CS231N FINAL PROJECT — MAY 2017

DenResNet: Ensembling DenseNet and ResNet

Victor Cheung

Stanford University
Computer Science
Department

Goal

To achieve high accuracy on Tiny ImageNet with transfer learning and ensembles with Residual Networks and Dense Networks.

Data

200 classes.
100,000 training images.
10,000 validation images.
10,000 test images.

Given limited resurces and training time, how best to produce results on new dataset?

1. *Transfer Learning*
Millions of parameters cannot be easily trained from scratch. Leverage existing pretrained models as feature extractors and finetune!

2. *Shallower Networks*
Experiments show going much deeper than "deep" produces only marginal gains. Marginal benefit small but constant cost of training.

3. *Ensemble Averaging*
Well-studied in statistics. Theoretically enables low bias and low variance predictions! Different model learn different features.

Approach

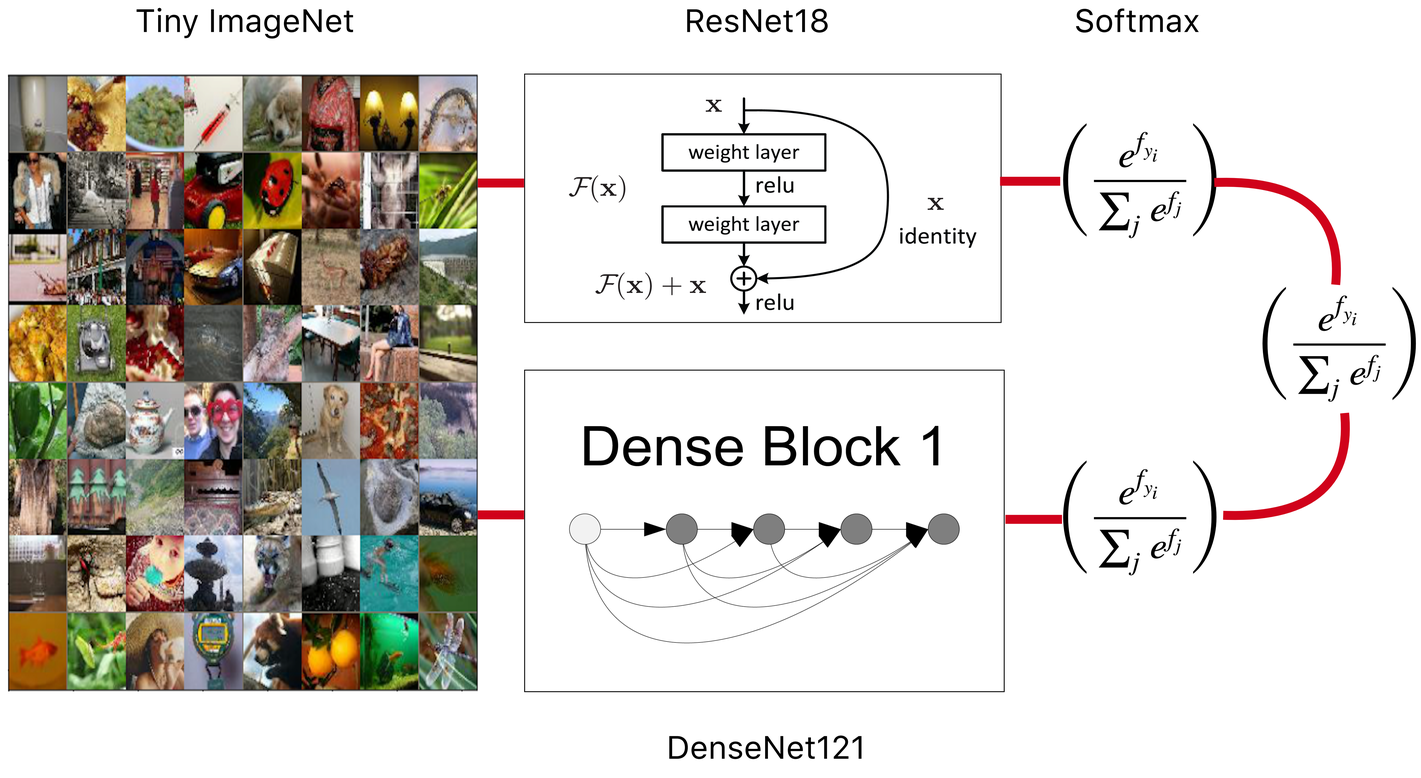

DenseNet121

Training

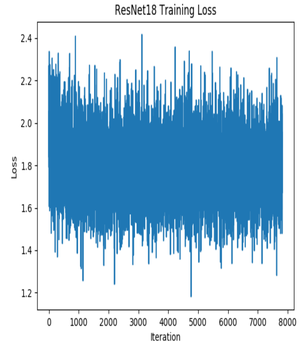

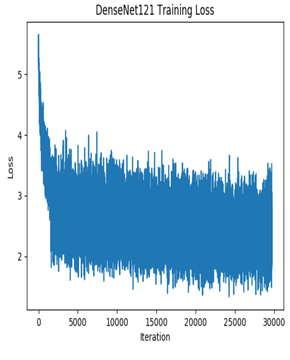

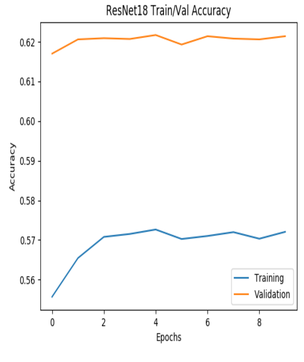

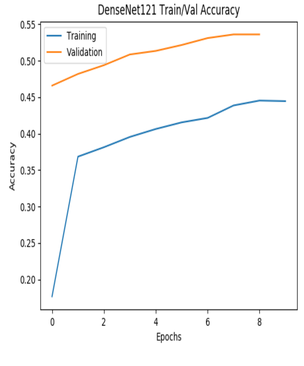

Discussion

~65% Accuracy!

Why softmax? Scores from last affine layer not guranteed to be on same scale across different models. Need to normalize to probability before applying softmax again to get final ensemble output.

Need to train models separately to ensure variance in features emphasized by model. However, only have access to one GPU! Need to train sequentially with handtuning on hyperparameters.

Original emphasis on Bayesian Optimization of hyperparameters did not pay off. Marginal gain in accuracy from updating with Bayesian Posterior.

Train errors due to
1. Incorrect object focus.
2. Similar colors.
3. Noise.

References

1. Deep Residual Networks for Image Recognition.
2. Densely Connected Convolutional Networks.
3. Practical Bayesian Optimization of Neural Networks.
4. PyTorch. Torchvision.