Transfer Learning with Feedback Networks

Max Spero
Stanford University

Introduction

Feedback networks are an alternative tool that find significantly different representations by running a representation of an image through a feedback loop and re-classifying at every iteration.

By applying a range of transfer learning techniques from feedforward networks to feedback networks, we show that transfer learning works in feedback networks as well. We show the additional advantages of early prediction at any iteration still hold through transfer learning.

Datasets

I used three datasets:
- CIFAR-10: ten classes of 32x32 images. Subjects are generally centered
- CIFAR-100: 100 classes, similar to CIFAR-10 but with 1/10 of the training samples per class. Class labels have almost no intersection with CIFAR-10
- Pascal VOC 2012: A 20-class image segmentation challenge with various image sizes and varying subject focus. We used this as a multilabel image classification task.

Preprocessing:
- Resize to 32x32, normalize to mean and standard deviation 0.5

Architecture

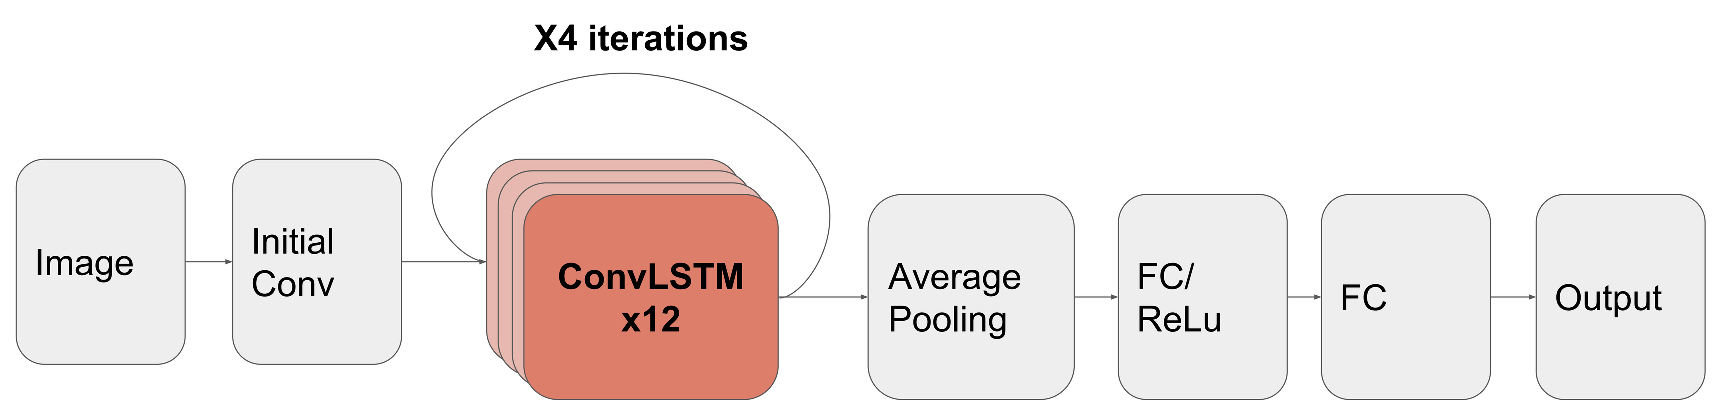

ConvLSTM+Feedback Unrolled

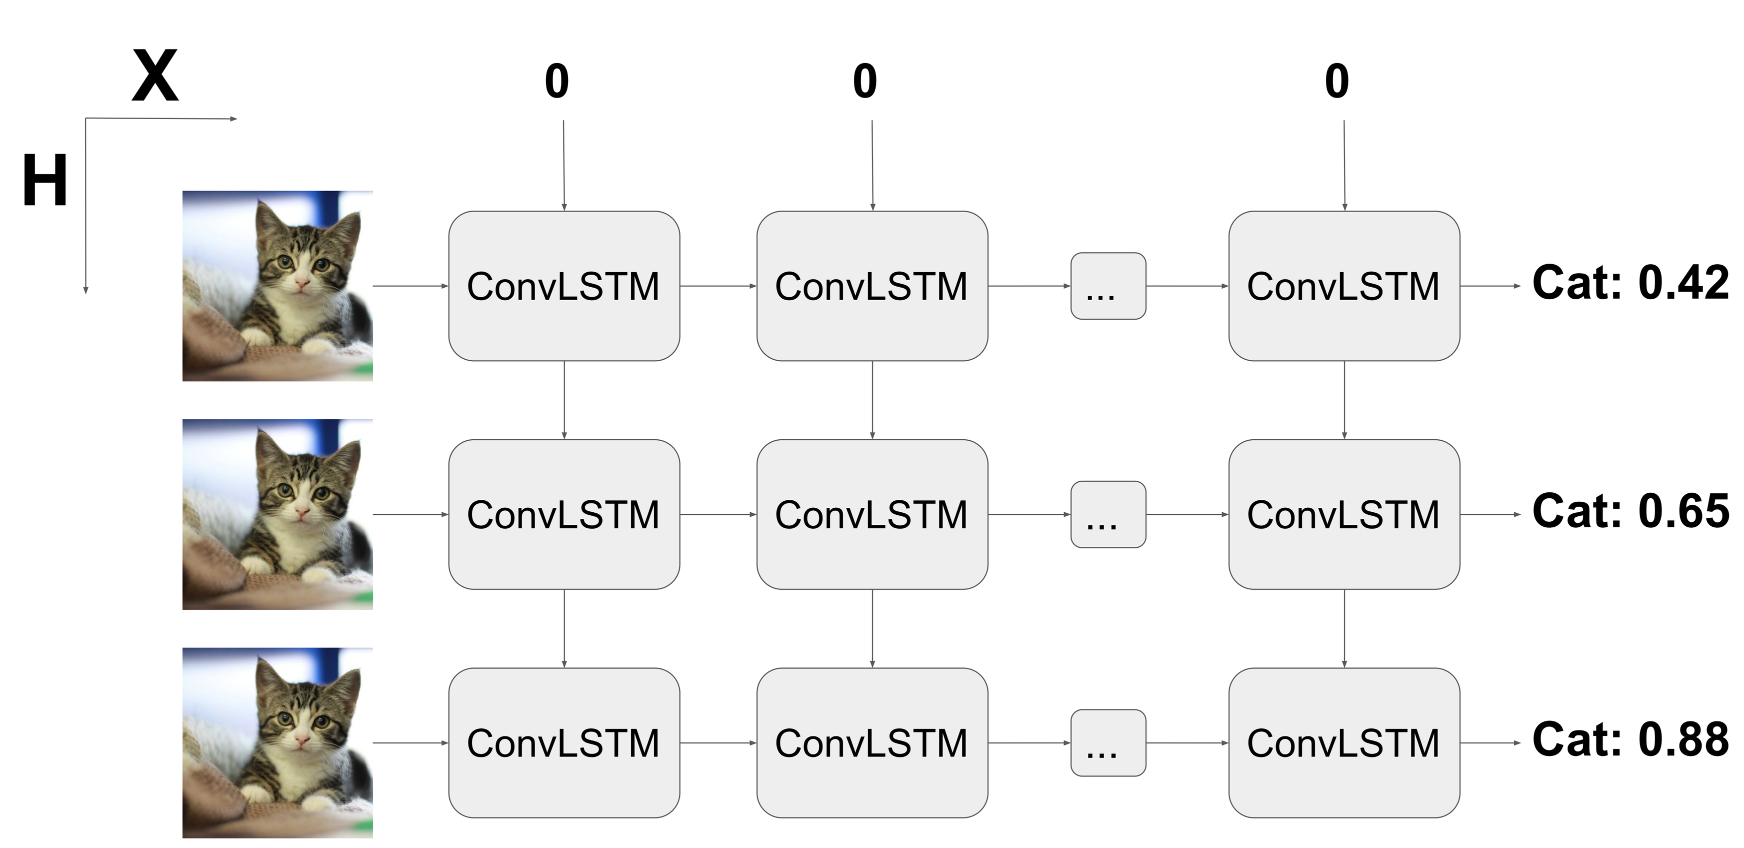

Initial Training

We trained a feedback network separately on three different datasets, yielding the validation accuracies:

	CIFAR-10	CIFAR-100	Pascal VOC (20 labels, multilabel)
Iteration 1	78.70%	33.25%	35.05%
Iteration 2	80.63%	37.22%	36.29%
Iteration 3	81.21%	39.18%	37.80%
Iteration 4	81.36%	38.85%	37.90%

Fine Tuning

By fine-tuning final layers of the networks pre-trained on CIFAR-10 and CIFAR-100, we are able to obtain better results on Pascal VOC.

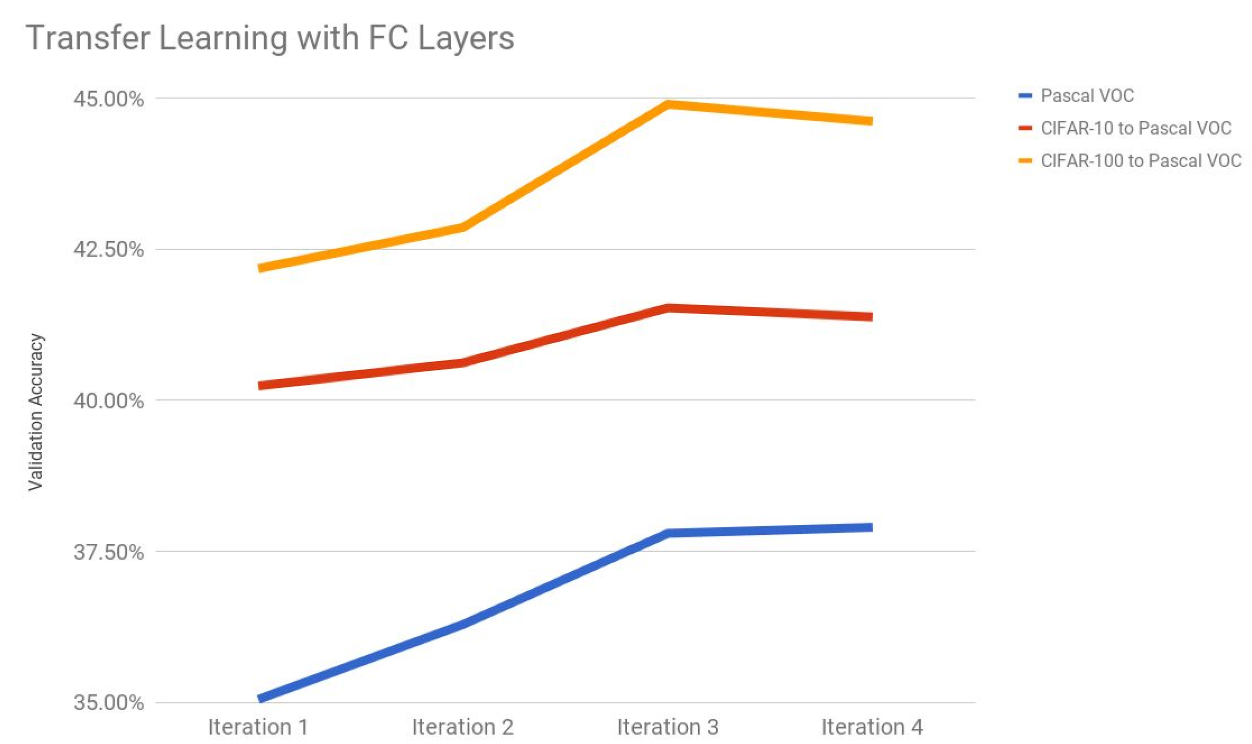

Random Forests

We used the output of the feedback networks from CIFAR-10 and -100 as a feature embedding to train a random forest classifier on Pascal VOC, but it didn't perform better than the fully supervised network.

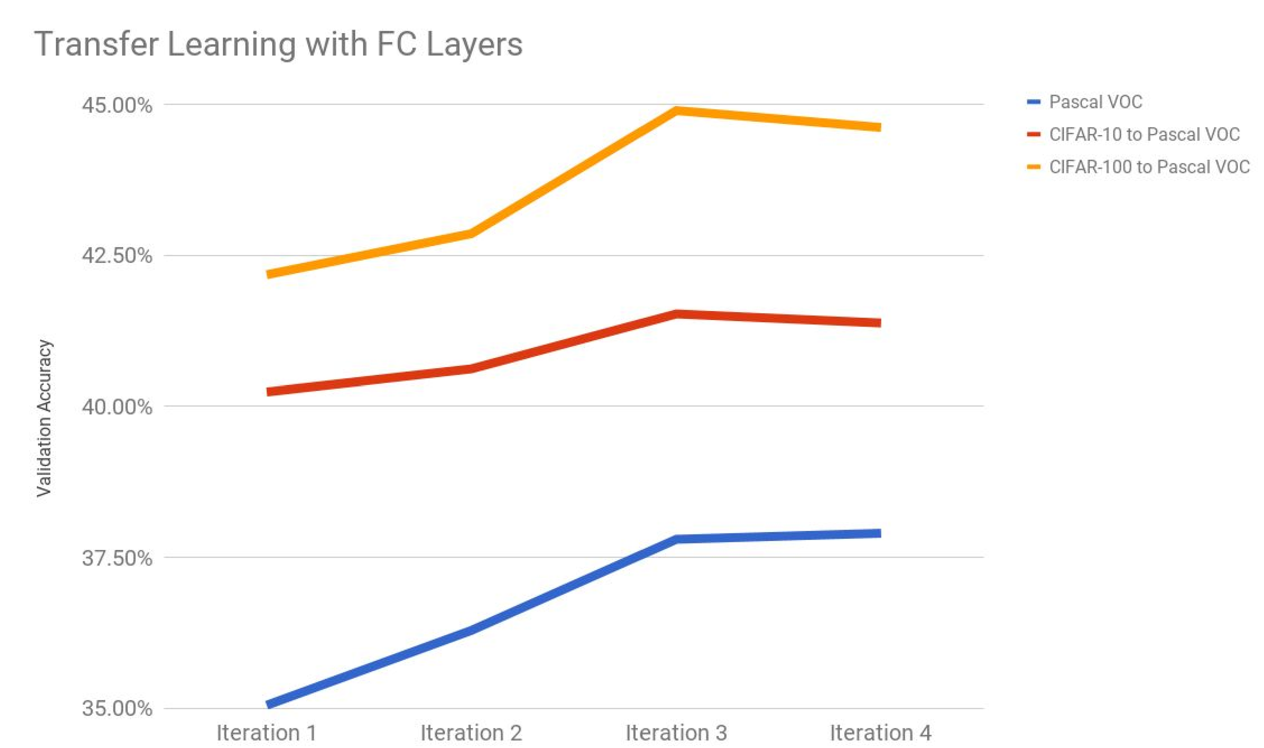

Which iterations are best?

We also ran experiments to see which iteration it is best to train on, by using feedback outputs as a feature embedding for logistic regression. Unsurprisingly, training on all outputs helped reduce overfitting and gave best validation accuracy.

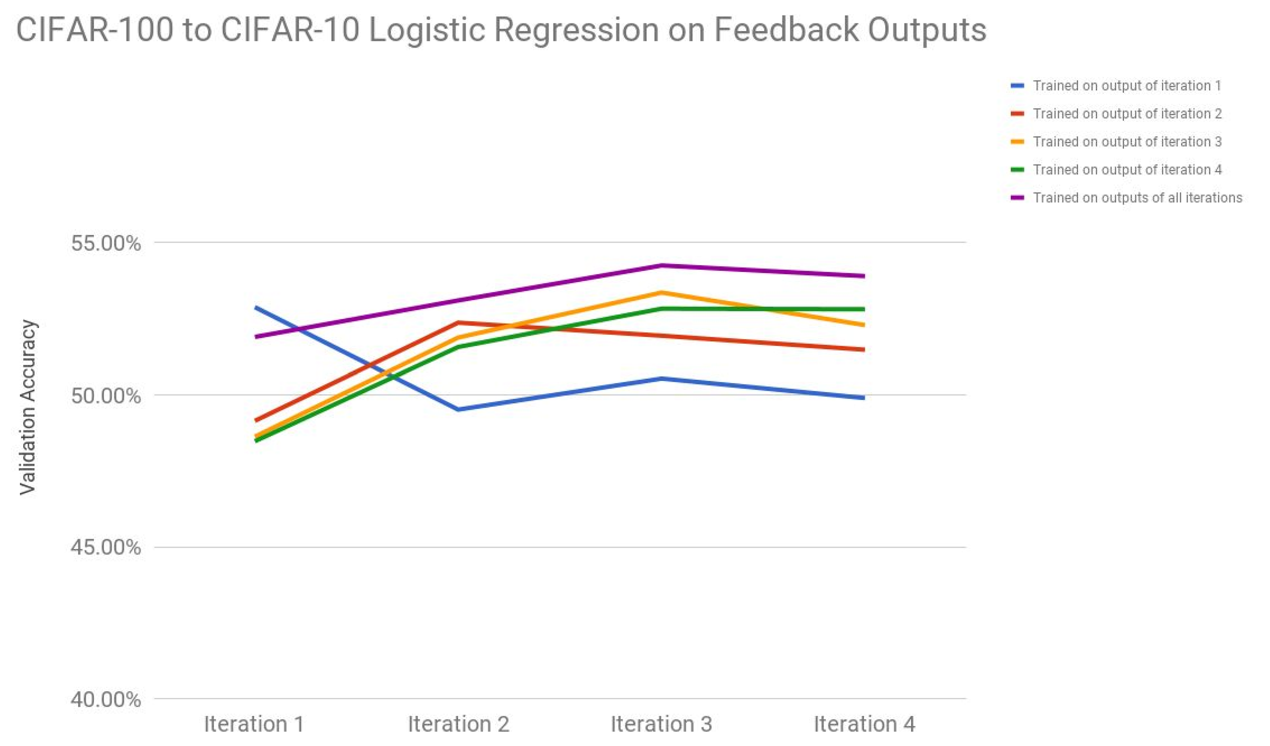

Discussion

Conclusions:
- Fine-tuning the final fully-connected layers of feedback networks provides good transfer learning results.
- Feedback networks are effective in early prediction
- Feedback output is effective as a feature embedding that improves with each iteration
- Training on the output of all iterations helps reduce overfitting and provides better results across the board.

Future Work

- Train a better Feedback Network
  - Wasn't able to obtain amazing results from this simple implementation
- See how feedback networks perform applied across different tasks (for example, detection or semantic segmentation)

References

[1] L. Breiman. Random forests. Machine learning, 45(1):5–32, 2001. 2
[2] S. Xingjian, Z. Chen, H. Wang, D.-Y. Yeung, W.-K. Wong, and W.-c. Woo. Convolutional lstm network: A machine learning approach for precipitation nowcasting. In Advances in neural information processing systems, pages 802–810,
2015. 2
[3] J. Yosinski, J. Clune, Y. Bengio, and H. Lipson. How transferable are features in deep neural networks? In Advances in neural information processing systems, pages 3320–3328, 2014. 2
[4] A. R. Zamir, T.-L.Wu, L. Sun,W. B. Shen, B. E. Shi, J. Malik, and S. Savarese. Feedback networks. In 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 1808–1817. IEEE, 2017. 2