

---

# meProp: Sparsified Back Propagation for Accelerated Deep Learning with Reduced Overfitting

---

Xu Sun<sup>1,2</sup> Xuancheng Ren<sup>1,2</sup> Shuming Ma<sup>1,2</sup> Houfeng Wang<sup>1,2</sup>

## Abstract

We propose a simple yet effective technique for neural network learning. The forward propagation is computed as usual. In back propagation, only a small subset of the full gradient is computed to update the model parameters. The gradient vectors are sparsified in such a way that only the top- $k$  elements (in terms of magnitude) are kept. As a result, only  $k$  rows or columns (depending on the layout) of the weight matrix are modified, leading to a linear reduction ( $k$  divided by the vector dimension) in the computational cost. Surprisingly, experimental results demonstrate that we can update only 1–4% of the weights at each back propagation pass. This does not result in a larger number of training iterations. More interestingly, the accuracy of the resulting models is actually improved rather than degraded, and a detailed analysis is given.

## 1. Introduction

Neural network learning is typically slow, where back propagation usually dominates the computational cost during the learning process. Back propagation entails a high computational cost because it needs to compute full gradients and update all model parameters in each learning step. It is not uncommon for a neural network to have a massive number of model parameters.

In this study, we propose a *minimal effort* back propagation method, which we call *meProp*, for neural network learning. The idea is that we compute only a very small but critical portion of the gradient information, and update only the corresponding minimal portion of the parameters in each learning step. This leads to sparsified gradients,

---

<sup>1</sup>School of Electronics Engineering and Computer Science, Peking University, China <sup>2</sup>MOE Key Laboratory of Computational Linguistics, Peking University, China. Correspondence to: Xu Sun <xusun@pku.edu.cn>.

such that only highly relevant parameters are updated and other parameters stay untouched. The sparsified back propagation leads to a linear reduction in the computational cost.

To realize our approach, we need to answer two questions. The first question is how to find the highly relevant subset of the parameters from the current sample in stochastic learning. We propose a top- $k$  search method to find the most important parameters. Interestingly, experimental results demonstrate that we can update only 1–4% of the weights at each back propagation pass. This does not result in a larger number of training iterations. The proposed method is general-purpose and it is independent of specific models and specific optimizers (e.g., Adam and AdaGrad).

The second question is whether or not this minimal effort back propagation strategy will hurt the accuracy of the trained models. We show that our strategy does not degrade the accuracy of the trained model, even when a very small portion of the parameters is updated. More interestingly, our experimental results reveal that our strategy actually improves the model accuracy in most cases. Based on our experiments, we find that it is probably because the minimal effort update does not modify weakly relevant parameters in each update, which makes overfitting less likely, similar to the *dropout* effect.

The contributions of this work are as follows:

- We propose a sparsified back propagation technique for neural network learning, in which only a small subset of the full gradient is computed to update the model parameters. Experimental results demonstrate that we can update only 1–4% of the weights at each back propagation pass. This does not result in a larger number of training iterations.
- Surprisingly, our experimental results reveal that the accuracy of the resulting models is actually improved, rather than degraded. We demonstrate this effect by conducting experiments on different deep learning models (LSTM and MLP), various optimization methods (Adam and AdaGrad), and diverse tasks (natural language processing and image recognition).




Figure 1. An illustration of meProp.

## 2. Proposed Method

We propose a simple yet effective technique for neural network learning. The forward propagation is computed as usual. During back propagation, only a small subset of the full gradient is computed to update the model parameters. The gradient vectors are “quantized” so that only the top- $k$  components in terms of magnitude are kept. We first present the proposed method and then describe the implementation details.

### 2.1. meProp

Forward propagation of neural network models, including feedforward neural networks, RNN, LSTM, consists of linear transformations and non-linear transformations. For simplicity, we take a computation unit with one linear transformation and one non-linear transformation as an example:

$$y = Wx \quad (1)$$

$$z = \sigma(y) \quad (2)$$

where  $W \in R^{n \times m}$ ,  $x \in R^m$ ,  $y \in R^n$ ,  $z \in R^n$ ,  $m$  is the dimension of the input vector,  $n$  is the dimension of the output vector, and  $\sigma$  is a non-linear function (e.g., *relu*, *tanh*, and *sigmoid*). During back propagation, we need to compute the gradient of the parameter matrix  $W$  and the input vector  $x$ :

$$\frac{\partial z}{\partial W_{ij}} = \sigma'_i x_j^T \quad (1 \leq i \leq n, 1 \leq j \leq m) \quad (3)$$

$$\frac{\partial z}{\partial x_i} = \sum_j W_{ij}^T \sigma'_j \quad (1 \leq j \leq n, 1 \leq i \leq m) \quad (4)$$

where  $\sigma'_i \in R^n$  means  $\frac{\partial z_i}{\partial y_i}$ . We can see that the computational cost of back propagation is directly proportional to the dimension of output vector  $n$ .

The proposed meProp uses approximate gradients by keeping only top- $k$  elements based on the *magnitude values*. That is, only the top- $k$  elements with the largest absolute values are kept. For example, suppose a vector  $v = \langle 1, 2, 3, -4 \rangle$ , then  $top_2(v) = \langle 0, 0, 3, -4 \rangle$ . We denote the indices of vector  $\sigma'(y)$ 's top- $k$  values as  $\{t_1, t_2, \dots, t_k\}$  ( $1 \leq k \leq n$ ), and the approximate gradient of the parameter matrix  $W$  and input vector  $x$  is:

$$\frac{\partial z}{\partial W_{ij}} \leftarrow \sigma'_i x_j^T \quad \text{if } i \in \{t_1, t_2, \dots, t_k\} \quad \text{else } 0 \quad (5)$$

$$\frac{\partial z}{\partial x_i} \leftarrow \sum_j W_{ij}^T \sigma'_j \quad \text{if } j \in \{t_1, t_2, \dots, t_k\} \quad \text{else } 0 \quad (6)$$

As a result, only  $k$  rows or columns (depending on the layout) of the weight matrix are modified, leading to a linear reduction ( $k$  divided by the vector dimension) in the computational cost.

Figure 1 is an illustration of meProp for a single computation unit of neural models. The original back propagation uses the full gradient of the output vectors to compute the gradient of the parameters. The proposed method selects the top- $k$  values of the gradient of the output vector, and backpropagates the loss through the corresponding subset of the total model parameters.

As for a complete neural network framework with a loss  $L$ , the original back propagation computes the gradient of the parameter matrix  $W$  as:

$$\frac{\partial L}{\partial W} = \frac{\partial L}{\partial y} \cdot \frac{\partial y}{\partial W} \quad (7)$$

while the gradient of the input vector  $x$  is:

$$\frac{\partial L}{\partial x} = \frac{\partial y}{\partial x} \cdot \frac{\partial L}{\partial y} \quad (8)$$

The proposed meProp selects top- $k$  elements of the gradient  $\frac{\partial L}{\partial y}$  to approximate the original gradient, and passes

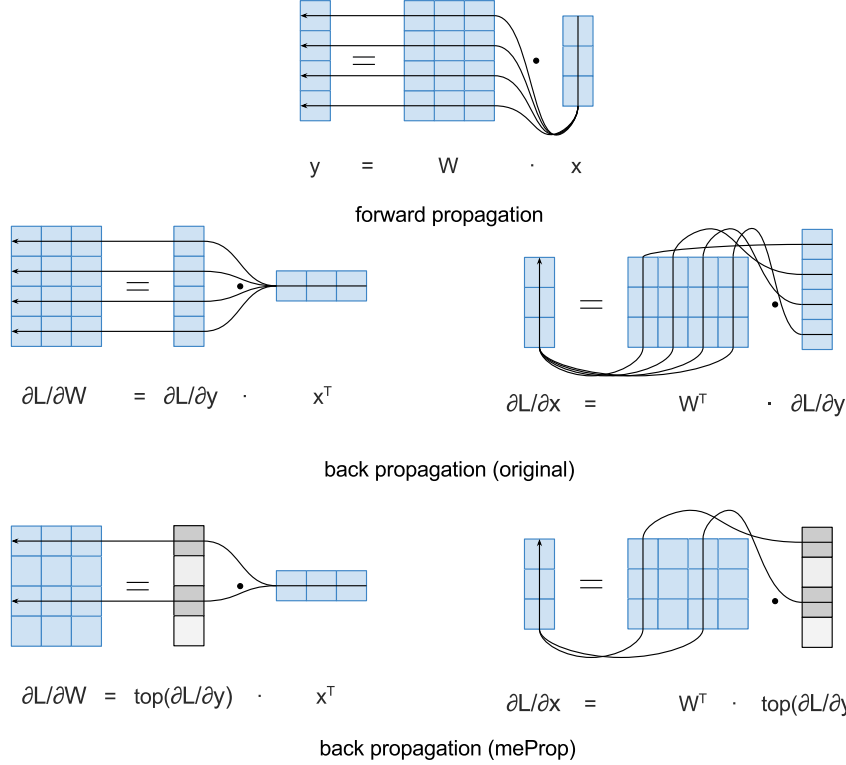


Figure 2. An illustration of the computational flow of meProp.

them through the gradient computation graph according to the chain rule. Hence, the gradient of  $W$  goes to:

$$\frac{\partial L}{\partial W} \leftarrow \text{top}_k\left(\frac{\partial L}{\partial y}\right) \cdot \frac{\partial y}{\partial W} \quad (9)$$

while the gradient of the vector  $x$  is:

$$\frac{\partial L}{\partial x} \leftarrow \frac{\partial y}{\partial x} \cdot \text{top}_k\left(\frac{\partial L}{\partial y}\right) \quad (10)$$

Figure 2 shows an illustration of the computational flow of meProp. The forward propagation is the same as traditional forward propagation, which computes the output vector via a matrix multiplication operation between two input tensors. The original back propagation computes the full gradient for the input vector and the weight matrix. For meProp, back propagation computes an approximate gradient by keeping top- $k$  values of the backward flowed gradient and masking the remaining values to 0.

Figure 3 further shows the computational flow of meProp for the mini-batch case.

## 2.2. Implementation

We have coded two neural network models, including an LSTM model for part-of-speech (POS) tagging, and a feed-forward NN model (MLP) for transition-based dependency

parsing and MNIST image recognition. We use the optimizers with automatically adaptive learning rates, including Adam (Kingma & Ba, 2014) and AdaGrad (Duchi et al., 2011). In our implementation, we make no modification to the optimizers, although there are many zero elements in the gradients.

Most of the experiments on CPU are conducted on the framework coded in C# on our own. This framework builds a dynamic computation graph of the model for each sample, making it suitable for data in variable lengths. A typical training procedure contains four parts: building the computation graph, forward propagation, back propagation, and parameter update. We also have an implementation based on the *PyTorch* framework for GPU based experiments.

### 2.2.1. WHERE TO APPLY MEPROP

The proposed method aims to reduce the complexity of the back propagation by reducing the elements in the computationally intensive operations. In our preliminary observations, matrix-matrix or matrix-vector multiplication consumed more than 90% of the time of back propagation. In our implementation, we apply meProp only to the back propagation from the output of the multiplication to its inputs. For other element-wise operations (e.g., activation functions), the original back propagation procedure is kept, be-

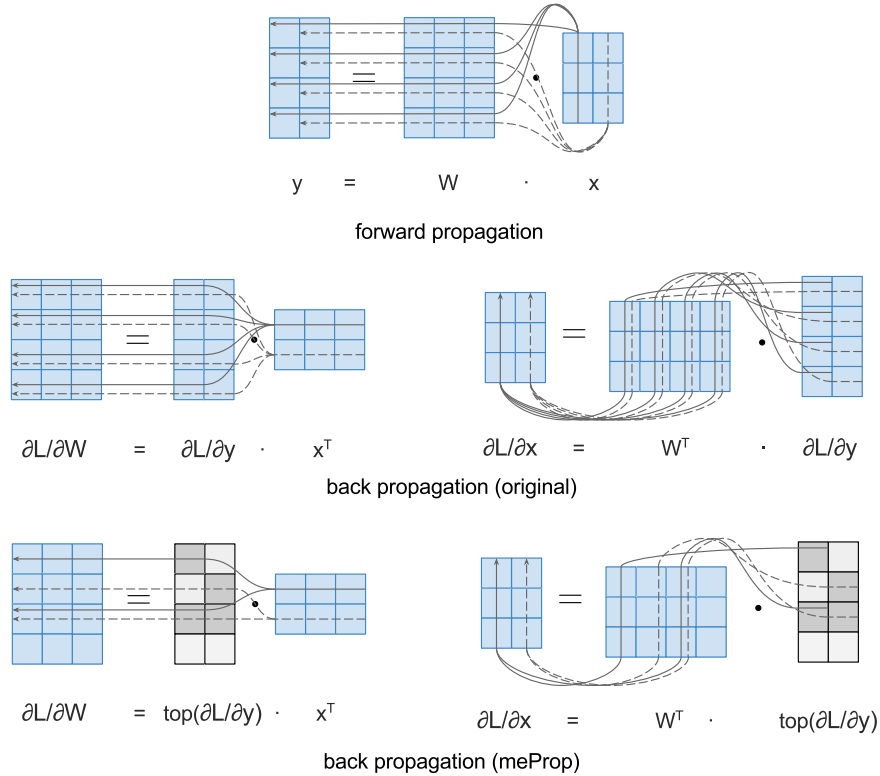


Figure 3. An illustration of the computational flow of meProp on a mini-batch learning setting.

cause those operations are already fast enough compared with matrix-matrix or matrix-vector multiplication operations.

If there are multiple hidden layers, the top- $k$  sparsification needs to be applied to every hidden layer, because the sparsified gradient will again be dense from one layer to another. That is, in meProp the gradients are sparsified with a top- $k$  operation at the output of every hidden layer.

While we apply meProp to all hidden layers using the same  $k$  of top- $k$ , usually the  $k$  for the output layer could be different from the  $k$  for the hidden layers, because the output layer typically has a very different dimension compared with the hidden layers. For example, there are 10 tags in the MNIST task, so the dimension of the output layer is 10, and we use an MLP with the hidden dimension of 500. Thus, the best  $k$  for the output layer could be different from that of the hidden layers.

### 2.2.2. CHOICE OF TOP- $k$ ALGORITHMS

Instead of sorting the entire vector, we use the well-known min-heap based top- $k$  selection method, which is slightly changed to focus on memory reuse. The algorithm has a time complexity of  $O(n \log k)$  and a space complexity of  $O(k)$ .

## 3. Related Work

Riedmiller and Braun (1993) proposed a direct adaptive method for fast learning, which performs a local adaptation of the weight update according to the behavior of the error function. Tollenaere (1990) also proposed an adaptive acceleration strategy for back propagation. Dropout (Srivastava et al., 2014) is proposed to improve training speed and reduce the risk of overfitting. Sparse coding is a class of unsupervised methods for learning sets of over-complete bases to represent data efficiently (Olshausen & Field, 1996). Ranzato et al. (2006) proposed a sparse autoencoder model for learning sparse over-complete features. The proposed method is quite different compared with those prior studies on back propagation, dropout, and sparse coding.

The *sampled-output-loss* methods (Jean et al., 2015) are limited to the softmax layer (output layer) and are only based on random sampling, while our method does not have those limitations. The sparsely-gated mixture-of-experts (Shazeer et al., 2017) only sparsifies the mixture-of-experts gated layer and it is limited to the specific setting of mixture-of-experts, while our method does not have those limitations. There are also prior studies focusing on reducing the communication cost in distributed systems (Seide et al., 2014; Dryden et al., 2016), by quanti-

Table 1. Results based on LSTM/MLP models and AdaGrad/Adam optimizers. Time means averaged time per iteration. Iter means the number of iterations to reach the optimal score on development data. The model of this iteration is then used to obtain the test score.

<b>POS-Tag (AdaGrad)</b>	Iter	Backprop time (s)	Dev Acc (%)	Test Acc (%)
LSTM (h=500)	4	17,534.4	96.89	96.93
meProp (k=5)	4	<b>253.2 (69.2x)</b>	97.18 (+0.29)	<b>97.25 (+0.32)</b>
<b>Parsing (AdaGrad)</b>	Iter	Backprop time (s)	Dev UAS (%)	Test UAS (%)
MLP (h=500)	11	8,899.7	89.07	88.92
meProp (k=20)	8	<b>492.3 (18.1x)</b>	89.17 (+0.10)	<b>88.95 (+0.03)</b>
<b>MNIST (AdaGrad)</b>	Iter	Backprop time (s)	Dev Acc (%)	Test Acc (%)
MLP (h=500)	8	171.0	98.20	97.52
meProp (k=10)	16	<b>4.1 (41.7x)</b>	98.20 (+0.00)	<b>98.00 (+0.48)</b>
<b>POS-Tag (Adam)</b>	Iter	Backprop time (s)	Dev Acc (%)	Test Acc (%)
LSTM (h=500)	2	16,167.3	97.18	97.07
meProp (k=5)	5	<b>247.2 (65.4x)</b>	97.14 (-0.04)	<b>97.12 (+0.05)</b>
<b>Parsing (Adam)</b>	Iter	Backprop time (s)	Dev UAS (%)	Test UAS (%)
MLP (h=500)	14	9,077.7	90.12	89.84
meProp (k=20)	6	<b>488.7 (18.6x)</b>	90.02 (-0.10)	<b>90.01 (+0.17)</b>
<b>MNIST (Adam)</b>	Iter	Backprop time (s)	Dev Acc (%)	Test Acc (%)
MLP (h=500)	17	169.5	98.32	97.82
meProp (k=20)	15	<b>7.9 (21.4x)</b>	98.22 (-0.10)	<b>98.01 (+0.19)</b>

Table 2. Overall forward propagation time vs. overall back propagation time. Time means averaged time per iteration. FP means forward propagation. BP means back propagation. Ov. time means overall training time (FP + BP).

<b>POS-Tag (Adam)</b>	Ov. FP time	Ov. BP time	Ov. time
LSTM (h=500)	7,334s	16,522s	23,856s
meProp (k=5)	7,362s	<b>540s (30.5x)</b>	<b>7,903s</b>
<b>Parsing (Adam)</b>	Ov. FP time	Ov. BP time	Ov. time
MLP (h=500)	3,906s	9,114s	13,020s
meProp (k=20)	4,002s	<b>513s (17.7x)</b>	<b>4,516s</b>
<b>MNIST (Adam)</b>	Ov. FP time	Ov. BP time	Ov. time
MLP (h=500)	69s	171s	240s
meProp (k=20)	68s	<b>9s (18.4x)</b>	<b>77s</b>

zing each value of the gradient from 32-bit float to only 1-bit. Those settings are also different from ours.

## 4. Experiments

To demonstrate that the proposed method is general-purpose, we perform experiments on different models (LSTM/MLP), various training methods (Adam/AdaGrad), and diverse tasks.

**Part-of-Speech Tagging (POS-Tag):** We use the standard benchmark dataset in prior work (Collins, 2002), which is derived from the Penn Treebank corpus, and use sections 0-18 of the Wall Street Journal (WSJ) for training (38,219 examples), and sections 22-24 for testing (5,462 examples). The evaluation metric is per-word accuracy. A popular model for this task is the LSTM model (Hochreiter & Schmidhuber, 1997),<sup>1</sup> which is used as our baseline.

<sup>1</sup>In this work, we use the bi-directional LSTM (Bi-LSTM) as the implementation of LSTM.

**Transition-based Dependency Parsing (Parsing):** Following prior work, we use English Penn TreeBank (PTB) (Marcus et al., 1993) for evaluation. We follow the standard split of the corpus and use sections 2-21 as the training set (39,832 sentences, 1,900,056 transition examples),<sup>2</sup> section 22 as the development set (1,700 sentences, 80,234 transition examples) and section 23 as the final test set (2,416 sentences, 113,368 transition examples). The evaluation metric is *unlabeled attachment score* (UAS). We implement a parser using MLP following Chen and Manning (2014), which is used as our baseline.

**MNIST Image Recognition (MNIST):** We use the MNIST handwritten digit dataset (LeCun et al., 1998) for evaluation. MNIST consists of 60,000 28×28 pixel training images and additional 10,000 test examples. Each image contains a single numerical digit (0-9). We select the first 5,000 images of the training images as the development set and the rest as the training set. The evaluation metric is per-image accuracy. We use the MLP model as the baseline.

### 4.1. Experimental Settings

We set the dimension of the hidden layers to 500 for all the tasks. For POS-Tag, the input dimension is 1 (word) × 50 (dim per word) + 7 (features) × 20 (dim per feature) = 190, and the output dimension is 45. For Parsing, the input dimension is 48 (features) × 50 (dim per feature) = 2400, and the output dimension is 25. For MNIST, the input dimension is 28 (pixels per row) × 28 (pixels per column) × 1 (dim per pixel) = 784, and the output dimension is 10. As discussed in Section 2, the optimal  $k$  of top- $k$  for the

<sup>2</sup>A transition example consists of a parsing context and its optimal transition action.




Figure 4. Accuracy vs. meProp’s backprop ratio (left). Results of top- $k$  meProp vs. random meProp (middle). Results of top- $k$  meProp vs. baseline with the hidden dimension  $h$  (right).

output layer could be different from the hidden layers, because their dimensions could be very different. For Parsing and MNIST, we find using the same  $k$  for the output and the hidden layers works well, and we simply do so. For another task, POS-Tag, we find the the output layer should use a different  $k$  from the hidden layers. For simplicity, we do not apply meProp to the output layer for POS-Tag, because in this task we find the computational cost of the output layer is almost negligible compared with other layers.

The hyper-parameters are tuned based on the development data. For the Adam optimization method, we find the default hyper-parameters work well on development sets, which are as follows: the learning rate  $\alpha = 0.001$ , and  $\beta_1 = 0.9, \beta_2 = 0.999, \epsilon = 1 \times 10^{-8}$ . For the AdaGrad learner, the learning rate is set to  $\alpha = 0.01, 0.01, 0.1$  for POS-Tag, Parsing, and MNIST, respectively, and  $\epsilon = 1 \times 10^{-6}$ . The experiments on CPU are conducted on a computer with the *INTEL(R) Xeon(R) 3.0GHz CPU*. The experiments on GPU are conducted on *NVIDIA GeForce GTX 1080*.

## 4.2. Experimental Results

In this experiment, the LSTM is based on one hidden layer and the MLP is based on two hidden layers (experiments on more hidden layers will be presented later). We conduct experiments on different optimization methods, including AdaGrad and Adam. Since meProp is applied to the linear transformations (which entail the major computational cost), we report the linear transformation related backprop time as *Backprop Time*. It does not include non-linear activations, which usually have only less than 2% computational cost. The total time of back propagation, including non-linear activations, is reported as *Overall Backprop Time*. Based on the development set and prior work, we set the mini-batch size to 1 (sentence), 10,000 (transition examples), and 10 (images) for POS-Tag, Parsing, and MNIST, respectively. Using 10,000 transition examples for Parsing follows Chen and Manning (2014).

Table 1 shows the results based on different models and

different optimization methods. In the table, *meProp* means applying meProp to the corresponding baseline model,  $h = 500$  means that the hidden layer dimension is 500, and  $k = 20$  means that meProp uses top-20 elements (among 500 in total) for back propagation. Note that, for fair comparisons, all experiments are first conducted on the development data and the test data is not observable. Then, the optimal number of iterations is decided based on the optimal score on development data, and the model of this iteration is used upon the test data to obtain the test scores.

As we can see, applying meProp can substantially speed up the back propagation. It provides a linear reduction in the computational cost. Surprisingly, results demonstrate that we can update only 1–4% of the weights at each back propagation pass. This does not result in a larger number of training iterations. More surprisingly, the accuracy of the resulting models is actually improved rather than decreased. The main reason could be that the minimal effort update does not modify weakly relevant parameters, which makes overfitting less likely, similar to the *dropout* effect.

Table 2 shows the overall forward propagation time, the overall back propagation time, and the training time by summing up forward and backward propagation time. As we can see, back propagation has the major computational cost in training LSTM/MLP.

The results are consistent among AdaGrad and Adam. The results demonstrate that meProp is independent of specific optimization methods. For simplicity, in the following experiments the optimizer is based on Adam.

## 4.3. Varying Backprop Ratio

In Figure 4 (left), we vary the  $k$  of top- $k$  meProp to compare the test accuracy on different ratios of meProp backprop. For example, when  $k=5$ , it means that the backprop ratio is  $5/500=1\%$ . The optimizer is Adam. As we can see, meProp achieves consistently better accuracy than the baseline. The best test accuracy of meProp, 98.15% (+0.33),

Table 3. Results based on the same  $k$  and  $h$ .

<b>POS-Tag (Adam)</b>	Iter	Test Acc (%)
LSTM (h=5)	7	96.40
meProp (k=5)	5	<b>97.12 (+0.72)</b>
<b>Parsing (Adam)</b>	Iter	Test UAS (%)
MLP (h=20)	18	88.37
meProp (k=20)	6	<b>90.01 (+1.64)</b>
<b>MNIST (Adam)</b>	Iter	Test Acc (%)
MLP (h=20)	15	95.77
meProp (k=20)	17	<b>98.01 (+2.24)</b>

 Table 4. Adding the *dropout* technique.

<b>POS-Tag (Adam)</b>	Dropout	Test Acc (%)
LSTM (h=500)	0.5	97.20
meProp (k=20)	0.5	<b>97.31 (+0.11)</b>
<b>Parsing (Adam)</b>	Dropout	Test UAS (%)
MLP (h=500)	0.5	91.53
meProp (k=40)	0.5	<b>91.99 (+0.46)</b>
<b>MNIST (Adam)</b>	Dropout	Test Acc (%)
MLP (h=500)	0.2	98.09
meProp (k=25)	0.2	<b>98.32 (+0.23)</b>

is actually better than the one reported in Table 1.

#### 4.4. Top-k vs. Random

It will be interesting to check the role of top- $k$  elements. Figure 4 (middle) shows the results of top- $k$  meProp vs. random meProp. The *random meProp* means that random elements (instead of top- $k$  ones) are selected for back propagation. As we can see, the top- $k$  version works better than the random version. It suggests that top- $k$  elements contain the most important information of the gradients.

#### 4.5. Varying Hidden Dimension

We still have a question: does the top- $k$  meProp work well simply because the original model does not require that big dimension of the hidden layers? For example, the meProp (topk=5) works simply because the LSTM works well with the hidden dimension of 5, and there is no need to use the hidden dimension of 500. To examine this, we perform experiments on using the same hidden dimension as  $k$ , and the results are shown in Table 3. As we can see, however, the results of the small hidden dimensions are much worse than those of meProp.

In addition, Figure 4 (right) shows more detailed curves by varying the value of  $k$ . In the figure, different  $k$  gives different backprop ratio for meProp and different hidden dimension ratio for LSTM/MLP. As we can see, the answer to that question is negative: meProp does not rely on redundant hidden layer elements.

Table 5. Varying the number of hidden layers on the MNIST task. The optimizer is Adam. Layers: the number of hidden layers.

Layers	Method	Test Acc (%)
2	MLP (h=500)	98.10
	meProp (k=25)	<b>98.20 (+0.10)</b>
3	MLP (h=500)	98.21
	meProp (k=25)	<b>98.37 (+0.16)</b>
4	MLP (h=500)	98.10
	meProp (k=25)	<b>98.15 (+0.05)</b>
5	MLP (h=500)	98.05
	meProp (k=25)	<b>98.21 (+0.16)</b>

 Table 6. Results of *simple unified top-k* meProp based on a whole mini-batch (i.e., unified sparse patterns). The optimizer is Adam. Mini-batch Size is 50.

Layers	Method	Test Acc (%)
2	MLP (h=500)	97.97
	meProp (k=30)	<b>98.08 (+0.11)</b>
5	MLP (h=500)	98.09
	meProp (k=50)	<b>98.36 (+0.27)</b>

#### 4.6. Adding Dropout

Since we have observed that meProp can reduce overfitting of deep learning, a natural question is that if meProp is reducing the same type of overfitting risk as dropout. Thus, we use development data to find a proper value of the dropout rate on those tasks, and then further add meProp to check if further improvement is possible.

Table 4 shows the results. As we can see, meProp can achieve further improvement over dropout. In particular, meProp has an improvement of 0.46 UAS on Parsing. The results suggest that the type of overfitting that meProp reduces is probably different from that of dropout. Thus, a model should be able to take advantage of both meProp and dropout to reduce overfitting.

#### 4.7. Adding More Hidden Layers

Another question is whether or not meProp relies on shallow models with only a few hidden layers. To answer this question, we also perform experiments on more hidden layers, from 2 hidden layers to 5 hidden layers. We find setting the dropout rate to 0.1 works well for most cases of different numbers of layers. For simplicity of comparison, we set the same dropout rate to 0.1 in this experiment. Table 5 shows that adding the number of hidden layers does not hurt the performance of meProp.

#### 4.8. Speedup on GPU

For implementing meProp on GPU, the simplest solution is to treat the entire mini-batch as a “big training example”, where the top- $k$  operation is based on the averaged values of all examples in the mini-batch. In this way, the big sparse matrix of the mini-batch will have consistent sparse

Table 7. Acceleration results on the matrix multiplication synthetic data using GPU. The batch size is 1024.

Method	Backprop time (ms)
Baseline (h=8192)	308.00
meProp (k=8)	8.37 (36.8x)
meProp (k=16)	9.16 (33.6x)
meProp (k=32)	11.20 (27.5x)
meProp (k=64)	14.38 (21.4x)
meProp (k=128)	21.28 (14.5x)
meProp (k=256)	38.57 (8.0x)
meProp (k=512)	69.95 (4.4x)

Table 8. Acceleration results on MNIST using GPU.

Method	Overall backprop time (ms)
MLP (h=8192)	17,696.2
meProp (k=8)	1,501.5 (11.8x)
meProp (k=16)	1,542.8 (11.5x)
meProp (k=32)	1,656.9 (10.7x)
meProp (k=64)	1,828.3 (9.7x)
meProp (k=128)	2,200.0 (8.0x)
meProp (k=256)	3,149.6 (5.6x)
meProp (k=512)	4,874.1 (3.6x)

patterns among examples, and this consistent sparse matrix can be transformed into a small dense matrix by removing the zero values. We call this implementation as *simple uni-fied top-k*. This experiment is based on *PyTorch*.

Despite its simplicity, Table 6 shows the good performance of this implementation, which is based on the mini-batch size of 50. We also find the speedup on GPU is less significant when the hidden dimension is low. The reason is that our GPU’s computational power is not fully consumed by the baseline (with small hidden layers), so that the normal back propagation is already fast enough, making it hard for meProp to achieve substantial speedup. For example, supposing a GPU can finish 1000 operations in one cycle, there could be no speed difference between a method with 100 and a method with 10 operations. Indeed, we find MLP (h=64) and MLP (h=512) have almost the same GPU speed even on *forward* propagation (i.e., without meProp), while theoretically there should be an 8x difference. With GPU, the *forward* propagation time of MLP (h=64) and MLP (h=512) is 572ms and 644ms, respectively. This provides evidence for our hypothesis that our GPU is not fully consumed with the small hidden dimensions.

Thus, the speedup test on GPU is more meaningful for the heavy models, such that the baseline can at least fully consume the GPU’s computational power. To check this, we test the GPU speedup on synthetic data of matrix multiplication with a larger hidden dimension. Indeed, Table 7 shows that meProp achieves much higher speed than the traditional backprop with the large hidden dimension. Furthermore, we test the GPU speedup on MLP with the large hidden dimension (Dryden et al., 2016). Table 8 shows that

meProp also has substantial GPU speedup on MNIST with the large hidden dimension. In this experiment, the speedup is based on *Overall Backprop Time* (see the prior definition). Those results demonstrate that meProp can achieve good speedup on GPU when it is applied to heavy models.

Finally, there are potentially other implementation choices of meProp on GPU. For example, another natural solution is to use a big sparse matrix to represent the sparsified gradient of the output of a mini-batch. Then, the sparse matrix multiplication library can be used to accelerate the computation. This could be an interesting direction of future work.

#### 4.9. Related Systems on the Tasks

The POS tagging task is a well-known benchmark task, with the accuracy reports from 97.2% to 97.4% (Toutanova et al., 2003; Sun, 2014; Shen et al., 2007; Tsuruoka et al., 2011; Collobert et al., 2011; Huang et al., 2015). Our method achieves 97.31% (Table 4).

For the transition-based dependency parsing task, existing approaches typically can achieve the UAS score from 91.4 to 91.5 (Zhang & Clark, 2008; Nivre et al., 2007; Huang & Sagae, 2010). As one of the most popular transition-based parsers, MaltParser (Nivre et al., 2007) has 91.5 UAS. Chen and Manning (2014) achieves 92.0 UAS using neural networks. Our method achieves 91.99 UAS (Table 4).

For MNIST, the MLP based approaches can achieve 98–99% accuracy, often around 98.3% (LeCun et al., 1998; Simard et al., 2003; Ciresan et al., 2010). Our method achieves 98.37% (Table 5). With the help from convolutional layers and other techniques, the accuracy can be improved to over 99% (Jarrett et al., 2009; Ciresan et al., 2012). Our method can also be improved with those additional techniques, which, however, are not the focus of this paper.

## 5. Conclusions

The back propagation in deep learning tries to modify all parameters in each stochastic update, which is inefficient and may even lead to overfitting due to unnecessary modification of many weakly relevant parameters. We propose a *minimal effort* back propagation method (*meProp*), in which we compute only a very small but critical portion of the gradient, and modify only the corresponding small portion of the parameters in each update. This leads to very sparsified gradients to modify only highly relevant parameters for the given training sample. The proposed meProp is independent of the optimization method. Experiments show that meProp can reduce the computational cost of back propagation by one to two orders of magnitude via updating only 1–4% parameters, and yet improve the model accuracy in most cases.



## Acknowledgements

The authors would like to thank the anonymous reviewers for insightful comments and suggestions on this paper. This work was supported in part by National Natural Science Foundation of China (No. 61673028), National High Technology Research and Development Program of China (863 Program, No. 2015AA015404), and an Okawa Research Grant (2016).

## References

- Chen, Danqi and Manning, Christopher D. A fast and accurate dependency parser using neural networks. In *Proceedings of EMNLP'14*, pp. 740–750, 2014.
- Ciresan, Dan C., Meier, Ueli, Gambardella, Luca Maria, and Schmidhuber, Jürgen. Deep, big, simple neural nets for handwritten digit recognition. *Neural Computation*, 22(12):3207–3220, 2010.
- Ciresan, Dan C., Meier, Ueli, and Schmidhuber, Jürgen. Multi-column deep neural networks for image classification. In *Proceedings of CVPR'12*, pp. 3642–3649, 2012.
- Collins, Michael. Discriminative training methods for hidden markov models: Theory and experiments with perceptron algorithms. In *Proceedings of EMNLP'02*, pp. 1–8, 2002.
- Collobert, Ronan, Weston, Jason, Bottou, Léon, Karlen, Michael, Kavukcuoglu, Koray, and Kuksa, Pavel P. Natural language processing (almost) from scratch. *Journal of Machine Learning Research*, 12:2493–2537, 2011.
- Dryden, Nikoli, Moon, Tim, Jacobs, Sam Ade, and Esen, Brian Van. Communication quantization for data-parallel training of deep neural networks. In *Proceedings of the 2nd Workshop on Machine Learning in HPC Environments*, pp. 1–8, 2016.
- Duchi, John C., Hazan, Elad, and Singer, Yoram. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12:2121–2159, 2011.
- Hochreiter, Sepp and Schmidhuber, Jürgen. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- Huang, Liang and Sagae, Kenji. Dynamic programming for linear-time incremental parsing. In *Proceedings of ACL'10*, pp. 1077–1086, 2010.
- Huang, Zhiheng, Xu, Wei, and Yu, Kai. Bidirectional lstm-crf models for sequence tagging. *CoRR*, abs/1508.01991, 2015.
- Jarrett, Kevin, Kavukcuoglu, Koray, Ranzato, Marc'Aurelio, and LeCun, Yann. What is the best multi-stage architecture for object recognition? In *Proceeding of ICCV'09*, pp. 2146–2153, 2009.
- Jean, Sébastien, Cho, KyungHyun, Memisevic, Roland, and Bengio, Yoshua. On using very large target vocabulary for neural machine translation. In *Proceedings of ACL/IJCNLP'15*, pp. 1–10, 2015.
- Kingma, Diederik P. and Ba, Jimmy. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2014.
- LeCun, Yann, Bottou, Léon, Bengio, Yoshua, and Haffner, Patrick. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- Marcus, Mitchell P, Marcinkiewicz, Mary Ann, and Santorini, Beatrice. Building a large annotated corpus of english: The penn treebank. *Computational linguistics*, 19(2):313–330, 1993.
- Nivre, Joakim, Hall, Johan, Nilsson, Jens, Chanev, Atanas, Eryigit, Gülsen, Kübler, Sandra, Marinov, Svetoslav, and Marsi, Erwin. Maltparser: A language-independent system for data-driven dependency parsing. *Natural Language Engineering*, 13(2):95–135, 2007.
- Olshausen, Bruno A and Field, David J. Natural image statistics and efficient coding. *Network: Computation in Neural Systems*, 7(2):333–339, 1996.
- Ranzato, Marc'Aurelio, Poultney, Christopher S., Chopra, Sumit, and LeCun, Yann. Efficient learning of sparse representations with an energy-based model. In *NIPS'06*, pp. 1137–1144, 2006.
- Riedmiller, Martin and Braun, Heinrich. A direct adaptive method for faster backpropagation learning: The rprop algorithm. In *Proceedings of IEEE International Conference on Neural Networks 1993*, pp. 586–591, 1993.
- Seide, Frank, Fu, Hao, Droppo, Jasha, Li, Gang, and Yu, Dong. 1-bit stochastic gradient descent and its application to data-parallel distributed training of speech dnns. In *Proceedings of INTERSPEECH'14*, pp. 1058–1062, 2014.
- Shazeer, Noam, Mirhoseini, Azalia, Maziarz, Krzysztof, Davis, Andy, Le, Quoc V., Hinton, Geoffrey E., and Dean, Jeff. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. *CoRR*, abs/1701.06538, 2017.
- Shen, Libin, Satta, Giorgio, and Joshi, Aravind K. Guided learning for bidirectional sequence classification. In *Proceedings of ACL'07*, pp. 760–767, 2007.

- Simard, Patrice Y., Steinkraus, Dave, and Platt, John C. Best practices for convolutional neural networks applied to visual document analysis. In *Proceedings of ICDR'03*, pp. 958–962, 2003.
- Srivastava, Nitish, Hinton, Geoffrey E., Krizhevsky, Alex, Sutskever, Ilya, and Salakhutdinov, Ruslan. Dropout: a simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(1): 1929–1958, 2014.
- Sun, Xu. Structure regularization for structured prediction. In *NIPS'14*, pp. 2402–2410. 2014.
- Tollenaere, Tom. Supersab: fast adaptive back propagation with good scaling properties. *Neural networks*, 3 (5):561–573, 1990.
- Toutanova, Kristina, Klein, Dan, Manning, Christopher D., and Singer, Yoram. Feature-rich part-of-speech tagging with a cyclic dependency network. In *Proceedings of HLT-NAACL'03*, pp. 173–180, 2003.
- Tsuruoka, Yoshimasa, Miyao, Yusuke, and Kazama, Jun'ichi. Learning with lookahead: Can history-based models rival globally optimized models? In *Proceedings of CoNLL'11*, pp. 238–246, 2011.
- Zhang, Yue and Clark, Stephen. A tale of two parsers: Investigating and combining graph-based and transition-based dependency parsing. In *Proceedings of EMNLP'08*, pp. 562–571, 2008.