \mathcal{L}}{\partial \mathbf{h}}, \quad \frac{\partial \mathcal{L}}{\partial \boldsymbol{\theta}} = \frac{\partial \mathcal{F}}{\partial \boldsymbol{\theta}} \left(\mathbf{I} - \frac{\partial \mathcal{F}}{\partial \mathbf{h}} \right)^{-1} \frac{\partial \mathcal{L}}{\partial \mathbf{h}}. \quad (4)$$

75 The gradients in Eq. (4) are symmetric w.r.t. \mathbf{u} and $\boldsymbol{\theta}$. Without loss of generality, we only discuss the
 76 gradient w.r.t. $\boldsymbol{\theta}$ in the following sections.

77 The most intriguing part lies in the Jacobian-inverse term, i.e., $(\mathbf{I} - \partial \mathcal{F} / \partial \mathbf{h})^{-1}$. Because of the
 78 inverse operation, a natural question arises from the numerical aspect. *Is it well-conditioned?* If the
 79 absolute value of the eigenvalue of $\partial \mathcal{F} / \partial \mathbf{h}$ is close to 1, the Jacobian-inverse will be numerically
 80 unstable. This spurs us to rethink the necessity of the Jacobian-inverse term in the standard implicit

81 differentiation. Note that the Jacobian-inverse is calculated to update model parameters with their
 82 gradients, but the exact gradient is not always optimal in model training. For example, previous
 83 research has instead used a moderate gradient noise as a regularization approach [15], which have
 84 been shown to play a central role in escaping poor local minima and improving generalization ability
 85 [16, 17, 18]. Moreover, as computing the inverse of a large matrix (e.g., $\partial\mathcal{F}/\partial\mathbf{h}$ is more than
 86 $10^5 \times 10^5$ on ImageNet) is intractable, prior works [3] proposes to iteratively solve a linear system
 87 involving a Jacobian-vector product instead, which makes the actual backward pass slow.

88 These observations motivate us to design an inexact, but theoretically sound and practically efficient
 89 gradient for training implicit models. Here, suppose the Jacobian $\partial\mathbf{h}^*/\partial\boldsymbol{\theta}$ is replaced with a matrix
 90 \mathbf{A} , and the corresponding *phantom gradient* is defined as

$$\frac{\widehat{\partial\mathcal{L}}}{\partial\boldsymbol{\theta}} := \mathbf{A} \frac{\partial\mathcal{L}}{\partial\mathbf{h}}. \quad (5)$$

91 Next, we give the general condition on \mathbf{A} so that a descent property of the phantom gradient can
 92 be guaranteed (Sec. 2.2), and provide two concrete instantiations of \mathbf{A} based on either fixed-point
 93 unrolling or the Neumann series (Sec. 2.3).

94 2.2 General Condition on the Phantom Gradient

95 The following theorem formulates a sufficient condition that the phantom gradient gives a descent
 96 direction of the loss landscape. Please refer to the appendix for the proof.

97 **Theorem 1.** *Suppose the exact gradient and the phantom gradient are given by Eq. (4) and Eq. (5),
 98 respectively. Let σ_{\max} and σ_{\min} be the maximal and minimal singular value of $\partial\mathcal{F}/\partial\boldsymbol{\theta}$. If*

$$\left\| \mathbf{A} \left(\mathbf{I} - \frac{\partial\mathcal{F}}{\partial\mathbf{h}} \right) - \frac{\partial\mathcal{F}}{\partial\boldsymbol{\theta}} \right\| \leq \frac{\sigma_{\min}^2}{\sigma_{\max}}, \quad (6)$$

99 *then the phantom gradient provides a descent direction of the function \mathcal{L} , i.e.,*

$$\left\langle \frac{\widehat{\partial\mathcal{L}}}{\partial\boldsymbol{\theta}}, \frac{\partial\mathcal{L}}{\partial\boldsymbol{\theta}} \right\rangle \geq 0. \quad (7)$$

100

101 **Remark 1.** Suppose only the $(\mathbf{I} - \partial\mathcal{F}/\partial\mathbf{h})^{-1}$ term is replaced with a matrix \mathbf{D} , namely, $\mathbf{A} =$
 102 $(\partial\mathcal{F}/\partial\boldsymbol{\theta}) \mathbf{D}$. Then, the condition in (6) can be reduced into

$$\left\| \mathbf{D} \left(\mathbf{I} - \frac{\partial\mathcal{F}}{\partial\mathbf{h}} \right) - \mathbf{I} \right\| \leq \frac{1}{\kappa^2}, \quad (8)$$

103 where κ is the condition number of $\partial\mathcal{F}/\partial\boldsymbol{\theta}$. The derivation can be found in the appendix.

104 **Remark 2.** The singular value σ and the condition number κ of $\partial\mathcal{F}/\partial\boldsymbol{\theta}$ make it tricky to ensure
 105 the condition in (6) or (8). However, with $\mathbf{J} = \partial\mathcal{F}/\partial\boldsymbol{\theta}$, $\boldsymbol{\Theta} = \mathbf{J}^\top \mathbf{J}$ is exactly the *neural tangent*
 106 *kernel* (NTK) [19] corresponding to the module \mathcal{F} . If \mathcal{F} is a multi-layer neural network, its NTK $\boldsymbol{\Theta}$
 107 converges in probability to a scalar matrix in the infinite-width limit, i.e.,

$$\boldsymbol{\Theta} \xrightarrow{P} s\mathbf{I}, \quad \text{for some } s \in \mathbb{R}_+, \text{ as width } \rightarrow \infty. \quad (9)$$

108 This conclusion holds both at initialization and during the training process [19, Theorem 1 & 2]¹,
 109 which implies that if \mathcal{F} is sufficiently wide, all singular values of \mathbf{J} and thus its condition number κ
 110 are close to 1 in the entire training stage. This property makes the threshold in (6) and (8) computable.

111 2.3 Instantiations of the Phantom Gradient

112 In this section, we present two practical instantiations of the phantom gradient. We also verify that
 113 the general condition in Theorem 1 can be satisfied if the hyperparameters in our instantiations are
 114 wisely selected.

¹A concrete bound of the approximation error for finite-width networks can be found in [20, Theorem 3.1 & 3.2].

115 Suppose we hope to differentiate through an implicit dynamic, *e.g.*, either a root-solving process or
 116 an optimization problem. Previous solutions towards this include differentiating through the unrolled
 117 steps of the dynamics [21] and using the Neumann series [7]. In our case, if we solve the root of
 118 Eq. (1) via fixed-point iteration:

$$\mathbf{h}_{t+1} = \mathcal{F}(\mathbf{h}_t, \mathbf{z}), \quad t = 0, 1, \dots, T-1, \quad (10)$$

119 then by differentiating through the unrolled steps of Eq. (10), we have

$$\frac{\partial \mathbf{h}_T}{\partial \boldsymbol{\theta}} = \sum_{t=0}^{T-1} \frac{\partial \mathcal{F}}{\partial \boldsymbol{\theta}} \Big|_{\mathbf{h}_t} \prod_{s=t+1}^{T-1} \frac{\partial \mathcal{F}}{\partial \mathbf{h}} \Big|_{\mathbf{h}_s}. \quad (11)$$

120 Besides, the Neumann series of the Jacobian-inverse $(\mathbf{I} - \partial \mathcal{F} / \partial \mathbf{h})^{-1}$ is

$$\mathbf{I} + \frac{\partial \mathcal{F}}{\partial \mathbf{h}} + \left(\frac{\partial \mathcal{F}}{\partial \mathbf{h}} \right)^2 + \left(\frac{\partial \mathcal{F}}{\partial \mathbf{h}} \right)^3 + \dots. \quad (12)$$

121 Notably, computing the Jacobian $\partial \mathbf{h}^* / \partial \boldsymbol{\theta}$ using the Neumann series in (12) is equivalent to differenti-
 122 ating through the unrolled steps of Eq. (10) at the exact solution point \mathbf{h}^* and taking the limit of
 123 infinite steps [7].

124 Without altering the root of Eq. (1), we consider a damped variant of the fixed-point iteration:

$$\mathbf{h}_{t+1} = \lambda \mathcal{F}(\mathbf{h}_t, \mathbf{z}) + (1 - \lambda) \mathbf{h}_t, \quad t = 0, 1, \dots, T-1. \quad (13)$$

125 Differentiating through the unrolled steps of Eq. (13), Eq. (11) is adapted as

$$\frac{\partial \mathbf{h}_T}{\partial \boldsymbol{\theta}} = \lambda \sum_{t=0}^{T-1} \frac{\partial \mathcal{F}}{\partial \boldsymbol{\theta}} \Big|_{\mathbf{h}_t} \prod_{s=t+1}^{T-1} \left(\lambda \frac{\partial \mathcal{F}}{\partial \mathbf{h}} \Big|_{\mathbf{h}_s} + (1 - \lambda) \mathbf{I} \right). \quad (14)$$

126 The Neumann series of $(\mathbf{I} - \partial \mathcal{F} / \partial \mathbf{h})^{-1}$ is correspondingly adapted as

$$\lambda (\mathbf{I} + \mathbf{B} + \mathbf{B}^2 + \mathbf{B}^3 + \dots), \quad \text{where } \mathbf{B} = \lambda \frac{\partial \mathcal{F}}{\partial \mathbf{h}} + (1 - \lambda) \mathbf{I}. \quad (15)$$

127 The next theorem shows that under mild conditions, the Jacobian in Eq. (14) converges to the exact
 128 Jacobian and the Neumann series in (15) converges to the Jacobian-inverse $(\mathbf{I} - \partial \mathcal{F} / \partial \mathbf{h})^{-1}$.

129 **Theorem 2.** *Suppose the matrix $\partial \mathcal{F} / \partial \mathbf{h}$ is a contractive mapping. Then,*

- 130 (i) *the Neumann series in (15) converges to the Jacobian-inverse $(\mathbf{I} - \partial \mathcal{F} / \partial \mathbf{h})^{-1}$; and*
 131 (ii) *if the function \mathcal{F} is continuously differentiable w.r.t. both \mathbf{h} and $\boldsymbol{\theta}$, the sequence in Eq. (14)*
 132 *converges to the exact Jacobian $\partial \mathbf{h}^* / \partial \boldsymbol{\theta}$ as $T \rightarrow \infty$, i.e.,*

$$\lim_{T \rightarrow \infty} \frac{\partial \mathbf{h}_T}{\partial \boldsymbol{\theta}} = \frac{\partial \mathcal{F}}{\partial \boldsymbol{\theta}} \Big|_{\mathbf{h}^*} \left(\mathbf{I} - \frac{\partial \mathcal{F}}{\partial \mathbf{h}} \Big|_{\mathbf{h}^*} \right)^{-1}. \quad (16)$$

133 However, as discussed in Sec. 2.1, it is unnecessary to compute the exact gradient with infinite
 134 terms. In the following context, we introduced two instantiations of the phantom gradient based on a
 135 finite-term truncation of Eq. (14) or (15).

136 **Unrolling-based Phantom Gradient.** In the unrolling form, the matrix \mathbf{A} is defined as

$$\mathbf{A}_{k,\lambda}^{\text{unr}} = \lambda \sum_{t=0}^{k-1} \frac{\partial \mathcal{F}}{\partial \boldsymbol{\theta}} \Big|_{\mathbf{h}_t} \prod_{s=t+1}^{k-1} \left(\lambda \frac{\partial \mathcal{F}}{\partial \mathbf{h}} \Big|_{\mathbf{h}_s} + (1 - \lambda) \mathbf{I} \right). \quad (17)$$

137 **Neumann-series-based Phantom Gradient.** In the Neumann form, the matrix \mathbf{A} is defined as

$$\mathbf{A}_{k,\lambda}^{\text{neu}} = \lambda \frac{\partial \mathcal{F}}{\partial \boldsymbol{\theta}} \Big|_{\mathbf{h}^*} (\mathbf{I} + \mathbf{B} + \mathbf{B}^2 + \dots + \mathbf{B}^{k-1}), \quad \text{where } \mathbf{B} = \lambda \frac{\partial \mathcal{F}}{\partial \mathbf{h}} \Big|_{\mathbf{h}^*} + (1 - \lambda) \mathbf{I}. \quad (18)$$

138 According to Theorem 2, the matrix \mathbf{A} defined by either Eq. (17) or Eq. (18) converges to the exact
 139 Jacobian $\partial \mathbf{h}^* / \partial \boldsymbol{\theta}$ as $k \rightarrow \infty$ for any $\lambda \in (0, 1]$. Therefore, by Theorem 2, the condition in (6) can
 140 be satisfied if a sufficiently large step k is selected, since

$$\left\| \mathbf{A} \left(\mathbf{I} - \frac{\partial \mathcal{F}}{\partial \mathbf{h}} \right) - \frac{\partial \mathcal{F}}{\partial \boldsymbol{\theta}} \right\| \leq (1 + L_h) \left\| \mathbf{A} - \frac{\partial \mathcal{F}}{\partial \boldsymbol{\theta}} \left(\mathbf{I} - \frac{\partial \mathcal{F}}{\partial \mathbf{h}} \right)^{-1} \right\|. \quad (19)$$

141 Next, we characterize the impact of the two hyperparameters, *i.e.*, k and λ , on the precision and
 142 conditioning of \mathbf{A} . Take the Neumann-series-based phantom gradient (Eq. (18)) as an example.

- 143 (i) On the precision of the phantom gradient,
- 144 • a large k makes the gradient estimate more accurate, as higher order terms of the
 - 145 Neumann series are included; while
 - 146 • a small λ slows down the convergence of the Neumann series because of the larger
 - 147 norm of \mathbf{B} with the decrease of λ .
- 148 (ii) On the conditioning of the phantom gradient,
- 149 • a large k impairs the conditioning of \mathbf{A} since the condition number of \mathbf{B}^k grows
 - 150 exponentially as k increases; while
 - 151 • a small λ helps maintain a small condition number of \mathbf{A} because the singular values of
 - 152 $\partial \mathcal{F} / \partial \mathbf{h}$ are “smoothed” by the identity matrix.

153 In a word, a large k is preferable for a more accurate \mathbf{A} , while a small λ contributes to the well-
 154 conditioning of \mathbf{A} . Practically, these hyperparameters should be selected in consideration of a
 155 balanced trade-off between the precision and conditioning of \mathbf{A} . See Sec. 3 for experimental results.

156 2.4 Convergence Theory

157 In this section, we provide convergence guarantee of the SGD algorithm using the phantom gradient.
 158 We prove that under mild conditions, if the approximation error of the phantom gradient is sufficiently
 159 small, the SGD algorithm converges to an ϵ -approximate stationary point in the expectation sense.
 160 Please refer to the appendix for the proof, where we also discuss the feasibility of our assumptions.

161 **Theorem 3.** *Suppose the loss function \mathcal{R} in Eq. (3) is ℓ -smooth, lower-bounded, and has bounded*
 162 *gradient almost surely in the training process. Besides, assume the gradient in Eq. (4) is an*
 163 *unbiased estimator of $\nabla \mathcal{R}(\boldsymbol{\theta})$ with a bounded covariance. If the phantom gradient in Eq. (5) is an*
 164 *ϵ -approximation to the gradient in Eq. (4), *i.e.*,*

$$\left\| \widehat{\frac{\partial \mathcal{L}}{\partial \boldsymbol{\theta}}} - \frac{\partial \mathcal{L}}{\partial \boldsymbol{\theta}} \right\| \leq \epsilon, \quad \text{almost surely}, \quad (20)$$

165 *then using Eq. (5) as a stochastic first-order oracle with a step size of $\eta_\tau = O(1/\sqrt{\tau})$ to update $\boldsymbol{\theta}$*
 166 *with gradient descent, it follows after T iterations that*

$$\mathbb{E} \left[\frac{\sum_{\tau=1}^T \eta_\tau \|\nabla \mathcal{R}(\boldsymbol{\theta}_\tau)\|^2}{\sum_{\tau=1}^T \eta_\tau} \right] \leq O \left(\epsilon + \frac{\log T}{\sqrt{T}} \right). \quad (21)$$

167 **Remark 3.** Consider the condition in (20):
 168

$$\left\| \widehat{\frac{\partial \mathcal{L}}{\partial \boldsymbol{\theta}}} - \frac{\partial \mathcal{L}}{\partial \boldsymbol{\theta}} \right\| \leq \left\| \mathbf{A} - \frac{\partial \mathcal{F}}{\partial \boldsymbol{\theta}} \left(\mathbf{I} - \frac{\partial \mathcal{F}}{\partial \mathbf{h}} \right)^{-1} \right\| \left\| \frac{\partial \mathcal{L}}{\partial \mathbf{h}} \right\|. \quad (22)$$

169 Suppose the norm of $\partial \mathcal{L} / \partial \mathbf{h}$ are almost-surely bounded. By Theorem 2, the condition in (20) can be
 170 guaranteed as long as a sufficiently large k is selected.

171 3 Experiments

172 In this section, we aim to answer the following questions via empirical results: (1) Does the phantom
 173 gradient form a descent direction in the practical scenario? (2) What is the difference between the

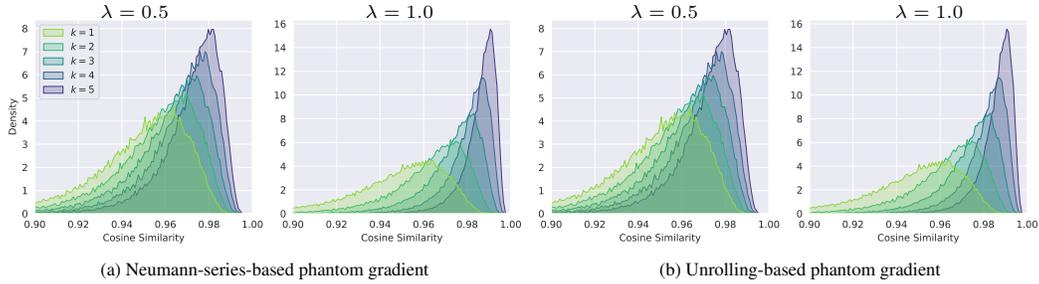


Figure 1: Cosine similarity between the phantom and the exact gradients in the synthetic setting.

174 phantom gradients in the unrolling form and in the Neumann form? (3) How is the phantom gradient
 175 influenced by the hyperparameters k and λ ? (4) How about the memory and computation cost of
 176 the phantom gradient compared with implicit differentiation? (5) Can the phantom gradient work at
 177 large-scale settings?

178 We have provided some theoretical analysis and intuitions to (1), (2), and (3) in Sec. 2.3. Now we
 179 answer (1) and (2) and demonstrate the performance curves under different hyperparameters k and λ
 180 on the CIFAR-10 dataset [10]. Besides, we also study other factors that have potential influences
 181 on the training process of the state-of-the-art implicit models [3, 4]. For (4) and (5), we conduct
 182 experiments on the large-scale datasets, including the CIFAR-10 and ImageNet [22] datasets.

183 We start by introducing two settings of experiments. On the CIFAR-10 dataset, we first use the
 184 MDEQ-Tiny [4] model (170K parameters) as the backbone model in an *ablation setting*. Additionally,
 185 we adopt a single-layer neural network with spectral normalization [23] as the function \mathcal{F} and the
 186 fixed-point iteration as solver of \mathbf{h}^* , which is the *synthetic setting*. Besides, the unrolling-based and
 187 Neumann-series-based phantom gradients are abbreviated to UPG and NPG, respectively.

188 **Precision of the Phantom Gradient.** The precision of the phantom gradient is measured by its
 189 angle (or cosine similarity) against the exact gradient. We discuss its precision in both the synthetic
 190 setting and the ablation setting.

191 In the synthetic setting, the function \mathcal{F} is restricted to be a contractive mapping. Specifically, we
 192 directly set the Lipschitz constant of \mathcal{F} as $L_{\mathbf{h}} = 0.9$, and use 100 fixed-point iterations to solve
 193 the root \mathbf{h}^* of Eq. (1) until the relative error satisfies $\|\mathbf{h} - \mathcal{F}(\mathbf{h}, \mathbf{z})\|/\|\mathbf{h}\| < 10^{-5}$. Here, the exact
 194 gradient is estimated by backpropagation through those fix point iterations, and cross-validated by
 195 implicit differentiation solved with 20 iterations of the Broyden’s method [24]. In our experiment, the
 196 cosine similarity between these two gradient estimates consistently succeeds 0.9999, indicating the
 197 gradient estimate is quite accurate. The cosine similarity between the phantom gradient and the exact
 198 gradient is shown in Fig. 1. It can be seen that the cosine similarity tends to increase as k grows, and
 199 that a small λ tends to slows down the convergence of the phantom gradient, allowing it to explore in
 200 a wider range with regard to its angle against the exact gradient.

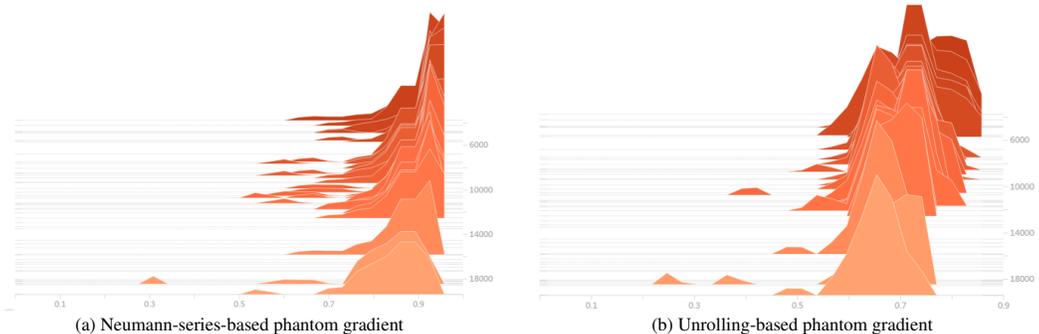


Figure 2: Cosine similarity between the phantom gradient and the exact gradient in the real scenario. The x-axis corresponds to the cosine similarity, and the y-axis to the training steps. This figure characterizes the evolution of the phantom gradient’s precision during the training process.

201 In the ablation setting, the precision of the phantom gradient during the training process is shown in
 202 Fig. 2. The model is trained by implicit differentiation under its official schedule². It can be seen that
 203 in real scenarios, the phantom gradient still provides a descent direction, as indicated by the large
 204 cosine similarity against the exact gradient.

205 **To Pretrain, or not to Pretrain?** To understand the components of training implicit models
 206 via implicit differentiation, we first show a detailed ablation study of the baseline models. All
 207 performances are based on 6 independent runs in the ablation setting, and the average and best
 208 accuracy are reported in Tab. 1.

209 The MDEQ model employs a pretraining stage in
 210 which the model \mathcal{F} is unrolled as a recurrent net-
 211 work. We study the impacts of the pretraining stage,
 212 the Dropout [25] operation, and the optimizer, re-
 213 spectively. It can be seen that the unrolled pretrain-
 214 ing stabilizes the training of the MDEQ model. Re-
 215 moving the pretraining stage leads to a large perfor-
 216 mance drop and apparent training instability among
 217 different runs because the solver cannot obtain an
 218 accurate fixed point \mathbf{h}^* when the model is not ade-
 219 quately trained. This also suggests that the MDEQ
 220 model is a strong baseline for our method to com-
 221 pare with.

222 However, pretraining is not always indispensable
 223 for training implicit models. It introduces a hyperparameter that how many steps should be used in
 224 pretraining. In the later paragraph, we discuss that how the unrolling-based phantom gradient can
 225 circumvent this issue.

226 **Trade-offs between Unrolling and Neumann.** For an exact fixed point \mathbf{h}^* , *i.e.*, $\mathbf{h}^* = \mathcal{F}(\mathbf{h}^*, \mathbf{z})$,
 227 there is no difference between the unrolling-based phantom gradient and Neumann-series-based
 228 one. However, when the numerical error exists in solving \mathbf{h}^* , *i.e.*, $\|\mathbf{h}^* - \mathcal{F}(\mathbf{h}^*, \mathbf{z})\| > 0$, phantom
 229 gradients in the two forms can have different behaviors. As in Fig. 1, the unrolling-based phantom
 230 gradient demonstrates greater robustness and higher tolerance to numerical errors in the backward
 231 computation.

232 We note that a particular benefit of the unrolling-
 233 based phantom gradient is its ability to automati-
 234 cally switch between the pretraining and training
 235 stages for the MDEQ model. When the model is
 236 not sufficiently trained and the solver possibly con-
 237 verges poorly (see [4]), the unrolling-based phantom
 238 gradient defines a forward computational graph that
 239 is essentially equivalent to a shallow recurrent net-
 240 work. In this stage, the phantom gradient serves as
 241 a backpropagation through time (BPTT) algorithm
 242 and hence behaves as in the pretraining stage. Then, as training progresses, the solver becomes more
 243 stable and converges to the fixed point \mathbf{h}^* better. This makes the unrolling-based phantom gradient
 244 behave more like the Neumann-series-based counterpart. Therefore, the unrolled pretraining is
 245 gradually transited into the regular training based on implicit differentiation, and the hyperparameter
 246 tuning of pretraining steps can be waived. We argue that such an ability to adaptively switch training
 247 stages is crucial to the implicit training protocol, which is also supported by the performance gain in
 248 Tab. 1.

249 Although the unrolling-based phantom gradient requires higher memory than implicit differentiation
 250 or the Neumann-series-based one, it does not surpass the peak memory usage in the entire training
 251 process of implicit differentiation due to the pretraining stage. In the ablation setting, the MDEQ
 252 [4] model employs a 10-layer unrolling for pretraining, which actually consumes double memory
 253 compared with a 5-step unrolling scheme (*e.g.*, $\mathcal{A}_{5,0.5}$ in Tab. 1).

Table 1: Ablation settings on CIFAR-10.

Method	Acc(%)
MDEQ-Tiny + Implicit	85.0(85.3)
w/o Pretraining	82.3(84.6)
w/o Dropout	83.7(84.0)
Adam \rightarrow SGD	84.4(84.8)
SGD w/o Pretrainig	81.9(85.6)
UPG ($\mathcal{A}_{5,0.5}$, w/o Dropout)	85.8(86.6)
NPG ($\mathcal{A}_{5,0.5}$, w/o Dropout)	85.6(86.1)
UPG ($\mathcal{A}_{9,0.5}$, w/ Dropout)	86.1(87.3)

Table 2: Complexity comparison. **Mem** means memory cost. $K \gg k \approx 1$, where K corresponds to the solver’s steps and k denotes the unrolling/Neumann steps.

Method	Time	Mem	Peak Mem
Implicit	$\mathcal{O}(K)$	$\mathcal{O}(1)$	$\mathcal{O}(k)$
UPG	$\mathcal{O}(k)$	$\mathcal{O}(k)$	$\mathcal{O}(k)$
NPG	$\mathcal{O}(k)$	$\mathcal{O}(1)$	$\mathcal{O}(1)$

²Code available at <https://github.com/locuslab/mdeq>.

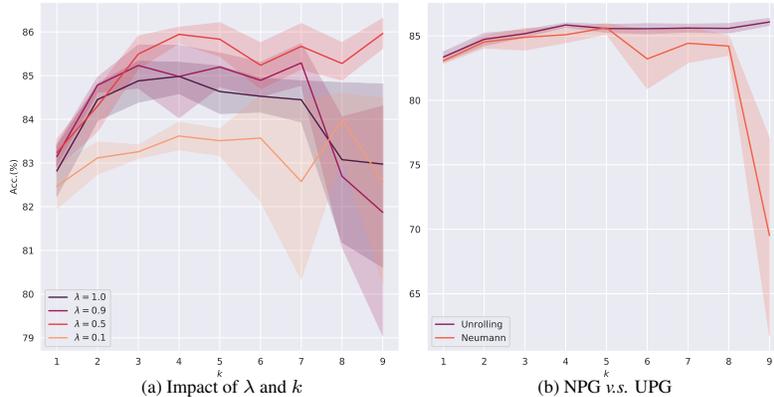


Figure 3: Ablation studies on (a) the hyperparameters λ and k , and (b) two forms of phantom gradient.

254 The performance curves in Fig. 3 further validates the advantages of the unrolling-based phantom
 255 gradient over the Neumann-series-based one and shows the influences of λ and k . In Tab. 2, we
 256 also demonstrate the time and memory complexity for implicit differentiation and the two forms of
 257 phantom gradient. In the following paragraph, we adopt the unrolling-based phantom gradient in the
 258 large-scale experiments.

259 **Phantom Gradient at Scale.** We conduct large-scale experiments to verify the advantages of the
 260 phantom gradient on CIFAR-10 and ImageNet classification benchmark. The results are illustrated
 261 in Tab. 3. Our method matches or surpasses the implicit differentiation training protocol on the
 262 state-of-the-art implicit models with a visible reduction in the training time and comparable or less
 263 peak memory usage.

Table 3: Large-scale experiments on CIFAR-10 and ImageNet classifications. Using phantom gradients, we are able to achieve comparable or better performance in these high-dimensional settings, while being much faster at training.

Task	Method	Params	Acc(%)	Speed	Peak Mem
CIFAR-10	MDEQ + Implicit	10M	93.8	1×	1×
CIFAR-10	MDEQ + UPG $\mathcal{A}_{5,0.5}$	10M	95.0	1.4×	0.5×
ImageNet	MDEQ + Implicit	18M	75.3	1×	1×
ImageNet	MDEQ + UPG $\mathcal{A}_{5,0.5}$	18M	75.1	1.7×	1×

264 4 Related Work

265 **Implicit Models.** Implicit models generalize the recursive forward/backward rules of neural net-
 266 works and characterize their internal mechanism by some pre-specified dynamics. Based on the
 267 dynamics, the implicit models can be broadly categorized into three classes: ODE-based [1, 2], root-
 268 solving-based [3, 4, 8, 5], and optimization-based [26, 27, 28, 29] implicit models. The ODE-based
 269 implicit models [1, 2] treat the iterative update rules of residual networks as the Euler discretization
 270 of an ODE, which could be solved by any black-box ODE solver. The gradient of the differential
 271 equation is calculated using the *adjoint method* [30], in which the adjoint state is obtained by solv-
 272 ing another ODE. The root-solving-based implicit models [3, 4, 8, 5] characterize layers of neural
 273 networks by the process of fixed-point equation solving. The equations are solved by either the
 274 black-box root-finding solver [3, 4] or the fixed-point iteration [5]. The optimization-based implicit
 275 models [26, 27, 28, 29] leverage the optimization programs as layers of neural networks. Previous
 276 work has studied differentiable layers of quadratic programming [26], submodular optimization
 277 [27], and maximum satisfiability (MAXSAT) problems [28]. As for the backward passes, implicit
 278 differentiation is applied to the problem-defining equations of the root-solving-based models [3, 4] or
 279 the KKT conditions of the optimization-based models [26]. As such, the gradient can be obtained by
 280 solving a linear system.

281 In this work, we focus on the root-solving-based implicit models. We differ from previous work in
 282 that we look into the theoretical aspect of the gradient-based algorithm in training implicit models. We
 283 show that besides the precision of the gradient estimate, its condition is also of great significance for

284 the training stability. With these considerations, we show that implicit models of the same architecture
285 could enjoy faster convergence and better generalization ability in practical applications.

286 **Non-End-to-End Optimization in Deep Learning.** Non-end-to-end optimization aims to replace
287 the standard gradient-based training of deep architectures with modular or weakly modular training
288 without the entire forward and backward passes. Currently, there are mainly three research directions
289 in this field, namely, the auxiliary variable methods [31, 32, 33, 34, 35, 36, 37], target propagation
290 [38, 39, 40], and synthetic gradient [41, 42, 43]. The auxiliary variable methods [31, 32, 33, 34,
291 35, 36, 37] formulate the optimization of neural networks as constrained optimization problems, in
292 which the layer-wise activations are considered as trainable auxiliary variables. Then, the equality
293 constraints are relaxed as penalty terms added to the objectives so that the parameters and auxiliary
294 variables can be divided into blocks and thus optimized in parallel. The target propagation method
295 [38, 39, 40] trains each module by having its activations regress to the pre-assigned targets, which are
296 propagated backwards from the downstream modules. Specifically, the auto-encoder architecture is
297 used to reconstruct targets at each layer. Finally, the synthetic gradient method [41, 42, 43] estimates
298 the local gradient of neural networks using auxiliary models, and employ the synthetic gradient in
299 place of the exact gradient to perform parameter update. In this way, the forward and backward
300 passes are decoupled and can be executed in an asynchronous manner.

301 Our work is in line with the non-end-to-end optimization research since we also aims to decouple
302 the forward and backward passes of neural networks. However, we show that finding a reasonable
303 “target” or a precise gradient estimate is not always the first principle in training deep architectures.
304 Our paper paves a path that an inexact but well-conditioned gradient estimate can contribute to both
305 training and generalization of implicit models.

306 **Differentiation through Implicit Dynamics.** Differentiation through certain implicit dynamics is
307 an important aspect in a wide range of research fields, including bilevel optimization [21, 7], meta-
308 learning [44, 45, 46], and sensitivity analysis [47]. Since the gradient (or Jacobian) usually cannot
309 be computed analytically, researchers have to implicitly differentiate the dynamics at the converged
310 point. The formula of the gradient typically contains a term of Jacobian-inverse (or Hessian-inverse),
311 which is computationally prohibitive for large-scale models. (See Eq. (2) in our case.) Herein, several
312 techniques have been developed to approximate the matrix inverse in the previous literature.

313 An intuitive solution is to differentiate through the unrolled steps of a numerical solver of the dynamics
314 [48, 49, 6]. In particular, if a single step is unrolled, it reduces to the well-known *one-step gradient*
315 [50, 44, 51, 46, 52], in which the Jacobian-inverse is simply approximated by an identity matrix.
316 On the contrary, unrolling a small number of steps may induce a bias [7], while the memory and
317 computational cost grows linearly as the number of unrolled steps increases. Towards this issue,
318 Shaban *et al.* [21] propose to truncate the long-term dependencies and differentiate through only the
319 last L steps. Furthermore, if the dynamics has converged to a stationary point, the approximation
320 in Shaban *et al.* [21] is exactly the Neumann approximation of the Jacobian-inverse with the first L
321 terms. Based on this, Lorraine *et al.* [7] choose to directly use the truncated Neumann series as an
322 approximation of the Jacobian-inverse. Besides the unrolling-based methods, optimization-based
323 approaches [53, 45] have also been studied in this field. Since the Jacobian-inverse-vector product
324 can be viewed as solution of a linear system, algorithms like the conjugate gradient method can be
325 used to solve it.

326 5 Conclusion

327 In this work, we question the necessity of rigorously estimating the exact gradient for training implicit
328 models. To back up our claim, we systematically analyze the general condition of a gradient estimate
329 so that the implicit models can be guaranteed to converge to an approximate stationary point of
330 the loss function. Specifically, we give a sufficient condition under which a first-order oracle could
331 always find a descent direction of the loss landscape in the training process. Moreover, we introduce
332 two instantiations of the proposed phantom gradient, based on either the fixed-point unrolling or the
333 Neumann series. The proposed method shows $1.4 \sim 1.7\times$ accelerations with comparable or better
334 performances on large-scale benchmarks. Overall, this paper provides an interesting and practical
335 perspective on training implicit models with theoretical guarantees.

References

- 336
- 337 [1] Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural Ordinary
338 Differential Equations. In *Neural Information Processing Systems (NeurIPS)*, 2018. 1, 8
- 339 [2] Emilien Dupont, Arnaud Doucet, and Yee Whye Teh. Augmented Neural ODEs. In *Neural*
340 *Information Processing Systems (NeurIPS)*, 2019. 1, 8
- 341 [3] Shaojie Bai, J. Zico Kolter, and Vladlen Koltun. Deep Equilibrium Models. In *Neural*
342 *Information Processing Systems (NeurIPS)*, 2019. 1, 2, 3, 6, 8
- 343 [4] Shaojie Bai, Vladlen Koltun, and J. Zico Kolter. Multiscale Deep Equilibrium Models. In
344 *Neural Information Processing Systems (NeurIPS)*, pages 5238–5250, 2020. 1, 6, 7, 8
- 345 [5] Samy Wu Fung, Howard Heaton, Qiuwei Li, Daniel McKenzie, Stanley J. Osher, and Wotao Yin.
346 Fixed Point Networks: Implicit Depth Models with Jacobian-Free Backprop. *arXiv preprint*
347 *arXiv:2103.12803*, 2021. 1, 8
- 348 [6] Luca Franceschi, Michele Donini, Paolo Frasconi, and Massimiliano Pontil. Forward and
349 Reverse Gradient-Based Hyperparameter Optimization. In *International Conference on Machine*
350 *Learning (ICML)*, pages 1165–1173, 2017. 1, 9
- 351 [7] Jonathan Lorraine, Paul Vicol, and David Duvenaud. Optimizing Millions of Hyperparameters
352 by Implicit Differentiation. In *International Conference on Artificial Intelligence and Statistics*
353 *(AISTATS)*, pages 1540–1552, 2020. 1, 4, 9
- 354 [8] Ezra Winston and J. Zico Kolter. Monotone operator equilibrium networks. In *Neural Informa-*
355 *tion Processing Systems (NeurIPS)*, pages 10718–10728, 2020. 1, 8
- 356 [9] Cheng Lu, Jianfei Chen, Chongxuan Li, Qiuhaio Wang, and Jun Zhu. Implicit normalizing flows.
357 In *International Conference on Learning Representations (ICLR)*, 2021. 1
- 358 [10] Alex Krizhevsky, Geoffrey Hinton, et al. Learning Multiple Layers of Features from Tiny
359 Images. Technical report, Citeseer, 2009. 2, 6
- 360 [11] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale
361 hierarchical image database. In *CVPR*, pages 248–255. Ieee, 2009. 2
- 362 [12] Samy Wu Fung, Howard Heaton, Qiuwei Li, Daniel Mckenzie, S. Osher, and W. Yin. Fixed
363 point networks: Implicit depth models with jacobian-free backprop. *ArXiv*, abs/2103.12803,
364 2021. 2
- 365 [13] Chirag Pabbaraju, Ezra Winston, and J. Zico Kolter. Estimating lipschitz constants of monotone
366 deep equilibrium models. In *International Conference on Learning Representations (ICLR)*,
367 2021. 2
- 368 [14] Max Revay, Ruigang Wang, and Ian R Manchester. Lipschitz bounded equilibrium networks.
369 *arXiv:2010.01732*, 2020. 2
- 370 [15] Xavier Gastaldi. Shake-Shake regularization of 3-branch residual networks. In *International*
371 *Conference on Learning Representations Workshop Track*, 2017. 3
- 372 [16] Guozhong An. The Effects of Adding Noise During Backpropagation Training on a Generaliza-
373 tion Performance. *Neural Computation*, 8(3):643–674, 1996. 3
- 374 [17] Zhanxing Zhu, Jingfeng Wu, Bing Yu, Lei Wu, and Jinwen Ma. The Anisotropic Noise in
375 Stochastic Gradient Descent: Its Behavior of Escaping from Sharp Minima and Regularization
376 Effects. In *International Conference on Machine Learning (ICML)*, pages 7654–7663, 2019. 3
- 377 [18] Jingfeng Wu, Wenqing Hu, Haoyi Xiong, Jun Huan, Vladimir Braverman, and Zhanxing Zhu.
378 On the Noisy Gradient Descent that Generalizes as SGD. In *International Conference on*
379 *Machine Learning (ICML)*, pages 10367–10376, 2020. 3
- 380 [19] Arthur Jacot, Franck Gabriel, and Clement Hongler. Neural Tangent Kernel: Convergence
381 and Generalization in Neural Networks. In *Neural Information Processing Systems (NeurIPS)*,
382 pages 8571–8580, 2018. 3

- 383 [20] Sanjeev Arora, Simon S Du, Wei Hu, Zhiyuan Li, Russ R Salakhutdinov, and Ruosong Wang.
384 On Exact Computation with an Infinitely Wide Neural Net. In *Neural Information Processing*
385 *Systems (NeurIPS)*, pages 8139–8148, 2019. 3
- 386 [21] Amirreza Shaban, Ching-An Cheng, Nathan Hatch, and Byron Boots. Truncated Back-
387 propagation for Bilevel Optimization. In *International Conference on Artificial Intelligence and*
388 *Statistics (AISTATS)*, pages 1723–1732, 2019. 4, 9
- 389 [22] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng
390 Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei.
391 ImageNet Large Scale Visual Recognition Challenge. *International Journal on Computer Vision*
392 *(IJCV)*, 115(3):211–252, 2015. 6
- 393 [23] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral Normaliza-
394 tion for Generative Adversarial Networks. In *International Conference on Learning Representations*
395 *(ICLR)*, 2018. 6
- 396 [24] Charles G Broyden. A Class of Methods for Solving Nonlinear Simultaneous Equations.
397 *Mathematics of computation*, 19(92):577–593, 1965. 6
- 398 [25] Geoffrey E Hinton, Nitish Srivastava, Alex Krizhevsky, Ilya Sutskever, and Ruslan R Salakhut-
399 dinov. Improving neural networks by preventing co-adaptation of feature detectors. *arXiv*
400 *preprint arXiv:1207.0580*, 2012. 7
- 401 [26] Brandon Amos and J. Zico Kolter. OptNet: Differentiable Optimization as a Layer in Neural
402 Networks. In *International Conference on Machine Learning (ICML)*, pages 136–145, 2017. 8
- 403 [27] Josip Djolonga and Andreas Krause. Differentiable Learning of Submodular Models. *Neural*
404 *Information Processing Systems (NeurIPS)*, pages 1013–1023, 2017. 8
- 405 [28] Po-Wei Wang, Priya Donti, Bryan Wilder, and Zico Kolter. SATNet: Bridging deep learning
406 and logical reasoning using a differentiable satisfiability solver. In *International Conference on*
407 *Machine Learning (ICML)*, pages 6545–6554, 2019. 8
- 408 [29] Marin Vlastelica, Anselm Paulus, Vít Musil, Georg Martius, and Michal Rolínek. Differentiation
409 of Blackbox Combinatorial Solvers. In *International Conference on Learning Representations*
410 *(ICLR)*, 2020. 8
- 411 [30] VG Boltyanskiy, Revaz V Gamkrelidze, YEF Mishchenko, and LS Pontryagin. Mathematical
412 theory of optimal processes. 1962. 8
- 413 [31] Miguel Carreira-Perpinan and Weiran Wang. Distributed optimization of deeply nested systems.
414 In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 10–19,
415 2014. 9
- 416 [32] Gavin Taylor, Ryan Burmeister, Zheng Xu, Bharat Singh, Ankit Patel, and Tom Goldstein.
417 Training Neural Networks Without Gradients: A Scalable ADMM Approach. In *International*
418 *Conference on Machine Learning (ICML)*, pages 2722–2731, 2016. 9
- 419 [33] Ziming Zhang, Yuting Chen, and Venkatesh Saligrama. Efficient Training of Very Deep Neural
420 Networks for Supervised Hashing. In *IEEE Conference on Computer Vision and Pattern*
421 *Recognition (CVPR)*, pages 1487–1495, 2016. 9
- 422 [34] Ziming Zhang and Matthew Brand. Convergent Block Coordinate Descent for Training
423 Tikhonov Regularized Deep Neural Networks. In *Neural Information Processing Systems*
424 *(NeurIPS)*, 2017. 9
- 425 [35] Jinshan Zeng, Shikang Ouyang, Tim Tsz-Kit Lau, Shaobo Lin, and Yuan Yao. Global con-
426 vergence in deep learning with variable splitting via the Kurdyka-Iojasiewicz property. *arXiv*
427 *preprint arXiv:1803.00225*, 2018. 9
- 428 [36] Jia Li, Cong Fang, and Zhouchen Lin. Lifted Proximal Operator Machines. In *Association for*
429 *the Advancement of Artificial Intelligence (AAAI)*, pages 4181–4188, 2019. 9

- 430 [37] Fangda Gu, Armin Askari, and Laurent El Ghaoui. Fenchel Lifted Networks: A Lagrange
431 Relaxation of Neural Network Training. In *International Conference on Artificial Intelligence
432 and Statistics (AISTATS)*, pages 3362–3371, 2020. [9](#)
- 433 [38] Yoshua Bengio. How Auto-Encoders Could Provide Credit Assignment in Deep Networks via
434 Target Propagation. *arXiv preprint arXiv:1407.7906*, 2014. [9](#)
- 435 [39] Dong-Hyun Lee, Saizheng Zhang, Asja Fischer, and Yoshua Bengio. Difference Target Propa-
436 gation. In *International Conference on Learning Representations (ICLR)*, page 498–515, 2015.
437 [9](#)
- 438 [40] Alexander Meulemans, Francesco Carzaniga, Johan Suykens, João Sacramento, and Benjamin F.
439 Grewe. A Theoretical Framework for Target Propagation. In *Neural Information Processing
440 Systems (NeurIPS)*, pages 20024–20036, 2020. [9](#)
- 441 [41] Max Jaderberg, Wojciech Marian Czarnecki, Simon Osindero, Oriol Vinyals, Alex Graves,
442 David Silver, and Koray Kavukcuoglu. Decoupled Neural Interfaces using Synthetic Gradients.
443 In *International Conference on Machine Learning (ICML)*, pages 1627–1635, 2017. [9](#)
- 444 [42] Wojciech Marian Czarnecki, Grzegorz Świrszcz, Max Jaderberg, Simon Osindero, Oriol Vinyals,
445 and Koray Kavukcuoglu. Understanding Synthetic Gradients and Decoupled Neural Interfaces.
446 In *International Conference on Machine Learning (ICML)*, pages 904–912, 2017. [9](#)
- 447 [43] Benjamin James Lansdell, Prashanth Ravi Prakash, and Konrad Paul Kording. Learning to
448 solve the credit assignment problem. In *International Conference on Learning Representations
449 (ICLR)*, 2020. [9](#)
- 450 [44] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-Agnostic Meta-Learning for Fast
451 Adaptation of Deep Networks. In *International Conference on Machine Learning (ICML)*,
452 pages 1126–1135, 2017. [9](#)
- 453 [45] Aravind Rajeswaran, Chelsea Finn, Sham M Kakade, and Sergey Levine. Meta-Learning with
454 Implicit Gradients. In *Neural Information Processing Systems (NeurIPS)*, 2019. [9](#)
- 455 [46] Xin-Yu Zhang, Taihong Xiao, Haolin Jia, Ming-Ming Cheng, and Ming-Hsuan Yang. Semi-
456 Supervised Learning with Meta-Gradient. In *International Conference on Artificial Intelligence
457 and Statistics (AISTATS)*, pages 73–81, 2021. [9](#)
- 458 [47] J Frédéric Bonnans and Alexander Shapiro. *Perturbation analysis of optimization problems*.
459 Springer Science & Business Media, 2013. [9](#)
- 460 [48] Justin Domke. Generic Methods for Optimization-Based Modeling. In *International Conference
461 on Artificial Intelligence and Statistics (AISTATS)*, pages 318–326, 2012. [9](#)
- 462 [49] Dougal Maclaurin, David Duvenaud, and Ryan Adams. Gradient-based Hyperparameter
463 Optimization through Reversible Learning. In *International Conference on Machine Learning
464 (ICML)*, pages 2113–2122, 2015. [9](#)
- 465 [50] Jelena Luketina, Mathias Berglund, Klaus Greff, and Tapani Raiko. Scalable Gradient-Based
466 Tuning of Continuous Regularization Hyperparameters. In *International Conference on Machine
467 Learning (ICML)*, pages 2952–2960, 2016. [9](#)
- 468 [51] Hanxiao Liu, Karen Simonyan, and Yiming Yang. DARTS: Differentiable Architecture Search.
469 *International Conference on Learning Representations (ICLR)*, 2018. [9](#)
- 470 [52] Zhengyang Geng, Meng-Hao Guo, Hongxu Chen, Xia Li, Ke Wei, and Zhouchen Lin. Is
471 Attention Better Than Matrix Decomposition? In *International Conference on Learning
472 Representations (ICLR)*, 2021. [9](#)
- 473 [53] Fabian Pedregosa. Hyperparameter Optimization with Approximate Gradient. In *International
474 Conference on Machine Learning (ICML)*, pages 737–746, 2016. [9](#)

475 **Checklist**

- 476 1. For all authors...
- 477 (a) Do the main claims made in the abstract and introduction accurately reflect the paper's
478 contributions and scope? [Yes]
- 479 (b) Did you describe the limitations of your work? [No]
- 480 (c) Did you discuss any potential negative societal impacts of your work? [N/A] Not
481 applicable.
- 482 (d) Have you read the ethics review guidelines and ensured that your paper conforms to
483 them? [Yes]
- 484 2. If you are including theoretical results...
- 485 (a) Did you state the full set of assumptions of all theoretical results? [Yes] See the
486 assumptions in Theorem 1 to 3.
- 487 (b) Did you include complete proofs of all theoretical results? [Yes] See the appendix.
- 488 3. If you ran experiments...
- 489 (a) Did you include the code, data, and instructions needed to reproduce the main experi-
490 mental results (either in the supplemental material or as a URL)? [Yes]
- 491 (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they
492 were chosen)? [Yes]
- 493 (c) Did you report error bars (e.g., with respect to the random seed after running experi-
494 ments multiple times)? [Yes]
- 495 (d) Did you include the total amount of compute and the type of resources used (e.g., type
496 of GPUs, internal cluster, or cloud provider)? [Yes]
- 497 4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
- 498 (a) If your work uses existing assets, did you cite the creators? [Yes] See Sec. 4.
- 499 (b) Did you mention the license of the assets? [Yes]
- 500 (c) Did you include any new assets either in the supplemental material or as a URL? [Yes]
501 See supplemental materials.
- 502 (d) Did you discuss whether and how consent was obtained from people whose data you're
503 using/curating? [N/A]
- 504 (e) Did you discuss whether the data you are using/curating contains personally identifiable
505 information or offensive content? [N/A]
- 506 5. If you used crowdsourcing or conducted research with human subjects...
- 507 (a) Did you include the full text of instructions given to participants and screenshots, if
508 applicable? [N/A]
- 509 (b) Did you describe any potential participant risks, with links to Institutional Review
510 Board (IRB) approvals, if applicable? [N/A]
- 511 (c) Did you include the estimated hourly wage paid to participants and the total amount
512 spent on participant compensation? [N/A]