GluonCV: Image Classification

Tong He
Applied Scientist, Amazon Web Services.
2018.12.17

Introduction

What is classification?

Tell you what is in the image

Applications

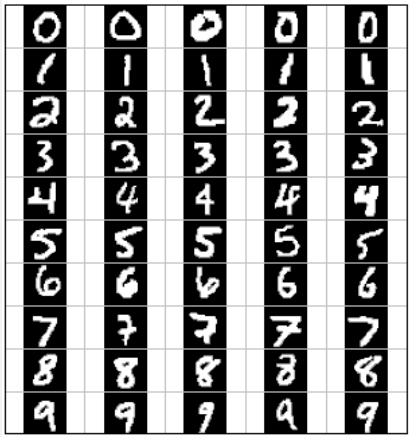

Digit Recognition

Car Model Detection

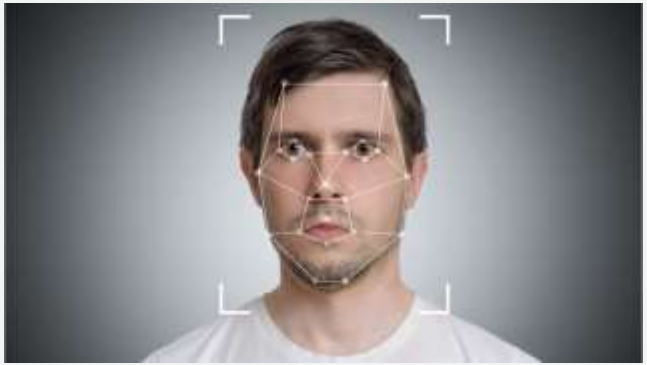

Facial Recognition

aws

Image Classification with Deep Learning

Classification with Deep Learning

MNIST: The "Hello World" dataset

- 28 x 28 input
- Grayscale
- 50000 training
- 10000 test

aws

Classification with Deep Learning

Fully-connected Layer

- Expensive
- 2D insensitive

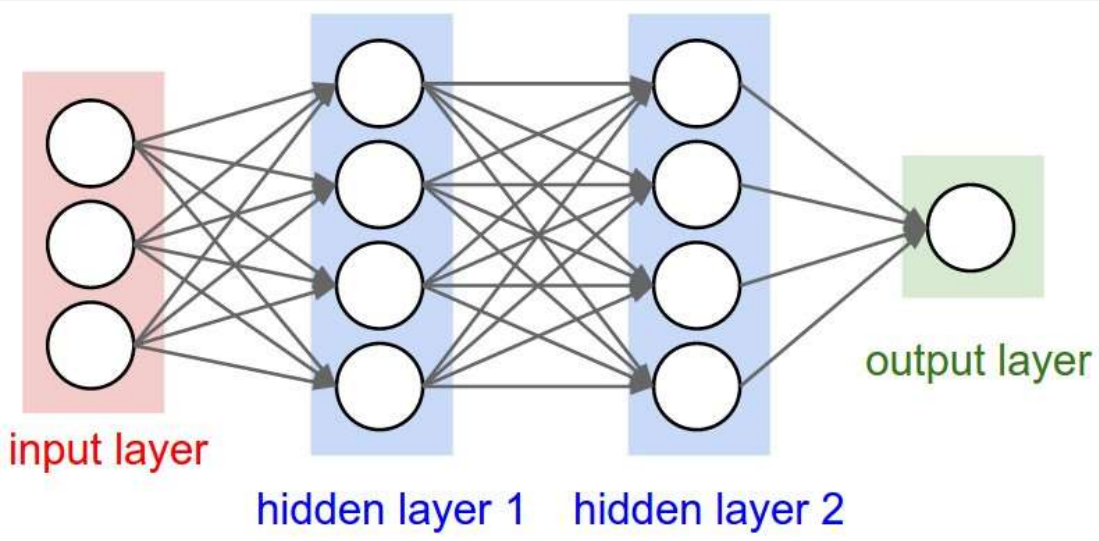

aws

Classification with Deep Learning

Convolution Layer

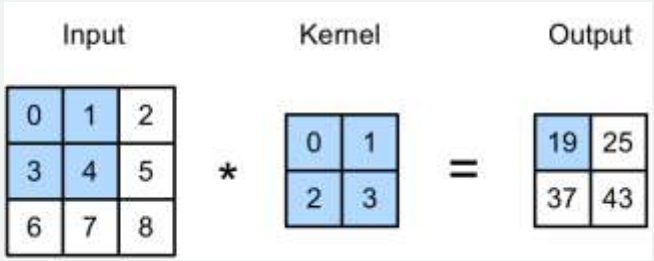

0*0 + 1*1 + 3*2 + 4*3 = 19

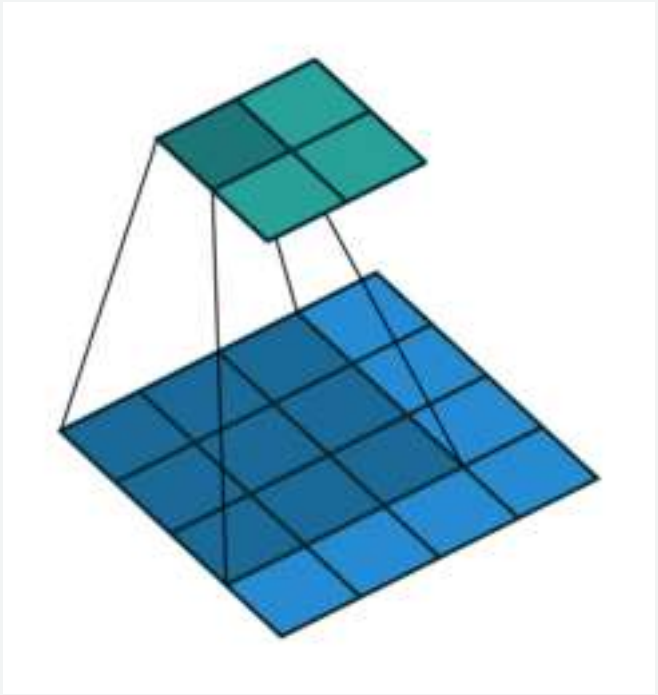

aws

Classification with Deep Learning

Pooling Layer

- Demonstration Video

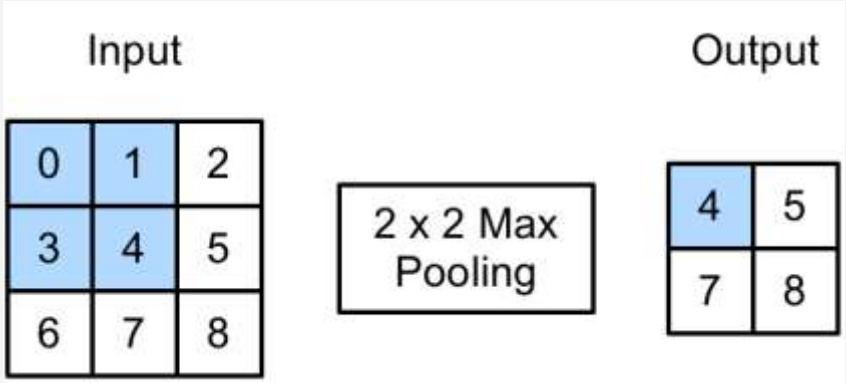

Classification with Deep Learning

Activation Layer

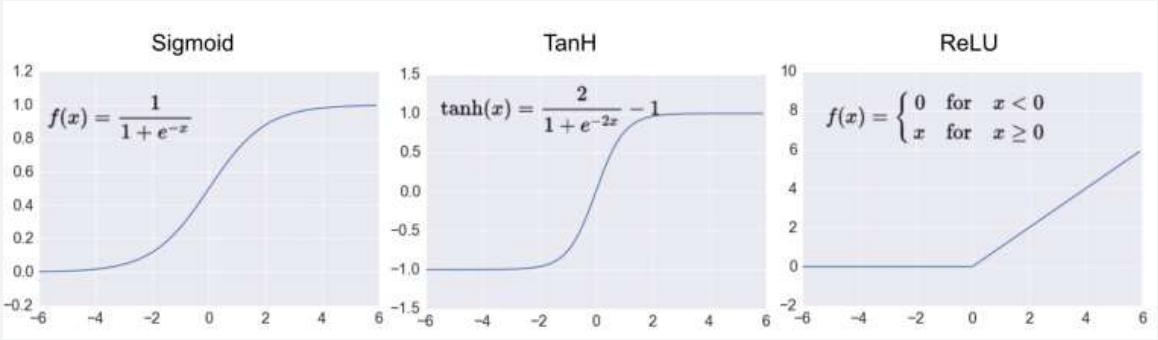

Classification with Deep Learning

LeNet

- Demonstration WebSite

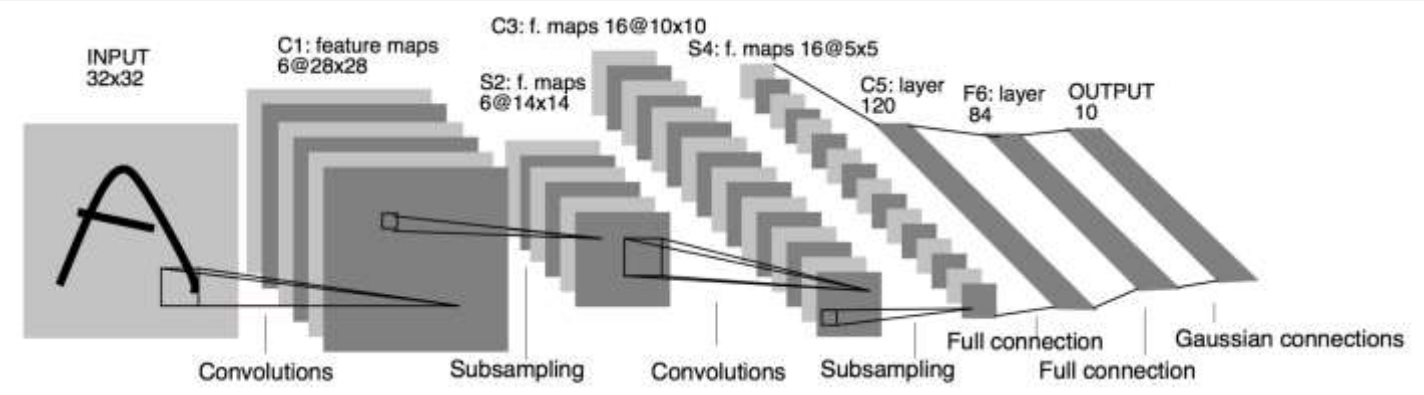

Classification with Deep Learning

ImageNet Challenge

- Natural images
- 1000 classes
- 1.2 million images

aws

Classification with Deep Learning

Difficulties

- LeNet is too small
- Hard to process huge amount of data

Classification with Deep Learning

AlexNet

- GPU Accelerated
- Deep (8 layers)

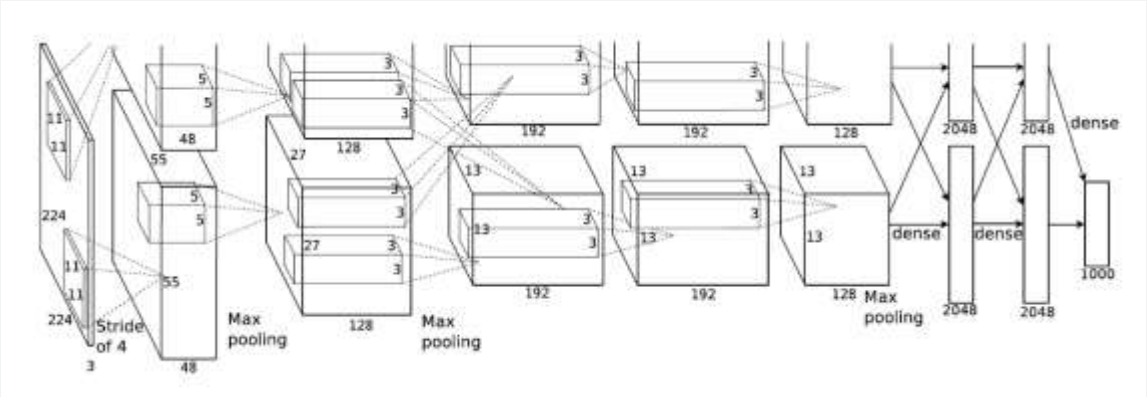

aws

Classification with Deep Learning

VGG

- Deeper
- Better architecture

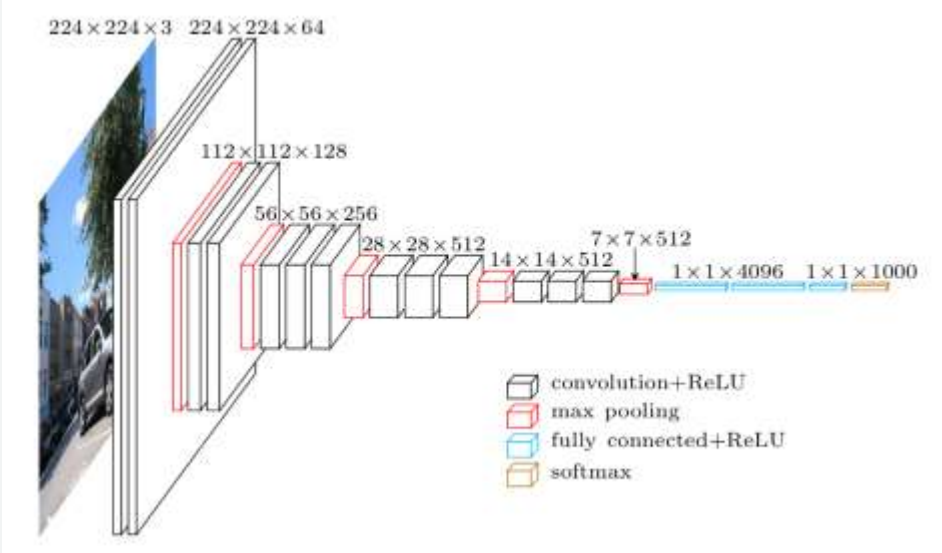

Classification with Deep Learning

Can we go deeper?

- Vanishing/Exploding Gradient
- Chain rule: $\frac{\partial f}{\partial w} = \frac{\partial f}{\partial g} \cdot \frac{\partial g}{\partial w}$

$$\frac{\partial f}{\partial g} = 0.1 \quad \frac{\partial g}{\partial w} = 0.1 \quad \Big\} \quad \frac{\partial f}{\partial w} = 0.01$$

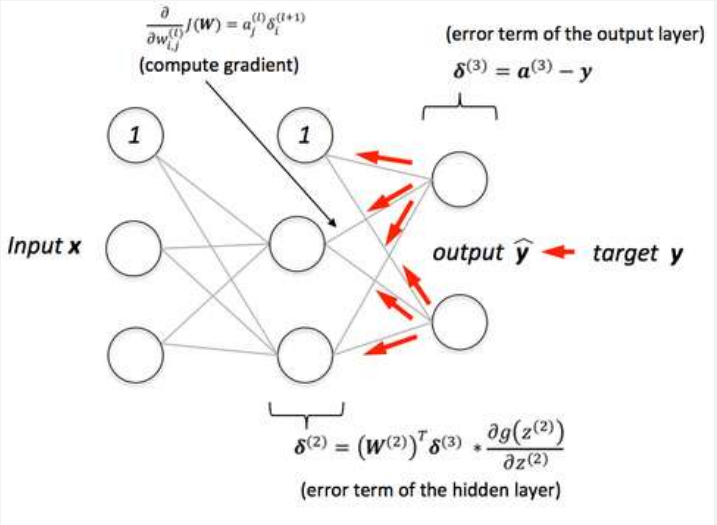

aws

Classification with Deep Learning

ResNet

- Residual
  - 152 layers!
- Modularized

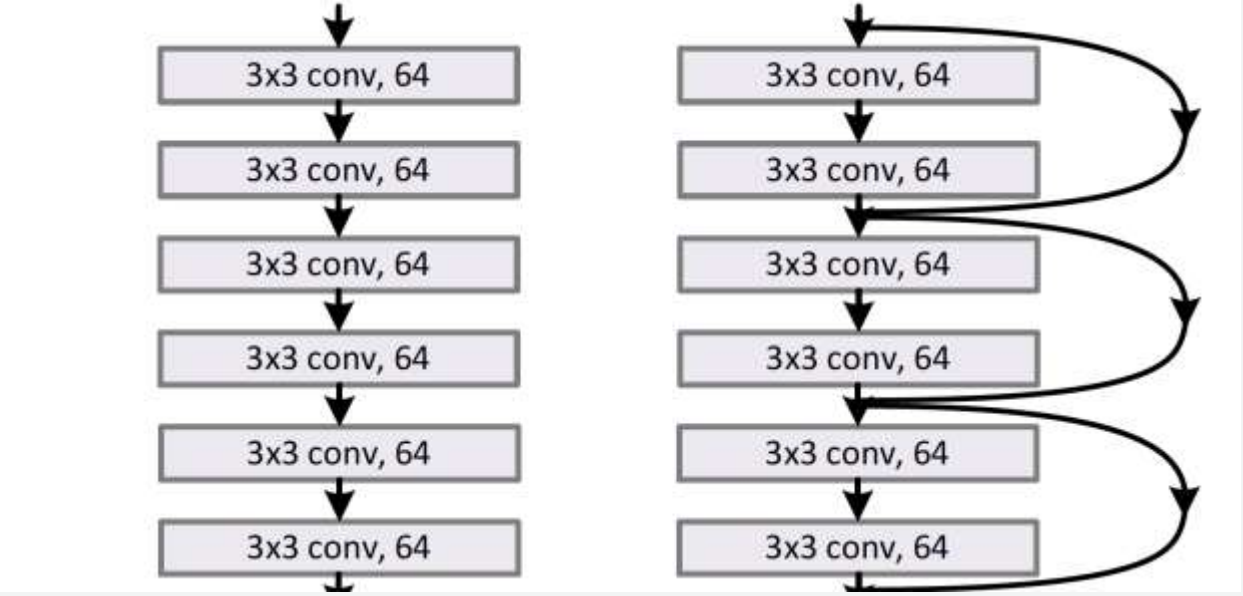

Plain network

Network with residual

aws

Classification with Deep Learning

ImageNet Results

- Top-5 error: 25% to 5%
- Hundreds of new models

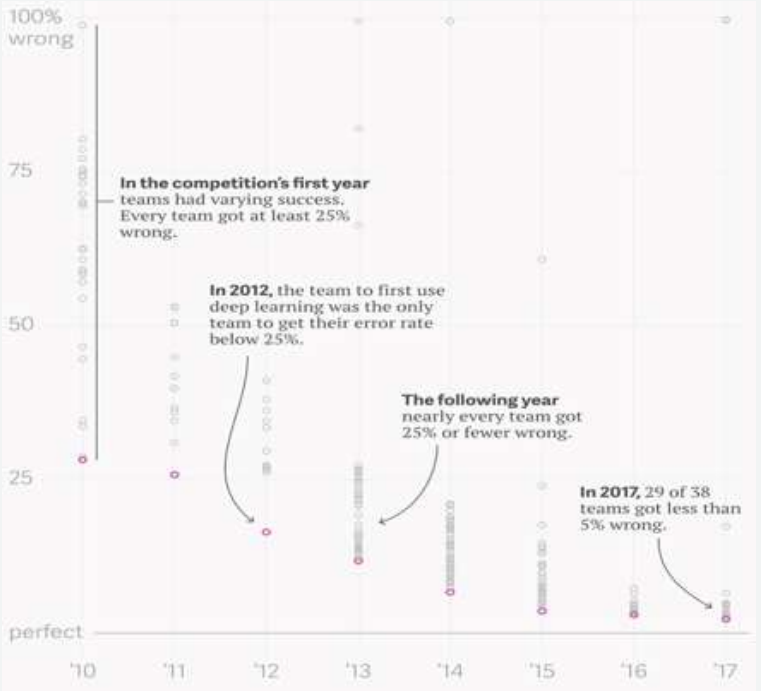

Classification with Deep Learning

MobileNet

- Depth-wise Convolution
- Parameterized Size
- Fast Inference on devices

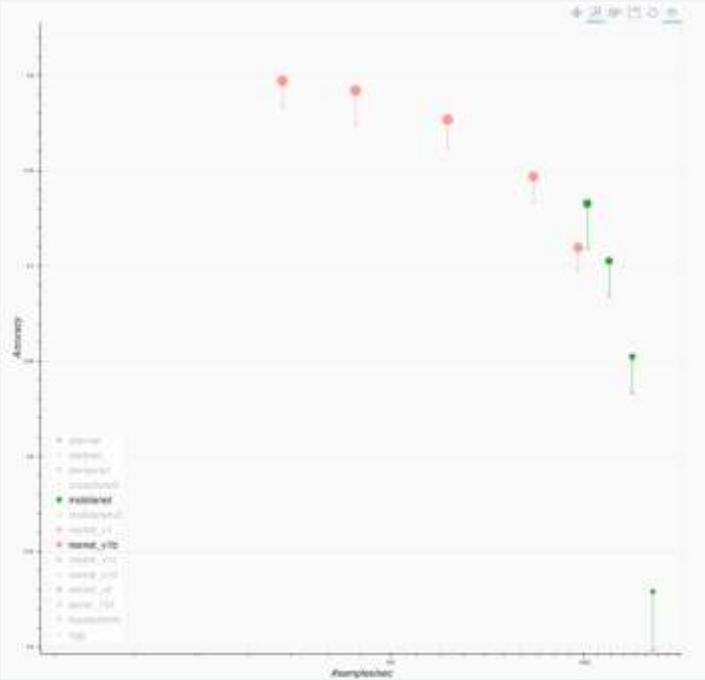

Classification with Deep Learning

NASNet

- Automatic Search
- Accurate and heavy

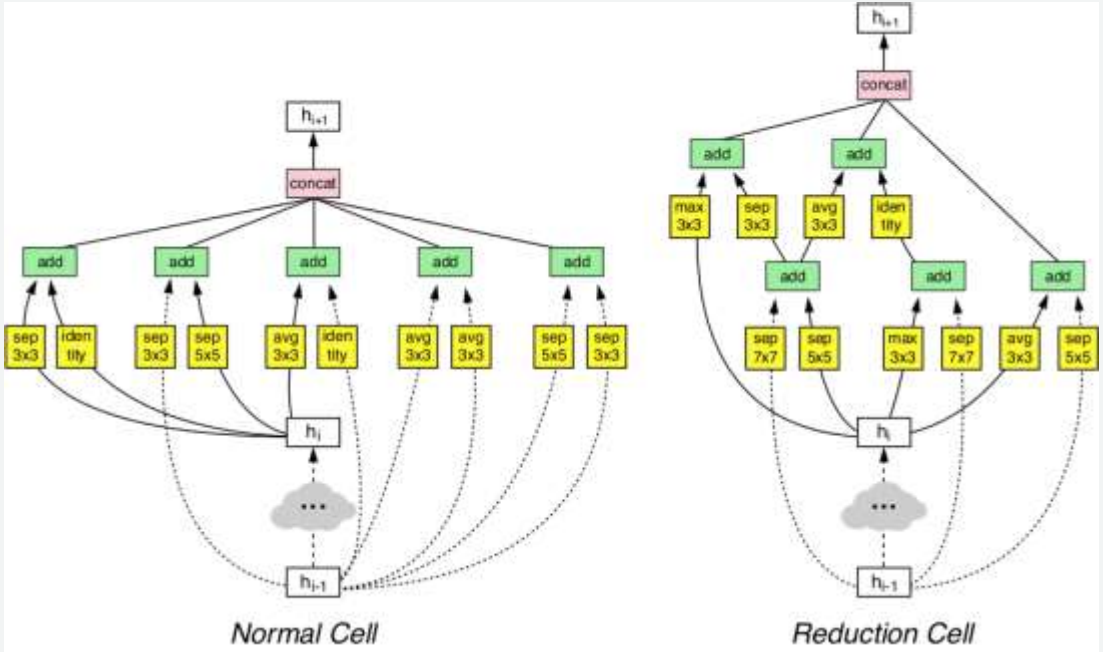

Model Summary

ResNet:

- Balanced accuracy and speed
- Well modularized

MobileNet:

- Small size
- Fast inference

NAS:

- Ongoing research topic

Classification with GluonCV

Classification with GluonCV

Model Zoo

- Pre-trained models
- Can be transferred or directly applied

Classification with GluonCV

GluonCV Model Zoo

- Comprehensive selection
  - AlexNet
  - VGG
  - ResNet
  - MobileNet
  - NASNet
  - …
- One of the most accurate open-sourced libraries
- Reproducible

aws

Classification with GluonCV

Training Essentials

- Data Preprocessing
- Network architecture definition
- Optimizer
- Loss
- Metric
- GPU Acceleration

Classification with GluonCV

Data Preprocessing

Classification with GluonCV

Optimizers

- SGD
- Adam
- RMSProp
- …

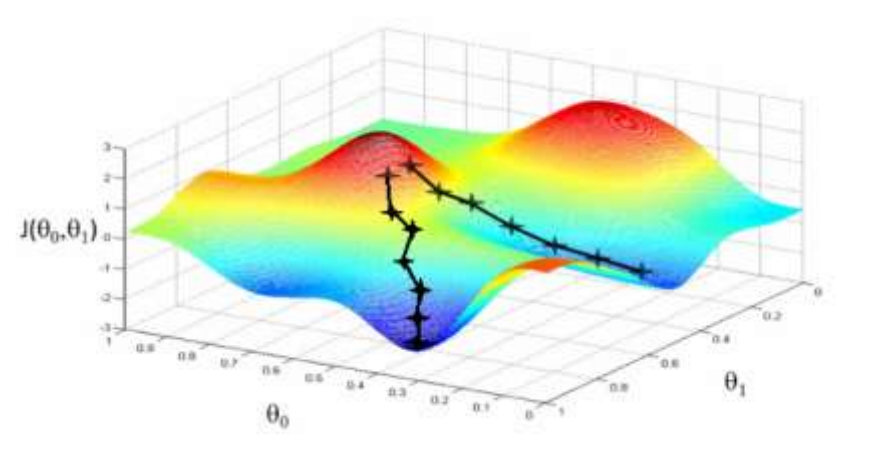

aws

Classification with GluonCV

Advanced Tricks

- Label smoothing
- Learning rate schedule
- Mix-Up
- Knowledge Distillation

Classification with GluonCV

Label Smoothing

- One hot: (0, 1, 0, 0, 0)
- Smoothed: (0.01, 0.96, 0.01, 0.01, 0.01)

- Prevent overfitting!

aws

Classification with GluonCV

Learning Rate Schedule

- Step
- Cosine
- Poly

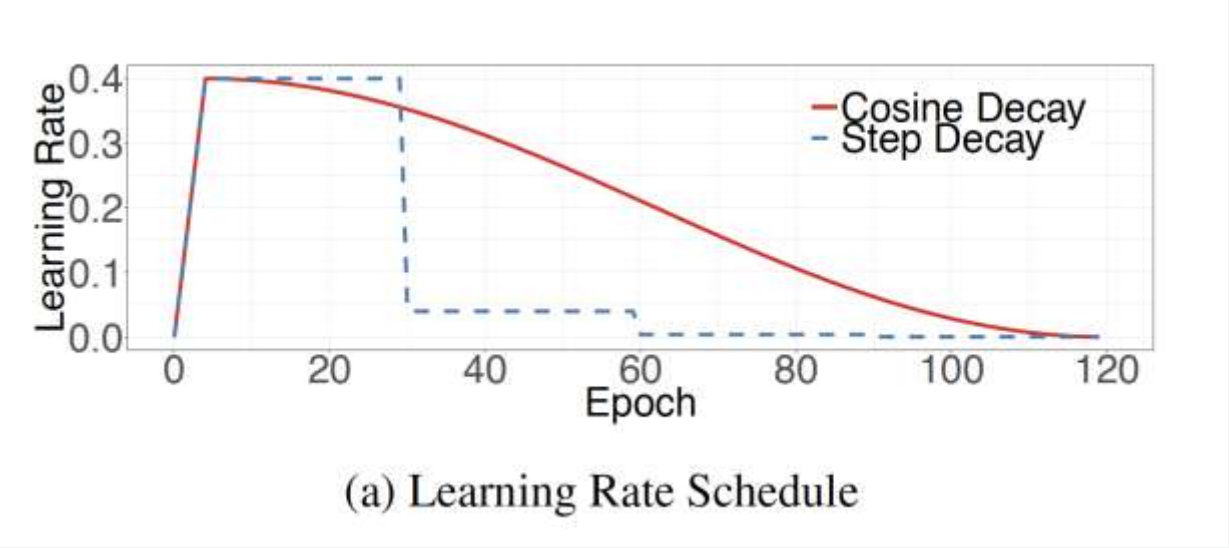

(a) Learning Rate Schedule

aws

Classification with GluonCV

Mix-Up

- Linear mapping
- $f(ax_i + bx_j) = af(x_i) + bf(x_j)$

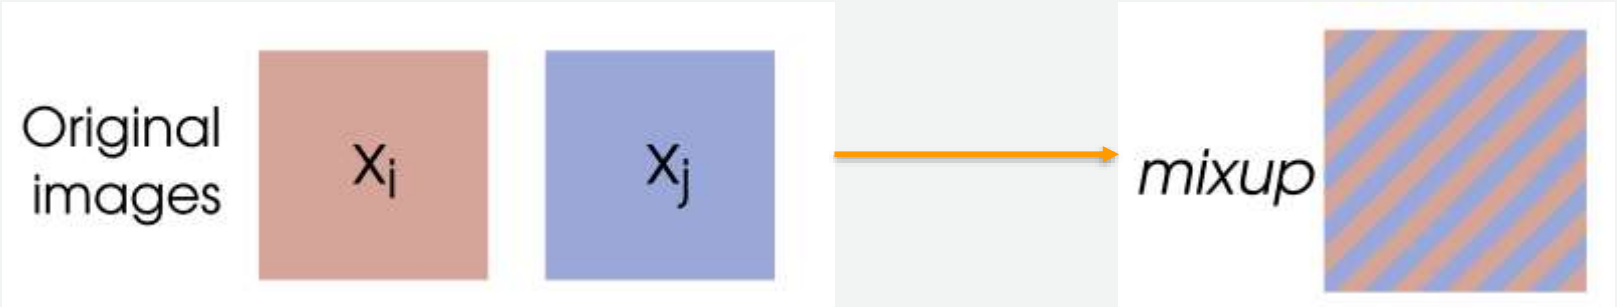

aws

Classification with GluonCV

Knowledge Distillation

- Dark Knowledge
  - Dog vs Cat
  - Dog vs Car

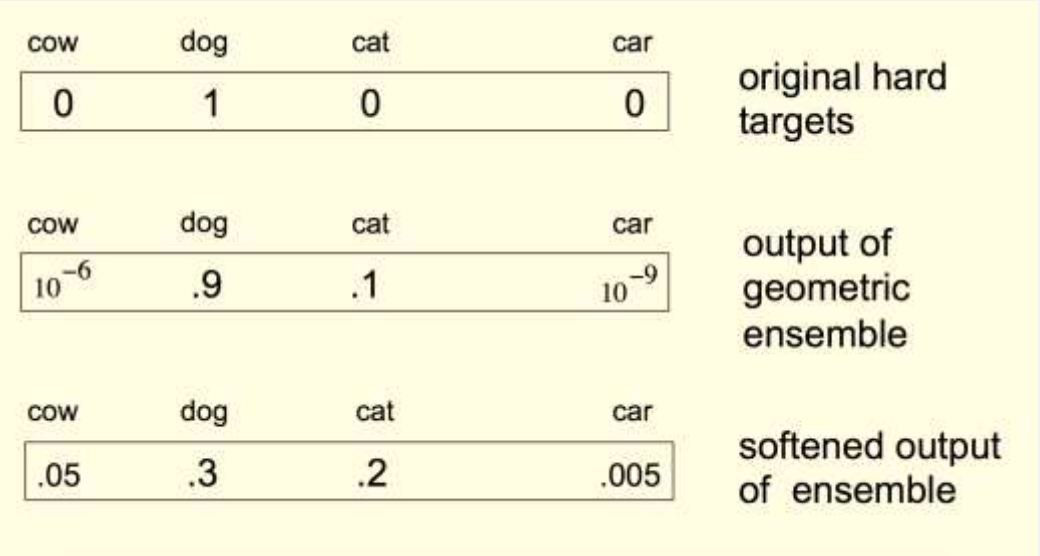

aws

Classification with GluonCV

Knowledge Distillation

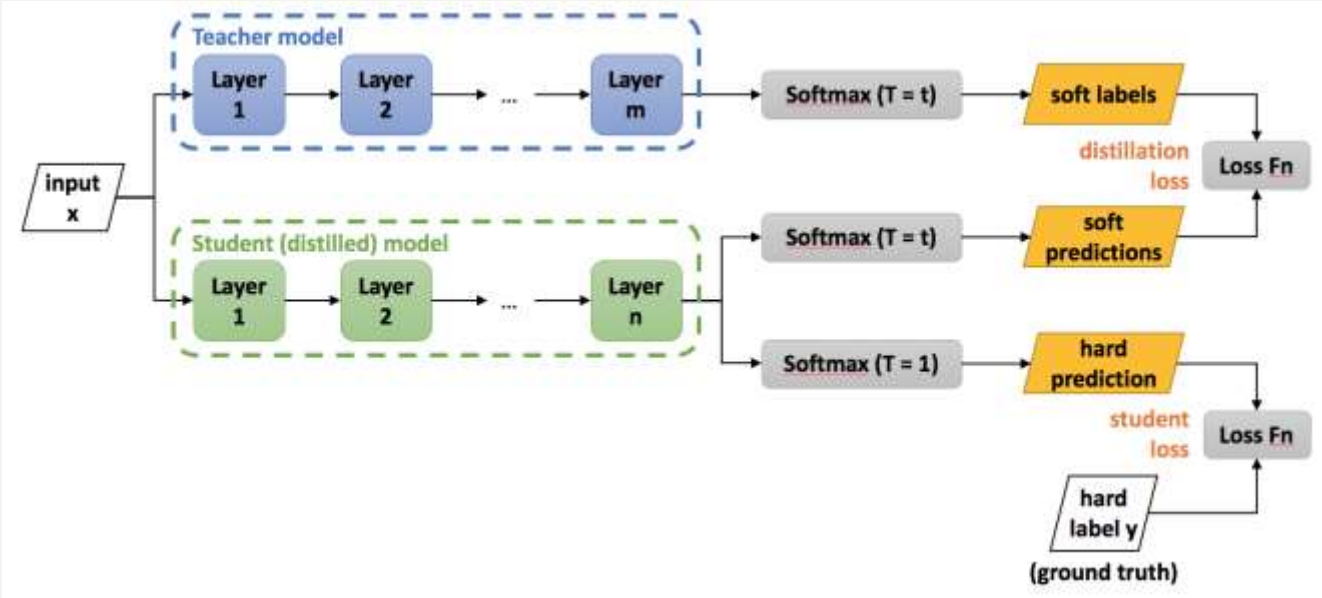

aws

Classification with GluonCV

GluonCV Model Zoo

He, Tong, et al. "Bag of Tricks for Image Classification with Convolutional Neural Networks" arXiv preprint arXiv:1812.01187 (2018).

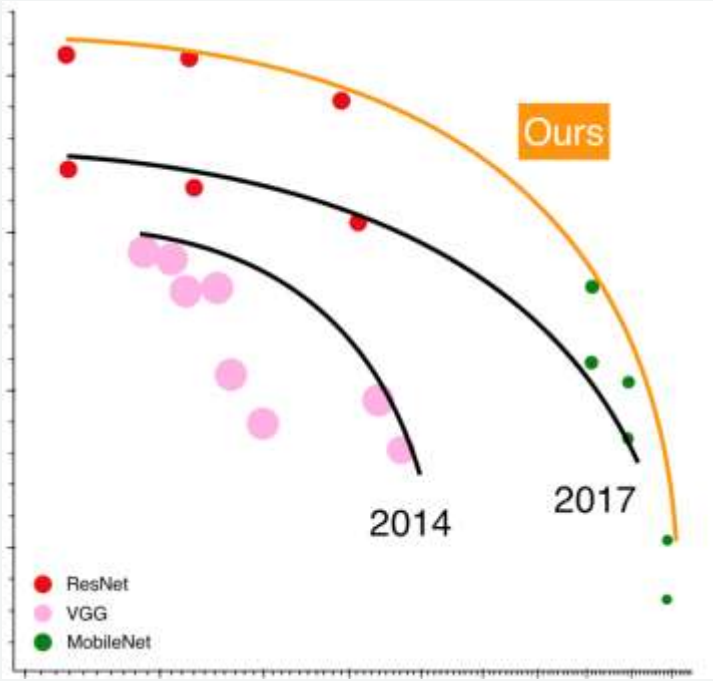

aws

Classification with GluonCV

Transfer learning

- Based on a pre-trained model
- Re-define the output layer

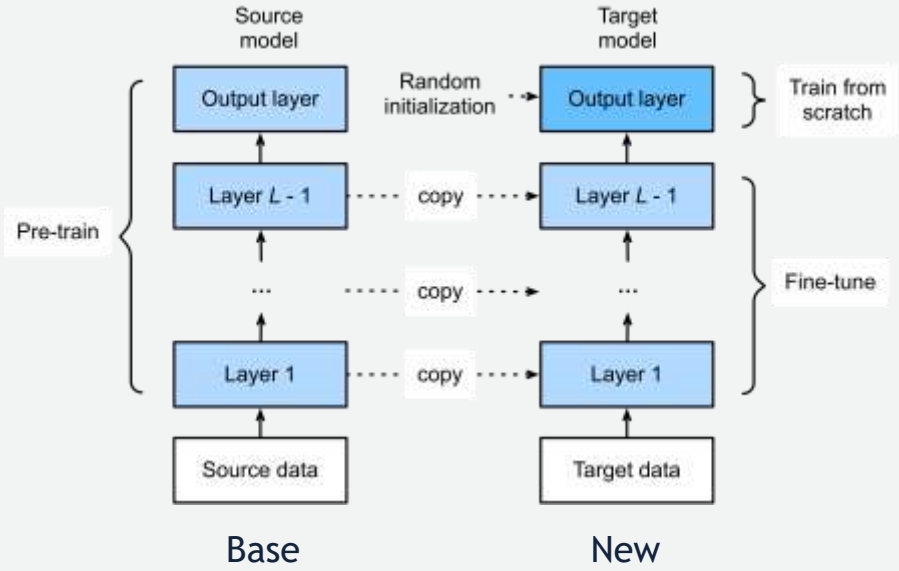

Classification with GluonCV

Resources:

- Model Zoo: https://gluon-cv.mxnet.io/model_zoo/classification.html
- Tutorials: https://gluon-cv.mxnet.io/build/examples_classification/index.html
- Deep Learning Book: http://diveintodeeplearning.org/

aws

Classification with GluonCV

Hands on!