
LocoProp: Enhancing BackProp via Local Loss Optimization

Ehsan Amid*

Google Research, Brain Team

Rohan Anil*

Google Research, Brain Team

Manfred K. Warmuth

Google Research, NY

Abstract

Second-order methods have shown state-of-the-art performance for optimizing deep neural networks. Nonetheless, their large memory requirement and high computational complexity, compared to first-order methods, hinder their versatility in a typical low-budget setup. This paper introduces a general framework of layerwise loss construction for multilayer neural networks that achieves a performance closer to second-order methods while utilizing first-order optimizers only. Our methodology lies upon a three-component loss, target, and regularizer combination, for which altering each component results in a new update rule. We provide examples using squared loss and layerwise Bregman divergences induced by the convex integral functions of various transfer functions. Our experiments on benchmark models and datasets validate the efficacy of our new approach, reducing the gap between first-order and second-order optimizers.

1 INTRODUCTION

Backpropagation (or *BackProp* for short) (Rumelhart et al., 1986) has been the prominent technique for training neural networks. BackProp is simply an expansion of the chain rule for calculating the derivative of the output *loss function* with respect to the weights in each layer. The BackProp update involves a forward pass to calculate the network’s activations given the input batch of training data. After the forward pass, the gradients of the loss function with respect to the network weights are *backpropagated* from the output layer all the way to the input and the weights are updated by applying a single gradient step.

Stochastic gradient descent is the most basic update rule, which involves taking a step in the direction of the negative backpropagated gradients. The more advanced first-order techniques such as AdaGrad (Duchi et al., 2011), RMSprop (Tieleman and Hinton, 2012), and Adam (Kingma and Ba, 2014) involve preconditioning the gradient by a diagonal matrix, as well as incorporating momentum. In contrast, second-order optimizers such as Shampoo (Gupta et al., 2018; Anil et al., 2020) and K-FAC (Heskes, 2000; Martens and Grosse, 2015; Ba et al., 2017) use Kronecker products to approximately form a full-matrix preconditioner, i.e., Full Matrix AdaGrad and Natural Gradient (Amari, 1998), respectively. Shampoo and K-FAC achieve significantly faster convergence than the first-order methods in terms of the number of steps and wall-time. However, large memory requirements (Anil et al., 2019; Shazeer and Stern, 2018) and the high computational cost of calculating matrix inverses (inverse- p th roots, in the case of Shampoo) make these approaches prohibitive for larger models. These methods require further extensions such as block diagonalization to make them scalable, and demand high precision arithmetic for computing the inverses. Moreover, they require more resources and are difficult to parallelize (Anil et al., 2020; Ba et al., 2017; Osawa et al., 2018). The question remains whether it is possible to achieve a similar performance to second-order methods with first-order optimizers without any extra forward-backward passes per batch or explicitly forming the preconditioners.

In this paper, we propose a simple local loss construction and minimization approach to enhance the efficacy of every BackProp step that is easily scalable to modern large-scale architectures. Our method involves forming a local loss for each layer by fixing a target and then minimizing this loss iteratively. The updates for each layer are entirely decoupled and run in parallel. A single iteration on the local loss always recovers a single BackProp step. Thus, any further iterations on the local loss work towards enhancing the initial BackProp step. Our approach is comparable to first-order methods in terms of memory and computation requirements: It involves a single forward-backward pass on each batch, followed by a few parallelizable lo-

Proceedings of the 25th International Conference on Artificial Intelligence and Statistics (AISTATS) 2022, Valencia, Spain. PMLR: Volume 151. Copyright 2022 by the author(s). *Equal contributions.

cal iterations. Our approach significantly improves the convergence of first-order methods, thus reducing the gap between first-order and second-order optimization techniques.

1.1 Related Work

Alternative update rules or extensions to BackProp have been proposed over the years (Carreira-Perpinan and Wang, 2014; Carreira-Perpinán and Alizadeh, 2016; Taylor et al., 2016; Zhang et al., 2018; Gotmare et al., 2018; Zach and Estellers, 2019; Zeng et al., 2019; Askari et al., 2018; Li et al., 2019; Gu et al., 2020). For instance, the *Difference Target Propagation* method (Lee et al., 2015) involves estimating a target value in each layer to approximate the BackProp gradient. The targets are formed via a separate network which performs as the inverse function. The main goal of such approaches is decoupled training (Jaderberg et al., 2017) by approximating the gradients or introducing auxiliary variables. Such decoupled training approaches are not the focus of our current work.

In this paper, we mainly focus on improving the performance of first-order optimizers (Nesterov, 1983; Duchi et al., 2011; Tieleman and Hinton, 2012; Kingma and Ba, 2014) to match the performance of their second-order contenders, namely K-FAC (Heskes, 2000; Martens and Grosse, 2015) and Shampoo (Gupta et al., 2018; Anil et al., 2020), via a local loss construction approach. The local problems are solved approximately using first-order optimizers, while we show that the exact solutions to these problems recover second-order update rules such as K-FAC. Our construction also recovers Proximal Backpropagation (Frerix et al., 2018), which builds upon the idea of local squared loss minimization, as a special case (analysis relegated to the appendix). Closely related is the work of Johnson and Zhang (2020) which proposes finding a guide function that is ahead of the current model with respect to the loss minimization and pushing the model towards the guide. They show that this construction improves training efficiency and is related to self-distillation.

1.2 Notation

We adopt the following notation throughout the paper. We use (\mathbf{x}, \mathbf{y}) to denote the input instance and target label pair. For an M -layer neural network and at a given layer $m \in [M]$, we use $\hat{\mathbf{a}}_m$ (respectively, $\hat{\mathbf{y}}_m$) and \mathbf{a}_m (respectively, \mathbf{y}_m) for the predicted and target pre (post)-activations, respectively. Note that $\hat{\mathbf{a}}_m = \mathbf{W}_m \hat{\mathbf{y}}_{m-1}$ and $\hat{\mathbf{y}}_m = f_m(\hat{\mathbf{a}}_m)$ where \mathbf{W}_m is the weight matrix¹ and f_m is an elementwise non-decreasing transfer function. (Also using this notation,

¹We assume that biases are incorporated in the weights.




Figure 1: The matching loss L_f of an increasing transfer function f as an integral under the curve f which is always convex in \hat{a} .

we have $\hat{\mathbf{y}}_0 = \mathbf{x}$ and $\hat{\mathbf{y}} = \hat{\mathbf{y}}_M$, but we distinguish between \mathbf{y}_M and \mathbf{y} .) We denote the loss of the network in the final layer (which is used for training with standard BackProp) as $L(\mathbf{y}, \hat{\mathbf{y}}) = L(\mathbf{y}, \hat{\mathbf{y}}_M)$. We make no assumption on the final loss L other than differentiability.

Lastly, \odot and \otimes denote Hadamard product and Kronecker product, respectively, and \mathbb{I} denotes the indicator function.

1.3 Our Contributions

- We introduce a layerwise loss construction framework that relies on a three-component loss, target, and regularizer combination. Our construction allows flexibility in setting each component to recover multiple different update rules.
- We discuss the motivation behind several combinations and show the connection to implicit gradient updates, Proximal Backpropagation, and second-order methods such as K-FAC. In each case, we show that the combination is set such that a single iteration on the local objective recovers BackProp (or a more advanced update such as natural gradient descent (Amari, 1998)), while applying further iterations recovers a second-order update.
- Experimentally, we show that performing additional iterations significantly improves convergence of the first-order methods on a benchmark optimization problem, thus reducing the gap between first-order and second-order techniques. After computing the BackProp gradients, our method is embarrassingly parallelizable across layers (Nielsen, 2016).

In particular, we will discuss in detail two variants of our local loss construction approach. The first method relies on using the squared loss as the local loss. Except for setting the targets, the squared loss construction approach uses no information about the local transfer

function. Instead, the second variant, based on the *matching loss*, adapts the local loss function to the transfer function of the layer.

Matching loss was initially introduced as a line integral (Helmhold et al., 1999) and was later motivated via a convex duality argument involving Bregman divergences (Kivinen and Warmuth, 2001; Amid et al., 2019). Specifically, consider an elementwise strictly increasing transfer function $f: \mathbb{R}^d \rightarrow \mathbb{R}^d$. Given input $\mathbf{x} \in \mathbb{R}^n$ and target label $\mathbf{y} \in \mathbb{R}^d$, let $\hat{\mathbf{a}} \in \mathbb{R}^d$ be the predicted pre (transfer function) activation produced by the model (for instance, the linear activation before softmax in a single-layer logistic regression classifier). The (post-activation) prediction of the model is given by $\hat{\mathbf{y}} = f(\hat{\mathbf{a}})$ (in this example, the softmax probabilities). The goal of training is to make the model prediction $\hat{\mathbf{y}}$ closer to the target label \mathbf{y} by means of minimizing a measure of discrepancy, commonly known as a loss function. In particular, the *matching loss* of the transfer function f between the target \mathbf{y} and the prediction $\hat{\mathbf{y}} = f(\hat{\mathbf{a}})$ is defined as the following line integral of f ,

$$L_f(\mathbf{y}, \hat{\mathbf{y}}(\hat{\mathbf{a}})) \doteq \int_{\mathbf{a}}^{\hat{\mathbf{a}}} (f(\mathbf{z}) - f(\mathbf{a}))^\top d\mathbf{z}, \quad (1)$$

where $\mathbf{a} = f^{-1}(\mathbf{y})$ is the target pre-activation. Figure 1 depicts this integral when $d = 1$. By definition, the gradient of the matching loss with respect to $\hat{\mathbf{a}}$ admits a simple form in terms of the difference between the prediction and the target:

$$\nabla_{\hat{\mathbf{a}}} L_f(\mathbf{y}, \hat{\mathbf{y}}(\hat{\mathbf{a}})) = f(\hat{\mathbf{a}}) - f(\mathbf{a}) = \hat{\mathbf{y}} - \mathbf{y}. \quad (2)$$

Consequently, when f is elementwise strictly increasing, then the local matching loss is strictly convex with respect to $\hat{\mathbf{a}}$. Note that the majority of the non-linear transfer functions used in practice (e.g., leaky ReLU, tanh, softmax, etc.) are indeed (elementwise) strictly increasing and thus their matching losses are strictly convex.²

As we shall see, the logistic loss is one of the main examples of a matching loss. When the logistic loss (i.e., softmax transfer function and KL (or cross entropy) divergence) is used for multilayer neural networks, then the overall loss is strictly convex in the weights of the last layer. Nonetheless, the convexity does not necessarily extend to the weights of the layers below. These weights are then updated by backpropagating the gradient of the *same* final loss. As one of our main contributions, we extend this construction to layerwise convex losses.

²The integral of the softmax function, known as “log-sum-exp”, is strictly convex in $\mathbb{R}^d - \{\pm c\mathbf{1}_d, c \in \mathbb{R}_+\}$.

2 LOCAL LOSS OPTIMIZATION

Given an input example³ \mathbf{x} , our local loss construction framework splits an M -layer feedforward neural network into independent layerwise (i.e., single neuron) problems. In layer $m \in [M]$, we fix the input to the layer $\hat{\mathbf{y}}_{m-1}$ and define a target \mathbf{a}_m (or $\mathbf{y}_m = f_m(\mathbf{a}_m)$), which lies in the pre (or post) activation domain. The update proceeds by minimizing a loss between the current prediction of the layer and a target \mathbf{a}_m , plus a regularizer term on the weights,

$$\mathbf{W}^{\text{new}} = \arg \min_{\tilde{\mathbf{W}}} \left\{ \underbrace{D(\tilde{\mathbf{W}} \hat{\mathbf{y}}_{m-1}, \mathbf{a}_m)}_{\text{loss}} + \underbrace{R(\tilde{\mathbf{W}}, \mathbf{W})}_{\text{regularizer}} \right\},$$

where D denotes a divergence function and R is a regularizer (e.g., a squared L_2 -norm). This formulation allows flexibility in term of setting the loss (i.e., divergence function), the target, and the regularizer term. In the next section, we will show how changing each component results in a different update rule, one of which recovers the K-FAC method. In each case, we pick the combination such that a single fixed-point iteration on the objective recovers standard gradient descent (i.e., BackProp) or a more advanced update such as natural gradient descent (Amari, 1998). We will refer to our **Local Loss Optimization** framework as **LocoProp**.

2.1 Local Squared Loss

We now motivate the more basic approach for constructing the local problems using a squared loss. Given the current pre-activation $\hat{\mathbf{a}}_m = \mathbf{W}_m \hat{\mathbf{y}}_{m-1}$ at layer $m \in [M]$, we first define the *gradient descent (GD) target* of the layer as $\mathbf{a}_m = \hat{\mathbf{a}}_m - \gamma \nabla_{\hat{\mathbf{a}}_m} L(\mathbf{y}, \hat{\mathbf{y}})$ where $\gamma > 0$ is an activation step size. That is, our target corresponds to a GD step on the current pre-activation using the gradient of the final loss with respect to the pre-activation. Keeping the input to the layer $\hat{\mathbf{y}}_{m-1}$ fixed, our local optimization problem at layer $m \in [M]$ consists of minimizing the squared loss between the new pre-activation $\tilde{\mathbf{W}} \hat{\mathbf{y}}_{m-1}$ and the GD target \mathbf{a}_m , plus a squared L_2 -norm regularizer which keeps the updated weights close to the current weights,

$$1/2 \|\tilde{\mathbf{W}} \hat{\mathbf{y}}_{m-1} - \mathbf{a}_m\|^2 + 1/2\eta \|\tilde{\mathbf{W}} - \mathbf{W}_m\|^2. \quad (3)$$

Here, $\eta > 0$ controls the trade-off between minimizing the loss and the regularizer. Setting the derivative of the objective to zero, we can write $\mathbf{W}_m^{\text{new}}$ as the solution of a fixed point iteration,

$$\mathbf{W}_m^{\text{new}} = \mathbf{W}_m - \eta (\mathbf{W}_m^{\text{new}} \hat{\mathbf{y}}_{m-1} - \mathbf{a}_m) \hat{\mathbf{y}}_{m-1}^\top. \quad (4)$$

³The construction generalizes to a batch of examples trivially.

Interestingly, one iteration over Eq. (4) by replacing $\mathbf{W}_m^{\text{new}}$ with \mathbf{W}_m on the r.h.s. yields,

$$\begin{aligned} \mathbf{W}_m^{\text{new}} &\approx \mathbf{W}_m - \eta (\mathbf{W}_m \hat{\mathbf{y}}_{m-1} - \mathbf{a}_m) \hat{\mathbf{y}}_{m-1}^\top \\ &= \mathbf{W}_m - \eta (\mathbf{W}_m \hat{\mathbf{y}}_{m-1} \\ &\quad - (\mathbf{W}_m \hat{\mathbf{y}}_{m-1} - \gamma \nabla_{\hat{\mathbf{a}}_m} L(\mathbf{y}, \hat{\mathbf{y}})) \hat{\mathbf{y}}_{m-1}^\top) \\ &= \mathbf{W}_m - \eta \gamma \nabla_{\hat{\mathbf{a}}_m} L(\mathbf{y}, \hat{\mathbf{y}}) \hat{\mathbf{y}}_{m-1}^\top \\ &= \mathbf{W}_m - \eta_e \frac{\partial L(\mathbf{y}, \hat{\mathbf{y}})}{\partial \hat{\mathbf{a}}_m} \frac{\partial \hat{\mathbf{a}}_m}{\partial \mathbf{W}_m}. \end{aligned} \quad (\text{BackProp})$$

Thus, a single iteration on the objective in Eq (3) recovers BackProp, with an effective learning rate of $\eta_e \doteq \eta \gamma > 0$. Nonetheless, Eq. (4) can in fact be solved in a closed-form as,

$$\mathbf{W}_m^{\text{new}} = \mathbf{W}_m - \eta_e \nabla_{\mathbf{W}_m} L(\mathbf{y}, \hat{\mathbf{y}}) (\mathbf{I} + \eta \hat{\mathbf{y}}_{m-1} \hat{\mathbf{y}}_{m-1}^\top)^{-1},$$

which corresponds to an *implicit gradient* update (Hasibi et al., 1996; Kivinen et al., 2006). Implicit updates can be motivated as the backward Euler approximation of the continuous-time gradient flow (Amid and Warmuth, 2020b) and generally provide faster convergence compared to their approximate (explicit) counterpart (Amid and Warmuth, 2020a). The matrix preconditioner which trails the BackProp gradient also corresponds to the right preconditioner that appears in the K-FAC update rule (see the appendix). In fact, the following proposition states that we can also recover the full K-FAC update as a special case.

Proposition 1. *The update that minimizes Eq. (3) with a natural gradient descent target $\mathbf{a}_m = \hat{\mathbf{a}}_m - \gamma \mathbf{F}_m^{-1} \nabla_{\hat{\mathbf{a}}_m} L(\mathbf{y}, \hat{\mathbf{y}})$, in which \mathbf{F}_m is the Fisher Information matrix treating the pre-activations as parameters, recovers the K-FAC update.*

Instead of forming the matrix inverse, which is computationally expensive for large layers, we apply the fixed-point update in Eq. (4) for a certain number of iterations that suits the computational budgets. We refer to our local squared loss construction approach with GD targets as *LocoProp-S* for short.

2.2 Matching Loss

We now extend the local loss construction in the previous section by replacing the squared loss between the pre-activations and the targets with a Bregman divergence which is tailored to the transfer function of each layer. We start by reviewing the idea of a matching loss of a transfer function and show that when using a matching loss, the local problem remains convex with respect to the weights in each layer. Additional LocoProp variants based on dual of the matching loss as well as those obtained by using different weight regularizers are given in the appendix.

2.2.1 Bregman Divergence

The integral in Eq. (1) expands to a *Bregman divergence* (Bregman, 1967) induced by the strictly convex integral function $F : \mathbb{R}^d \rightarrow \mathbb{R}$ such that $f = \nabla F$:

$$\begin{aligned} L_f(\mathbf{y}, \hat{\mathbf{y}}(\hat{\mathbf{a}})) &= \left(F(\mathbf{z}) - f(\mathbf{a})^\top \mathbf{z} \right) \Big|_{\mathbf{a}}^{\hat{\mathbf{a}}} \\ &= F(\hat{\mathbf{a}}) - F(\mathbf{a}) - f(\mathbf{a})^\top (\hat{\mathbf{a}} - \mathbf{a}) \doteq D_F(\hat{\mathbf{a}}, \mathbf{a}). \end{aligned} \quad (5)$$

Bregman divergences are non-negative distance measures that satisfy many desirable properties including: **(I) Convexity:** $D_F(\hat{\mathbf{a}}, \mathbf{a})$ is always convex in the first argument, but not necessarily in the second argument. **(II) Duality:** $D_F(\hat{\mathbf{a}}, \mathbf{a}) = D_{F^*}(\mathbf{y}, \hat{\mathbf{y}})$ where $F^*(\mathbf{y}) \doteq \sup_{\mathbf{z}} \{\mathbf{y} \cdot \mathbf{z} - F(\mathbf{z})\}$ is the Fenchel dual (Hiriart-Urruty and Lemaréchal, 2001) of the convex function F , and $(\hat{\mathbf{a}}, \hat{\mathbf{y}})$ and (\mathbf{a}, \mathbf{y}) are pairs of dual points, $\hat{\mathbf{y}} = f(\hat{\mathbf{a}})$, $\hat{\mathbf{a}} = f^*(\hat{\mathbf{y}})$, $\mathbf{y} = f(\mathbf{a})$, $\mathbf{a} = f^*(\mathbf{y})$, where $f^* = \nabla F^* = f^{-1}$. **(III) Strict non-negativity:** $D_F(\hat{\mathbf{a}}, \mathbf{a}) \geq 0$ for all $\hat{\mathbf{a}}, \mathbf{a} \in \text{dom}(F)$ and $D_F(\hat{\mathbf{a}}, \mathbf{a}) = 0$ iff $\hat{\mathbf{a}} = \mathbf{a}$.

Since Bregman divergence are convex with respect to the first argument (Property (I)), the rewrite (5) of the matching loss as a Bregman divergence implies that this loss is convex in the pre-activation $\hat{\mathbf{a}}$. Additionally, using the duality argument (Property (II)), we have $L_f(\mathbf{y}, \hat{\mathbf{y}}) = D_{F^*}(\mathbf{y}, \hat{\mathbf{y}})$. Consequently, Property (III) ensures that $L_f(\mathbf{y}, \hat{\mathbf{y}}) \geq 0$ and when f is (elementwise) strictly increasing, $L_f(\mathbf{y}, \hat{\mathbf{y}}) = 0$ iff $\mathbf{y} = \hat{\mathbf{y}}$ (or $\hat{\mathbf{a}} = \mathbf{a}$).

A classical example of a matching loss is the commonly used logistic loss for classification. The logistic loss amounts to the relative entropy divergence (a.k.a. KL divergence)⁴:

$$\text{KL}(\mathbf{y}, \hat{\mathbf{y}}) \doteq \sum_i y_i \log \frac{y_i}{\hat{y}_i} - y_i + \hat{y}_i, \quad (6)$$

between the target \mathbf{y} and the softmax probabilities, $\hat{\mathbf{y}} = f_{\text{SM}}(\hat{\mathbf{a}}) = \text{softmax}(\hat{\mathbf{a}}) \doteq \frac{\exp(\hat{\mathbf{a}})}{\sum_i \exp \hat{a}_i}$. The integral of the softmax function is the so called ‘‘log-sum-exp’’ function, $F_{\text{SM}}(\mathbf{a}) = \log \sum_i \exp(a_i)$. KL divergence is a Bregman divergence induced by the negative Shannon entropy function $F_{\text{SM}}^*(\mathbf{y}) = \sum_i (y_i \log y_i - y_i)$, which is in fact the Fenchel dual of F_{SM} (restricted to the unit simplex (Amid, 2020)). Thus, by Property (II), the KL divergence in Eq. (6) between the label \mathbf{y} and the output $\hat{\mathbf{y}}$ is equal to the Bregman divergence induced by the convex function F_{SM} formed between the pre-activation of the final layer $\hat{\mathbf{a}}$ and the target $\mathbf{a} = f_{\text{SM}}^{-1}(\mathbf{y}) = \log \mathbf{y} - \frac{1}{d} \sum_i \log y_i \mathbf{1}$,

$$D_{F_{\text{SM}}}(\hat{\mathbf{a}}, \mathbf{a}) = \log \frac{\sum_i \exp(\hat{a}_i)}{\sum_i \exp(a_i)} - \sum_i \frac{\exp(a_i) (\hat{a}_i - a_i)}{\sum_j \exp(a_j)}.$$

⁴The last two terms cancel when \mathbf{y} and $\hat{\mathbf{y}}$ are probability distributions.

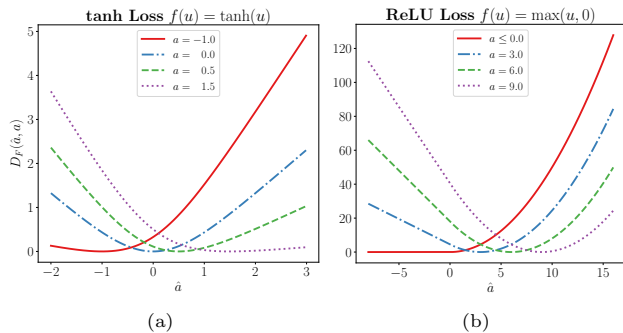


Figure 2: Matching losses induced by the (a) tanh and (b) ReLU transfer functions. For strictly increasing transfer functions such as tanh, the matching loss is equal to zero iff $\hat{\mathbf{a}} = \mathbf{a}$. On the contrary, non-decreasing (but not strictly increasing) transfer functions such as ReLU only satisfy the if condition, thus inducing matching losses with flat regions. (For instance for ReLU, $D_F(\hat{\mathbf{a}}, \mathbf{a}) = 0$ for all $\hat{\mathbf{a}}, \mathbf{a} \leq 0$.)

Consequently, neural networks trained with logistic loss (i.e., softmax transfer function paired with KL divergence) aim to minimize the matching loss of the softmax transfer function⁵.

$$L_{f_{\text{SM}}}(\mathbf{y}, \hat{\mathbf{y}}) = \text{KL}(\mathbf{y}, \hat{\mathbf{y}}) = D_{F_{\text{SM}}}(\hat{\mathbf{a}}, \mathbf{a}).$$

Also, the matching loss ensures convexity with respect to the weights of the last layer since $\hat{\mathbf{a}}$ is linear in the weights, i.e., $\hat{\mathbf{a}} = \hat{\mathbf{a}}_M = \mathbf{W}\hat{\mathbf{y}}_{M-1}$ where $\hat{\mathbf{y}}_{M-1}$ is the input to the last layer.

2.2.2 Common Transfer Functions

Interestingly, the majority of the commonly used transfer functions in the modern neural network architectures are elementwise non-decreasing. Examples of such transfer functions include (leaky) ReLU, hyperbolic tan, softplus, softmax (sigmoid), etc. (see the appendix for an extensive list). To simplify the construction, we refer to the Bregman divergence $D_F(\hat{\mathbf{a}}, \mathbf{a})$ induced by the convex integral function F as the matching loss of the transfer function f . Although in some cases (such as softplus) the matching loss has a sophisticated form, in practice, the only requirement is to calculate the gradient with respect to the first argument, i.e., $\nabla_{\hat{\mathbf{a}}} D_F(\hat{\mathbf{a}}, \mathbf{a})$. Favorably, Bregman divergences facilitate this calculation by providing a general form for the derivative with respect to the first argument; regardless of the form, the derivative only requires evaluation of the transfer function, as in Eq. (2).

Lastly, transfer functions such as step function and ReLU are non-decreasing but not strictly increasing.

⁵Although the target pre-activation \mathbf{a} may be unbounded in the case of a one-hot \mathbf{y} vector, the construction still carries out.

Thus, the integral function does not induce a strictly convex function. As a result, $D_F(\hat{\mathbf{a}}, \mathbf{a}) = 0$ does not necessarily imply $\hat{\mathbf{a}} = \mathbf{a}$. Also, the inverses $\mathbf{a} = f^{-1}(\mathbf{y})$ are not uniquely defined in this case. Figure 2 illustrates an example. Nonetheless, our construction still applies for non-decreasing transfer functions as we discuss in the next section. Thus, we will loosely refer to the divergences induced by non-decreasing transfer functions as Bregman divergences.

2.2.3 Local Matching Loss

We now discuss another variant of LocoProp using the matching loss of each layer. Given the current post-activation $\hat{\mathbf{y}}_m = f_m(\mathbf{W}_m \hat{\mathbf{y}}_{m-1})$ at layer $m \in [M]$ with a non-decreasing transfer function f_m , we define the *Mirror Descent (MD) target* of the layer as $\mathbf{y}_m = \hat{\mathbf{y}}_m - \gamma \nabla_{\hat{\mathbf{a}}_m} L(\mathbf{y}, \hat{\mathbf{y}})$ where $\gamma > 0$ is again the activation step size. The MD target corresponds to a MD step (Nemirovsky and Yudin, 1983) on the current post-activations using the gradient of the final loss with respect to the pre-activation. Keeping the input to the layer $\hat{\mathbf{y}}_{m-1}$ fixed, our local optimization problem at layer $m \in [M]$ now consists of minimizing the matching loss between the new post-activation $f_m(\tilde{\mathbf{W}} \hat{\mathbf{y}}_{m-1})$ and the MD target, plus a similar squared L₂-norm regularizer,

$$D_{F_m}(\mathbf{y}_m, f_m(\tilde{\mathbf{W}} \hat{\mathbf{y}}_{m-1})) + 1/2\eta \|\tilde{\mathbf{W}} - \mathbf{W}_m\|^2. \quad (7)$$

Similarly, $\eta > 0$ controls the trade-off between minimizing the loss and the regularizer. The following proposition shows the convexity of the problem (7) in \mathbf{W}_m in every layer $m \in [M]$ for a fixed input $\hat{\mathbf{y}}_{m-1}$ and MD target \mathbf{y}_m .

Proposition 2. *Given fixed input $\hat{\mathbf{y}}_{m-1}$ and MD target \mathbf{y}_m , the optimization problem (7) is convex in \mathbf{W}_m in every layer $m \in [M]$.*

Setting the derivative of the objective to zero, we can write $\mathbf{W}_m^{\text{new}}$ as the solution of a fixed point iteration,

$$\mathbf{W}_m^{\text{new}} = \mathbf{W}_m - \eta (f_m(\mathbf{W}_m^{\text{new}} \hat{\mathbf{y}}_{m-1}) - \mathbf{y}_m) \hat{\mathbf{y}}_{m-1}^\top. \quad (8)$$

Notably, calculating the gradient in Eq. (8) only requires the value of the input to the layer along with the difference between the post-activations. Unlike LocoProp-S, the fixed-point iteration of LocoProp-M in Eq. (8) does not yield a closed-form solution in general. However, the following proposition provides the approximate preconditioned update form.

Proposition 3. *The LocoProp-M update in Eq. (8) can approximately be written in the vectorized form $\mathbf{w}_m = \text{vec } \mathbf{W}_m$ as $\mathbf{w}_m^{\text{new}} = \mathbf{w}_m + \delta_m$,*

$$\delta_m = -\eta_e \mathbf{C}_m^{-1} (\mathbf{I} \otimes \mathbf{H}_{F_m}^{-1}) \text{vec}(\nabla_{\mathbf{w}_m} L(\mathbf{y}, \hat{\mathbf{y}})),$$

where $\mathbf{H}_{F_m} = \nabla^2 F_m$ is the Hessian matrix and

$$\mathbf{C}_m = \mathbf{I} \otimes \mathbf{H}_{F_m}^{-1} + \eta \hat{\mathbf{y}}_{m-1} \hat{\mathbf{y}}_{m-1}^\top \otimes \mathbf{I},$$

Algorithm 1 LocoProp-S: LocoProp Using Squared Loss

Input weights $\{\mathbf{W}_m\}$ where $m \in [M]$ for an M -layer network, activation step size γ , weight learning rate η

repeat

- perform a **forward pass** and fix the *inputs* $\{\hat{\mathbf{y}}_{m-1}\}$
- perform a **backward pass** and set the *GD targets*

$$\mathbf{a}_m = \hat{\mathbf{a}}_m - \gamma \nabla_{\hat{\mathbf{a}}_m} L(\mathbf{y}, \hat{\mathbf{y}})$$

for each layer $m \in [M]$ **in parallel do**

for T iterations **do**

$$\mathbf{W}_m \leftarrow \mathbf{W}_m - \eta (\mathbf{W}_m \hat{\mathbf{y}}_{m-1} - \mathbf{a}_m) \hat{\mathbf{y}}_{m-1}^\top$$

until $\{\mathbf{W}_m\}$ not converged

Algorithm 2 LocoProp-M: LocoProp Using Matching Loss

Input weights $\{\mathbf{W}_m\}$ where $m \in [M]$ for an M -layer network, activation step size γ , weight learning rate η

repeat

- perform a **forward pass** fix the *inputs* $\{\hat{\mathbf{y}}_{m-1}\}$
- perform a **backward pass** and set the *MD targets*

$$\mathbf{y}_m = \hat{\mathbf{y}}_m - \gamma \nabla_{\hat{\mathbf{a}}_m} L(\mathbf{y}, \hat{\mathbf{y}})$$

for each layer $m \in [M]$ **in parallel do**

for T iterations **do**

$$\mathbf{W}_m \leftarrow \mathbf{W}_m - \eta (f_m(\mathbf{W}_m \hat{\mathbf{y}}_{m-1}) - \mathbf{y}_m) \hat{\mathbf{y}}_{m-1}^\top$$

until $\{\mathbf{W}_m\}$ not converged

is the preconditioner matrix.

Interestingly, similar to LocoProp-S, the first iteration of the procedure always recovers BackProp,

$$\begin{aligned} \mathbf{W}_m^{\text{new}} &\approx \mathbf{W}_m - \eta (f_m(\mathbf{W}_m \hat{\mathbf{y}}_{m-1}) - \mathbf{y}_m) \hat{\mathbf{y}}_{m-1}^\top \\ &= \mathbf{W}_m - \eta_e \nabla_{\hat{\mathbf{a}}_m} L(\mathbf{y}, \hat{\mathbf{y}}) \hat{\mathbf{y}}_{m-1}^\top. \quad (\text{BackProp}) \end{aligned}$$

Thus, any further iterations enhances the initial BackProp update towards the fixed-point solution of Eq. (8). Also interestingly for a network which already utilizes a matching loss in the last layer, the objective in Eq. (7) with MD target and $\gamma = 1$ corresponds to directly minimizing the final loss.

Proposition 4. *For networks with a matching loss $L(\mathbf{y}, \hat{\mathbf{y}}) = L_{f_M}(\mathbf{y}, \hat{\mathbf{y}})$ in the last layer M , the LocoProp-M objective in Eq. (7) with $\gamma = 1$ corresponds to the loss of the network plus the regularizer term on the weights.*

Proposition 4 shows why layerwise matching loss is a more natural choice for defining the local problems, as the last layer simply minimizes the final loss with respect to the weights of the last layer. The LocoProp-S and LocoProp-M variants are given in Algorithm 1 and Algorithm 2, respectively.

3 EXPERIMENTS

We perform an extensive study on optimizing a deep auto-encoder on three standard datasets: MNIST (LeCun and Cortes, 2010), Fashion MNIST (Xiao et al., 2017), and CURVES⁶. Our emphasis on the deep auto-encoder task is due to it being a standard benchmark when studying new second-order methods for optimization (Martens, 2010; Goldfarb et al., 2020). Due to the limited size (only a few million parameters), it allows answering a plethora of questions with rigor on the effectiveness of our local loss construction. With our extensive tuning for the experiments, we find that adaptive first-order methods work far better than what is reported in the existing literature (Goldfarb et al., 2020).

⁶Downloadable at www.cs.toronto.edu/~jmartens/digs3pts_1.mat

For all experiments, the batch size is set to 1000 and the model is trained for 100 epochs with a learning rate schedule that includes a linear warmup for 5 epochs followed by a linear decay towards zero. The local iterations in LocoProp variants apply an additional linear decay schedule on the given learning rate. There are three configurations of autoencoders used in training: (a) Standard sized (2.72M parameters) with layer sizes: [1000, 500, 250, 30, 250, 500, 1000] (b) Deep (6.72M parameters) [1000, 500 \times 8 times, 250, 30, 250, 500 \times 8 times, 1000] and (c) Wide (26M parameters) [4000, 2000, 1000, 120, 1000, 2000, 4000]. We conduct several ablation studies on choices such as (a) transfer functions (ReLU, tanh), (b) model sizes, (c) datasets, (d) number of local iterations, and (e) inner optimizer for local iterations. Our implementation is in TensorFlow (Abadi et al., 2015) and all walltime measurements are made on a V100 GPU. For the second-order method implementation, K-FAC computes statistics based on the sampled gradients from the model distribution (Kunstner et al., 2019). For tuning, we use a Bayesian optimization package for each of our experiments. We search for hyper-parameters ($\eta, \beta_1, \beta_2, \epsilon$) for over hundreds (for larger scale models) to several thousand trials (for standard sized models). The search space, list of hyper-parameters, and further ablations across transfer functions, regularization, and batch sizes are available in the supplementary material, and the code to reproduce the experiments is available at <https://github.com/google-research/google-research/tree/master/locoprop>. We focus on our results on the standard, deep, and wide autoencoder variants with tanh transfer function on the MNIST dataset and relegate the additional results to the appendix.

3.1 Tuned Results for First-order Methods

We tune baseline first-order optimizers with a Bayesian optimization package for thousands of trials. We use standard implementations of optimizers in TensorFlow and tune all relevant hyper-parameters. The search space and further information are given in the supplementary material. Firstly, we observe that the RMSProp optimizer works remarkably well, as seen




Figure 3: Results on the MNIST dataset with tanh transfer function. Comparisons on the standard autoencoder: (a) first-order methods, (b) LocoProp-S vs. LocoProp-M, (c) effect of inner optimizer on LocoProp-M.

Table 1: Additional computational and memory cost for a layer with dimensions $\{d_i\}$ and batch size b . The number of local iterations for LocoProp is denoted by T .

Algorithm	Memory	Computation
Adam	$\mathcal{O}(\prod_i d_i)$	$\mathcal{O}(\prod_i d_i)$
Shampoo	$\mathcal{O}(\sum_i d_i^2)$	$\mathcal{O}(\sum_i d_i^3)$
K-FAC	$\mathcal{O}(\sum_i d_i^2)$	$\mathcal{O}(b \sum_i d_i^2) + \mathcal{O}(\sum_i d_i^3) + \mathcal{O}(b \prod_i d_i)$
LocoProp	$\mathcal{O}(\prod_i d_i) + \mathcal{O}(b \sum_i d_i)$	$\mathcal{O}(bT \prod_i d_i)$

in Figure 3(a). A difference to note is that RMSProp in TensorFlow includes an option for enabling momentum. We utilize this option, which is very similar to Adam: The former uses heavy ball momentum whereas the latter uses exponential moving averages. We discover that the performance gap between first-order and second-order methods, while still significant, is smaller than what has been previously reported in the literature. We attribute this to our exhaustive hyperparameter tuning as well as the fact that we made use of best practices for training neural networks such as learning rate warmup and decay.

3.2 Comparison between LocoProp-S and LocoProp-M Variants

We compare the performance of (a) LocoProp-S (using GD targets) and (b) LocoProp-M (using MD targets) on the standard sized deep encoder model on the MNIST dataset. We notice that the matching loss variant consistently performs better as seen in Figure 3(b). A rationale is that the matching losses are more adaptive to the local non-linear transfer functions (tanh, in this case) than squared loss.

3.3 Number of Local Iterations and Choice of Optimizer for LocoProp-M

The number of local gradient descent steps is a hyperparameter for LocoProp. As described earlier, a single

local iteration is equivalent to BackProp. Intuitively, further iterations are expected to improve on minimizing the local losses and make the update closer to a preconditioned gradient form. Thus, we expect to see more improvements as we increase the number of local iterations. However, additional improvements should diminish with the number of iterations. This is aligned with our observation in Figure 3(b), where increasing the number of local iterations from 5 to 10 and 50 consistently improves the performance. However, the performance gain after 10 local iterations is marginal, and thus, for all further experiments, we fix the number of local iterations to 10. Consequently, we refer to LocoProp-M with 10 local iterations as LocoProp. The LocoProp iterations minimize the local losses for each layer by gradient descent. We ablate across various choices of optimizers in Figure 3(c). Again, we find that RMSProp works quite well.

3.4 Computational and Memory Complexity

Here we provide a detailed account of the additional computational and memory complexity of LocoProp along with standard first-order and second-order methods in Table 1. For notations, we use b for batch size, and $\{d_i\}$ denotes the dimension of the weights of the layer, for example, $\{d_1, d_2\}$ for a fully connected network, and T denotes the number of local iterations. K-FAC has an additional computational complexity




Figure 4: Comparisons for standard, wide, and deep autoencoders on MNIST with tanh transfer function. All algorithms are trained for 100 epochs. @k indicates the interval for carrying out inverse (p th root) operation.

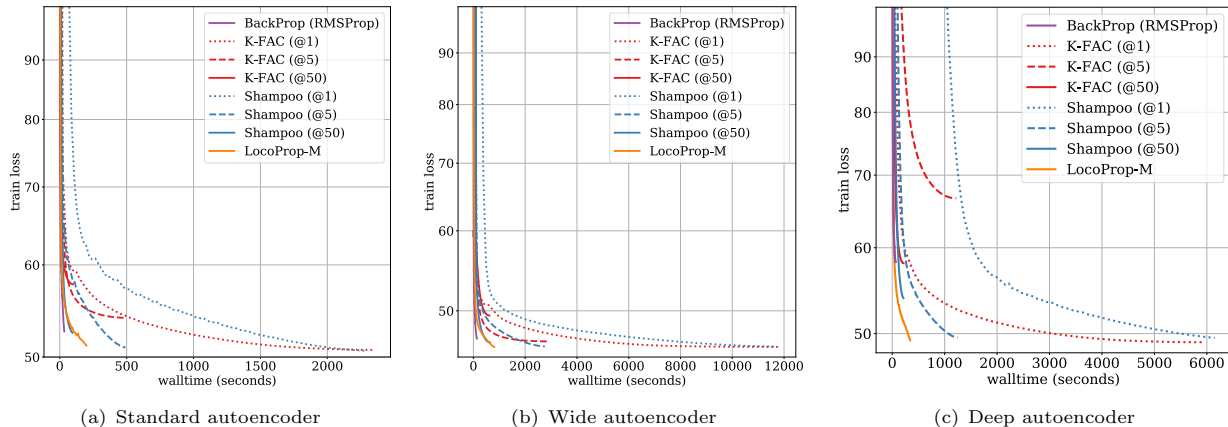


Figure 5: Walltime comparisons for standard, wide, and deep autoencoders on MNIST with tanh transfer function. LocoProp is much faster than optimized second-order methods on a single V100 GPU (limited parallelism). @k indicates the interval for carrying out inverse (p th root) operation.

compared to Shampoo due to calculating sampled gradients for the Fisher approximation. LocoProp has an advantage over KFAC and Shampoo both in terms of memory and computational complexity when batch size is comparable or smaller than the dimensions of the layers. Indeed we observe this in our experiments in Section 3.5

3.5 Comparison to Second-order Methods

We perform an end to end wall time comparison against second-order methods. LocoProp has similar convergence behavior to second-order methods while having significantly less overhead both in terms of memory and computation time. Second-order methods can be made faster by reducing the amount of computation by running the most expensive part of the step (i.e., matrix inverse or root) only every k steps. Note this comes with degradation in overall quality which we notice especially for K-FAC. The effectiveness of delayed inverse calculation depends on task specific param-

eters such as batch size and number of examples. For larger batch sizes, every step makes large progress and the curvature estimates need to be updated more frequently. Several investigations along these lines have studied the effect of stale preconditioners (Anil et al., 2020; Osawa et al., 2018; Ba et al., 2017). We find that LocoProp has a significant wall-time advantage compared to second-order methods. Second-order methods barely match the walltime performance of LocoProp by running the inverse p -th roots every 50 steps for Shampoo whereas for K-FAC delayed preconditioning configuration degrades the solution quality as seen in Figure 5 and the corresponding steps to convergence in Figure 4

4 Conclusions

We presented a simple yet effective enhancement to BackProp via a local loss construction that reduces the gap between first-order and second-order optimizers.

We perform a critical study of its effectiveness across a wide range of model sizes and dataset choices. The construction is embarrassingly parallel, allowing it to have a wallclock time advantage. Finally, we also have shown that our variant with squared (matching) loss is (approximately) an iterative version of a preconditioned update. Future work involves scaling the method up to much larger architectures across tasks.

Broader Impact and Limitations. This paper introduces a technique that could be used to accelerate the training of neural networks. This could have a positive downstream implication on reducing the energy usage of training large models. LocoProp is a new technique; it still remains to be seen how well the method generally works across tasks.

Acknowledgments

We would like to thank Frank Nielsen and our anonymous reviewers for their valuable feedback.

References

- Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., and et al. (2015). TensorFlow: Large-scale machine learning on heterogeneous systems. Software available from tensorflow.org.
- Amari, S. (1998). Natural gradient works efficiently in learning. *Neural Computation*, 10(2):251–276.
- Amid, E. (2020). *Tempered Bregman Divergence for Continuous and Discrete Time Mirror Descent and Robust Classification*. PhD thesis, University of California, Santa Cruz.
- Amid, E. and Warmuth, M. K. (2020a). An implicit form of Krasulina’s k-PCA update without the orthonormality constraint. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 3179–3186.
- Amid, E., Warmuth, M. K., Anil, R., and Koren, T. (2019). Robust bi-tempered logistic loss based on Bregman divergences. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, NeurIPS.
- Amid, E. and Warmuth, M. K. (2020b). Reparameterizing mirror descent as gradient descent. In *Advances in Neural Information Processing Systems*, volume 33, pages 8430–8439. Curran Associates, Inc.
- Anil, R., Gupta, V., Koren, T., Regan, K., and Singer, Y. (2020). Scalable Second Order Optimization for Deep Learning. *arXiv e-prints*, page arXiv:2002.09018.
- Anil, R., Gupta, V., Koren, T., and Singer, Y. (2019). Memory efficient adaptive optimization. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Askari, A., Negiar, G., Sambharya, R., and Ghaoui, L. E. (2018). Lifted neural networks. *arXiv preprint arXiv:1805.01532*.
- Ba, J., Grosse, R., and Martens, J. (2017). Distributed second-order optimization using Kronecker-factored approximations. *International Conference on Learning Representations (ICLR)*.
- Birkhoff, G. and Mac Lane, S. (2017). *A survey of modern algebra*. AK Peters/CRC Press.
- Bregman, L. M. (1967). The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming. *USSR computational mathematics and mathematical physics*, 7(3):200–217.
- Carreira-Perpinan, M. and Wang, W. (2014). Distributed optimization of deeply nested systems. In *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Statistics*, pages 10–19.
- Carreira-Perpinán, M. A. and Alizadeh, M. (2016). PAR-MAC: Distributed optimisation of nested functions, with application to learning binary autoencoders. *arXiv preprint arXiv:1605.09114*.
- Duchi, J., Hazan, E., and Singer, Y. (2011). Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12(Jul):2121–2159.
- Frerix, T., Möllenhoff, T., Moeller, M., and Cremers, D. (2018). Proximal Backpropagation. *International Conference on Learning Representations (ICLR)*.
- Goldfarb, D., Ren, Y., and Bahamou, A. (2020). Practical quasi-Newton methods for training deep neural networks. In *NeurIPS*.
- Gotmare, A. D., Thomas, V., Brea, J., and Jaggi, M. (2018). Decoupling Backpropagation using constrained optimization methods. In *ICML Workshop on Efficient Credit Assignment in Deep Learning and Deep Reinforcement Learning*.
- Gu, F., Askari, A., and El Ghaoui, L. (2020). Fenchel lifted networks: A Lagrange relaxation of neural network training. In *International Conference on Artificial Intelligence and Statistics*, pages 3362–3371. PMLR.
- Gupta, V., Koren, T., and Singer, Y. (2018). Shampoo: Preconditioned stochastic tensor optimization. In *International Conference on Machine Learning*, pages 1842–1850. PMLR.
- Hassibi, B., Sayed, A. H., and Kailath, T. (1996). H^∞ optimality of the LMS algorithm. *IEEE Transactions on Signal Processing*, 44(2):267–280.

- Helmbold, D. P., Kivinen, J., and Warmuth, M. K. (1999). Relative loss bounds for single neurons. *IEEE Transactions on Neural Networks*, 10(6):1291–1304.
- Heskes, T. (2000). On natural learning and pruning in multilayered perceptrons. *Neural Computation*, 12:881–901.
- Hiriart-Urruty, J.-B. and Lemaréchal, C. (2001). *Fundamentals of Convex Analysis*. Springer-Verlag Berlin Heidelberg, first edition.
- Jaderberg, M., Czarnecki, W. M., Osindero, S., Vinyals, O., Graves, A., Silver, D., and Kavukcuoglu, K. (2017). Decoupled neural interfaces using synthetic gradients. In *Proceedings of the 34th International Conference on Machine Learning*, pages 1627–1635.
- Johnson, R. and Zhang, T. (2020). Guided learning of nonconvex models through successive functional gradient optimization. In *International Conference on Machine Learning*, pages 4921–4930. PMLR.
- Kingma, D. P. and Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Kivinen, J. and Warmuth, M. K. (2001). Relative loss bounds for multidimensional regression problems. *Journal of Machine Learning*, 45(3):301–329.
- Kivinen, J., Warmuth, M. K., and Hassibi, B. (2006). The p-norm generalization of the LMS algorithm for adaptive filtering. *IEEE Transactions on Signal Processing*, 54(5):1782–1793.
- Kunstner, F., Balles, L., and Hennig, P. (2019). Limitations of the empirical Fisher approximation for natural gradient descent. *arXiv preprint arXiv:1905.12558*.
- LeCun, Y. and Cortes, C. (2010). MNIST handwritten digit database. *online*.
- Lee, D.-H., Zhang, S., Fischer, A., and Bengio, Y. (2015). Difference target propagation. In *Joint european conference on machine learning and knowledge discovery in databases*, pages 498–515. Springer.
- Li, J., Fang, C., and Lin, Z. (2019). Lifted proximal operator machines. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33:4181–4188.
- Martens, J. (2010). Deep learning via Hessian-free optimization. In *ICML*, volume 27, pages 735–742.
- Martens, J. and Grosse, R. (2015). Optimizing neural networks with Kronecker-factored approximate curvature. In *International conference on machine learning*, pages 2408–2417. PMLR.
- Moreau, J.-J. (1965). Proximité et dualité dans un espace hilbertien. *Bulletin de la Société mathématique de France*, 93:273–299.
- Nemirovsky, A. and Yudin, D. (1983). *Problem Complexity and Method Efficiency in Optimization*. Wiley & Sons, New York.
- Nesterov, Y. E. (1983). A method for solving the convex programming problem with convergence rate of $(1/k^2)$. *Dokl. akad. nauk Sssr*, 269:543–547.
- Nielsen, F. (2016). *Introduction to HPC with MPI for Data Science*. Springer.
- Osawa, K., Tsuji, Y., Ueno, Y., Naruse, A., Yokota, R., and Matsuoka, S. (2018). Second-order optimization method for large mini-batch: Training ResNet-50 on ImageNet in 35 epochs. *CoRR*, abs/1811.12019.
- Rumelhart, D. E., Hinton, G. E., and Williams, R. J. (1986). Learning representations by back-propagating errors. *nature*, 323(6088):533–536.
- Shazeer, N. and Stern, M. (2018). Adafactor: Adaptive learning rates with sublinear memory cost. In Dy, J. and Krause, A., editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 4596–4604. PMLR.
- Taylor, G., Burmeister, R., Xu, Z., Singh, B., Patel, A., and Goldstein, T. (2016). Training neural networks without gradients: A scalable ADMM approach. In *International conference on machine learning*, pages 2722–2731. PMLR.
- Tieleman, T. and Hinton, G. (2012). Lecture 6.5—RMSProp: Divide the gradient by a running average of its recent magnitude. COURSE: Neural Networks for Machine Learning.
- Xiao, H., Rasul, K., and Vollgraf, R. (2017). Fashion-MNIST: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:cs.LG/1708.07747*.
- Zach, C. and Estellers, V. (2019). Contrastive learning for lifted networks. *British Machine Vision Conference*.
- Zeng, J., Lau, T. T.-K., Lin, S., and Yao, Y. (2019). Global convergence of block coordinate descent in deep learning. In *International Conference on Machine Learning*, pages 7313–7323. PMLR.
- Zhang, H., Chen, W., and Liu, T.-Y. (2018). On the local Hessian in Back-propagation. *Advances in Neural Information Processing Systems*, 31:6520–6530.

Supplementary Material: LocoProp: Enhancing BackProp via Local Loss Optimization

A MATCHING LOSSES OF COMMON TRANSFER FUNCTIONS

We provide a more extensive list of convex integral functions of common transfer functions in Table 2.

Table 2: Example of elementwise non-decreasing transfer functions and their corresponding convex integral function. The Bregman divergence can be formed by plugging in the convex integral function into the definition: $D_F(\hat{\mathbf{a}}, \mathbf{a}) = F(\hat{\mathbf{a}}) - F(\mathbf{a}) - f(\mathbf{a})^\top (\hat{\mathbf{a}} - \mathbf{a})$. For instance, for “linear” activation, we have $F_{\text{Sq}}(\mathbf{a}) = 1/2 \|\mathbf{a}\|^2$ and the Bregman divergence is $D_{F_{\text{Sq}}}(\hat{\mathbf{a}}, \mathbf{a}) = 1/2 \|\hat{\mathbf{a}}\|^2 - 1/2 \|\mathbf{a}\|^2 - \mathbf{a}^\top (\hat{\mathbf{a}} - \mathbf{a}) = 1/2 \|\hat{\mathbf{a}} - \mathbf{a}\|^2$. For non-decreasing but not strictly increasing transfer functions such as “ReLU”, the induced divergence does not satisfy the necessity condition in Property (III): $D_F(\hat{\mathbf{a}}, \mathbf{a}) = 0$ if $\hat{\mathbf{a}} = \mathbf{a}$ (but not only if).

NAME	TRANSFER FUNCTION $f(\mathbf{a})$	CONVEX INTEGRAL FUNCTION $F(\mathbf{a})$	NOTE
STEP FUNCTION	$1/2 (1 + \text{sign}(\mathbf{a}))$	$\sum_i \max(a_i, 0)$	–
LINEAR	\mathbf{a}	$1/2 \ \mathbf{a}\ ^2$	–
(LEAKY) RELU	$\max(\mathbf{a}, 0) - \beta \max(-\mathbf{a}, 0)$	$1/2 \sum_i a_i (\max(a_i, 0) - \beta \max(-a_i, 0))$	$\beta \geq 0$
SIGMOID	$(1 + \exp(-\mathbf{a}))^{-1}$	$\sum_i (a_i + \log(1 + \exp(-a_i)))$	–
SOFTMAX	$\exp(\mathbf{a}) / \sum_i \exp(a_i)$	$\log \sum_i \exp(a_i)$	–
HYPERBOLIC TAN	$\tanh(\mathbf{a})$	$\sum_i \log \cosh(a_i)$	–
ARC TAN	$\arctan(\mathbf{a})$	$\sum_i (a_i \arctan(a_i) - \log \sqrt{1 + a_i^2})$	–
SOFTPLUS	$\log(1 + \exp(\mathbf{a}))$	$-\sum_i \text{Li}_2(-\exp(a_i))$	$\text{Li}_2 := \text{SPENCE'S FUNC.}$
ELU	$[f(\mathbf{a})]_i = \begin{cases} a_i & a_i \geq 0 \\ \beta(\exp a_i - 1) & \text{OTHERWISE} \end{cases}$	$\sum_i (a_i^2/2 \mathbb{I}(a_i \geq 0) + \beta(\exp a_i - a_i - 1) \mathbb{I}(a_i < 0))$	$\beta \geq 0$

B MISSING PROOFS

Proposition 2. Given fixed input $\hat{\mathbf{y}}_{m-1}$ and MD target \mathbf{y}_m , the optimization problem (7) is convex in \mathbf{W}_m in every layer $m \in [M]$.

Proof. The proof simply follows from the construction of the matching loss; that is, the convexity of Bregman divergence $D_{F_m}(\hat{\mathbf{a}}_m, \mathbf{a}_m) = L_{f_m}(\mathbf{W}_m \hat{\mathbf{y}}_{m-1}, f_m^{-1}(\mathbf{y}_m))$ in the first argument $\hat{\mathbf{a}}_m$ and the fact that $\hat{\mathbf{a}}_m$ is a linear function of \mathbf{W}_m , i.e., $\hat{\mathbf{a}}_m = \mathbf{W}_m \hat{\mathbf{y}}_{m-1}$. \square

Proposition 3. The LocoProp-M update in Eq. (8) can approximately be written in the vectorized form $\mathbf{w}_m = \text{vec } \mathbf{W}_m$ as $\mathbf{w}_m^{\text{new}} = \mathbf{w}_m + \delta_m$,

$$\delta_m = -\eta_e \mathbf{C}_m^{-1} (\mathbf{I} \otimes \mathbf{H}_{F_m}^{-1}) \text{vec}(\nabla_{\mathbf{W}_m} L(\mathbf{y}, \hat{\mathbf{y}})),$$

where $\mathbf{H}_{F_m} = \nabla^2 F_m$ is the Hessian matrix and

$$\mathbf{C}_m = \mathbf{I} \otimes \mathbf{H}_{F_m}^{-1} + \eta \hat{\mathbf{y}}_{m-1} \hat{\mathbf{y}}_{m-1}^\top \otimes \mathbf{I},$$

is the preconditioner matrix.

Proof. Consider the LocoProp-M update:

$$\mathbf{W}_m^{\text{new}} = \mathbf{W}_m - \eta (f_m(\mathbf{W}_m^{\text{new}} \hat{\mathbf{y}}_{m-1}) - \mathbf{y}_m) \hat{\mathbf{y}}_{m-1}^\top.$$

Writing $\mathbf{W}_m^{\text{new}} = \mathbf{W}_m + \Delta_m$, we have

$$\Delta_m = -\eta (f_m((\mathbf{W}_m + \Delta_m)\hat{\mathbf{y}}_{m-1}) - \mathbf{y}_m) \hat{\mathbf{y}}_{m-1}^\top \quad (9)$$

Assuming that $\|\Delta_m\| \ll \|\mathbf{W}_m\|$, we can approximate the rhs as

$$\begin{aligned} \Delta_m &\approx -\eta (f_m(\mathbf{W}_m \hat{\mathbf{y}}_{m-1}) + \mathbf{H}_{F_m} \Delta_m \hat{\mathbf{y}}_{m-1} - \mathbf{y}_m) \hat{\mathbf{y}}_{m-1}^\top \\ &= -\eta \mathbf{H}_{F_m} \Delta_m \hat{\mathbf{y}}_{m-1} \hat{\mathbf{y}}_{m-1}^\top - \eta_e \nabla_{\hat{\mathbf{a}}_m} L(\mathbf{y}, \hat{\mathbf{y}}) \hat{\mathbf{y}}_{m-1}^\top \end{aligned}$$

where $\mathbf{H}_{F_m} := \nabla^2 F_m$ is the Hessian. Rearranging the terms, we can write

$$\Delta_m + \eta \mathbf{H}_{F_m} \Delta_m \hat{\mathbf{y}}_{m-1} \hat{\mathbf{y}}_{m-1}^\top = -\eta_e \nabla_{\mathbf{W}_m} L(\mathbf{y}, \hat{\mathbf{y}})$$

Or alternatively

$$\mathbf{H}_{F_m}^{-1} \Delta_m + \eta \Delta_m \hat{\mathbf{y}}_{m-1} \hat{\mathbf{y}}_{m-1}^\top = -\eta_e \mathbf{H}_{F_m}^{-1} \nabla_{\mathbf{W}_m} L(\mathbf{y}, \hat{\mathbf{y}}) \quad (10)$$

Eq. (10) corresponds to a Sylvester equation (Birkhoff and Mac Lane, 2017). The solution to this equation can be written in the vectorized form as

$$\delta_m = \text{vec } \Delta_m = -\eta_e (\mathbf{I} \otimes \mathbf{H}_{F_m}^{-1} + \eta \hat{\mathbf{y}}_{m-1} \hat{\mathbf{y}}_{m-1}^\top \otimes \mathbf{I})^{-1} (\mathbf{I} \otimes \mathbf{H}_{F_m}^{-1}) \text{vec}(\nabla_{\mathbf{W}_m} L(\mathbf{y}, \hat{\mathbf{y}})) \quad (11)$$

Note that LocoProp-S assumes a flat geometry for which $\mathbf{H}_{F_m} = \mathbf{I}_d$ and Eq. (10) reduces to the implicit gradient update obtained from the closed-form solution of LocoProp-S. \square

Proposition 4. For networks with a matching loss $L(\mathbf{y}, \hat{\mathbf{y}}) = L_{f_M}(\mathbf{y}, \hat{\mathbf{y}})$ in the last layer M , the LocoProp- M objective in Eq. (7) with $\gamma = 1$ corresponds to the the loss of the network plus the regularizer term on the weights.

Proof. For a MD target with $\gamma = 1$ at the last layer M , we have

$$\mathbf{y}_M = \hat{\mathbf{y}} - \nabla_{\hat{\mathbf{a}}_M} L_{f_M}(\mathbf{y}, \hat{\mathbf{y}}) = \hat{\mathbf{y}} - \underbrace{(\hat{\mathbf{y}} - \mathbf{y})}_{\text{gradient of matching loss}} = \mathbf{y}.$$

Thus, the objective in Eq. (7) corresponds to minimizing the final loss $L_{f_M}(\mathbf{y}, \hat{\mathbf{y}})$ plus the regularizer term with respect to \mathbf{W}_M . \square

C LOCOPROP-S RECOVERS PROXPROP

ProxProp (Frerix et al., 2018) is motivated by approximating the proximal mapping (Moreau, 1965) of a function, which is commonly known as the implicit update (Hassibi et al., 1996; Kivinen et al., 2006). ProxProp is motivated by first forming updated pre and post-activations recursively starting from the layer $M-1$ (one before the last layer):

$$\hat{\mathbf{y}}_{M-1}^+ = \hat{\mathbf{y}}_{M-1} - \tau \nabla_{\hat{\mathbf{y}}_{M-1}} L(\mathbf{y}, \hat{\mathbf{y}}),$$

and for the layers below as,

$$\begin{aligned} \hat{\mathbf{a}}_m^+ &= \hat{\mathbf{a}}_m - \nabla_{\hat{\mathbf{a}}_m} f_m(\hat{\mathbf{a}}_m) (f_m(\hat{\mathbf{a}}_m) - \hat{\mathbf{y}}_{m-1}^+), \text{ for } m \in [M-1], \\ \hat{\mathbf{y}}_{m-1}^+ &= \hat{\mathbf{y}}_{m-1} - \nabla_{\hat{\mathbf{y}}_{m-1}} \left(\frac{1}{2} \|\mathbf{W}_m \hat{\mathbf{y}}_{m-1} - \hat{\mathbf{a}}_m^+\|^2 \right), \text{ for } m \in [M-1] - \{1\}. \end{aligned}$$

The ProxProp then proceeds by first updating the weights of the last layer using gradient descent,

$$\mathbf{W}_M^{\text{new}} = \mathbf{W}_M - \tau \nabla_{\mathbf{W}_M} L(\mathbf{y}, \hat{\mathbf{y}}).$$

The remaining weights are then updated by as minimizing a squared loss between the current pre-activations and the updated pre-activations plus a squared regularizer terms on the wights:

$$\mathbf{W}_m^{\text{new}} = \arg \min_{\widetilde{\mathbf{W}}} \left\{ 1/2 \|\widetilde{\mathbf{W}} \hat{\mathbf{y}}_{m-1} - \hat{\mathbf{a}}_m^+\|^2 + 1/2\tau_m \|\widetilde{\mathbf{W}} - \mathbf{W}_m\|^2 \right\}, m \in [M-1].$$

At first glance, due to the convoluted formulation of the updated activations, it is unclear what the objective of ProxProp is chasing after. Although not mentioned in the original ProxProp paper, after some simplification, we can rewrite the updated activations as a gradient step on the current values of the activations:

$$\begin{aligned}\hat{\mathbf{y}}_m^+ &= \hat{\mathbf{y}}_m - \tau \nabla_{\hat{\mathbf{y}}_M} L(\mathbf{y}, \hat{\mathbf{y}}), \\ \hat{\mathbf{a}}_m^+ &= \hat{\mathbf{a}}_m - \tau \nabla_{\hat{\mathbf{a}}_M} L(\mathbf{y}, \hat{\mathbf{y}}), \text{ for } m \in [M-1].\end{aligned}$$

With this simplification, we can recognize that the updated activations (i.e., the targets, in our terminology) can in fact be formed in parallel for all layers after the backward pass. Thus, the LocoProp-S and ProxProp formulations become equivalent except for updating the weights of the last layer; ProxProp applies vanilla gradient descent while LocoProp-S uses a local loss function for updating these weights as well, which recovers ProxProp when the number of local iterations is one.

This example further illustrates how our LocoProp formulation, which minimizes a loss to a target plus a regularizer on the weights, simplifies the construction and allows developing several extensions by adjusting each component.

D LOCOPROP-S CLOSED-FORM UPDATE RESEMBLES K-FAC PRECONDITIONERS

There are several types of preconditioners such as K-FAC and Shampoo that one could use to train neural networks. In this section, we make a connection to the K-FAC (Martens and Grosse, 2015) update rule, which applies a Kronecker factored approximation to the Fisher information matrix (Amari, 1998). Let \mathbf{w} be the vectorized form of matrix \mathbf{W} . Recall that the Fisher information matrix is defined as

$$\mathbf{F}(\mathbf{w}) = \mathbb{E}_{\substack{\mathbf{x} \sim \text{data} \\ \mathbf{y} \sim p(\mathbf{y}|\mathbf{x}, \mathbf{w})}} [\nabla_{\mathbf{w}} L(\mathbf{y}, \hat{\mathbf{y}}) \nabla_{\mathbf{w}} L(\mathbf{y}, \hat{\mathbf{y}})^\top], \quad (12)$$

where the expectation is over the data and the model’s predictive distribution. The natural gradient update (Amari, 1998) is then defined as

$$\mathbf{w}^+ = \mathbf{w} - \eta \mathbf{F}(\mathbf{w})^{-1} \nabla_{\mathbf{w}} L(\mathbf{y}, \hat{\mathbf{y}}). \quad (13)$$

For a single fully connected layer $\mathbf{W}_m \in \mathbb{R}^{d \times n}$ where $m \in [M]$, the gradient, denoted by $\mathbf{G}_m \in \mathbb{R}^{d \times n}$ can be obtained via the chain rule as $\mathbf{G}_m = \nabla_{\hat{\mathbf{a}}_m} L(\mathbf{y}, \hat{\mathbf{y}}) \hat{\mathbf{y}}_{m-1}^\top$, which in the vectorized form can be written as: $\nabla_{\hat{\mathbf{a}}_m} L(\mathbf{y}, \hat{\mathbf{y}}) \otimes \hat{\mathbf{y}}_{m-1}$. We can then write the Fisher information matrix as

$$\begin{aligned}\mathbf{F}(\mathbf{w}_m) &= \mathbb{E} [(\nabla_{\hat{\mathbf{a}}_m} L(\mathbf{y}, \hat{\mathbf{y}}) \otimes \hat{\mathbf{y}}_{m-1}) (\nabla_{\hat{\mathbf{a}}_m} L(\mathbf{y}, \hat{\mathbf{y}}) \otimes \hat{\mathbf{y}}_{m-1})^\top] \\ &= \mathbb{E} [(\nabla_{\hat{\mathbf{a}}_m} L(\mathbf{y}, \hat{\mathbf{y}}) \nabla_{\hat{\mathbf{a}}_m} L(\mathbf{y}, \hat{\mathbf{y}})^\top) \otimes (\hat{\mathbf{y}}_{m-1} \hat{\mathbf{y}}_{m-1}^\top)].\end{aligned}$$

Assuming independence between $\nabla_{\hat{\mathbf{a}}_m} L(\mathbf{y}, \hat{\mathbf{y}})$ and $\hat{\mathbf{y}}_{m-1}$, K-FAC rewrites the Fisher in a tractable form as

$$\mathbf{F}(\mathbf{w}_m) \approx \mathbb{E} [(\nabla_{\hat{\mathbf{a}}_m} L(\mathbf{y}, \hat{\mathbf{y}}) \nabla_{\hat{\mathbf{a}}_m} L(\mathbf{y}, \hat{\mathbf{y}})^\top)] \otimes \mathbb{E} [\hat{\mathbf{y}}_{m-1} \hat{\mathbf{y}}_{m-1}^\top].$$

Let $\mathbf{D} \doteq \mathbb{E} [(\nabla_{\hat{\mathbf{a}}_m} L(\mathbf{y}, \hat{\mathbf{y}}) \nabla_{\hat{\mathbf{a}}_m} L(\mathbf{y}, \hat{\mathbf{y}})^\top)]$ and $\mathbf{X} \doteq \mathbb{E} [\hat{\mathbf{y}}_{m-1} \hat{\mathbf{y}}_{m-1}^\top]$. Then the K-FAC update rule can be simplified as:

$$\mathbf{W}_m^+ \approx \mathbf{W}_m - \eta \mathbf{D}^{-1} \mathbf{G}_m \mathbf{X}^{-1}. \quad (14)$$

The prototypical implementation of K-FAC uses the moving average of the statistics over training batches. In Section 2.1, we noted that how the exact solution of the squared loss in Eq. (3) emerges as a preconditioner that is similar to the right preconditioner of K-FAC in Eq. (14), formed by the outer-product of the input activations. (In practice, constant diagonal entries are added to matrix \mathbf{X} to avoid numerical issues.) Interestingly, one could similarly recover the left preconditioner as well by setting the target pre-activations using a *natural gradient* step on the pre-activations:

$$\mathbf{a}_m = \hat{\mathbf{a}}_m - \gamma \mathbf{D}^{-1} \nabla_{\hat{\mathbf{a}}_m} L(\mathbf{y}, \hat{\mathbf{y}}),$$

where \mathbf{D} acts as the Fisher Information matrix treating the pre-activations as parameters, in which case the expectation is conditioned on \mathbf{x} and is taken over the model’s predictive distribution. Substituting for the targets yields the closed-form update rule for LocoProp-S with *natural gradient descent targets*:

$$\mathbf{W}_m^+ = \mathbf{W}_m - \eta_e \mathbf{D}^{-1} \nabla_{\mathbf{W}_m} L(\mathbf{y}, \hat{\mathbf{y}}) (\mathbf{I} + \eta \hat{\mathbf{y}}_{m-1} \hat{\mathbf{y}}_{m-1}^\top)^{-1}. \quad (15)$$

LocoProp: Enhancing BackProp via Local Loss Optimization

Algorithm 3 PocoProp-S: PocoProp Using Squared Loss

Input weights $\{\mathbf{W}_m\}$ where $m \in [M]$ for an M -layer network, activation step size γ , weight learning rate η

repeat

- perform a **forward pass** and fix the *inputs* $\{\hat{\mathbf{y}}_{m-1}\}$
- perform a **backward pass** and set the *post GD targets*

$$\mathbf{y}_m = \hat{\mathbf{y}}_m - \gamma \nabla_{\hat{\mathbf{y}}_m} L(\mathbf{y}, \hat{\mathbf{y}})$$

for each layer $m \in [M]$ **in parallel do**

for T iterations **do**

$$\mathbf{W}_m \leftarrow \mathbf{W}_m - \eta \mathbf{J}_{f_m}^\top (f_m(\mathbf{W}_m \hat{\mathbf{y}}_{m-1}) - \mathbf{y}_m) \hat{\mathbf{y}}_{m-1}^\top$$

until $\{\mathbf{W}_m\}$ not converged

Algorithm 4 PocoProp-M: PocoProp Using Matching Loss

Input weights $\{\mathbf{W}_m\}$ where $m \in [M]$ for an M -layer network, activation step size γ , weight learning rate η

repeat

- perform a **forward pass** fix the *inputs* $\{\hat{\mathbf{y}}_{m-1}\}$
- perform a **backward pass** and set the *dual MD targets*

$$\mathbf{a}_m = \hat{\mathbf{a}}_m - \gamma \nabla_{\hat{\mathbf{y}}_m} L(\mathbf{y}, \hat{\mathbf{y}})$$

for each layer $m \in [M]$ **in parallel do**

for T iterations **do**

$$\mathbf{W}_m \leftarrow \mathbf{W}_m - \eta \mathbf{J}_{f_m}^\top (\mathbf{W}_m \hat{\mathbf{y}}_{m-1} - \mathbf{a}_m) \hat{\mathbf{y}}_{m-1}^\top$$

until $\{\mathbf{W}_m\}$ not converged

Our exposition here is to make the connection which may aid us in analyzing the properties of LocoProp in the future.

E ALTERNATE LOCOPROP FORMULATIONS

We introduce two additional formulations of LocoProp based on the post-activations, for which we defer the experimental evaluation to future work. Similarly, the first formulation is based on the squared loss whereas the second formulation uses the dual form of the Bregman divergence used in LocoProp-M. In both cases, a single local iteration recovers BackProp. These two variants are given in Algorithm 3 and 4.

E.1 A Variant of LocoProp-S Based on Post-activations

We define a LocoProp-S variant by first setting the targets as $\mathbf{y}_m = \hat{\mathbf{y}}_m - \gamma \nabla_{\hat{\mathbf{y}}_m} L(\mathbf{y}, \hat{\mathbf{y}})$ for $m \in [M]$, where $\gamma > 0$ is the activation step size. We refer to these targets as *post GD targets*. Next, the local optimization problem is defined as

$$\mathbf{W}_m^{\text{new}} = \arg \min_{\tilde{\mathbf{W}}} \left\{ \frac{1}{2} \|f_m(\tilde{\mathbf{W}} \hat{\mathbf{y}}_{m-1}) - \mathbf{y}_m\|^2 + \frac{1}{2\eta} \|\tilde{\mathbf{W}} - \mathbf{W}_m\|^2 \right\}. \quad (16)$$

Similarly, $\eta > 0$ controls the trade-off between minimizing the loss and the regularizer. By setting the derivative of the objective to zero, the fixed point iteration for $\mathbf{W}_m^{\text{new}}$ can be written as,

$$\mathbf{W}_m^{\text{new}} = \mathbf{W}_m - \eta \mathbf{J}_{f_m} (\mathbf{W}_m^{\text{new}} \hat{\mathbf{y}}_{m-1})^\top (f_m(\mathbf{W}_m^{\text{new}} \hat{\mathbf{y}}_{m-1}) - \mathbf{y}_m) \hat{\mathbf{y}}_{m-1}^\top, \quad (17)$$

where $\mathbf{J}_{f_m}(\hat{\mathbf{a}}) = \frac{\partial f_m(\hat{\mathbf{a}})}{\partial \hat{\mathbf{a}}}$ is the Jacobian of the transfer function f_m (which in many cases is a diagonal matrix). Note that for the choice of post GD targets, a single fixed point iteration again recovers vanilla BackProp,

$$\begin{aligned} \mathbf{W}_m^{\text{new}} &\approx \mathbf{W}_m - \eta \mathbf{J}_{f_m} (\mathbf{W}_m \hat{\mathbf{y}}_{m-1})^\top (f_m(\mathbf{W}_m \hat{\mathbf{y}}_{m-1}) - \mathbf{y}_m) \hat{\mathbf{y}}_{m-1}^\top \\ &= \mathbf{W}_m - \eta \mathbf{J}_{f_m} (\mathbf{W}_m \hat{\mathbf{y}}_{m-1})^\top (f_m(\mathbf{W}_m \hat{\mathbf{y}}_{m-1}) - (f_m(\mathbf{W}_m \hat{\mathbf{y}}_{m-1}) - \gamma \nabla_{\hat{\mathbf{y}}_m} L(\mathbf{y}, \hat{\mathbf{y}}))) \hat{\mathbf{y}}_{m-1}^\top \\ &= \mathbf{W}_m - \eta \gamma \mathbf{J}_{f_m} (\mathbf{W}_m \hat{\mathbf{y}}_{m-1})^\top \nabla_{\hat{\mathbf{y}}_m} L(\mathbf{y}, \hat{\mathbf{y}}) \hat{\mathbf{y}}_{m-1}^\top \\ &= \mathbf{W}_m - \eta_e \frac{\partial L(\mathbf{y}, \hat{\mathbf{y}})}{\partial \hat{\mathbf{a}}_m} \frac{\partial \hat{\mathbf{a}}_m}{\partial \mathbf{W}_m}, \end{aligned} \quad (\text{BackProp})$$

with an effective learning rate of $\eta_e = \eta \gamma > 0$. However in this case, the fixed point iteration in Eq. (17) does not yield a closed-form solution in general and therefore, should be solved iteratively. We refer to this local loss construction variant as *Post-LoCoProp-S* for short. Other variants of PocoProp-S can be obtained by replacing the post GD targets with post natural GD targets with respect to the post activations. We defer the analysis of such variants to future work.

E.2 A Variant of LocoProp-M Based on Post-activations

Similarly, we define the *dual MD targets* as $\mathbf{a}_m = \hat{\mathbf{a}}_m - \gamma \nabla_{\hat{\mathbf{y}}_m} L(\mathbf{y}, \hat{\mathbf{y}})$ for $m \in [M]$, where $\gamma > 0$ is the activation step size. The local optimization problem is then defined using the dual form of the Bregman divergence used in Eq. (7), that is

$$\mathbf{W}_m^{\text{new}} = \arg \min_{\tilde{\mathbf{W}}} \left\{ D_{F_m^*} (f_m(\tilde{\mathbf{W}} \hat{\mathbf{y}}_{m-1}), f_m(\mathbf{a}_m)) + \frac{1}{2\eta} \|\tilde{\mathbf{W}} - \mathbf{W}_m\|^2 \right\} \quad (18)$$

Hyper-parameter	Range	Scaling
η	$[10^{-7}, 0.1]$	Log
$1 - \beta_1$	$[10^{-3}, 0.1]$	Log
$1 - \beta_2$	$[10^{-3}, 0.1]$	Log
ϵ	$[10^{-10}, 10^{-5}]$	Log

Table 3: Search space for comparing first-order optimizers. Note that the hyper-parameters are dropped when not applicable (e.g., β_2 for non-adaptive methods).

Hyper-parameter	Range	Scaling
η	$[10^{-7}, 10]$	Log
$1 - \beta_1$	$[10^{-3}, 0.1]$	Log
$1 - \beta_2$	$[10^{-3}, 0.1]$	Log
ϵ	$[10^{-10}, 10^{-1}]$	Log

Table 4: Search space for Shampoo and K-FAC.

Similarly, $\eta > 0$ controls the trade-off between minimizing the loss and the regularizer. Setting the derivative of the objective to zero, we can write \mathbf{W}_m^+ as the solution of a fixed point iteration,

$$\begin{aligned} \mathbf{W}_m^{\text{new}} &= \mathbf{W}_m - \eta \mathbf{J}_{f_m}(\mathbf{W}_m^{\text{new}} \hat{\mathbf{y}}_{m-1})^\top (f_m^*(f_m(\mathbf{W}_m^{\text{new}} \hat{\mathbf{y}}_{m-1})) - f_m^*(f_m(\mathbf{a}_m))) \hat{\mathbf{y}}_{m-1}^\top \\ &= \mathbf{W}_m - \eta \mathbf{J}_{f_m}(\mathbf{W}_m^{\text{new}} \hat{\mathbf{y}}_{m-1})^\top (\mathbf{W}_m^{\text{new}} \hat{\mathbf{y}}_{m-1} - \mathbf{a}_m) \hat{\mathbf{y}}_{m-1}^\top. \end{aligned} \quad (19)$$

Interestingly, the first iteration again recovers BackProp with an effective learning rate of $\eta_e = \eta \gamma$. We refer to this LocoProp-M variant with dual MD targets as *PocoProp-M*.

E.3 Other Possible Variants Based on the Regularizer Term

So far, we discussed LocoProp (and PocoProp) variants by considering different local loss functions as well as targets. A tangential approach to such constructions is to consider other possible options for the regularizer term on the weights. A number of possible alternatives are L_1 -regularizer, i.e., $\|\tilde{\mathbf{W}} - \mathbf{W}_m\|_1$ (for sparse updates), and a local Mahalanobis distance based on the Fisher matrix, that is

$$^{1/2}(\tilde{\mathbf{w}} - \mathbf{w}_m)^\top \mathbf{F}(\mathbf{w}_m)(\tilde{\mathbf{w}} - \mathbf{w}_m),$$

where $\tilde{\mathbf{w}}$ and \mathbf{w}_m are the vectorized forms of $\tilde{\mathbf{W}}$ and \mathbf{W}_m , respectively, and $\mathbf{F}(\mathbf{w}_m)$ is the Fisher information matrix, defined in Eq. (12). Note that a single iteration of LocoProp and PocoProp variants with this local Mahalanobis distance corresponds to a natural gradient descent step on \mathbf{W}_m (see Eq. (13)). Thus, computationally effective approximations of such variants may yield an improved convergence. We also defer the analysis of such variants to future work.

F TUNING PROTOCOL AND SEARCH SPACES

We setup a search space for each experiment as shown in Table 3, 5, and 4. We use a Bayesian optimization package for hyper-parameter search to optimize for training loss. We use decoupled weight decay of 10^{-5} across all experiments. We run up to 4096 trails for all experiments, except for the wider/deeper sized autoencoder where we run only up to 512 trials. Experiments are conducted on V100 GPUs. Our rough estimate for the total GPU time is $\sim 15\text{k}$ GPU hours.

G RESULTS ON RELU AND OTHER DATASETS

We provide additional results using ReLU transfer function on the MNIST dataset in Figure 6(a). Interestingly, K-FAC performs poorly on this problem and converges to a worse solution than RMSProp. LocoProp-M also performs well initially, but converges to a similar solution to RMSProp. Shampoo outperforms all other methods on this problem.

Hyper-parameter	Range	Scaling
γ	$[10, 1 \times 10^2, 5 \times 10^2, 10^3, 2 \times 10^3, 3 \times 10^3, 5 \times 10^3]$	Log
η	$[10^{-7}, 0.1]$	Log
$1 - \beta_1$	$[10^{-3}, 0.1]$	Log
$1 - \beta_2$	$[10^{-3}, 0.1]$	Log
ϵ	$[10^{-10}, 10^{-5}]$	Log

Table 5: Search space for LocoProp (both variants). Base optimizer is RMSProp.

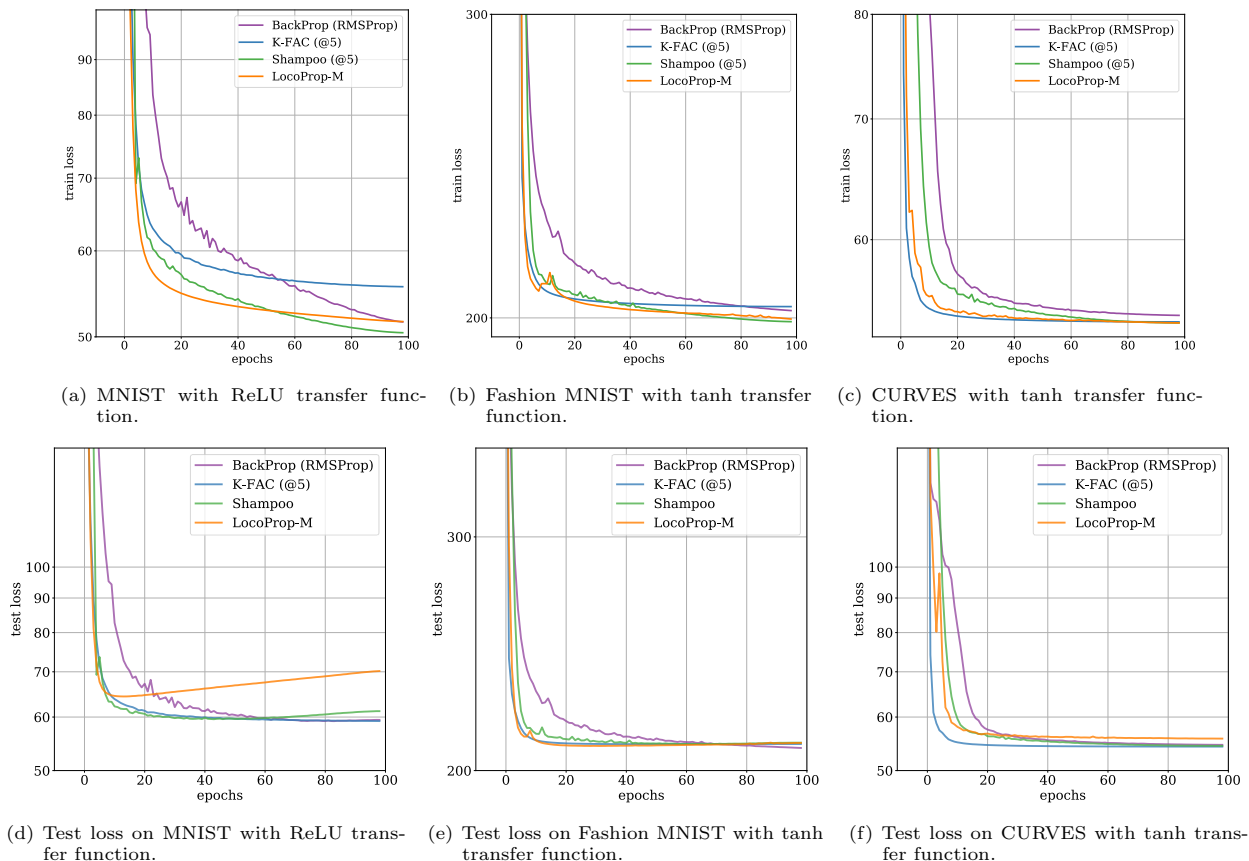
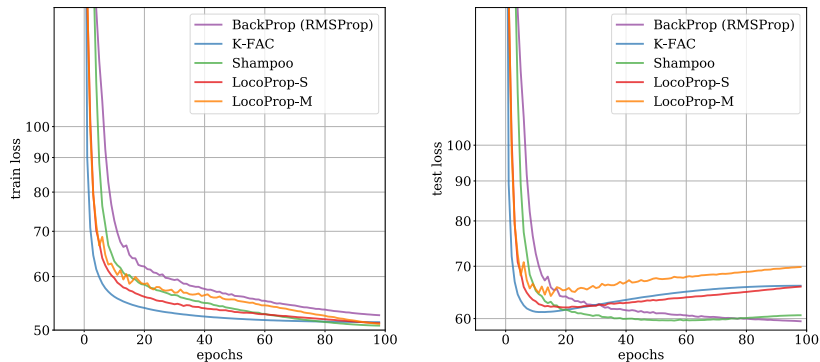


Figure 6: Results on different transfer functions and datasets. Comparisons on the standard autoencoder: (a) MNIST dataset and ReLU transfer function, (b) Fashion MNIST and (c) CURVES datasets using tanh transfer function. @k indicates the interval for carrying out inverse (p th root) operation.


We also show results on the Fashion MNIST and CURVES datasets (using the standard autoencoder with tanh transfer function) in Figure 6(b) and 6(c), respectively. Again, K-FAC performs poorly on Fashion MNIST, but works well on the CURVES dataset. LocoProp-M performs closely to the best performing second-order method.

Lastly, we also show the test loss in Figure 6(d)-(f) for these problems. The test loss results show signs of overfitting in all cases. The main reason for overfitting is that, in this work, we only focus on optimizing the fixed loss and thus, do not tune the weight decay regularizer term, which is set to $1e-5$ in the default autoencoder definition. Similar behavior on this problem has been observed in previous work using the default hyper-parameters (Goldfarb et al., 2020).



(a) LocoProp-S and LocoProp-M has comparable performance to second-order methods even with a smaller batch size. (b) Test loss on the same problem for different methods.

Figure 7: MNIST autoencoder with tanh transfer function trained with batch size of 100.



(a) Test loss on the standard autoencoder with tanh transfer function. (b) Test loss on the wide autoencoder with tanh transfer function. (c) Test loss on the deep autoencoder with tanh transfer function.

Figure 8: Test loss results on the MNIST dataset with tanh transfer function. Comparisons on: (a) standard, (b) wide, and (c) deep autoencoder variants.

H RESULTS WITH A SMALLER BATCH SIZE

We train the standard autoencoder on the MNIST dataset with a batch size of 100. Results are presented in Figure 7, which look similar to the earlier results at batch size 1000. One thing that stands out is that LocoProp-S works just as well in this setting as LocoProp-M.

I TEST LOSS FOR STANDARD, WIDE, AND DEEP VARIANTS

We provide the plots of the test loss for all variants of the autoencoder model in Figure 8. In general, the autoencoder model does not generalize well when trained with second-order methods or LocoProp. As discussed earlier, this behavior is due to an undertuned weight decay regularizer parameter for each case. In this work, we only focus on minimizing the training loss and defer the analysis of the generalization properties of LocoProp to future work.

J HYPER-PARAMETERS FOR EXPERIMENTS

The entire hyper-parameter sweeps along with the sensitivity analysis, training and test losses, and the code to reproduce the results is available at <https://github.com/google-research/google-research/tree/master/locoprop>.