



# Tiny ImageNet Challenge

Yinbin Ma, Stanford University

## Abstract

In this project we classify the images in the tiny ImageNet dataset. We train the residual network on GPU using PyTorch. We use ResNet-18 to achieve 53.9% error rate on the test dataset. The effect of weight initialization is studied.

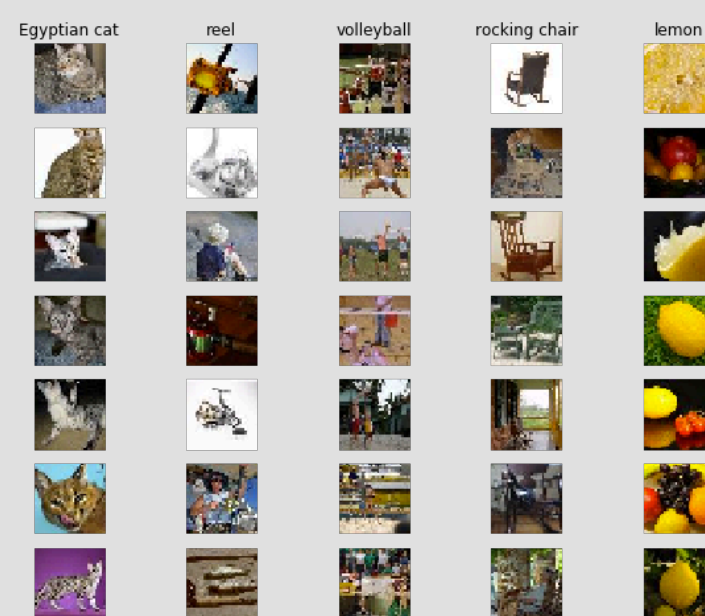
## Background and Introduction

- Convolution neural network (CNN) is a powerful tool for image classification, however it is hard to train deep network
- Residual networks allow us to train deeper network
- Weight initialization and the activation function affects the training process

## Problem Statement and Dataset

In this project we work on image classification problem using the tiny ImageNet dataset.

The tiny ImageNet dataset contains images with 200 different categories. The training set has  $10^5$  images (500 images per category), while the validation set and test set has  $10^4$  images.




Samples of images from the tiny ImageNet dataset

## References

- K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. CoRR, abs/1512.03385, 2015.
- K. He, X. Zhang, S. Ren, and J. Sun. Identity mappings in deep residual networks. In ECCV, 2016.
- X. Glorot and Y. Bengio. Understanding the difficulty of training deep feedforward neural networks. In Aistats, volume 9, pages 249–256, 2010.
- D. Clevert, T. Unterthiner, and S. Hochreiter. Fast and accurate deep network learning by exponential linear units (elus). CoRR, abs/1511.07289, 2015.
- D. Sussillo and L. F. Abbott. Random Walk Initialization for Training Very Deep Feedforward Networks. ArXiv e-prints, Dec. 2014.

## Residual network

We use ResNet-18 with the following configuration:




Convolution filters from the first layer. We run 30 iterations from the pre-trained model.


Left: ResNet-18 configuration. Pre-trained models exist for ImageNet with the last layer being fc-1000 instead of fc-200

## Overfit a small dataset

First we use ResNet-18 to overfit a small dataset of 100 images.



Overfit from pre-trained model.




Overfit from scratch.


The pre-trained model start from low accuracy but converge fast. We fail to overfit when starting from scratch.

## Train the residual network

We train the ResNet-18 with 5000 images from the training set. The best error rate on the test dataset is 53.9% so far



Filters from pre-trained model.




Starting from scratch.

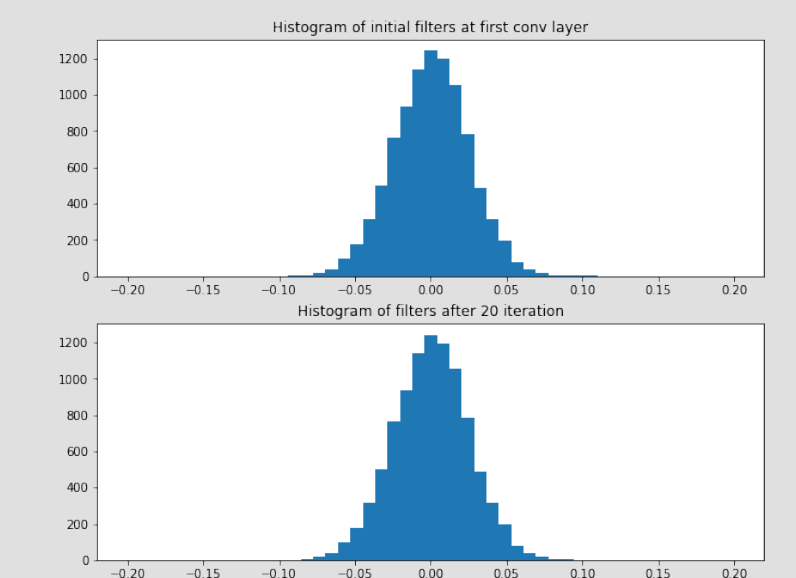
Training from scratch is still in progress

## Weight initialization

We train the ResNet with different weight initialization



Gaussian initialization mean=0, variance=0.1



Xavier initialization

## Conclusions and Future work

- We train the residual network on the tiny ImageNet dataset
- We use pre-trained model to achieve 53.9% error rate on the test dataset, while training from scratch is challenging
- The study of weight initialization and activation function is in progress