



PyCon Thailand 2019

First steps in Deep Learning with TensorFlow 2.0 :
CNNs



Martin Andrews

Google Developer Expert, Machine Learning

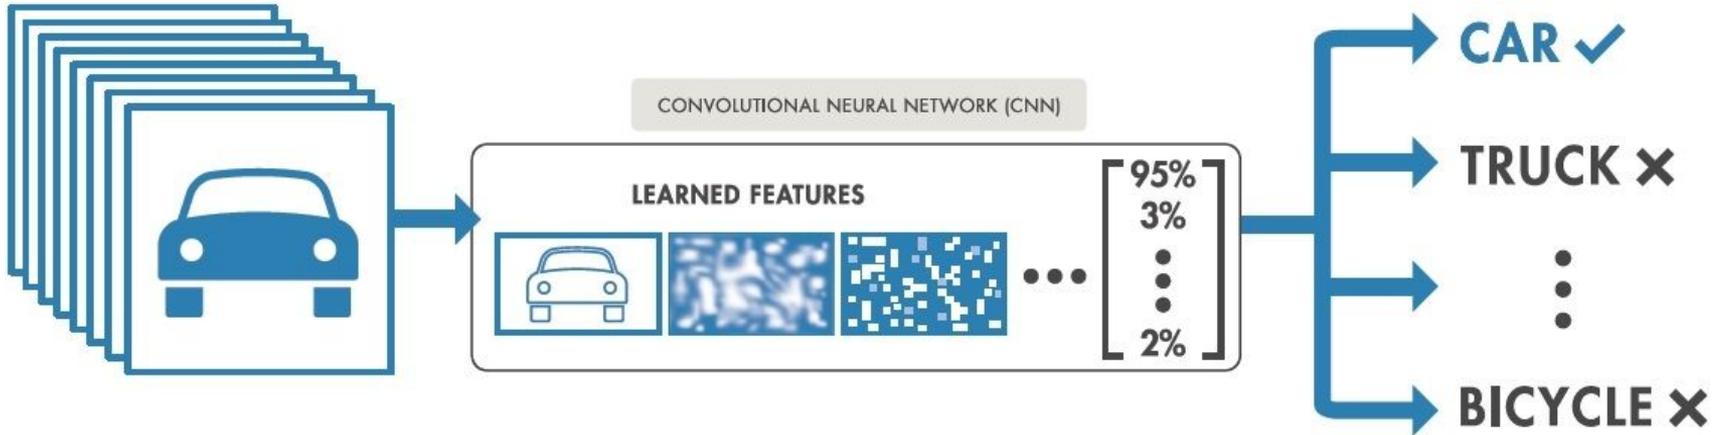
Red Dragon AI, Singapore

Outline

- Image classification : The problem
- What are CNNs?
 - Code-first explanation
- Why is a CNN model trainable?
- Let's train a model
 - Standard thing
 - Less standard thing
- Wrap-up

The Problem

Image Classification



What are CNNs?

Convolutional Neural Networks

- Translational invariance
 - ... yada yada yada ...
- Instead, let's understand what these things **do**

Code

We have the layers...

How does the model learn?

- We need a few more pieces :
 - List of Parameters
 - Training data (i.e. correct answers)
 - Optimisation
 - Testing

List of Parameters

- We've seen this already :
 - Each Conv2d has a kernel (+ bias term)
 - (ReLUs and SoftMax don't need parameters)
 - ... that's it

Training Data

- We'll see this soon :
 - Lots of images and their correct classes
- Split the data into 'training' and 'test' portions
 - Ensures that model works on unseen data

Optimisation

- This is a bit more involved...
- Key piece of mathematics :
 - Measure the model's error (= 'loss')
 - Want to reduce the 'loss' to improve model
 - But 'loss' depends on every parameter
 - Repeatedly nudge each parameter in correct direction
 - Model should improve (may be slow)

Testing

- We kept some testing data to one side :
 - (more images and their correct classes)

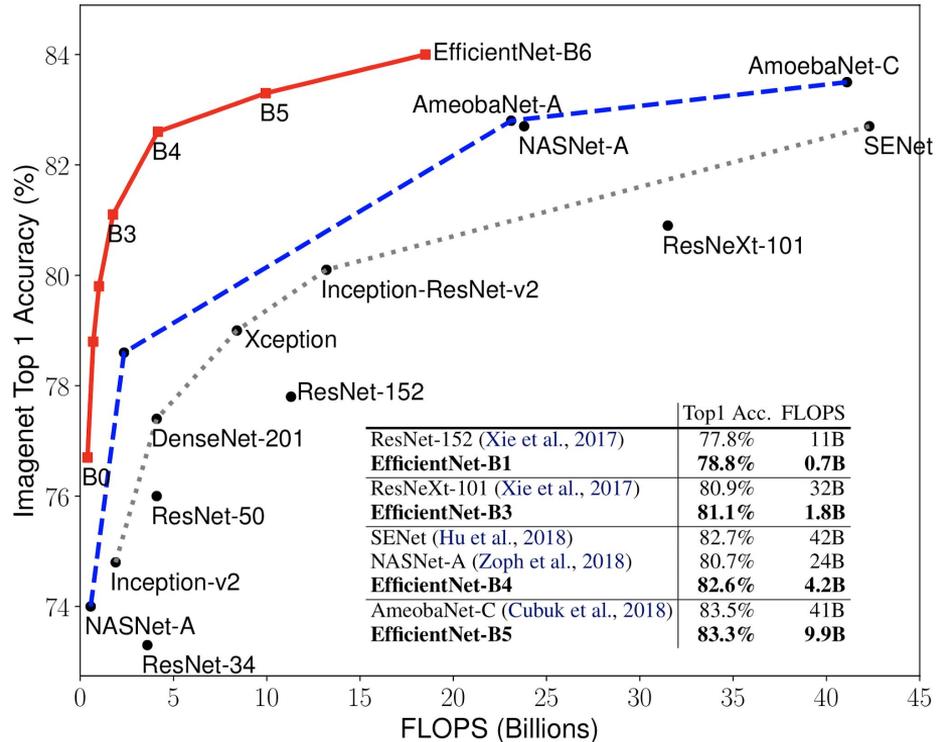
- We'll see whether the trained model works
 - (fingers crossed)

Code (again)

How to invent the architecture?

- Historically:
 - Make networks bigger/deeper
 - Modularising : Inception Blocks
 - Outside-the-box thinking : ResNets
- Recently:
 - Problem : Lots of possibilities
 - Use machines to design the network

CNN Efficient Frontier



Current Winner : EfficientNets

- Research published in May 2019
 - Models available in multiple frameworks
 - *But team-keras is not yet TF2-aware*



Code (last one)

Wrap-Up

- Basic CNN components shown
- Basic CNN model built
(Hands-waved about training)
- Model trained on 2 datasets
- Advanced, pretrained model used

- Success!

Deep Learning MeetUp Group

The Group :

- MeetUp.com / TensorFlow-and-Deep-Learning-Singapore
- > 3,800 members

The Meetings :

- Next = 6-July, hosted at Google
 - Something for Beginners
 - Something from the Bleeding Edge
 - Lightning Talks

Deep Learning Courses in Singapore

Jumpstart Course : July 4+5

- Hands-on with real model code
- Build your own Project

Other Modules:

- Advanced Computer Vision; Advanced NLP; Self-supervised ...

Each 'module' includes :

- In-depth instruction, by practitioners
- 70%-100% funding via IMDA for SG/PR

Deep Learning Everywhere...

Want to explore doing courses in Thailand

- ... but need to understand who is interested

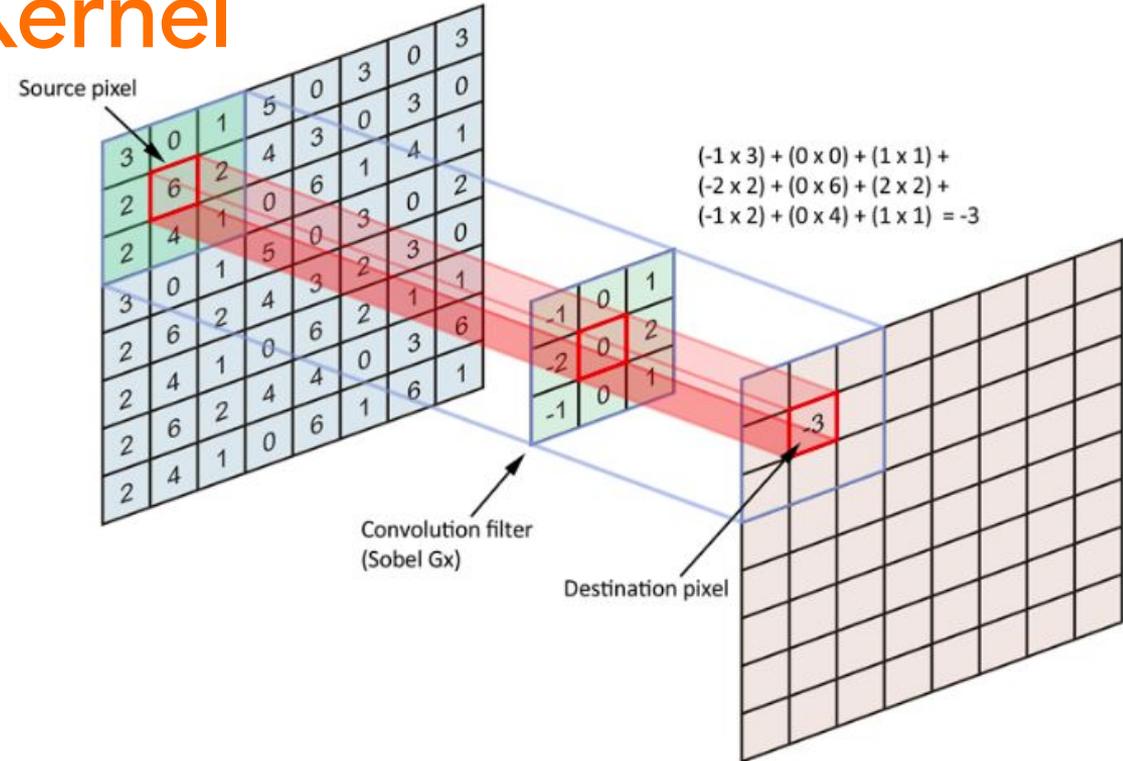
Action points :

- `http:// bit.ly / thai_courses`

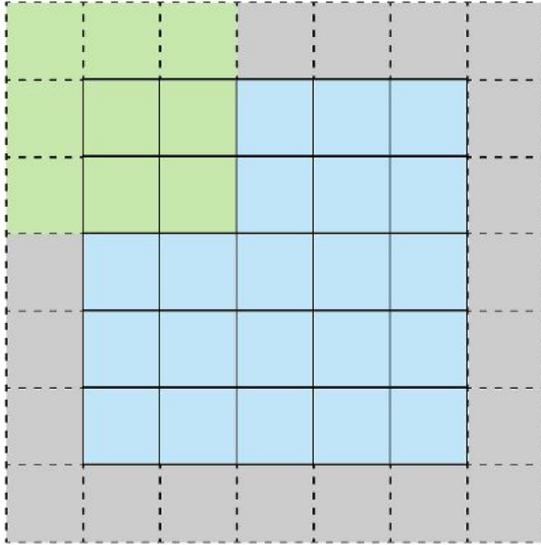
Questions?

Extras

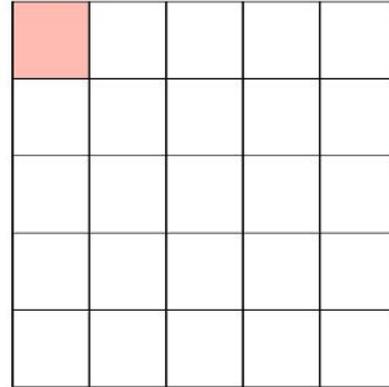
CNN 3x3 Kernel



CNN stride=1 (with padding)

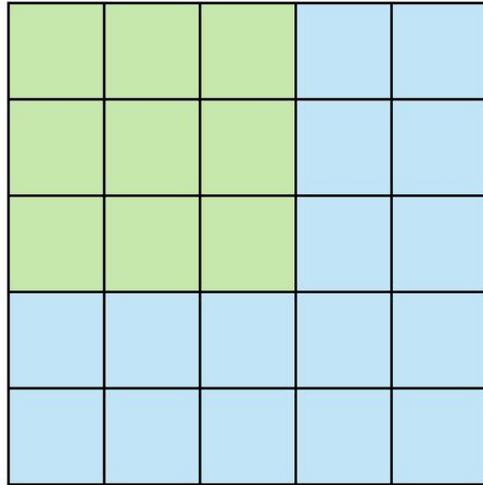


Stride 1 with Padding

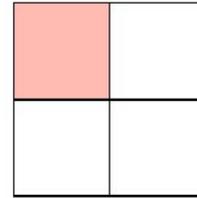


Feature Map

CNN stride=2 (no padding)



Stride 2



Feature Map