Tiny ImageNet Challenge

Jean-Baptiste Boin

Motivation: We want to achieve the best training error on the but also get more insight on how the different components of a convnet contribute to its performance.

Current data augmentation

Mirror

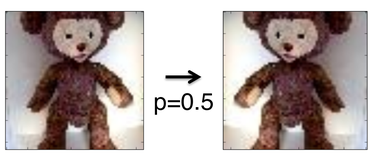

Crop

Disassembling a convnet

Starting with a simple architecture, we add features one after each other and train a network

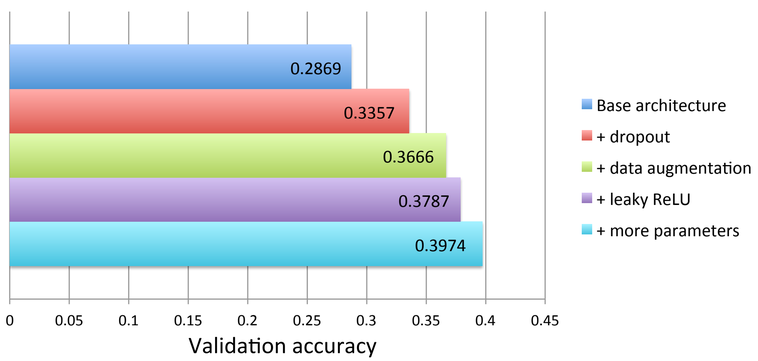

Next steps

Fighting against overfitting:

- Adding more data augmentation (more cropping, color jittering, …)
- Dropout after each stage: according to Srivastava et al., this should help, if done carefully (normalization, etc.)

Improving performance:

- Model ensembles
- Image preprocessing (different color space / PCA)

Architecture of best network

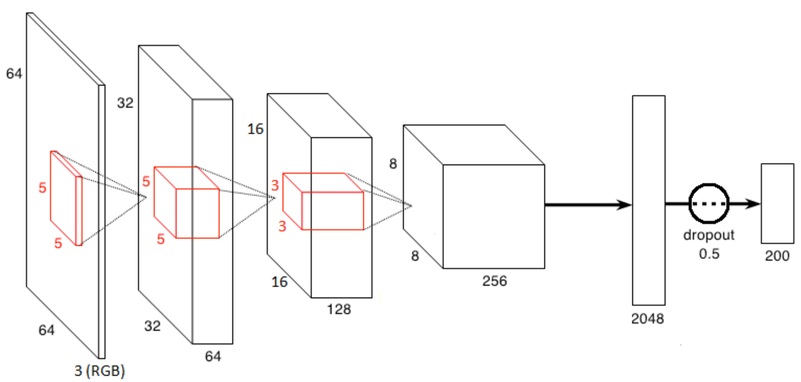

- Softmax
- Leaky ReLU (negative slope: 0.01)
- Max pool after each conv. Layer
- Gaussian initialization

Performance

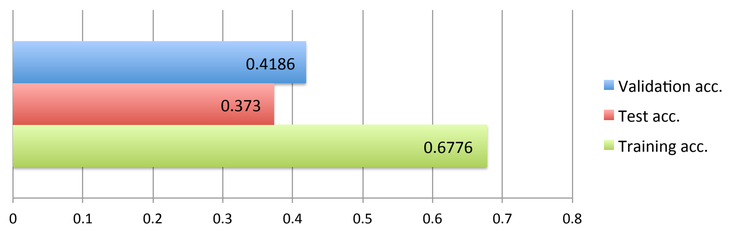

Visualization of first layer

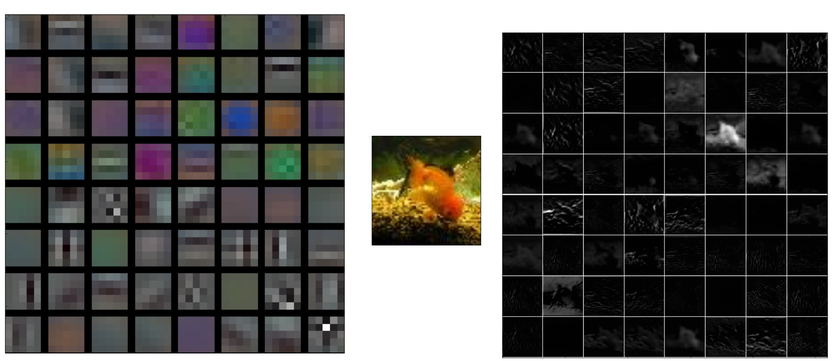

Analysis

Pros:

- Smooth and discriminative weights
- Architecture with relatively high capacity
- Fast training

Cons:

- Considerable overfitting, despite dropout
- Large filter size for the first two layers (tried to replace by several layers with smaller filters ; no improvement yet)