# On Training Implicit Models

**Zhengyang Geng**[1,2]* **Xin-Yu Zhang**[2]* **Shaojie Bai**[3] **Yisen Wang**[2] **Zhouchen Lin**[2,4]†

[1]Zhejiang Lab, China [2]Key Lab. of Machine Perception, School of EECS, Peking University
[3]Carnegie Mellon University [4]Pazhou Lab, China

## Abstract

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

## 1 Introduction

Conventional neural networks are typically constructed by explicitly stacking multiple linear and non-linear operators in a feed-forward manner. Recently, the implicitly-defined models [1, 2, 3, 4, 5] have attracted increasing attentions and are able to match the state-of-the-art results by explicit models on several vision [3, 6], language [2] and graph [4] tasks. These works treat the evolution of the intermediate hidden states as a certain form of dynamics, such as fixed-point equations [2, 3] or ordinary differential equations (ODEs) [1, 7], which represents infinite latent states. The forward passes of implicit models are therefore formulated as solving the underlying dynamics, by either black-box ODE solvers [1, 7] or root-finding algorithms [2, 3]. As for the backward passes, however, directly differentiating through the forward pass trajectories could induce a heavy memory overhead [8, 9]. To this end, researchers have developed memory-efficient backpropagation via implicit differentiation, such as solving a Jacobian-based linear fixed-point equation for the backward pass of deep equilibrium models (DEQs) [2], which eventually makes the backpropagation trajectories independent of the forward passes, allowing one to train implicit models with essentially constant memory consumption, as we only need to store the final output and the layer itself without saving any intermediate states. However, in order to estimate the exact gradient promised by the implicit differentiation, implicit models have to rely on black-box solvers for backward passes, *e.g.,* ODE solvers or root-solving algorithms, whose iterative nature usually makes the gradient computation very costly in practice, even involving over 30 steps for large-scale DEQ models.

This work investigates fast approximate gradients for training implicit models. We found that a first-order oracle that produces good gradient estimates is enough to efficiently and effectively train implicit models, circumventing laboriously computing the exact gradient as in prior arts [2, 3, 4,

*Equal contribution
†Corresponding author, zlin@pku.edu.cn

35th Conference on Neural Information Processing Systems (NeurIPS 2021).

10, 11]. We develop a framework in which a balanced trade-off is made between the precision and conditioning of the gradient estimate. Specifically, we provide the general condition under which the phantom gradient can provide an ascent direction of the loss landscape. We further propose two instantiations of phantom gradients in the context of DEQ models, which are based on the the damped fixed-point unrolling and the Neumann series, respectively. Importantly, we show that our proposed instantiations satisfy the theoretical condition, and that the stochastic gradient descent (SGD) algorithm based on the phantom gradient enjoys a sound convergence property as long as the relevant hyperparameters, *e.g.,* the damping factor, are wisely selected. Note that our proposed method only directly affects, and thus accelerates, the backward formulation of these implicit models, leaving the forward pass formulation (*i.e.,* the root-solving process) and inference-time behavior mostly intact. We conduct an extensive set of synthetic, ablation, and large-scale experiments to both analyze the theoretical properties of the phantom gradient and validate its performances on various tasks, such as ImageNet [12] classification and Wikitext-103 [13] language modeling.

Overall, our results suggest that: 1) the phantom gradient estimates an ascent direction; 2) it is applicable to large-scale tasks and is capable of achieving a strong performance which is comparable with or even better than using exact gradients; and 3) it significantly shortens the total training time needed for implicit models roughly by a factor of $1.4 \sim 1.7\times$, and even accelerates the backward passes by astonishing $12\times$ on ImageNet. We believe that our results provide strong evidence for effectively training implicit models with the lightweight phantom gradient.

## 2 Method

### 2.1 Inspection of Implicit Differentiation

In this work, we primarily focus on the formulation of implicit models based on root-solving, represented by the DEQ models [2]. The table of notations is arranged in the Appendix. Specifically, given an equilibrium module $\mathcal{F}$, the output of the implicit model is characterized as the solution $\boldsymbol{h}^*$ to the following fixed-point equation:

$$\boldsymbol{h}^* = \mathcal{F}(\boldsymbol{h}^*, \boldsymbol{z}), \tag{1}$$

where $\boldsymbol{z} \in \mathbb{R}^{d_{\boldsymbol{u}}+d_{\boldsymbol{\theta}}}$ is the union of the module's input $\boldsymbol{u} \in \mathbb{R}^{d_{\boldsymbol{u}}}$ and parameters $\boldsymbol{\theta} \in \mathbb{R}^{d_{\boldsymbol{\theta}}}$, *i.e.,* $\boldsymbol{z}^\top = [\boldsymbol{u}^\top, \boldsymbol{\theta}^\top]$. Here, $\boldsymbol{u}$ is usually the projection of the original data point $\boldsymbol{x} \in \mathbb{R}^{d_{\boldsymbol{x}}}$, *e.g.,* $\boldsymbol{u} = \mathcal{M}(\boldsymbol{x})$. In this section, we assume $\mathcal{F}$ is a contraction mapping *w.r.t.* $\boldsymbol{h}$ so that its Lipschitz constant $L_{\boldsymbol{h}}$ *w.r.t.* $\boldsymbol{h}$ is less than one, *i.e.,* $L_{\boldsymbol{h}} < 1$, a setting that has been analyzed in recent works [14, 15][1].

To backpropagate through the module induced by Eq. (1), we need to calculate the Jacobian matrix of $\boldsymbol{h}^*$ *w.r.t.* the input $\boldsymbol{z}$. By Implicit Function Theorem (IFT), we have

$$\frac{\partial \boldsymbol{h}^*}{\partial \boldsymbol{z}} = \frac{\partial \mathcal{F}}{\partial \boldsymbol{z}}\bigg|_{\boldsymbol{h}^*} \left(\boldsymbol{I} - \frac{\partial \mathcal{F}}{\partial \boldsymbol{h}}\bigg|_{\boldsymbol{h}^*}\right)^{-1} \tag{2}$$

The equilibrium point $\boldsymbol{h}^*$ of Eq. (1) is then passed to a post-processing function $\mathcal{G}$ to obtain a prediction $\hat{\boldsymbol{y}} = \mathcal{G}(\boldsymbol{h}^*)$. In the generic learning scenario, the training objective is the following expected loss:

$$\mathcal{R}(\boldsymbol{\theta}) = \mathbb{E}_{(\boldsymbol{x},\boldsymbol{y})\sim\mathcal{P}}\left[\mathcal{L}(\hat{\boldsymbol{y}}(\boldsymbol{x}, \boldsymbol{\theta}), \boldsymbol{y})\right], \tag{3}$$

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

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

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

[1] Note that the contraction condition could be largely relaxed to the one that the spectral radius of $\partial\mathcal{F}/\partial\boldsymbol{h}^*$ on the given data is less than 1, *i.e.,* $\rho(\partial\mathcal{F}/\partial\boldsymbol{h}^*) < 1$, as indicated by the well-posedness condition in [5].

### 2.2 Motivation

The most intriguing part lies in the Jacobian-inverse term, *i.e.,* $(\boldsymbol{I} - \partial \mathcal{F}/\partial \boldsymbol{h})^{-1}$. Computing the inverse term by brute force is intractable due to the $\mathcal{O}(n^3)$ complexity. Previous implicit models [2] approach this by solving a linear system involving a Jacobian-vector product iteratively via a gradient solver, introducing over 30 Broyden [16] iterations in the backward pass. However, the Jacobian matrix can exceed $10^6 \times 10^6$ in the real scenarios, leading to a prohibitive cost in computing the exact gradients. For example, training a small-scale state-of-the-art implicit model on ImageNet can consume weeks using 8 GPUs while training explicit models usually takes several days, demonstrating that pursuing the exact gradient severely slows down the training process of implicit models compared with explicit models.

Secondly, because of the inverse operation, we cast doubt on the conditioning of the gradient and further the stability of training process from the numerical aspect. The Jacobian-inverse can be numerically unstable when encountering the ill-conditioned issue. The conditioning problem might lead to the training stability issues studied in the recent work [17] as well.

Plus, the Jacobian-inverse is calculated to update model parameters with their gradients, but the exact gradient is not always optimal in model training. For example, previous research has used a moderate gradient noise as a regularization approach [18], which has been shown to play a central role in escaping poor local minima and improving generalization ability [19, 20, 21].

The concerns and observations motivate us to rethink the possibility of replacing the Jacobian-inverse term in the standard implicit differentiation with a cheaper and more stable counterpart. We believe that a precise gradient estimate is not always required, especially for a black-box layer like those in the implicit models. Hence this work designs an inexact, theoretically sound, and practically efficient gradient estimate for training implicit models under various settings. We name the proposed gradient estimate as the *phantom gradient*.

Suppose the Jacobian $\partial \boldsymbol{h}^*/\partial \boldsymbol{\theta}$ is replaced with a matrix $\boldsymbol{A}$, and the corresponding phantom gradient is defined as

$$\widehat{\frac{\partial \mathcal{L}}{\partial \boldsymbol{\theta}}} := \boldsymbol{A} \, \frac{\partial \mathcal{L}}{\partial \boldsymbol{h}}. \tag{5}$$

Next, we give the general condition on $\boldsymbol{A}$ so that the phantom gradient can be guaranteed valid for optimization (Sec. 2.3), and provide two concrete instantiations of $\boldsymbol{A}$ based on either damped fixed-point unrolling or the Neumann series (Sec. 2.4).

### 2.3 General Condition on the Phantom Gradient

The following theorem formulates a sufficient condition that the phantom gradient gives an ascent direction of the loss landscape. Please refer to the Appendix for the proof.

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

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

*then the phantom gradient provides an ascent direction of the function $\mathcal{L}$, i.e.,*

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

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

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

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

### 2.4 Instantiations of the Phantom Gradient

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

Suppose we hope to differentiate through an implicit dynamic, *e.g.,* either a root-solving process or an optimization problem. In the context of hyperparameter optimization (HO), previous solutions include differentiating through the unrolled steps of the dynamics [22] or employing the Neumann series of the Jacobian-inverse [9]. In our case, if we solve the root of Eq. (1) via the fixed-point iteration:

$$\boldsymbol{h}_{t+1} = \mathcal{F}(\boldsymbol{h}_t, \boldsymbol{z}), \quad t = 0, 1, \cdots, T-1, \tag{9}$$

then by differentiating through the unrolled steps of Eq. (9), we have

$$\frac{\partial \boldsymbol{h}_T}{\partial \boldsymbol{\theta}} = \sum_{t=0}^{T-1} \frac{\partial \mathcal{F}}{\partial \boldsymbol{\theta}}\bigg|_{\boldsymbol{h}_t} \prod_{s=t+1}^{T-1} \frac{\partial \mathcal{F}}{\partial \boldsymbol{h}}\bigg|_{\boldsymbol{h}_s}. \tag{10}$$

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

$$\boldsymbol{I} + \frac{\partial \mathcal{F}}{\partial \boldsymbol{h}} + \left(\frac{\partial \mathcal{F}}{\partial \boldsymbol{h}}\right)^2 + \left(\frac{\partial \mathcal{F}}{\partial \boldsymbol{h}}\right)^3 + \cdots. \tag{11}$$

Notably, computing the Jacobian $\partial\boldsymbol{h}^*/\partial\boldsymbol{\theta}$ using the Neumann series in (11) is equivalent to differentiating through the unrolled steps of Eq. (9) at the exact equilibrium point $\boldsymbol{h}^*$ and taking the limit of infinite steps [9].

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

$$\boldsymbol{h}_{t+1} = \mathcal{F}_\lambda(\boldsymbol{h}_t, \boldsymbol{z}) = \lambda\mathcal{F}(\boldsymbol{h}_t, \boldsymbol{z}) + (1-\lambda)\boldsymbol{h}_t, \quad t = 0, 1, \cdots, T-1. \tag{12}$$

Differentiating through the unrolled steps of Eq. (12), Eq. (10) is adapted as

$$\frac{\partial \boldsymbol{h}_T}{\partial \boldsymbol{\theta}} = \lambda \sum_{t=0}^{T-1} \frac{\partial \mathcal{F}}{\partial \boldsymbol{\theta}}\bigg|_{\boldsymbol{h}_t} \prod_{s=t+1}^{T-1} \left(\lambda \frac{\partial \mathcal{F}}{\partial \boldsymbol{h}}\bigg|_{\boldsymbol{h}_s} + (1-\lambda)\,\boldsymbol{I}\right). \tag{13}$$

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

$$\lambda\left(\boldsymbol{I} + \boldsymbol{B} + \boldsymbol{B}^2 + \boldsymbol{B}^3 + \cdots\right), \quad \text{where } \boldsymbol{B} = \lambda\frac{\partial \mathcal{F}}{\partial \boldsymbol{h}} + (1-\lambda)\boldsymbol{I}. \tag{14}$$

The next theorem shows that under mild conditions, the Jacobian from the damped unrolling in Eq. (13) converges to the exact Jacobian and the Neumann series in (14) converges to the Jacobian-inverse $(\boldsymbol{I} - \partial\mathcal{F}/\partial\boldsymbol{h})^{-1}$ as well. The proof can be found in the Appendix.

**Theorem 2.** *Suppose the Jacobian $\partial\mathcal{F}/\partial\boldsymbol{h}$ is a contraction mapping. Then,*

- *(i) the Neumann series in (14) converges to the Jacobian-inverse $(\boldsymbol{I} - \partial\mathcal{F}/\partial\boldsymbol{h})^{-1}$; and*
- *(ii) if the function $\mathcal{F}$ is continuously differentiable w.r.t. both $\boldsymbol{h}$ and $\boldsymbol{\theta}$, the sequence in Eq. (13) converges to the exact Jacobian $\partial\boldsymbol{h}^*/\partial\boldsymbol{\theta}$ as $T \to \infty$, i.e.,*

$$\lim_{T\to\infty} \frac{\partial \boldsymbol{h}_T}{\partial \boldsymbol{\theta}} = \frac{\partial \mathcal{F}}{\partial \boldsymbol{\theta}}\bigg|_{\boldsymbol{h}^*} \left(\boldsymbol{I} - \frac{\partial \mathcal{F}}{\partial \boldsymbol{h}}\bigg|_{\boldsymbol{h}^*}\right)^{-1}. \tag{15}$$

However, as discussed in Sec. 2.2, it is unnecessary to compute the exact gradient with infinite terms. In the following context, we introduce two instantiations of the phantom gradient based on the finite-term truncation of Eq. (13) or (14).

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

$$\boldsymbol{A}^{\text{unr}}_{k,\lambda} = \lambda \sum_{t=0}^{k-1} \frac{\partial \mathcal{F}}{\partial \boldsymbol{\theta}}\bigg|_{\boldsymbol{h}_t} \prod_{s=t+1}^{k-1} \left(\lambda \frac{\partial \mathcal{F}}{\partial \boldsymbol{h}}\bigg|_{\boldsymbol{h}_s} + (1-\lambda)\,\boldsymbol{I}\right). \tag{16}$$

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

$$\boldsymbol{A}_{k,\lambda}^{\text{neu}} = \lambda \left.\frac{\partial \mathcal{F}}{\partial \boldsymbol{\theta}}\right|_{\boldsymbol{h}^*} \left(\boldsymbol{I} + \boldsymbol{B} + \boldsymbol{B}^2 + \cdots + \boldsymbol{B}^{k-1}\right), \quad \text{where } \boldsymbol{B} = \lambda \left.\frac{\partial \mathcal{F}}{\partial \boldsymbol{h}}\right|_{\boldsymbol{h}^*} + (1-\lambda)\boldsymbol{I}. \tag{17}$$

Note that both the initial point of the fixed-point iteration (*i.e.,* $\boldsymbol{h}_0$ in Eq. (16)) and the point at which the Neumann series is evaluated (*i.e.,* $\boldsymbol{h}^*$ in (17)) are the solution of the root-finding solver.

According to Theorem 2, the matrix $\boldsymbol{A}$ defined by either Eq. (16) or (17) converges to the exact Jacobian $\partial \boldsymbol{h}^*/\partial \boldsymbol{\theta}$ as $k \to \infty$ for any $\lambda \in (0, 1]$. Therefore, by Theorem 2, the condition in (6) can be satisfied if a sufficiently large step $k$ is selected, since

$$\left\| \boldsymbol{A}\left(\boldsymbol{I} - \frac{\partial \mathcal{F}}{\partial \boldsymbol{h}}\right) - \frac{\partial \mathcal{F}}{\partial \boldsymbol{\theta}} \right\| \leq (1 + L_{\boldsymbol{h}}) \left\| \boldsymbol{A} - \frac{\partial \mathcal{F}}{\partial \boldsymbol{\theta}}\left(\boldsymbol{I} - \frac{\partial \mathcal{F}}{\partial \boldsymbol{h}}\right)^{-1} \right\|. \tag{18}$$

Next, we characterize the impact of the two hyperparameters, *i.e.,* $k$ and $\lambda$, on the precision and conditioning of $\boldsymbol{A}$. Take the Neumann-series-based phantom gradient (Eq. (17)) as an example.

(i) On the precision of the phantom gradient,
- a large $k$ makes the gradient estimate more accurate, as higher-order terms of the Neumann series are included; while
- a small $\lambda$ slows down the convergence of the Neumann series because the norm $\|\boldsymbol{B}\|$ increases as $\lambda$ decreases.

(ii) On the conditioning of the phantom gradient,
- a large $k$ impairs the conditioning of $\boldsymbol{A}$ since the condition number of $\boldsymbol{B}^k$ grows exponentially as $k$ increases; while
- a small $\lambda$ helps maintain a small condition number of $\boldsymbol{A}$ because the singular values of $\partial \mathcal{F}/\partial \boldsymbol{h}$ are "smoothed" by the identity matrix.

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

### 2.5 Convergence Theory

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

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

$$\left\| \widehat{\frac{\partial \mathcal{L}}{\partial \boldsymbol{\theta}}} - \frac{\partial \mathcal{L}}{\partial \boldsymbol{\theta}} \right\| \leq \epsilon, \quad \textit{almost surely}, \tag{19}$$

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

$$\mathbb{E}\left[\frac{\sum_{\tau=1}^{T} \eta_\tau \|\nabla \mathcal{R}(\boldsymbol{\theta}_\tau)\|^2}{\sum_{\tau=1}^{T} \eta_\tau}\right] \leq \mathcal{O}\left(\epsilon + \frac{\log T}{\sqrt{T}}\right). \tag{20}$$

**Remark 2.** Consider the condition in (19):

$$\left\| \widehat{\frac{\partial \mathcal{L}}{\partial \boldsymbol{\theta}}} - \frac{\partial \mathcal{L}}{\partial \boldsymbol{\theta}} \right\| \leq \left\| \boldsymbol{A} - \frac{\partial \mathcal{F}}{\partial \boldsymbol{\theta}}\left(\boldsymbol{I} - \frac{\partial \mathcal{F}}{\partial \boldsymbol{h}}\right)^{-1} \right\| \left\| \frac{\partial \mathcal{L}}{\partial \boldsymbol{h}} \right\|. \tag{21}$$

Suppose the gradient $\partial \mathcal{L}/\partial \boldsymbol{h}$ is almost surely bounded. By Theorem 2, the condition in (19) can be guaranteed as long as a sufficiently large $k$ is selected.

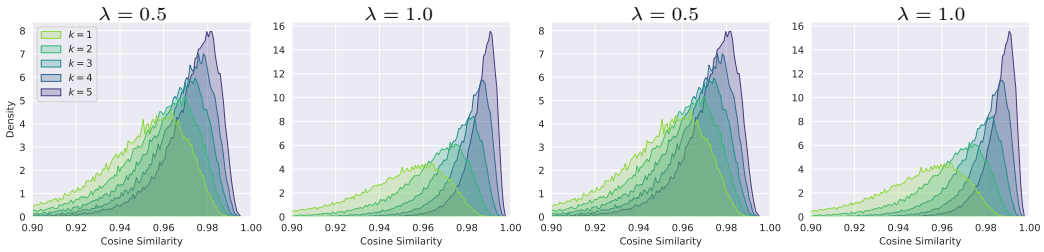


(a) Neumann-series-based phantom gradient.

(b) Unrolling-based phantom gradient.

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

## 3 Experiments

In this section, we aim to answer the following questions via empirical results: (1) How precise is the phantom gradient? (2) What is the difference between the unrolling-based and the Neumann-series-based phantom gradients? (3) How is the phantom gradient influenced by the hyperparameters $k$ and $\lambda$? (4) How about the computation cost of the phantom gradient compared with implicit differentiation? (5) Can the phantom gradient work at large-scale settings for various tasks?

We have provided some theoretical analysis and intuitions to (1), (2), and (3) in Sec. 2.4. Now we answer (1) and (2) and demonstrate the performance curves under different hyperparameters $k$ and $\lambda$ on CIFAR-10 [23]. Besides, we also study other factors that have potential influences on the training process of the state-of-the-art implicit models [2, 3] like pretraining. For (4) and (5), we conduct experiments on large-scale datasets to highlight the ultra-fast speed and competitive performances, including image classification on ImageNet [12] and language modeling on Wikitext-103 [13].

We start by introducing two experiment settings. We adopt a single-layer neural network with spectral normalization [24] as the function $\mathcal{F}$ and fixed-point iterations as the equilibrium solver, which is the *synthetic setting*. Additionally, on the CIFAR-10 dataset, we use the MDEQ-Tiny model [3] (170K parameters) as the backbone model, denoted as the *ablation setting*. Besides, the unrolling-based and the Neumann-series-based phantom gradients are abbreviated to UPG and NPG, respectively.

**Precision of the Phantom Gradient.** The precision of the phantom gradient is measured by its angle against the exact gradient, indicated by the cosine similarity between the two. We discuss its precision in both the synthetic setting and the ablation setting. The former is under the static and randomly generated weights, while the latter provides characterization of the training dynamics.

In the synthetic setting, the function $\mathcal{F}$ is restricted to be a contraction mapping. Specifically, we directly set the Lipschitz constant of $\mathcal{F}$ as $L_{\boldsymbol{h}} = 0.9$, and use 100 fixed-point iterations to solve the root $\boldsymbol{h}^*$ of Eq. (1) until the relative error satisfies $\|\boldsymbol{h} - \mathcal{F}(\boldsymbol{h}, \boldsymbol{z})\|/\|\boldsymbol{h}\| < 10^{-5}$. Here, the exact gradient is estimated by backpropagation through those fixed-point iterations, and cross-validated by implicit differentiation solved with 20 iterations of the Broyden's method [16]. In our experiment, the cosine similarity between these two gradient estimates consistently succeeds $0.9999$, indicating the gradient estimate is quite accurate when the relative errors of forward solver are minor. The cosine similarity between phantom gradients and exact gradients is shown in Fig. 1. It shows that the cosine similarity tends to increase as $k$ grows and that a small $\lambda$ tends to slow down the convergence of the phantom gradient, allowing it to explore in a wider range regarding the angle against the exact gradient.

In the ablation setting, the precision of the phantom gradient during the training process is shown in Fig. 2. The model is trained by implicit differentiation under the official schedule[2]. It shows that the phantom gradient still provides an ascent direction in the real training process, as indicated by the considerable cosine similarity against the exact gradient. Interestingly, the similarity seemingly slightly decays during the training process, which implies a possibility to construct an adaptive gradient solver for implicit models.

[2]Code available at https://github.com/locuslab/mdeq.

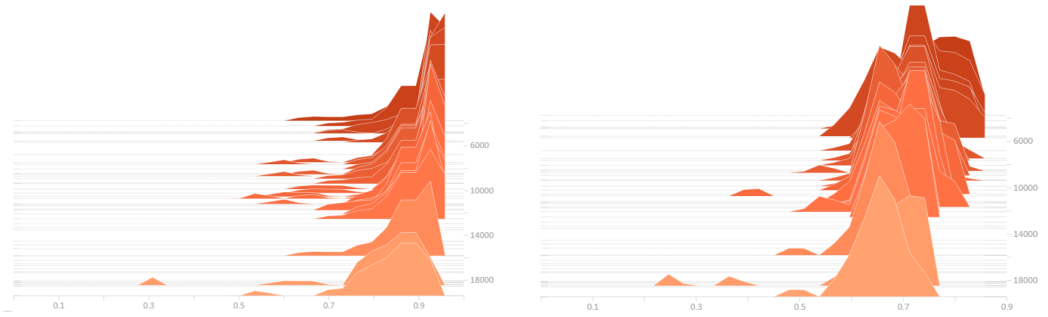

(a) Neumann-series-based phantom gradient. (b) Unrolling-based phantom gradient.

Figure 2: Cosine similarity between the phantom gradient and the exact gradient in the real scenario. The horizontal axis corresponds to the cosine similarity, and the vertical axis to the training step.

**To Pretrain, or not to Pretrain?** To better understand the components of the implicit models' training schedule, we first illustrate a detailed ablation study of the baseline model in the ablation setting. The average accuracies with standard deviations are reported in Tab. 1.

The MDEQ model employs a pretraining stage in which the model $\mathcal{F}$ is unrolled as a recurrent network. We study the impact of the pretraining stage, the Dropout [25] operation, and the optimizer separately. It can be seen that the unrolled pretraining stabilizes the training of the MDEQ model. Removing the pretraining stage leads to a severe performance drop and apparent training instability among different trials because the solver cannot obtain an accurate fixed point $\boldsymbol{h}^*$ when the model is not adequately trained. This ablation study also suggests that the MDEQ model is a strong baseline for our method to compare.

Table 1: Ablation settings on CIFAR-10.

| **Method** | **Acc. (%)** |
|---|---|
| Implicit Differentiation | $85.0 \pm 0.2$ |
| w/o Pretraining | $82.3 \pm 1.3$ |
| w/o Dropout | $83.7 \pm 0.1$ |
| Adam $\rightarrow$ SGD | $84.5 \pm 0.3$ |
| SGD w/o Pretraining | $82.9 \pm 1.5$ |
| UPG ($\boldsymbol{A}_{5,0.5}$, w/o Dropout) | $\mathbf{85.8} \pm 0.5$ |
| NPG ($\boldsymbol{A}_{5,0.5}$, w/o Dropout) | $85.6 \pm 0.5$ |
| UPG ($\boldsymbol{A}_{9,0.5}$, w/ Dropout) | $\mathbf{86.1} \pm 0.5$ |

However, pretraining is not always indispensable for training implicit models. It introduces an extra hyperparameter, *i.e.,* how many steps should be involved in the pretraining stage. Next, we discuss how the UPG could circumvent this issue.

**Trade-offs between Unrolling and Neumann.** For an exact fixed point $\boldsymbol{h}^*$, *i.e.,* $\boldsymbol{h}^* = \mathcal{F}(\boldsymbol{h}^*, \boldsymbol{z})$, there is no difference between UPG and NPG. However, when the numerical error exists in solving $\boldsymbol{h}^*$, *i.e.,* $\|\boldsymbol{h}^* - \mathcal{F}(\boldsymbol{h}^*, \boldsymbol{z})\| > 0$, these two instatiations of the phantom gradient can behave differently.

We note that a particular benefit of the UPG is its ability to automatically switch between the pretraining and training stages for implicit models. When the model is not sufficiently trained and the solver converges poorly (see [3]), the UPG defines a forward computational graph that is essentially equivalent to a shallow weight-tied network to refine the coarse equilibrium states. In this stage, the phantom gradient serves as a backpropagation through time (BPTT) algorithm and hence behaves as in the pretraining stage. Then, as training progresses, the solver becomes more stable and converges to the fixed point $\boldsymbol{h}^*$ better. This makes the UPG behave more like the NPG. Therefore, the unrolled pretraining is gradually transited into the regular training phase based on implicit differentiation, and the hyperparameter tuning of pretraining steps can be waived. We argue that such an ability to adaptively switch training stages is benign to the implicit models' training protocol, which is also supported by the performance gain in Tab. 1.

Table 2: Complexity comparison. Mem means the memory cost, and $K$ and $k$ denote the solver's steps and the unrolling/Neumann steps, respectively. Here, $K \gg k \approx 1$.

| **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 | $\mathcal{O}(k)$ | $\mathcal{O}(1)$ | $\mathcal{O}(1)$ |

Although the UPG requires higher memory than implicit differentiation or the NPG, it does not surpass the peak memory usage in the entire training process of implicit differentiation due to


(a) Impact of $\lambda$ and $k$ (b) NPG *v.s.* UPG

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

the pretraining stage. In the ablation setting, the MDEQ model employs a 10-layer unrolling for pretraining, which actually consumes double memory compared with a 5-step unrolling scheme, *e.g.,* $\boldsymbol{A}_{5,0.5}$ in Tab. 1. In Tab. 2, we also demonstrate the time and memory complexity for implicit differentiation and the two forms of phantom gradient.

In addition, leaving the context of approximate and exact gradients aside, we also develop insights into understanding the subtle differences between UPG and NPG in terms of the state-dependent and state-free gradients. Actually, UPG is state-dependent, which means it is the "exact" gradient of a computation sub-graph, although it serves as an approximation to the whole gradient. Both the NPG and the gradient solved by implicit differentiation, however, do not exactly match the gradient of a forward computation graph unless the numerical errors are entirely eliminated in both the forward and backward passes for implicit differentiation or the one-step gradient [26, 27, 28, 29] is used for the NPG. Interestingly, we observe that models trainde on the state-dependent gradients demonstrate an additional training stability regarding the Jacobian spectral radius, compared with the state-free ones *when numerical errors exist (IFT) or approximate gradients are applied (NPG)*. We empirically note that it can be seen as certain form of implicit Jacobian regularization for implicit models as the supplement to the explicit Jacobian regularization [17], indicated by the stable estimated Jacobian spectral radius along the training, *e.g.,* $\rho(\partial \mathcal{F}_\lambda / \partial \boldsymbol{h}) \approx 1$.

The experimental results also agree with our insight. The performance curves in Fig. 3 demonstrate the influences of $\lambda$ and $k$ and further validate that the UPG is more robust to a wide range of steps $k$ than the NPG. When the Jacobian spectral radius $\rho(\partial \mathcal{F} / \partial \boldsymbol{h})$ increases freely and lacks of regularization, the exploding high-order terms in NPG can ruin the training dynamics, leading to performance degradation in the large $k$ realms of Fig. 3(b), also detected by the surging estimated radius and gradient explosion in our experiments, while UPG can circumvent this issue through its implicit Jacobian regularization.

Table 3: Experiments using DEQ [2] and MDEQ [3] on vision and language tasks. Metrics stand for accuracy(%)↑ for image classification on CIFAR-10 and ImageNet, and perplexity↓ for language modeling on Wikitext-103. JR stands for Jacobian Regularization [17]. † indicates additional iterations in the forward equilibrium solver.

| **Datasets** | **Model** | **Method** | **Params** | **Metrics** | **Speed** |
|---|---|---|---|---|---|
| CIFAR-10 | MDEQ | Implicit | 10M | $93.8 \pm 0.17$ | $1.0\times$ |
| CIFAR-10 | MDEQ | UPG $\boldsymbol{A}_{5,0.5}$ | 10M | $95.0 \pm 0.16$ | $1.4\times$ |
| ImageNet | MDEQ | Implicit | 18M | 75.3 | $1.0\times$ |
| ImageNet | MDEQ | UPG $\boldsymbol{A}_{5,0.6}$ | 18M | 75.7 | $1.7\times$ |
| Wikitext-103 | DEQ (PostLN) | Implicit | 98M | 24.0 | $1.0\times$ |
| Wikitext-103 | DEQ (PostLN) | UPG $\boldsymbol{A}_{5,0.8}$ | 98M | 25.7 | $1.7\times$ |
| Wikitext-103 | DEQ (PreLN) | JR + Implicit | 98M | 24.5 | $1.7\times$ |
| Wikitext-103 | DEQ (PreLN) | JR + UPG $\boldsymbol{A}_{5,0.8}$ | 98M | 24.4 | $2.2\times$ |
| Wikitext-103 | DEQ (PreLN) | JR + UPG $\boldsymbol{A}_{5,0.8}$ | 98M | $24.0^\dagger$ | $1.7\times$ |

**Phantom Gradient at Scale.** We conduct large-scale experiments to verify the advantages of the phantom gradient on vision, graph, and language benchmarks. We adopt the UPG in the large-scale experiments. The results are illustrated in Tab. 3 and Tab. 4. Our method matches or surpasses the implicit differentiation training protocol on the state-of-the-art implicit models with a visible reduction on the training time. When only considering the backward pass, the acceleration for MDEQ can be remarkably $12\times$ on ImageNet classification.

Table 4: Experiments using IGNN [4] on graph tasks. Metrics stand for accuracy(%)↑ for graph classification on COX2 and PROTEINS, Micro-F1(%)↑ for node classification on PPI.

| Datasets | Model | Method | Params | Metrics (%) |
|---|---|---|---|---|
| COX2 | IGNN | Implicit | 38K | $84.1 \pm 2.9$ |
| COX2 | IGNN | UPG $\boldsymbol{A}_{5,0.5}$ | 38K | $83.9 \pm 3.0$ |
| COX2 | IGNN | UPG $\boldsymbol{A}_{5,0.8}$ | 38K | $83.9 \pm 2.7$ |
| COX2 | IGNN | UPG $\boldsymbol{A}_{5,1.0}$ | 38K | $83.0 \pm 2.9$ |
| PROTEINS | IGNN | Implicit | 34K | $78.6 \pm 4.1$ |
| PROTEINS | IGNN | UPG $\boldsymbol{A}_{5,0.5}$ | 34K | $78.4 \pm 4.2$ |
| PROTEINS | IGNN | UPG $\boldsymbol{A}_{5,0.8}$ | 34K | $78.6 \pm 4.2$ |
| PROTEINS | IGNN | UPG $\boldsymbol{A}_{5,1.0}$ | 34K | $78.8 \pm 4.2$ |
| PPI | IGNN | Implicit | 4.7M | 97.6 |
| PPI | IGNN | UPG $\boldsymbol{A}_{5,0.5}$ | 4.7M | 98.2 |
| PPI | IGNN | UPG $\boldsymbol{A}_{5,0.8}$ | 4.7M | 97.4 |
| PPI | IGNN | UPG $\boldsymbol{A}_{5,1.0}$ | 4.7M | 96.2 |

## 4 Related Work

**Implicit Models.** Implicit models generalize the recursive forward/backward rules of neural networks and characterize their internal mechanism by some pre-specified dynamics. Based on the dynamics, the implicit mechanisms can be broadly categorized into three classes: ODE-based [1, 7], root-solving-based [2, 3, 4, 5, 10], and optimization-based [30, 31, 32] implicit models.

The ODE-based implicit models [1, 7] treat the iterative update rules of residual networks as the Euler discretization of an ODE, which could be solved by any black-box ODE solver. The gradient of the differential equation is calculated using the *adjoint method* [33], in which the adjoint state is obtained by solving another ODE. The root-solving-based implicit models [2, 5, 3, 4, 6, 10, 29] characterize layers of neural networks by solving fixed-point equations. The equations are solved by either the black-box root-finding solver [2, 3] or the fixed-point iteration [4, 29]. The optimization-based implicit models [30, 31, 32, 34, 26, 27, 35, 36] leverage the optimization programs as layers of neural networks. Previous works have studied differentiable layers of quadratic programming [30], submodular optimization [31], maximum satisfiability (MAXSAT) problems [32], and structured decomposition [27]. As for the backward passes, implicit differentiation is applied to the problem-defining equations of the root-solving-based models [2, 3] or the KKT conditions of the optimization-based models [30]. As such, the gradient can be obtained from solving the backward linear system.

In this work, we focus on the root-solving-based implicit models. Theorectical works towards root-solving-based implicit models including the well-posedness [5, 4], monotone operators [10], global convergence [37, 36], and Lipschitz analysis [15]. We look into the theoretical aspect of the gradient-based algorithm in training implicit models and the efficient practice guidance. With these considerations, we show that implicit models of the same architecture could enjoy faster training speed and strong generalization ability in practical applications.

**Non-End-to-End Optimization in Deep Learning.** Non-end-to-end optimization aims to replace the standard gradient-based training of deep architectures with modular or weakly modular training without the entire forward and backward passes. Currently, there are mainly three research directions in this field, namely, the auxiliary variable methods [38, 39, 40, 41, 42, 43, 44], target propagation [45, 46, 47], and synthetic gradient [48, 49, 50]. The auxiliary variable methods [38, 39, 40, 41, 42, 43, 44] formulate the optimization of neural networks as constrained optimization problems, in which the layer-wise activations are considered as trainable auxiliary variables. Then, the equality constraints are relaxed as penalty terms added to the objectives so that the parameters and auxiliary variables can be divided into blocks and thus optimized in parallel. The target propagation method [45, 46, 47] trains each module by having its activations regress to the pre-assigned targets, which are

propagated backwards from the downstream modules. Specifically, the auto-encoder architecture is used to reconstruct targets at each layer. Finally, the synthetic gradient method [48, 49, 50] estimates the local gradient of neural networks using auxiliary models, and employ the synthetic gradient in place of the exact gradient to perform parameter update. In this way, the forward and backward passes are decoupled and can be executed in an asynchronous manner.

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

**Differentiation through Implicit Dynamics.** Differentiation through certain implicit dynamics is an important aspect in a wide range of research fields, including bilevel optimization [22, 9], meta-learning [51, 52, 28], and sensitivity analysis [53]. Since the gradient usually cannot be computed analytically, researchers have to implicitly differentiate the dynamics at the converged point. The formula of the gradient typically contains a term of Jacobian-inverse (or Hessian-inverse), which is computationally prohibitive for large-scale models. (See Eq. (2) in our case.) Herein, several techniques have been developed to approximate the matrix inverse in the previous literature.

An intuitive solution is to differentiate through the unrolled steps of a numerical solver of the dynamics [54, 55, 8]. In particular, if a single step is unrolled, it reduces to the well-known *one-step gradient* [56, 51, 57, 26, 28, 27], in which the inverse of Jacobian/Hessian is simply approximated by an identity matrix. On the contrary, unrolling a small number of steps may induce a bias [9], while the memory and computation cost grows linearly as the number of unrolled steps increases. Towards this issue, Shaban *et al.* [22] propose to truncate the long-term dependencies and differentiate through only the last $L$ steps. In fact, if the dynamics has converged to a stationary point, the finite-term truncation in Shaban *et al.* [22] is exactly the Neumann approximation of the Jacobian-inverse with the first $L$ terms. Based on this, Lorraine *et al.* [9] directly use the truncated Neumann series as an approximation of the Jacobian-inverse. Besides the unrolling-based methods, optimization-based approaches [58, 52] have been studied in this field as well. Since the Jacobian-inverse-vector product can be viewed as solution of a linear system, algorithms like the conjugate gradient method can be used to solve it.

## 5 Limitation and Future Work

The main limitation of this work lies in the hyperparameter tuning of phantom gradients, especially for the damping factor $\lambda$, which directly controls the gradient's precision and conditioning, the implicit Jacobian regularization for UPG, the stability for NPG, and the final generalization behaviors. However, it has not been a hindrance to the application of phantom gradients in training implicit models as one can tune the hyperparameter according to the validation loss in the early training stage.

Regarding future works, we would like to highlight the following aspects: (1) eliminating the structured bias in the current phantom gradient, (2) constructing an adaptive gradient solver for implicit models, (3) analyzing the damping factor to provide practical guidance, (4) investigating the implicit Jacobian regularization, and (5) understanding how different noises in the gradients can impact the training of implicit models under different loss landscapes.

## 6 Conclusion

In this work, we explore the possibility of training implicit models via efficient approximate phantom gradients. We systematically analyze the general condition of a gradient estimate so that the implicit model can be guaranteed to converge to an approximate stationary point of the loss function. Specifically, we give a sufficient condition under which a first-order oracle could always find an ascent direction of the loss landscape in the training process. Moreover, we introduce two instantiations of the proposed phantom gradient, based on either the damped fixed-point unrolling or the Neumann series. The proposed method shows a $1.4 \sim 1.7\times$ acceleration with comparable or better performances on large-scale benchmarks. Overall, this paper provides a practical perspective on training implicit models with theoretical guarantees.

**Acknowledgments**

Zhouchen Lin was supported by the NSF China under Grants 61625301 and 61731018, Project 2020BD006 supported by PKU-Baidu Fund, and Zhejiang Lab (grant no.s 2019KB0AC01 and 2019KB0AB02). Yisen Wang was partially supported by the National Natural Science Foundation of China under Grant 62006153, and Project 2020BD006 supported by PKU-Baidu Fund. Shaojie Bai was sponsored by a grant from the Bosch Center for Artificial Intelligence.

## References

[1] Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural Ordinary Differential Equations. In *Neural Information Processing Systems (NeurIPS)*, 2018. 1, 9

[2] Shaojie Bai, J. Zico Kolter, and Vladlen Koltun. Deep Equilibrium Models. In *Neural Information Processing Systems (NeurIPS)*, 2019. 1, 2, 3, 6, 8, 9

[3] Shaojie Bai, Vladlen Koltun, and J. Zico Kolter. Multiscale Deep Equilibrium Models. In *Neural Information Processing Systems (NeurIPS)*, pages 5238–5250, 2020. 1, 6, 7, 8, 9

[4] Fangda Gu, Heng Chang, Wenwu Zhu, Somayeh Sojoudi, and Laurent El Ghaoui. Implicit graph neural networks. In *Advances in Neural Information Processing Systems*, volume 33, pages 11984–11995. Curran Associates, Inc., 2020. 1, 9

[5] Laurent El Ghaoui, Fangda Gu, Bertrand Travacca, Armin Askari, and Alicia Tsai. Implicit Deep Learning. *SIAM Journal on Mathematics of Data Science*, 3(3):930–958, 2021. 1, 2, 9

[6] Tiancai Wang, Xiangyu Zhang, and Jian Sun. Implicit feature pyramid network for object detection. *arXiv preprint arXiv:2012.13563*, 2020. 1, 9

[7] Emilien Dupont, Arnaud Doucet, and Yee Whye Teh. Augmented Neural ODEs. In *Neural Information Processing Systems (NeurIPS)*, 2019. 1, 9

[8] Luca Franceschi, Michele Donini, Paolo Frasconi, and Massimiliano Pontil. Forward and Reverse Gradient-Based Hyperparameter Optimization. In *International Conference on Machine Learning (ICML)*, pages 1165–1173, 2017. 1, 10

[9] Jonathan Lorraine, Paul Vicol, and David Duvenaud. Optimizing Millions of Hyperparameters by Implicit Differentiation. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 1540–1552, 2020. 1, 4, 10

[10] Ezra Winston and J. Zico Kolter. Monotone operator equilibrium networks. In *Neural Information Processing Systems (NeurIPS)*, pages 10718–10728, 2020. 2, 9

[11] Cheng Lu, Jianfei Chen, Chongxuan Li, Qiuhao Wang, and Jun Zhu. Implicit Normalizing Flows. In *International Conference on Learning Representations (ICLR)*, 2021. 2

[12] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet: Large Scale Visual Recognition Challenge. *International Journal on Computer Vision (IJCV)*, 115(3):211–252, 2015. 2, 6

[13] Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. Pointer Sentinel Mixture Models. In *International Conference on Learning Representations (ICLR)*, 2017. 2, 6

[14] Chirag Pabbaraju, Ezra Winston, and J. Zico Kolter. Estimating Lipschitz constants of monotone deep equilibrium models. In *International Conference on Learning Representations (ICLR)*, 2021. 2

[15] Max Revay, Ruigang Wang, and Ian R Manchester. Lipschitz Bounded Equilibrium Networks. *arXiv:2010.01732*, 2020. 2, 9

[16] Charles G Broyden. A Class of Methods for Solving Nonlinear Simultaneous Equations. *Mathematics of computation*, 19(92):577–593, 1965. 3, 6

[17] Shaojie Bai, Vladlen Koltun, and J. Zico Kolter. Stabilizing equilibrium models by jacobian regularization. In *International Conference on Machine Learning (ICML)*, 2021. 3, 8

[18] Xavier Gastaldi. Shake-Shake Regularization. *arXiv preprint arXiv:1705.07485*, 2017. 3

[19] Guozhong An. The Effects of Adding Noise During Backpropagation Training on a Generalization Performance. *Neural Computation*, 8(3):643–674, 1996. 3

[20] Zhanxing Zhu, Jingfeng Wu, Bing Yu, Lei Wu, and Jinwen Ma. The Anisotropic Noise in Stochastic Gradient Descent: Its Behavior of Escaping from Sharp Minima and Regularization Effects. In *International Conference on Machine Learning (ICML)*, pages 7654–7663, 2019. 3

[21] Jingfeng Wu, Wenqing Hu, Haoyi Xiong, Jun Huan, Vladimir Braverman, and Zhanxing Zhu. On the Noisy Gradient Descent that Generalizes as SGD. In *International Conference on Machine Learning (ICML)*, pages 10367–10376, 2020. 3

[22] Amirreza Shaban, Ching-An Cheng, Nathan Hatch, and Byron Boots. Truncated Backpropagation for Bilevel Optimization. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 1723–1732, 2019. 4, 10

[23] Alex Krizhevsky, Geoffrey Hinton, et al. Learning Multiple Layers of Features from Tiny Images. Technical report, Citeseer, 2009. 6

[24] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral Normalization for Generative Adversarial Networks. In *International Conference on Learning Representations (ICLR)*, 2018. 6

[25] Geoffrey E Hinton, Nitish Srivastava, Alex Krizhevsky, Ilya Sutskever, and Ruslan R Salakhutdinov. Improving neural networks by preventing co-adaptation of feature detectors. *arXiv preprint arXiv:1207.0580*, 2012. 7

[26] Bryan Wilder, Eric Ewing, Bistra Dilkina, and Milind Tambe. End to end learning and optimization on graphs. *Advances in Neural Information Processing Systems*, 32:4672–4683, 2019. 8, 9, 10

[27] Zhengyang Geng, Meng-Hao Guo, Hongxu Chen, Xia Li, Ke Wei, and Zhouchen Lin. Is Attention Better Than Matrix Decomposition? In *International Conference on Learning Representations (ICLR)*, 2021. 8, 9, 10

[28] Xin-Yu Zhang, Taihong Xiao, Haolin Jia, Ming-Ming Cheng, and Ming-Hsuan Yang. Semi-Supervised Learning with Meta-Gradient. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 73–81, 2021. 8, 10

[29] Samy Wu Fung, Howard Heaton, Qiuwei Li, Daniel McKenzie, Stanley J. Osher, and Wotao Yin. Fixed Point Networks: Implicit Depth Models with Jacobian-Free Backprop. *arXiv preprint arXiv:2103.12803*, 2021. 8, 9

[30] Brandon Amos and J. Zico Kolter. OptNet: Differentiable Optimization as a Layer in Neural Networks. In *International Conference on Machine Learning (ICML)*, pages 136–145, 2017. 9

[31] Josip Djolonga and Andreas Krause. Differentiable Learning of Submodular Models. *Neural Information Processing Systems (NeurIPS)*, pages 1013–1023, 2017. 9

[32] Po-Wei Wang, Priya Donti, Bryan Wilder, and Zico Kolter. SATNet: Bridging deep learning and logical reasoning using a differentiable satisfiability solver. In *International Conference on Machine Learning (ICML)*, pages 6545–6554, 2019. 9

[33] VG Boltyanskiy, Revaz V Gamkrelidze, YEF Mishchenko, and LS Pontryagin. Mathematical theory of optimal processes. 1962. 9

[34] Marin Vlastelica, Anselm Paulus, Vít Musil, Georg Martius, and Michal Rolínek. Differentiation of Blackbox Combinatorial Solvers. In *International Conference on Learning Representations (ICLR)*, 2020. 9

[35] Stephen Gould, Richard Hartley, and Dylan Campbell. Deep declarative networks: A new hope. *arXiv preprint arXiv:1909.04866*, 2019. 9

[36] Xingyu Xie, Qiuhao Wang, Zenan Ling, Xia Li, Yisen Wang, Guangcan Liu, and Zhouchen Lin. Optimization induced equilibrium networks. *arXiv preprint arXiv:2105.13228*, 2021. 9

[37] Kenji Kawaguchi. On the theory of implicit deep learning: Global convergence with implicit layers. In *International Conference on Learning Representations*, 2020. 9

[38] Miguel Carreira-Perpinan and Weiran Wang. Distributed Optimization of Deeply Nested Systems. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 10–19, 2014. 9

[39] Gavin Taylor, Ryan Burmeister, Zheng Xu, Bharat Singh, Ankit Patel, and Tom Goldstein. Training Neural Networks Without Gradients: A Scalable ADMM Approach. In *International Conference on Machine Learning (ICML)*, pages 2722–2731, 2016. 9

[40] Ziming Zhang, Yuting Chen, and Venkatesh Saligrama. Efficient Training of Very Deep Neural Networks for Supervised Hashing. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1487–1495, 2016. 9

[41] Ziming Zhang and Matthew Brand. Convergent Block Coordinate Descent for Training Tikhonov Regularized Deep Neural Networks. In *Neural Information Processing Systems (NeurIPS)*, 2017. 9

[42] Jinshan Zeng, Shikang Ouyang, Tim Tsz-Kit Lau, Shaobo Lin, and Yuan Yao. Global Convergence in Deep Learning with Variable Splitting via the Kurdyka-Łojasiewicz Property. *arXiv preprint arXiv:1803.00225*, 2018. 9

[43] Jia Li, Cong Fang, and Zhouchen Lin. Lifted Proximal Operator Machines. In *Association for the Advancement of Artificial Intelligence (AAAI)*, pages 4181–4188, 2019. 9

[44] Fangda Gu, Armin Askari, and Laurent El Ghaoui. Fenchel Lifted Networks: A Lagrange Relaxation of Neural Network Training. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 3362–3371, 2020. 9

[45] Yoshua Bengio. How Auto-Encoders Could Provide Credit Assignment in Deep Networks via Target Propagation. *arXiv preprint arXiv:1407.7906*, 2014. 9

[46] Dong-Hyun Lee, Saizheng Zhang, Asja Fischer, and Yoshua Bengio. Difference Target Propagation. In *International Conference on Learning Representations (ICLR)*, page 498–515, 2015. 9

[47] Alexander Meulemans, Francesco Carzaniga, Johan Suykens, João Sacramento, and Benjamin F. Grewe. A Theoretical Framework for Target Propagation. In *Neural Information Processing Systems (NeurIPS)*, pages 20024–20036, 2020. 9

[48] Max Jaderberg, Wojciech Marian Czarnecki, Simon Osindero, Oriol Vinyals, Alex Graves, David Silver, and Koray Kavukcuoglu. Decoupled Neural Interfaces using Synthetic Gradients. In *International Conference on Machine Learning (ICML)*, pages 1627–1635, 2017. 9, 10

[49] Wojciech Marian Czarnecki, Grzegorz Świrszcz, Max Jaderberg, Simon Osindero, Oriol Vinyals, and Koray Kavukcuoglu. Understanding Synthetic Gradients and Decoupled Neural Interfaces. In *International Conference on Machine Learning (ICML)*, pages 904–912, 2017. 9, 10

[50] Benjamin James Lansdell, Prashanth Ravi Prakash, and Konrad Paul Kording. Learning to solve the credit assignment problem. In *International Conference on Learning Representations (ICLR)*, 2020. 9, 10

[51] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks. In *International Conference on Machine Learning (ICML)*, pages 1126–1135, 2017. 10

[52] Aravind Rajeswaran, Chelsea Finn, Sham M Kakade, and Sergey Levine. Meta-Learning with Implicit Gradients. In *Neural Information Processing Systems (NeurIPS)*, 2019. 10

[53] J Frédéric Bonnans and Alexander Shapiro. *Perturbation analysis of optimization problems*. Springer Science & Business Media, 2013. 10

[54] Justin Domke. Generic Methods for Optimization-Based Modeling. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 318–326, 2012. 10

[55] Dougal Maclaurin, David Duvenaud, and Ryan Adams. Gradient-based Hyperparameter Optimization through Reversible Learning. In *International Conference on Machine Learning (ICML)*, pages 2113–2122, 2015. 10

[56] Jelena Luketina, Mathias Berglund, Klaus Greff, and Tapani Raiko. Scalable Gradient-Based Tuning of Continuous Regularization Hyperparameters. In *International Conference on Machine Learning (ICML)*, pages 2952–2960, 2016. 10

[57] Hanxiao Liu, Karen Simonyan, and Yiming Yang. DARTS: Differentiable Architecture Search. *International Conference on Learning Representations (ICLR)*, 2018. 10

[58] Fabian Pedregosa. Hyperparameter Optimization with Approximate Gradient. In *International Conference on Machine Learning (ICML)*, pages 737–746, 2016. 10