Loopy Neural Nets

- Imitating Feedback Loops In The Human Brain -

Isaac Caswell, Lisa Wang, Chuanqi Shen

Stanford University

Stanford | ENGINEERING
Computer Science

Motivation

Biomimicry: Artificial neural networks are directed acyclic graphs (DAGs), but the network of neurons in our brains, in particular the two main visual object recognition systems ("what" and "where" pathways), contain many feedback loops[1]. We propose loopy neural networks (LNN) as a more faithful mimicry of actual neural net.

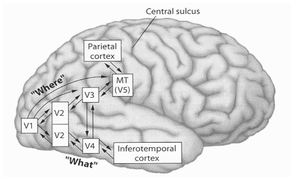

Practical: We are interested in exploring whether LNNs with fewer layers can perform as well, if not better than, deep networks. If true, then LNNs can become a more compact alternative when training neural net models.

Related Existing Architectures

1. **RNNs:** Unrolling of LNN is similar to that in RNN. However, RNNs accept new input at each time step, while LNNs accept the same input at each time-step. Furthermore, original input is elementwise composed with the loop input, rather than processed through matrix multiply first.

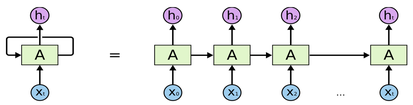

2. **ResNet:** Unrolled addition loops can be interpreted as skip connections.
3. **Attention Networks:** ANs digest input once to indicate where to heed in the input, and then digest input again. This is similar to a one-unroll LNN.
4. **Recurrent CNNs for Scene Labeling[2]:** Feed an image through a convnet, get a result, element-wise compose this with a lower-res version of the input, feed through the convnet again. Iterate. Special case of LNN.

Loopy Neural Net Architecture

We propose a new model of loopy neural networks that mimics the cyclic structures in the human brain by augmenting conventional neural networks with "loops" that allow information from deeper layers to be fed to the earlier layers. The loops can also contain parameters themselves, which we will refer to as "loop layers."

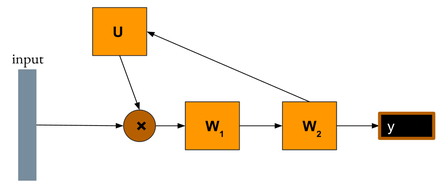

Example LNN with 1 loop

Semantics: the output of U is a layer that weights the importance of each input feature, given that we know what the higher level features produced by W1 are. U allows us to learn the relation between these higher and lower level features.

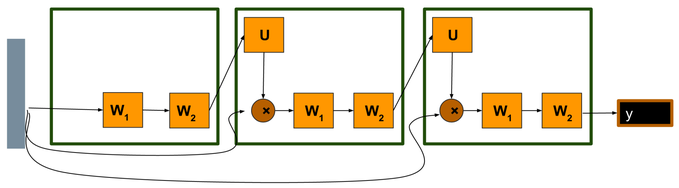

Example LNN unrolled 3 times

Models

1. Model Vanilla (5 convolutional layers, no loops):

input → {conv} X **5** → fc

2. Model Loopy (Model Vanilla with 1 sum loop around all 5 convolutional layers and unrolled 5 times):

input → {conv} X **5** → fc

3. Model Deep (25 convolutional layers and no loops, imitating the depth of the unrolled Model Loopy):

input → {conv} X **25** → fc

Results on CIFAR-10

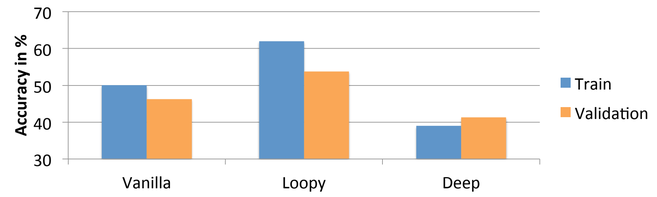

Introducing unrolls results in significantly better accuracy, implying that LNNs are more expressive models. Contrary to our expectation, the deep model performed much worse than the loopy model. This could be due to vanishing gradients.

Visualizations

To explore how loops influence the learning process of the networks, we used guided backpropagation to visualize the filters at different stages of unrolling.

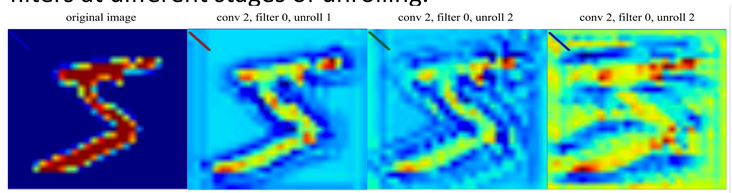

MNIST: Increasing the number of unrolls resulted in more non-linear reactions, a behaviour characteristic of deep layers, from the shallow layers of LNN.

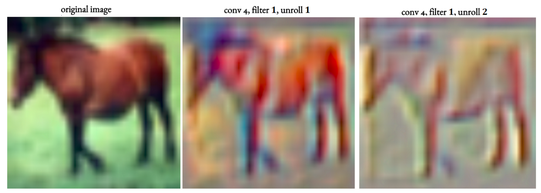

CIFAR-10: Unrolls allowed shallower layers to find features (a characteristic of deep layers) in the image.

[1] M. A. Goodale and A. D. Milner. Separate visual pathways for perception and action. Trends in neurosciences, 15(1):20–25, 1992.
[2] Pinheiro and Collobert. Recurrent Convolutional Neural Networks for Scene Labeling, ICML 2014.