Feedback alignment for training neural networks

Backpropagation

 $\mathbf{h}^{(0)} \in \mathbb{R}^{N}$ is the input to the network.

Forward computation

For each layer i:

$$\mathbf{a}^{(i)} = \mathbf{W}^{(i)}\mathbf{h}^{(i-1)} + \mathbf{b}^{(i)},$$
$$\mathbf{h}^{(i)} = f(\mathbf{a}^{(i)})$$



The weight matrices are used in the Computation of the gradients.

$$\boldsymbol{\delta}^{(k-1)} = \frac{\partial \mathcal{L}(\mathbf{y}, \hat{\mathbf{y}})}{\partial \mathbf{a}^{(k)}} = \mathbf{W}^{(k)^{T}} \boldsymbol{\delta}^{(k)} f'(\mathbf{a}^{(k)})$$
$$\boldsymbol{\delta}^{(3)} = \frac{\partial \mathcal{L}(\mathbf{y}, \hat{\mathbf{y}})}{\partial \hat{\mathbf{y}}} \frac{\partial \hat{\mathbf{y}}}{\partial \mathbf{a}^{(3)}} = \hat{\mathbf{y}} - \mathbf{y}$$

Backpropagation biologically plausible?

$$\Delta \mathbf{W}^{(2)^{T}} \propto \mathbf{W}^{(3)^{T}} (\hat{\mathbf{y}} - \mathbf{y}) f'(\mathbf{a}^{(3)}) \mathbf{h}^{(1)^{T}}$$

 Requires symmetric weights -> bi-directional synapses. Complete knowledge of downstream weights is necessary.

T. P. Lillicrap, D. Cownden, D.B. Tweed, C. J. Akerman, Random feedback weights support learning in deep neural networks

Feedback alignment

Forward computation

For each layer i:

$$\mathbf{a}^{(i)} = \mathbf{W}^{(i)}\mathbf{h}^{(i-1)} + \mathbf{b}^{(i)}$$
$$\mathbf{h}^{(i)} = f(\mathbf{a}^{(i)})$$



Feedback alignment questions

- To which solutions does it converge and why?
- How does their performance compare to back propagation?
- Can it somehow prevent over-fitting?

•

Results on MNIST

-200-80-10 Architecture, tanh units hidden layers, softmax output

-Networks randomly initialized at the same weights in both experiments

-random normal feedback weights N(0, 1)

-Learning rate 0.01



Here feedback alignment mitigates over-fitting

Why would it work?

• *Alignment* of the forward weights with the *feedback* weights.



Lillicrap, Cownden, Tweed, Akerman, Nat Comm. 2016

Results MNIST Feedback Alignment

-200-80-10 Architecture, tanh units hidden layers, softmax output -Random normal feedback weights N(0, 1) -Learning rate 0.01



-1000-10 Architecture, tanh units hidden layers, softmax output

-Random normal feedback weights N(0, 1) -Learning rate 0.01





Epoch average alignment

Test error 3.23%

-1000-10 Architecture, tanh units hidden layers, softmax output-Random normal feedback weights N(0, 1)

-Learning rate 0.01



Alignment of second layer vectors with Feedback weights after each epoch.

Two-layer student teacher with sigmoidal units





During learning keep track of the weight vector overlaps:

 $R_{in} = \boldsymbol{w}_i^{(1)} \cdot \widetilde{\boldsymbol{w}}_n^{(1)}$ $Q_{ik} = \boldsymbol{w}_i^{(1)} \cdot \boldsymbol{w}_k^{(1)}$

In addition, we keep track of $F = w^{(2)} \cdot b^{(2)}$, the overlap between the second layer weights and the feedback vector.

- Number of inputs N=500
- Start from same weight initialization in first layer with $R_{in} = 0, Q_{ik} = 0.2\delta_{ik}, T_{nm} =$, for $1 \le i, k, m, n, \le 5$

Second layer teacher weights δ_{nm}



Student-Teacher FA





Second layer weights align with the feedback vector.

Student-Teacher BP





Student-teacher (K=5) FA vs. BP generalization error



MNIST experiment



Direct Feedback Alignment



A. Nøkland, Direct Feedback Alignment Provides Learning in Deep Neural Networks

Results MNIST Direct Feedback Alignment

-200-200-100-10 Architecture, tanh units hidden layers, softmax output

-Random normal feedback weights N(0, 1)

-Learning rate 0.01



Test error: 4.31%

References

[1] T. P. Lillicrap, D. Cownden, D.B. Tweed, C. J. Akerman, Random feedback weights support learning in deep neural networks

[2] A. Nøkland, Direct Feedback Alignment Provides Learning in Deep Neural Networks