# TensorNet: putting neural networks on a Tensor Train 

A. Novikov<br>D. Podoprikhin<br>A. Osokin<br>D. Vetrov

March 6, 2015

## Outline

## (1) Neural networks

## (2) Tensor train

(4) Experiments

## Neural networks



First layer
Second layer

## Motivation

Why do we care about memory?

- State-of-the-art deep networks doesn't fit to mobile devices;
- Up to $95 \%$ percent of parameters are in the fully connected layers;
- Shallow networks with huge fully connected layers can achieve almost the same accuracy, as ensemble of deep CNNs (Ba and Caruana 2014).


## Matrix rank decomposition

Lets consider an $M \times N$ matrix $\boldsymbol{W}$ with the rank equals $r$. We can use $(M+N) r$ memory instead of $M N$ :

$$
\underbrace{\boldsymbol{W}}_{M \times N}=\underbrace{\boldsymbol{A}}_{M \times r r \times N} \underbrace{\boldsymbol{B}}
$$

## Drawbacks of the rank decomposition

## Problems:

(1) The low rank compression rate is limited (we want more);
(2) There is no practical way to train low rank shallow networks.

## (1) Neural networks

(2) Tensor train

## (3) TensorNet

(4) Experiments

## Tensor Train summary

Tensor Train (TT) decomposition:

- Compact representation for vectors, matrices and tensors;
- Allows for efficient application of linear algebra operations.


## Mapping example: vector

Build a mapping from the vector $\boldsymbol{b}$ indices to tensor's elements: $x \leftrightarrow \boldsymbol{i}=\left(i_{1}, \ldots, i_{d}\right)$

Example (Matlab reshape):

$$
\begin{aligned}
\boldsymbol{B}(1,1,1)=\boldsymbol{b}(x(1,1,1)) & =\boldsymbol{b}(1) \\
\boldsymbol{B}(2,1,1)=\boldsymbol{b}(x(2,1,1)) & =\boldsymbol{b}(2) \\
\ldots & \\
\boldsymbol{B}(2,3,3)=\boldsymbol{b}(\times(2,3,3)) & =\boldsymbol{b}(18) .
\end{aligned}
$$

## Matrices in the TT-format

Build a mapping from row / column indices of matrix $\boldsymbol{W}=[W(x, y)]$ to vectors $\boldsymbol{i}$ and $\boldsymbol{j}: x \leftrightarrow \boldsymbol{i}=\left(i_{1}, \ldots, i_{d}\right)$ and $y \leftrightarrow \boldsymbol{j}=\left(j_{1}, \ldots, j_{d}\right)$.

TT-format for matrix $\boldsymbol{W}$ :
$\boldsymbol{W}\left(i_{1}, \ldots, i_{d} ; j_{1}, \ldots, j_{d}\right)=\boldsymbol{W}(x(\boldsymbol{i}), y(\boldsymbol{j}))=\underbrace{\boldsymbol{G}_{1}\left[i_{1}, j_{1}\right]}_{1 \times r} \underbrace{\boldsymbol{G}_{2}\left[i_{2}, j_{2}\right]}_{r \times r} \ldots \underbrace{\boldsymbol{G}_{d}\left[i_{d}, j_{d}\right]}_{r \times 1}$
Notation \& terminology:

- $\boldsymbol{W} \in \mathbb{R}^{M \times N}, M=m^{d}, N=n^{d}$;
- $i_{k} \in\{1, \ldots, m\}, \quad j_{k} \in\{1, \ldots, n\}$;
- $\boldsymbol{G}_{k}$ - TT-cores;
- r - TT-rank;

TT-format exists for any matrix $\boldsymbol{W}$ and uses $O\left(d m n r^{2}\right)$ memory to store $O\left(m^{d} n^{d}\right)$ elements. Efficient only if TT-rank is small.

## (1) Neural networks

(2) Tensor train
(3) TensorNet

## Tensor Train layer: feedforward

Input is a $N \times 1$ vector $\boldsymbol{x}$, output is a $M \times 1$ vector $\boldsymbol{y}$ :

$$
\boldsymbol{y}=\boldsymbol{W} \boldsymbol{x}+\boldsymbol{b}
$$

$\boldsymbol{W}$ is represented in the TT-format:

$$
\boldsymbol{y}\left(i_{1}, \ldots, i_{d}\right)=\sum_{j_{1}, \ldots, j_{d}} \boldsymbol{G}_{1}\left[i_{1}, j_{1}\right] \ldots \boldsymbol{G}_{d}\left[i_{d}, j_{d}\right] \boldsymbol{x}\left(j_{1}, \ldots, j_{d}\right)+\boldsymbol{b}(i)
$$

The parameters are the vector $\boldsymbol{b}$ and the TT-cores $\left\{\boldsymbol{G}_{k}\right\}_{k=1}^{d}$

## Backpropagation

$$
\begin{aligned}
L & =\frac{1}{2} \sum_{s=1}^{S}\left\|\boldsymbol{y}_{2}^{s}-\boldsymbol{y}^{s}\right\|_{2}^{2} \\
\frac{\partial L}{\partial \boldsymbol{y}_{2}} & =\sum_{s=1}^{S}\left(\boldsymbol{y}_{2}^{s}-\boldsymbol{y}^{s}\right) \\
\frac{\partial L}{\partial \boldsymbol{x}_{2}} & =\sum_{i} \frac{\partial L}{\partial \boldsymbol{y}_{2}(i)} \frac{\partial \boldsymbol{y}_{2}(i)}{\partial \boldsymbol{x}_{2}} \\
& =\frac{\partial f_{2}}{\partial \boldsymbol{x}_{2}} \frac{\partial \boldsymbol{L}}{\partial \boldsymbol{y}_{2}} \\
\frac{\partial L}{\partial \boldsymbol{\theta}_{2}} & =\frac{\partial f_{2}}{\partial \boldsymbol{\theta}_{2}} \frac{\partial L}{\partial \boldsymbol{y}_{2}}
\end{aligned}
$$

## Backpropagation cont'd

From each layer we need only this:

$$
y_{1}=x_{2}
$$



$$
\begin{aligned}
& \frac{\partial L}{\partial \boldsymbol{x}_{k}}=g_{k}^{x}\left(\boldsymbol{x}_{k}, \frac{\partial L}{\partial \boldsymbol{y}_{k}}\right) \\
& \frac{\partial L}{\partial \boldsymbol{\theta}_{k}}=g_{k}^{\theta}\left(\boldsymbol{x}_{k}, \frac{\partial L}{\partial \boldsymbol{y}_{k}}\right)
\end{aligned}
$$

$$
f_{1}\left(x_{1}, \boldsymbol{\theta}_{1}\right) \quad f_{2}\left(x_{2}, \boldsymbol{\theta}_{2}\right)
$$

## Tensor Train layer: backpropagation

Input: vectors $\frac{\partial L}{\partial \boldsymbol{y}} \in \mathbb{R}^{M}$ and $\boldsymbol{x} \in \mathbb{R}^{N}$.
Output: $\frac{\partial L}{\partial \boldsymbol{x}}, \frac{\partial L}{\partial b}$ and $\frac{\partial L}{\partial \boldsymbol{G}_{k}\left[i_{k}, j_{k}\right]}$.

## Tensor Train layer: backpropagation

Input: vectors $\frac{\partial L}{\partial \boldsymbol{y}} \in \mathbb{R}^{M}$ and $\boldsymbol{x} \in \mathbb{R}^{N}$.
Output: $\frac{\partial L}{\partial \boldsymbol{x}}, \frac{\partial L}{\partial b}$ and $\frac{\partial L}{\partial \boldsymbol{G}_{k}\left[\dot{j}_{k}, j_{k}\right]}$.

$$
\begin{gathered}
\frac{\partial L}{\partial \boldsymbol{x}}=\boldsymbol{W}^{\top} \frac{\partial L}{\partial \boldsymbol{y}} \\
\frac{\partial L}{\partial \boldsymbol{b}}=\frac{\partial L}{\partial \boldsymbol{y}} . \\
\underbrace{\frac{\partial L}{\partial \boldsymbol{G}_{k}\left[i_{k}, j_{k}\right]}}_{r \times r}=\sum_{\boldsymbol{i}{ }^{\backslash k}} \frac{\partial L}{\partial \boldsymbol{y}(\boldsymbol{i})} \frac{\partial \boldsymbol{y}(\boldsymbol{i})}{\partial \boldsymbol{G}_{k}\left[i_{k}, j_{k}\right]}
\end{gathered}
$$

## Tensor Train layer: Jacobian

We want to differentiate the following expression:

$$
\boldsymbol{y}(i)=\sum_{j} \boldsymbol{G}_{1}\left[i_{1}, j_{1}\right] \ldots \boldsymbol{G}_{k}\left[i_{k}, j_{k}\right] \ldots \boldsymbol{G}_{d}\left[i_{d}, j_{d}\right] \boldsymbol{x}(j)+\boldsymbol{b}(\boldsymbol{i})
$$

## Tensor Train layer: Jacobian

We want to differentiate the following expression:

$$
\begin{aligned}
& \boldsymbol{y}(\boldsymbol{i})=\sum_{j} \boldsymbol{G}_{1}\left[i_{1}, j_{1}\right] \ldots \boldsymbol{G}_{k}\left[i_{k}, j_{k}\right] \ldots \boldsymbol{G}_{d}\left[i_{d}, j_{d}\right] \boldsymbol{x}(\boldsymbol{j})+\boldsymbol{b}(\boldsymbol{i}) .
\end{aligned}
$$

$$
\begin{aligned}
& \left.\sum_{j^{k, d}} \boldsymbol{G}_{1}\left[i_{1}, j_{1}\right] \ldots \boldsymbol{G}_{k} \mid i_{k, 1 k}\right] \ldots \boldsymbol{G}_{d-1}\left[i_{d-1}, j_{d-1}\right] \\
& \underbrace{\sum_{j_{d}} G_{d}\left[i_{d}, j_{d}\right] x(j)}_{r \times m n^{d-1}}
\end{aligned}
$$

## Intuition

$$
\boldsymbol{W}\left(i_{1}, i_{2}, i_{3} ; j_{1}, j_{2}, j_{3}\right)=\underbrace{G_{1}\left[i_{1}, j_{1}\right]}_{\in \mathbb{R}} \underbrace{G_{2}\left[i_{2}, j_{2}\right]}_{\in \mathbb{R}} \underbrace{G_{3}\left[i_{3}, j_{3}\right]}_{\in \mathbb{R}}
$$

## $\boldsymbol{W} \in \mathbb{R}^{64 \times 64}$

Input $\boldsymbol{x}$ and output $\boldsymbol{y}$ are reshaped to $4 \times 4 \times 4$ tensor.

To vanish all dashed line weights, set $G_{3}\left[i_{3}=2, j_{3}=4\right]=0$.


Hidden units

Input image

## (1) Neural networks

(2) Tensor train
(3) TensorNet
(4) Experiments

## Mnist

Mnist dataset, two layered neural network. The input $28 \times 28$ image is reshaped to $2 \times 2 \times 7 \times 2 \times 2 \times 7$ tensor.

| Network | Error | Neurons | I layer params | II layer |
| :--- | :---: | :---: | :---: | :---: |
| Baseline | $2.34 \%$ | 500 | 400000 | 5000 |
| TensorNet (one TT-layer) | $2.26 \%$ | 46656 | 420 | 466560 |
| TensorNet (two TT-layer) | $2.07 \%$ | 15625 | 3360 | 1380 |
| TensorNet | $2.6 \%$ | 15625 | 350 | 156250 |
| TensorNet (random order) | $3.5 \%$ | 15625 | 350 | 156250 |

## Mnist cont'd

Mnist dataset, two layered neural network. The input $28 \times 28$ image is reshaped to $4 \times 7 \times 4 \times 7$ tensor.

| Network | Error | Neurons | I layer params | II layer |
| :--- | :---: | :---: | :---: | :---: |
| Baseline | $2.34 \%$ | 500 | 400000 | 5000 |
| TensorNet (one TT-layer) | $1.68 \%$ | 4096 | 1760 | 40960 |

## Cifar

Deep convolutional neural network for CIFAR-10 (image classification). We compressed the last two fully connected layers $\times 11$ (the error increased from $23.25 \%$ to $23.74 \%$ ).

## References I

Ba, Jimmy and Rich Caruana (2014). "Do Deep Nets Really Need to be Deep?" In: Advances in Neural Information Processing Systems 27. Ed. by Z. Ghahramani et al. Curran Associates, Inc., pp. 2654-2662.

