Proximal Backpropagation

Thomas Frerix[*]
Technical University of Munich
thomas.frerix@tum.de

Thomas Möllenhoff[*]
Technical University of Munich
moellenh@in.tum.de

Michael Moeller[*]
University of Siegen
michael.moeller@uni-siegen.de

Daniel Cremers
Technical University of Munich
cremers@tum.de

Abstract

We offer a generalized point of view on the backpropagation algorithm, currently the most common technique to train neural networks via stochastic gradient descent and variants thereof. Specifically, we show that backpropagation of a prediction error is equivalent to sequential gradient descent steps on a quadratic penalty energy. This energy comprises the network activations as variables of the optimization and couples them to the network parameters. Based on this viewpoint, we illustrate the limitations on step sizes when optimizing a nested function with gradient descent. Rather than taking explicit gradient steps, where step size restrictions are an impediment for optimization, we propose proximal backpropagation (ProxProp) as a novel algorithm that takes *implicit gradient steps* to update the network parameters. We experimentally demonstrate that our algorithm is robust in the sense that it decreases the objective function for a wide range of parameter values. In a systematic quantitative analysis, we compare to related approaches on a supervised visual learning task (CIFAR-10) for fully connected as well as convolutional neural networks and for an unsupervised autoencoder (USPS dataset). We demonstrate that ProxProp leads to a significant speed up in training performance.

1   Introduction

In recent years neural networks have gained considerable attention in solving difficult correlation tasks such as classification in computer vision [8] or sequence learning [20] and as building blocks of larger learning systems [17]. Training neural networks is accomplished by optimizing a non-convex, possibly nonsmooth, nested function of the network parameters. Since the introduction of stochastic gradient descent (SGD) [15, 1], several more sophisticated optimization methods have been studied. One such class is that of quasi-Newton methods, as for example the comparison of L-BFGS with SGD in [9], Hessian-free approaches [10], and the Sum of Functions Optimizer in [18]. The latter uses BFGS to approximate the Hessian of every subfunction of a sum of differentiable functions, combining the computational efficiency of SGD with tractable second-order information. Several works consider specific properties of energy landscapes of certain deep learning models such as frequent saddle points [4] and well-generalizable local optima [3]. Among the most popular optimization methods in currently used deep learning frameworks are momentum based improvements of classical SGD, notably Nesterov's Accelerated Gradient [13, 19], and the Adam optimizer [7], which uses estimates of first and second order moments of the gradients for parameter updates.

Nevertheless, the optimization of these models remains challenging, as learning with SGD and its variants requires careful weight initialization and a sufficiently small learning rate in order to yield

---

[*]contributed equally.

a stable and convergent algorithm. Moreover, SGD often has difficulties in propagating a learning signal deeply into a network, commonly referred to as the vanishing gradient problem [6].

Recently, the authors of [2] have tackled the problem of optimizing the nested objective function by explicitly introducing the network activations as variables of the optimization, also known as the method of auxiliary coordinates (MAC). They use layer-wise constraints to couple the activation variables with the network parameters and propose a quadratic penalty method to solve the constrained problem. Closely related to the previous approach, Taylor et al. [21] introduce additional auxiliary variables to further split linear and nonlinear transfer between layers.

In this work we draw a connection between the penalty formulations [2, 21] and the classical backpropagation algorithm [16]. In particular, we show that the backpropagation algorithm can be interpreted as a method alternating between two steps. First, a forward pass that resets the auxiliary activation variables in accordance with the previously computed weights. Secondly, an ordered sequence of gradient descent steps on a quadratic penalty energy.

Interestingly, for many common network architectures the updates arising in the second step still have a closed form solution if the explicit gradient descent step is replaced by an implicit gradient step (also known as proximal step, for the definition see (8), Section 4.1). We investigate different combinations of explicit and implicit optimization subproblems and demonstrate that implicit gradient steps can significantly improve optimization progress.

2 Notation and model

We denote the Euclidean norm for vectors and the Frobenius norm for matrices by $|| \cdot ||$, induced by an inner product $\langle \cdot, \cdot \rangle$. We use the gradient symbol ∇ to denote the transpose of the Jacobian matrix, such that the chain rule applies in the form "inner derivative times outer derivative". For all computations involving matrix-valued functions and their gradient/Jacobian, we uniquely identify all involved quantities with their vectorized form by flattening matrices in a column-first order.

The method we propose in this work is suitable for optimizing a parametrized function composition:

$$f(\boldsymbol{\theta}) = f_K \left(\theta_K, f_{K-1} \left(\theta_{K-1}, \cdots (\theta_2, f_1 (\theta_1)) \cdots \right) \right). \tag{1}$$

Note that almost any common neural network architecture can be phrased in the form of (1).

In this paper we consider fully connected feedforward neural networks as well as convolutional neural networks. For the sake of notational simplicity we will limit the following discussion to the example of a fully connected L-layer network with nonlinear activation functions before the $(L-2)$ hidden layers and an affine transformation at the last layer. We will remark the difference for convolutional architectures where necessary. For now we denote the input batch as $X \in \mathbb{R}^{n_0 \times N}$, where N is the batch size and n_0 is the number of input neurons. The subsequent dimensions of the layers are denoted as $n_1, \ldots, n_{L-1}$. The network parameters $\boldsymbol{\theta} = (\boldsymbol{W}, \boldsymbol{b})$ consist of weights $\boldsymbol{W} := (W_1, \ldots, W_{L-1})$ and network biases $\boldsymbol{b} := (b_1, \ldots, b_{L-1})$, with $W_l \in \mathbb{R}^{n_l \times n_{l-1}}$ and $b_l \in \mathbb{R}^{n_l}$. The targets for the batch are denoted as $y \in \mathbb{R}^{n_{L-1} \times N}$ and we use the convention $a_0 := X$.
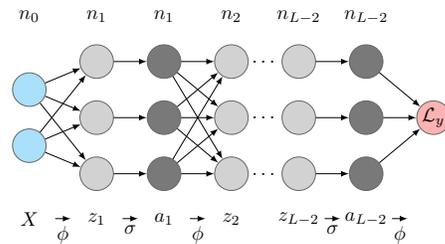

Figure 1: Fully connected L-layer feedforward neural network with explicit layerwise activation variables a, z and affine and nonlinear transfer functions ϕ and σ.

Our nonlinearities are either rectified linear units (ReLU), $\sigma(x) = \max(0, x)$, or sigmoid units, $\sigma(x) = 1/(1 + \exp(-x))$. For convolutional neural networks we also combine the ReLU with a max pooling nonlinearity and we make our choice explicit where needed. When writing $\sigma(z_l)$ for some matrix $z_l \in \mathbb{R}^{n_l \times N}$ it is to be understood elementwise. The final objective J is a nested function consisting of the prediction loss $\mathcal{L}_y : \mathbb{R}^{n_{L-1} \times N} \to \mathbb{R}$ evaluated at layer L. More precisely, let us define the function $\phi(W, b, a) = Wa + b\mathbf{1}$, for matrices W and a, a column vector b, and a row vector $\mathbf{1} \in \mathbb{R}^{1 \times N}$. Note that a convolution is also a linear operator which is why it can be written as

2

Algorithm 1 - Penalty formulation of backprop.	**Algorithm 2** - Generalization of backprop.
Input: Current parameters $(\boldsymbol{W}, \boldsymbol{b})^k$.	**Input:** Current parameters $(\boldsymbol{W}, \boldsymbol{b})^k$.
// Forward pass.	*// Forward pass.*
for $l = 1$ **to** $L - 2$ **do**	**for** $l = 1$ **to** $L - 2$ **do**
$\quad z_l^{(k)} = \phi(W_l^{(k)}, b_l^{(k)}, a_{l-1}^{(k)}), \qquad$ *// $a_0 = X$.*	$\quad z_l^{(k)} = \phi(W_l^{(k)}, b_l^{(k)}, a_{l-1}^{(k)}), \qquad$ *// $a_0 = X$.*
$\quad a_l^{(k)} = \sigma(z_l^{(k)})$.	$\quad a_l^{(k)} = \sigma(z_l^{(k)})$.
end for	**end for**
// Backward pass by τ-gradient steps on (4).	*// Perform minimization steps on* (4).
ⓐ grad. step on E wrt. $(W_{L-1}, b_{L-1}, a_{L-2})$.	ⓐ min. step on E wrt. $(W_{L-1}, b_{L-1}, a_{L-2})$.
for $l = L - 2$ **to** 1 **do**	**for** $l = L - 2$ **to** 1 **do**
$\quad$ⓑ grad. step on E wrt. z_l,	$\quad$ⓑ min. step on E wrt. z_l,
$\quad$ⓒ grad. step on E wrt. (W_l, b_l, a_{l-1}).	$\quad$ⓒ min. step on E wrt. (W_l, b_l, a_{l-1}).
end for	**end for**
Output: New parameters $(\boldsymbol{W}, \boldsymbol{b})^{k+1}$.	**Output:** New parameters $(\boldsymbol{W}, \boldsymbol{b})^{k+1}$.

a matrix-vector product. Therefore, except for the matrix W being constrained to be of a particular structure, the form of the problem remains the same for convolutional layers.

Using the above notation we can express the final objective of the fully connected feed-forward network analogous to (1) as

$$J(\boldsymbol{W}, \boldsymbol{b}; X, y) = \mathcal{L}_y(\phi(W_{L-1}, b_{L-1}, \sigma(\phi(\cdots, \sigma(\phi(W_1, b_1, X))\cdots)))). \tag{2}$$

Formally, the functions σ and ϕ map between spaces of different dimensions depending on the layer. However, to keep the presentation clean we do not state this dependence explicitly. Figure 1 illustrates our notation for the fully-connected network architecture.

3   Penalty formulation of backpropagation

The gradient descent iteration on a nested function $J(\boldsymbol{W}, \boldsymbol{b}; X, y)$ like (2),

$$(\boldsymbol{W}, \boldsymbol{b})^{k+1} = (\boldsymbol{W}, \boldsymbol{b})^k - \tau \nabla J(\boldsymbol{W}^k, \boldsymbol{b}^k; X, y), \tag{3}$$

is commonly implemented using the backpropagation algorithm [16]. As the basis for the proposed proximal optimization method, we will now derive a connection between the classical backpropagation algorithm and quadratic penalty functions of the form

$$E(\boldsymbol{W}, \boldsymbol{b}, \boldsymbol{a}, \boldsymbol{z}) = \mathcal{L}_y(\phi(W_{L-1}, b_{L-1}, a_{L-2})) + \sum_{l=1}^{L-2} \frac{\gamma}{2} \|\sigma(z_l) - a_l\|^2 + \frac{\rho}{2} \|\phi(W_l, b_l, a_{l-1}) - z_l\|^2. \tag{4}$$

Under mild conditions the limit $\rho, \gamma \to \infty$ leads to the convergence of the sequence of minimizers of E to the minimizer of J, see [14, Theorem 17.1], also [2, Theorem B.1-B.4].

One of the main insights of this paper, which motivates a novel class of optimization algorithms, is that the iteration shown in Algorithm 1 of forward passes followed by a sequential gradient descent on the penalty function E is equivalent to the classical gradient descent iteration.

Proposition 1. *For $\rho = \gamma = 1/\tau$ and $(\boldsymbol{W}, \boldsymbol{b})^k$ as the input to Algorithm 1, its output $(\boldsymbol{W}, \boldsymbol{b})^{k+1}$ satisfies* (3), *i.e., Algorithm 1 computes one gradient descent iteration on J.*

Proof. See appendix. $\square$

Observe that although the above proposition refers to our specific setting no assumptions on σ and ϕ (except their differentiability) are made in the proof of Proposition 1.

4 Generalizing backpropagation

The interpretation of Proposition 1 leads to the natural idea of replacing the explicit gradient steps ⓐ, ⓑ and ⓒ in Algorithm 1 with other – possibly more powerful – minimization steps. For example, for $l = 1$ step ⓒ of Algorithm 1 is

$$(W_1, b_1)^{k+1} = (W_1, b_1)^k - \tau\rho\nabla f((W_1, b_1)^k) \tag{5}$$

for $f(W, b) = \frac{1}{2}\|\phi(W, b, X) - z_1\|^2$. As we will discuss in the next subsection, replacing Equation (5) by an *implicit gradient step*

$$(W_1, b_1)^{k+1} = (W_1, b_1)^k - \tau\rho\nabla f((W_1, b_1)^{\mathbf{k+1}}) \tag{6}$$

has several advantageous properties and therefore motivates the general concept presented in Algorithm 2, i.e., alternating between forward passes of the network and sequential minimization steps on the penalty function (4). In the following section we explore novel minimization strategies for ⓐ, ⓑ and ⓒ in Algorithm 2 by replacing the explicit gradient descent steps with implicit steps known as *proximal mappings*.

4.1 Implicit gradient steps

The proximal mapping of a function $f : \mathbb{R}^n \to \mathbb{R}$ is defined as the following [12]:

$$\text{prox}_{\tau f}(y) := \underset{x \in \mathbb{R}^n}{\text{argmin}}\ f(x) + \frac{1}{2\tau}\|x - y\|^2. \tag{7}$$

By rearranging the optimality conditions to (7) and taking $y = x^k$, it can be interpreted as an *implicit gradient step* evaluated at the new point x^{k+1}:

$$x^{k+1} := \underset{x \in \mathbb{R}^n}{\text{argmin}}\ f(x) + \frac{1}{2\tau}\|x - x^k\|^2 = x^k - \tau\nabla f(x^{k+1}). \tag{8}$$

The advantage of such proximal steps is that in contrast to explicit gradient descent, the update scheme (8), also known as the proximal point algorithm [11], is *unconditionally stable*. Unconditional stability means that (8) monotonically decreases the energy f for any $\tau > 0$, as by definition of x^{k+1} it holds that $f(x^{k+1}) + \frac{1}{2\tau}\|x^{k+1} - x^k\|^2 \leq f(x^k)$. Within our model, we can recover the proximal point algorithm (8) as proximal optimization of the prediction loss (4) for a 2-layer network.

Note that explicit gradient steps pose severe restrictions on the allowed step size τ: Even for a convex, twice continuously differentiable, $\mathscr{L}$-smooth function $f : \mathbb{R}^n \to \mathbb{R}$, the convergence of the gradient descent algorithm can only be guaranteed for step sizes $0 < \tau < 2/\mathscr{L}$. The Lipschitz constant $\mathscr{L}$ of the gradient ∇f is in this case equal to the largest eigenvalue of the Hessian H. For the example of the first layer shown in (5), the Hessian is $H = \begin{bmatrix} X & \mathbf{1} \end{bmatrix}\begin{bmatrix} X & \mathbf{1} \end{bmatrix}^\top$, which for the CIFAR-10 dataset has a largest eigenvalue of $6.7 \cdot 10^6$. Similar problems also arise in other layers where poorly conditioned matrices a_l pose severe restrictions for guaranteeing the energy of the subproblem to decrease.

Note that the instability of the substeps does not need to imply the instability of the gradient descent algorithm. Nevertheless, proximal mappings yield unconditionally stable subproblems, which motivates us to use proximal steps in Algorithm 2.

In many cases, the full proximity operator in steps ⓐ and ⓒ is difficult to compute exactly due to the nonconvex product of W and a in the function ϕ. Therefore, we use one iteration of coordinate descent in the two blocks (W, b) and a as an approximate solution. For the two block coordinate descent steps, we have the option to take proximal steps, or make the updates explicit. We discuss the proximal steps one can possibly take on each of the variables (W, b), a, and z in the next subsection.

Of particular importance for our proposed scheme is the proximity operator of the squared ℓ^2 norm

$$\text{prox}_{\frac{\varsigma}{2}\|T\cdot -c\|_2^2}(d) = \arg\min_x \frac{1}{2}\|x - d\|^2 + \frac{\zeta}{2}\|T \cdot -c\|_2^2 = (I + \zeta T^*T)^{-1}(d + \zeta T^*c), \tag{9}$$

for different linear operators T. In the above T^* denotes the adjoint operator (which would be the transpose in the case of matrices). For any linear operator T, equation (9) merely requires the solution of a linear system of equations.

4.2 Computing the proximal steps

In this subsection we discuss the proximal operators for taking implicit gradient steps with respect to our variables W, b, a, and z. For the sake of simplicity, we limit the discussion to the fully connected case. Please note that convolutional layers can be treated in a similar fashion, which we detail in the appendix.

4.2.1 Proximal steps in (W, b)

Last layer In the last layer one needs to solve

$$(W_{L-1}^{(k+1)}, b_{L-1}^{(k+1)}) = \underset{W,b}{\operatorname{argmin}} \ \mathcal{L}_y(W a_{L-2}^{(k)} + b\mathbf{1}) + \frac{1}{2\tau}||W - W_{L-1}^{(k)}||^2 + \frac{1}{2\tau}||b - b_{L-1}^{(k)}||^2.$$
(10)

Depending on the choice of the loss function $\mathcal{L}_y$, this leads to different subproblems.

For the square loss $\mathcal{L}_y(z) = \frac{1}{2}||z - y||^2$, the proximity operator has a closed form solution given by (9). We detail the update formulas in the appendix.

For $\mathcal{L}_y(z) = -\sum_{i=1}^{N} \sum_{j=1}^{n_{L-1}} y_{j,i} z_{j,i} - \log \sum_{j=1}^{n_{L-1}} \exp(z_{j,i})$, the proximity operator does not have a simple, closed form solution. While in principle it can be solved using algorithms from convex optimization, we found that it is more efficient to simply take an explicit gradient step in this case.

First and hidden layers In all other layers, one needs to determine

$$(W_l^{(k+1)}, b_l^{(k+1)}) = \underset{W,b}{\operatorname{argmin}} \ \frac{\rho}{2}||W a_{l-1}^{(k)} + b\mathbf{1} - z_l^{(k+1)}||^2 + \frac{1}{2\tau}||W - W_l^{(k)}||^2 + \frac{1}{2\tau}||b - b_l^{(k)}||^2$$
(11)

with the convention $a_0 = X$, which again resembles a quadratic proximity operator (9).

4.2.2 Proximal steps in a

Last layer Similar to the above case of (W, b) one needs to solve

$$a_{L-2}^{(k+1)} = \underset{a}{\operatorname{argmin}} \ \mathcal{L}_y(W_{L-1}^{(k+1)} a + b_{L-1}^{(k+1)} \mathbf{1}) + \frac{\gamma}{2}||a - a_{L-2}^{(k)}||^2 + \frac{1}{2\tau}||a - a_{L-2}^{(k)}||^2,$$
(12)

where we substituted $\sigma(z_{L-2}^{(k)}) = a_{L-2}^{(k)}$ in the γ-penalty term. This is a valid substitution due to the forward pass. Again, for a square loss the above problem becomes an instance of (9), and we propose to take explicit steps for a softmax loss.

First and hidden layers All layers except the last result in the problem

$$a_{l-1}^{(k+1)} = \underset{a}{\operatorname{argmin}} \ \frac{1}{2}||W_l^{(k+1)} a + b_l^{(k+1)} \mathbf{1}||_2^2 + \frac{\gamma}{2}||a - a_{l-1}^{(k)}||^2 + \frac{1}{2\tau}||a - a_{l-1}^{(k)}||^2,$$
(13)

and again amount to an instance of (9).

4.2.3 Proximal steps in z

Taking a proximal step on E with respect to z_l, i.e. taking implicit steps in ⓑ, amounts to the nonconvex minimization problem

$$z_l^{(k+1)} = \underset{z}{\operatorname{argmin}} \ \frac{\gamma}{2}||\sigma(z) - a_l^{(k+1)}||^2 + \frac{\rho}{2}||z_l^{(k)} - z||^2 + \frac{1}{2\tau}||z - z_l^{(k)}||^2,$$
(14)

where we substituted the variable $z_l^{(k)}$ from the forward pass. Problem (14) decomposes into independent univariate problems which can be solved to global optimality for $\sigma(x) = \max(0, x)$ despite its nonconvexity. We state the explicit solution of (14) in the appendix.
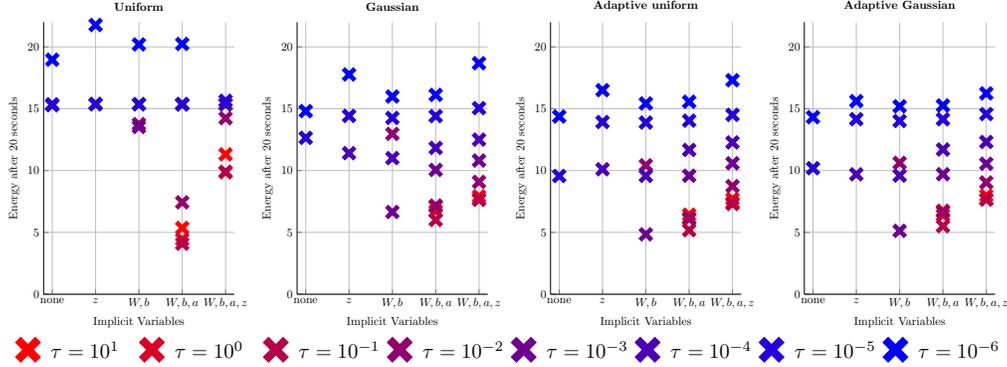
Figure 2: Final energies for implicit and explicit optimization of subproblems with different step sizes and initializations after a runtime of 20s. The color of the data points represents the step size and for missing step sizes the method did not decrease the initial energy.

4.3 Explicit vs. implicit steps

Proposition 1 shows that Algorithm 2 with explicit gradient descent steps and $\gamma = \rho = 1/\tau$ yields the gradient descent iteration. Motivated by the more general scheme of Algorithm 2, we consider the following variants: 1. Explicit steps on all variables, 2. implicit steps on z_l, explicit steps on (W_l, b_l, a_{l-1}), 3. explicit steps on (z_l, a_{l-1}), explicit steps on (W_l, b_l), 4. explicit steps on z_l, implicit steps on (W_l, b_l, a_{l-1}), and 5. implicit steps on all variables.

To investigate their stability, we train a fully connected $256 - 300 - 100 - 20 - 100 - 300 - 256$ autoencoder with ReLU nonlinearities and square loss using the above update equations. We train on the USPS dataset with 2000 grayscale images of size 16×16 and set $\gamma = \rho = 1/\tau$ for all five variants.

Each method is run with stepsizes $\tau \in \{10, 1, 10^{-1}, 10^{-2}, 10^{-3}, 10^{-4}, 10^{-5}, 10^{-6}\}$ and for four different initializations of the weights: A Gaussian distribution with variance 0.1, a uniform distribution of values in $[-0.1, 0.1]$ as well as the adaptive initialization scheme suggested in [5], denoted as adaptive uniform and adaptive Gaussian respectively. Figure 2 illustrates the final energies of each combination after training for 20 seconds. For the case of adaptive uniform weight initialization, Figure 3 furthermore shows the decay of the training energy over the iterations using the optimal step size for each method.

While the more implicit methods yield decaying energies for a much wider range of different step sizes, implicit steps do not necessarily yield the fastest decay of the energy: Taking implicit steps in (W, b, a) and explicit ones in z usually yields lower energies than the fully implicit method.

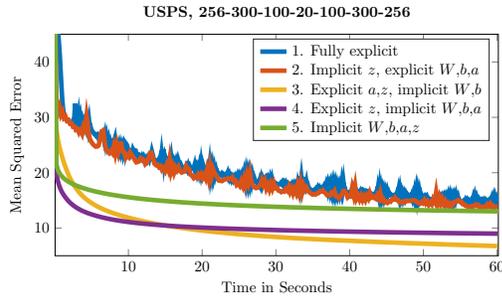

Figure 3: Comparison of implicit with explicit steps for the subproblems. The combination of implicit steps on (W, b) and explicit steps on (a, z) achieves most optimization progress per time and is further evaluated in Section 6.2.

Moreover, we can see that for adaptive initializations the method using implicit steps in (W, b) and explicit steps in (a, z) experienced the fastest convergence. Altogether, we therefore focus on this method in our further evaluation in Section 6.

5 Proximal backpropagation as a first-order oracle

Algorithm 2 can be used as a gradient oracle for first-order methods (such as Adam [7]) or quasi-Newton methods (such as SFO [18]), by using it to compute a direction $\boldsymbol{d}^{k+1} = (\boldsymbol{W}, \boldsymbol{b})^{k+1} - (\boldsymbol{W}, \boldsymbol{b})^k$ instead of the usual gradient. Indeed, as seen in Proposition 1, for $\rho = \gamma = \frac{1}{\tau}$ and explicit
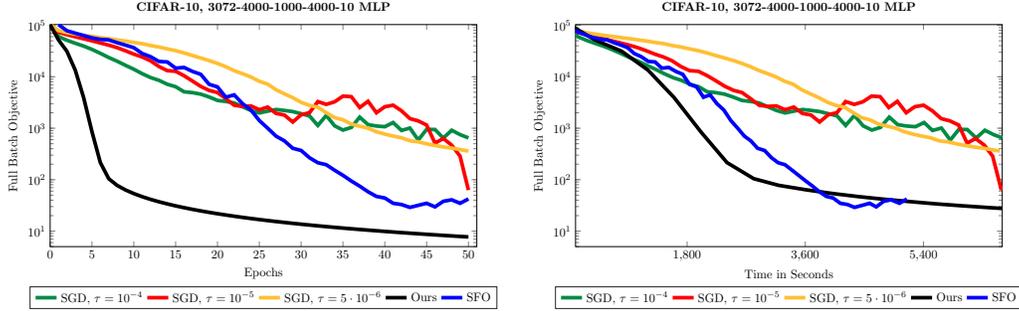
Figure 4: Comparison of different optimization algorithms evaluated on the CIFAR-10 full batch supervised learning objective function shown on a log-scale. Our method clearly outperforms SGD, and performs comparably to the recent work SFO [18] while being more memory efficient.

gradient steps, one obtains the direction of steepest descent

$$\boldsymbol{d}^{k+1} = -\tau \nabla_{(\boldsymbol{W}, \boldsymbol{b})} J(\boldsymbol{W}^k, \boldsymbol{b}^k; X, y). \tag{15}$$

Note that in general it is not clear if $\boldsymbol{d}$ is a descent direction for the energy J, i.e., whether $\langle \nabla J, \boldsymbol{d} \rangle < 0$. This property is often assumed in convergence proofs for gradient-based methods. Nevertheless, in practice we observed convergence over a wide range of parameters. For the experiments in this paper, we stick to SGD with Nesterov momentum μ, which is given as:

$$\boldsymbol{v}^{k+1} = \mu \boldsymbol{v}^k + \boldsymbol{d}^{k+1}, \qquad (\boldsymbol{W}, \boldsymbol{b})^{k+1} = (\boldsymbol{W}, \boldsymbol{b})^k + \boldsymbol{d}^{k+1} + \mu \boldsymbol{v}^{k+1}. \tag{16}$$

In the following section, we will refer to SGD with Nesterov momentum as **SGD**. **Ours** will be (16), where $\boldsymbol{d}^{k+1}$ is computed using ProxProp (Algorithm 2 with explicit updates in the network activations a, z and proximal steps in the network parameters W, b).

6   Numerical evaluation

We evaluate our method in two representative experiments, a supervised visual learning problem on the CIFAR-10 dataset and an autoencoder on the USPS handwritten digit dataset.

6.1   CIFAR-10 supervised learning

Fully connected network   We trained a fully connected network with architecture $3072 - 4000 - 1000 - 4000 - 10$, ReLU nonlinearities and cross-entropy + softmax loss. The weights are initialized according to the uniform distribution scheme described in [5] and a weight-decay of 10^{-6} was chosen. In Fig. 4, we compare SGD to the proposed method (ours). The momentum was set to $\mu = 0.95$ and the batch size chosen as 250 for both methods. For the proposed method, we picked $\tau = 5 \cdot 10^{-3}$ and the penalty parameters were chosen as $\rho = 10/\tau, \gamma = 2.5/\tau$. Since the proximal step for the softmax loss does not have a closed form solution (see Section 4.2.1), we performed a single explicit τ-gradient step to approximately solve the subproblem in (13). We further compare the proposed method to SFO [18] and used the publicly available MATLAB implementation[2] with 20 subfunctions (minibatch size of 2250). As the memory requirement of SFO is linear in the minibatches, a larger number of subfunctions exceeded the available memory. All methods were implemented in MATLAB to have a comparable performance. From Fig. 4 we infer that the proposed method leads to lower energies in less time and iterations, and also to the overall lowest energy. The test accuracies reached by the different methods were 54% for SFO, 54% - 57% for SGD and 56% for ours.

Convolutional neural network   Our ProxProp algorithm can be used with any architecture with a linear transfer function to which we apply the proximal updates. To demonstrate our algorithm on a

---

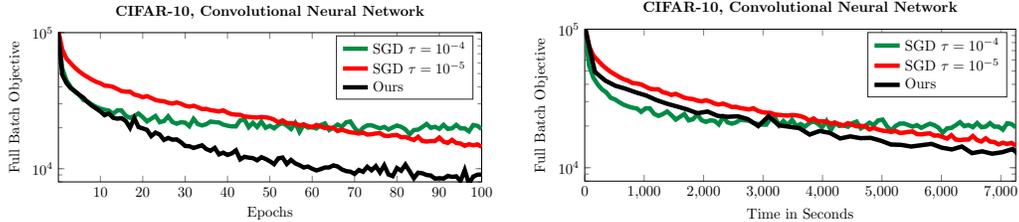[2]`https://github.com/Sohl-Dickstein/Sum-of-Functions-Optimizer`

7

Figure 5: Comparison of our Lasagne/Theano ProxProp optimizer to Lasagne's Nesterov SGD evaluated on the CIFAR-10 full batch supervised learning objective function. ProxProp reaches a lower energy, while achieving a comparable generalization accuracy. Note that the objective value is shown on a log-scale.
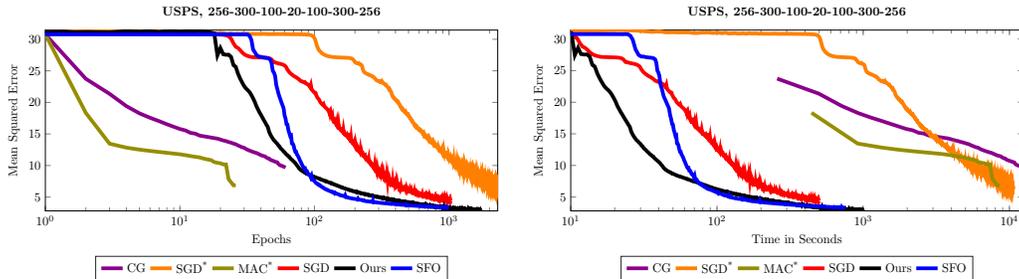

Figure 6: Comparison of different optimization algorithms with respect to a USPS autoencoder mean squared error. Note that the time and iteration axes are shown on a log-scale. Our method outperforms MAC [2] by several orders of magnitude (w.r.t. runtime) and also minimizes the energy significantly faster also than SGD. The methods denoted with an asterisk are based on code kindly provided by [2].

generic architecture with layers commonly used in practice, we trained on the convolutional neural network of the form:

$$\text{Conv}[16 \times 32 \times 32] \rightarrow \text{ReLU} \rightarrow \text{Pool}[16 \times 16 \times 16] \rightarrow \text{Conv}[20 \times 16 \times 16] \rightarrow \text{ReLU}$$

$$\rightarrow \text{Pool}[20 \times 8 \times 8] \rightarrow \text{Conv}[20 \times 8 \times 8] \rightarrow \text{ReLU} \rightarrow \text{Pool}[20 \times 4 \times 4] \rightarrow \text{FC} + \text{Softmax}[10 \times 1 \times 1]$$

Here, the first dimension denotes the respective number of filters with kernel size 5×5 and max pooling downsamples its input by a factor of two. For this experiment, we chose $\tau = 5 \cdot 10^{-3}, \rho = 10/\tau, \gamma = 2.5/\tau, \mu = 0.9$ and a batch size of 250. We compared with a Nesterov SGD optimizer with parameters $\tau = 1 \cdot 10^{-4}/\tau = 1 \cdot 10^{-5}, \mu = 0.95$. To this end, we have implemented our algorithm as an optimizer in Lasagne/Theano [22] and compared to Lasagne's Nesterov SGD optimizer. Training was executed on the CPU of a machine with 8 Intel Xeon CPU E5-2637 with 3.50GHz. The results are shown in Fig. 5. The ProxProp optimizer achieves a lower energy on this nonconvex full batch training objective, while arriving at a comparable generalization accuracy on the test set, namely 65% for ProxProp, 66% for SGD with step size $\tau = 1 \cdot 10^{-4}$ and 68% for SGD with step size $\tau = 1 \cdot 10^{-5}$. Note that we have focused on the conceptual difference in parameter updates and have not performed extensive hyperparameter tuning for either of the methods.

6.2 USPS autoencoder

To compare with MAC [2], we trained a fully connected autoencoder with the same architecture as for our experiments in Section 4.3, used sigmoid activation functions, and trained on the USPS dataset with 5000 grayscale images of size 16×16 using minibatches of size 250. We used the same initialization as [2] with the initial weights in layer l being uniformly sampled from $[-1/\sqrt{n_{l-1}}, 1/\sqrt{n_{l-1}}]$. The authors of [2] kindly provided us with their MATLAB code to reproduce the results of their parallel MAC method, as well as the SGD and CG methods they compared to. We are therefore able to reproduce Figure 2 in [2] without the minibatch MAC on our computer with an Intel Core i7-6700HQ CPU with 2.60GHz,and MATLABs parpool parallelizing MAC on 4 cores.

Figure 6 shows the proposed semi-implicit method along with our implementation of a minibatch-SGD method with momentum, as well as the SFO algorithm. We used $\tau = 5 \cdot 10^{-5}$ and $\tau = 5 \cdot 10^{-4}$ as a stepsize for the SGD and our semi-implicit algorithm respectively, used a momentum of $\mu = 0.99$ for both methods, and set $\rho = \gamma = 1/\tau$. As we can see in Fig. 6 the resulting advantage of our approach to explicit SGD steps and competing approaches like SFO and MAC becomes significant, particularly in the earlier phase of training.

7 Conclusion

We have proposed ProxProp as an efficient novel method for training neural networks. To this end, we first showed the equivalence of the classical backprop algorithm with an algorithm that alternates between sequential gradient steps on a quadratic penalty function and forward passes through the network. Subsequently, we developed a generalization of backprop, which replaces explicit gradient steps with implicit (proximal) steps. Numerical experiments demonstrate that ProxProp yields a stable decrease of the objective function on a much wider range of step sizes than classical backprop. A performance comparison to related minimization methods shows that our algorithm exhibits faster convergence and often yields overall lower nonconvex objective values. We believe that the proposed framework creates an important bridge between ideas from convex optimization and deep learning.

Appendix

A   Theoretical results

Proof of Proposition 1. We first take a gradient step on

$$
E(\boldsymbol{W}, \boldsymbol{b}, \boldsymbol{a}, \boldsymbol{z}) = \mathcal{L}_y(\phi(W_{L-1}, b_{L-1}, a_{L-2}))
$$
$$
+ \frac{\gamma}{2} \sum_{l=1}^{L-2} \|\sigma(z_l) - a_l\|^2 + \frac{\rho}{2} \sum_{l=1}^{L-2} \|\phi(W_l, b_l, a_{l-1}) - z_l\|^2, \tag{17}
$$

with respect to $(W_{L-1}, b_{L-1}, a_{L-2})$. The gradient step with respect to (W_{L-1}, b_{L-1}) is the same as in the gradient descent update,

$$
(\boldsymbol{W}, \boldsymbol{b})^{k+1} = (\boldsymbol{W}, \boldsymbol{b})^k - \tau \nabla J(\boldsymbol{W}^k, \boldsymbol{b}^k; X, y), \tag{18}
$$

since J depends on (W_{L-1}, b_{L-1}) only via $\mathcal{L}_y \circ \phi$.

The gradient descent step on a_{L-2} in ⓐ yields

$$
a_{L-2}^{(k+1/2)} = a_{L-2}^{(k)} - \tau \nabla_a \phi(W_{L-1}^{(k)}, b_{L-1}^{(k)}, a_{L-2}^{(k)}) \cdot \nabla_\phi \mathcal{L}_y(\phi(W_{L-1}^{(k)}, b_{L-1}^{(k)}, a_{L-2}^{(k)})), \tag{19}
$$

where we use $a_{L-2}^{(k+1/2)}$ to denote the updated variable a_{L-2} before the forward pass of the next iteration.

For all layers $l \leq L - 2$ note that due to the forward pass in Algorithm 1 we have

$$
\sigma(z_l^{(k)}) = a_l^{(k)}, \quad \phi(W_l^{(k)}, b_l^{(k)}, a_{l-1}^{(k)}) = z_l^{(k)}
$$

and we therefore get the following update equations in the gradient step ⓑ

$$
z_l^{(k+1/2)} = z_l^{(k)} - \tau \gamma \nabla \sigma(z_l^{(k)}) \left(\sigma(z_l^{(k)}) - a_l^{(k+1/2)} \right) = z_l^{(k)} - \nabla \sigma(z_l^{(k)}) \left(a_l^{(k)} - a_l^{(k+1/2)} \right), \tag{20}
$$

and in the gradient step ⓒ w.r.t. a_{l-1},

$$
\begin{aligned}
a_{l-1}^{(k+1/2)} &= a_l^{(k)} - \tau \rho \nabla_a \phi(W_l^{(k)}, b_l^{(k)}, a_{l-1}^{(k)}) \cdot \left(\phi(W_l^{(k)}, b_l^{(k)}, a_{l-1}^{(k)}) - z_l^{(k+1/2)} \right) \\
&= a_l^{(k)} - \nabla_a \phi(W_l^{(k)}, b_l^{(k)}, a_{l-1}^{(k)}) \cdot \left(z_l^{(k)} - z_l^{(k+1/2)} \right).
\end{aligned} \tag{21}
$$

Equations (20) and (21) can be combined to obtain:

$$
z_l^{(k)} - z_l^{(k+1/2)} = \nabla \sigma(z_l^{(k)}) \nabla_a \phi(W_{l+1}^{(k)}, b_{l+1}^{(k)}, a_l^{(k)}) \cdot \left(z_{l+1}^{(k)} - z_{l+1}^{(k+1/2)} \right). \tag{22}
$$

9

The above formula allows us to backtrack the differences of the old $z_l^{(k)}$ and the updated $z_l^{(k+1/2)}$ up to layer $L - 2$, where we can use equations (20) and (19) to relate the difference to the loss. Altogether, we obtain

$$z_l^{(k)} - z_l^{(k+1/2)} = \tau \left(\prod_{q=l}^{L-2} \nabla \sigma(z_q^{(k)}) \nabla_a \phi(W_{q+1}^{(k)}, b_{q+1}^{(k)}, a_q^{(k)}) \right) \cdot \nabla_\phi \mathcal{L}_y(\phi(W_{L-1}^{(k)}, b_{L-1}^{(k)}, a_{L-2}^{(k)})).$$

(23)

By inserting (23) into the gradient descent update equation with respect to (W_l, b_l) in Ⓒ,

$$(W_l, b_l)^{k+1} = (W_l, b_l)^k - \nabla_{(W,b)} \phi(W_l^{(k)}, b_l^{(k)}, a_{l-1}^{(k)}) \cdot \left(z_l^{(k)} - z_l^{(k+1/2)} \right),$$

(24)

we obtain the chain rule for update (18). $\qquad\square$

B Solution of the proximal subproblems

B.1 Proximal operator for the square loss

The solution of the problem

$$(W_{L-1}^{(k+1)}, b_{L-1}^{(k+1)}) = \underset{W,b}{\operatorname{argmin}} \; \mathcal{L}_y(W a_{L-2}^{(k)} + b\mathbf{1}) + \frac{1}{2\tau} ||W - W_{L-1}^{(k)}||^2 + \frac{1}{2\tau} ||b - b_{L-1}^{(k)}||^2.$$

(25)

for the choice $\mathcal{L}_y(z) = \frac{1}{2}||z - y||^2$ is given by

$$(W_{L-1}, b_{L-1})^{k+1} = \left[ya^\top + \frac{1}{\tau} W \quad y\mathbf{1}^\top + \frac{1}{\tau} b \right] \left(\begin{bmatrix} aa^\top & a\mathbf{1}^\top \\ \mathbf{1}a^\top & \mathbf{1}\mathbf{1}^\top \end{bmatrix} + \frac{1}{\tau} I \right)^{-1},$$

(26)

where $a \leftarrow a_{L-2}^{(k)}$, $W \leftarrow W_{L-1}^{(k)}$, $b \leftarrow b_{L-1}^{(k)}$ and y is the ground-truth of the training data. Interestingly, since we have $a \in \mathbb{R}^{n_{L-2} \times N}$ and $\mathbf{1} \in \mathbb{R}^{1 \times N}$, the size of the linear system is independent of the batch size N.

B.2 Proximal operator for ReLU penalty term

It can be quickly verified by checking the individual cases, that the global solution of the one-dimensional nonconvex optimization problem

$$\widehat{x} = \underset{x \in \mathbb{R}}{\operatorname{argmin}} \; \frac{1}{2} \left(\max(0, x) - a \right)^2 + \frac{1}{2\alpha} \left(x - z \right)^2,$$

(27)

is given by the following:

$$\widehat{x} = \begin{cases} \frac{\alpha a + z}{\alpha + 1}, & \text{if } (z \geq -\alpha a) \wedge \left[(z \geq 0) \vee \left((z < 0) \wedge (\frac{(a-z)^2}{2(\alpha+1)} < \frac{a^2}{2}) \right) \right], \\ z, & \text{if } (z < 0) \wedge \left[(z < -\alpha a) \vee \left((z \geq -\alpha a) \wedge (\frac{(a-z)^2}{2(\alpha+1)} > \frac{a^2}{2}) \right) \right], \\ 0, & \text{otherwise.} \end{cases}$$

(28)

B.3 Proximal operator for convolutional layers

For some input activation a of shape (n_b, n_c, n_y, n_x) with batch size n_b, number of input channels n_c, sample dimensions $n_y \times n_x$, respectively, and a filterbank K with n_f filters and quadratic kernel size n_k of dimension (n_f, n_c, n_k, n_k), we denote the operation of the convolutional layer by $y = K * a$ (with bias absorbed in K). For a proximal update of the network parameters K we then have to solve the following problem for some parameter $\tau > 0$:

$$K^{(k+1)} = \underset{K}{\operatorname{argmin}} \; \frac{1}{2} ||K * a - y||^2 + \frac{1}{2\tau} ||K - K^{(k)}||^2.$$

(29)

To this end, we explicitly construct a dense matrix representation of the convolution operation. This matrix contains a flattened version of every input image shifted to all positions of the convolution

kernel, i.e. for a kernel of size 5×5, every image is represented by 25 shifted entries in that matrix. This representation is written to the first dimension and indexed by the kernel position in the second dimension. One can then perform the convolution by matrix vector multiplication with the flattened filterbank. To be concrete, denote this matrix by A, which, for the above case of dimensions, is of shape $(n_b \cdot n_x \cdot n_y, n_c \cdot n_k \cdot n_k)$. Problem (29) can then be equivalently formulated as

$$\hat{K}^{(k+1)} = \underset{\hat{K}}{\operatorname{argmin}} \frac{1}{2}||\hat{K}A^\top - y||^2 + \frac{1}{2\tau}||\hat{K} - \hat{K}^{(k)}||^2, \tag{30}$$

where $\hat{K}A^\top$ denotes the standard matrix-matrix multiplication for a matrix $\hat{K}$ of shape $(n_f, n_c \cdot n_k \cdot n_k)$. Note that in order to solve (30), one has to invert a matrix of the form $I + A^\top A \in \mathbb{R}^{n_c n_k n_k \times n_c n_k n_k}$, which is very small. Hence this proximal subproblem can be solved efficiently.

References

[1] Léon Bottou. Stochastic gradient learning in neural networks. *Proceedings of Neuro-Nımes*, 91(8), 1991.

[2] Miguel Á. Carreira-Perpiñán and Weiran Wang. Distributed optimization of deeply nested systems. In *Proceedings of the 17th International Conference on Artificial Intelligence and Statistics, AISTATS*, 2014.

[3] Pratik Chaudhari, Anna Choromanska, Stefano Soatto, and Yann LeCun. Entropy-SGD: Biasing gradient descent into wide valleys. *arXiv:1611.01838*, 2016.

[4] Yann N. Dauphin, Razvan Pascanu, Caglar Gulcehre, Kyunghyun Cho, Surya Ganguli, and Yoshua Bengio. Identifying and attacking the saddle point problem in high-dimensional non-convex optimization. In *Proceedings of the 27th International Conference on Neural Information Processing Systems, NIPS*, 2014.

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

[6] Sepp Hochreiter, Yoshua Bengio, and Paolo Frasconi. Gradient flow in recurrent nets: the difficulty of learning long-term dependencies. In *Field Guide to Dynamical Recurrent Networks*. 2001.

[7] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *International Conference on Learning Representations, ICLR*, 2015.

[8] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. ImageNet classification with deep convolutional neural networks. In *Proceedings of the 25th International Conference of Neural Information Processing Systems, NIPS*, 2012.

[9] Quoc V Le, Adam Coates, Bobby Prochnow, and Andrew Y Ng. On optimization methods for deep learning. *Proceedings of The 28th International Conference on Machine Learning, ICML*, 2011.

[10] James Martens. Deep learning via Hessian-free optimization. In *Proceedings of the 27th International Conference on Machine Learning, ICML*, 2010.

[11] B. Martinet. Régularisation d'inéquations variationnelles par approximations successives. *Rev. Francaise Inf. Rech. Oper.*, pages 154–159, 1970.

[12] Jean-Jacques Moreau. Proximité et dualité dans un espace hilbertien. *Bulletin de la Société mathématique de France*, 93:273–299, 1965.

[13] Yurii Nesterov. A method of solving a convex programming problem with convergence rate $O(1/k^2)$. *Soviet Mathematics Doklady*, 27(2):372–376, 1983.

[14] J. Nocedal and S. Wright. *Numerical Optimization*. Springer Series in Operations Research and Financial Engineering. 2006.

[15] Herbert Robbins and Sutton Monro. A stochastic approximation method. *The Annals of Mathematical Statistics*, 22(3):400–407, 1951.

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

[17] David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel, and Demis Hassabis. Mastering the game of go with deep neural networks and tree search. *Nature*, 529(7587):484–489, 2016.

[18] Jascha Sohl-Dickstein, Ben Poole, and Surya Ganguli. Fast large-scale optimization by unifying stochastic gradient and quasi-newton methods. In *Proceedings of The 31st International Conference on Machine Learning, ICML*, 2013.

[19] Ilya Sutskever, James Martens, George Dahl, and Geoffrey Hinton. On the importance of initialization and momentum in deep learning. In *Proceedings of the 30th International Conference on International Conference on Machine Learning, ICML*, 2013.

[20] Ilya Sutskever, Oriol Vinyals, and Quoc V Le. Sequence to sequence learning with neural networks. In *Proceedings of the 27th International Conference of Neural Information Processing Systems, NIPS*, 2014.

[21] Gavin Taylor, Ryan Burmeister, Zheng Xu, Bharat Singh, Ankit Patel, and Tom Goldstein. Training neural networks without gradients: A scalable ADMM approach. In *Proceedings of the 33rd International Conference on Machine Learning, ICML*, 2016.

[22] Theano Development Team. Theano: A Python framework for fast computation of mathematical expressions. *arXiv:1605.02688*, 2016.