
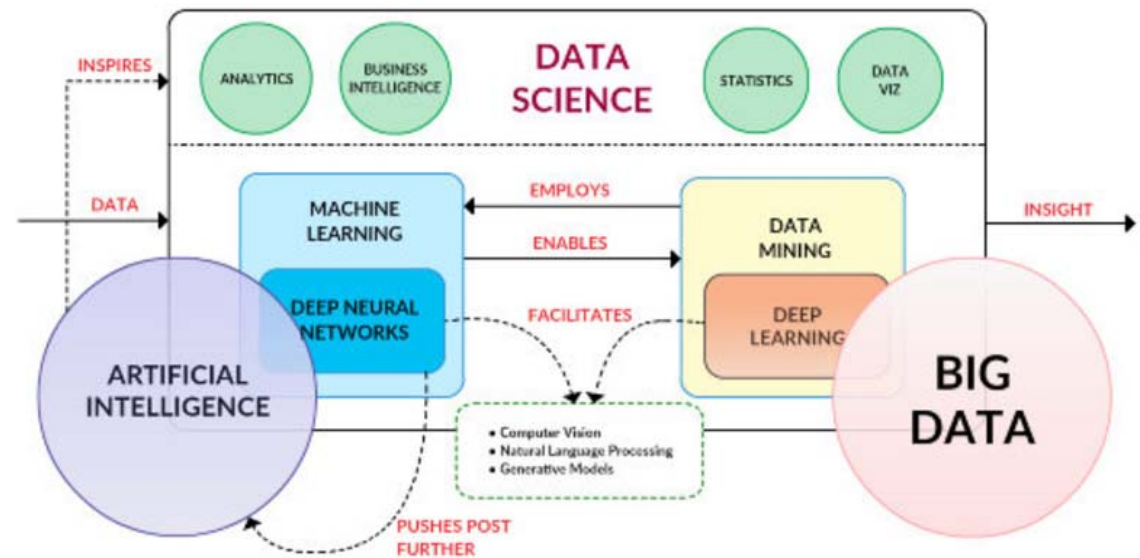



# Data Science @ Sioux LIME

A peek at a recent project...



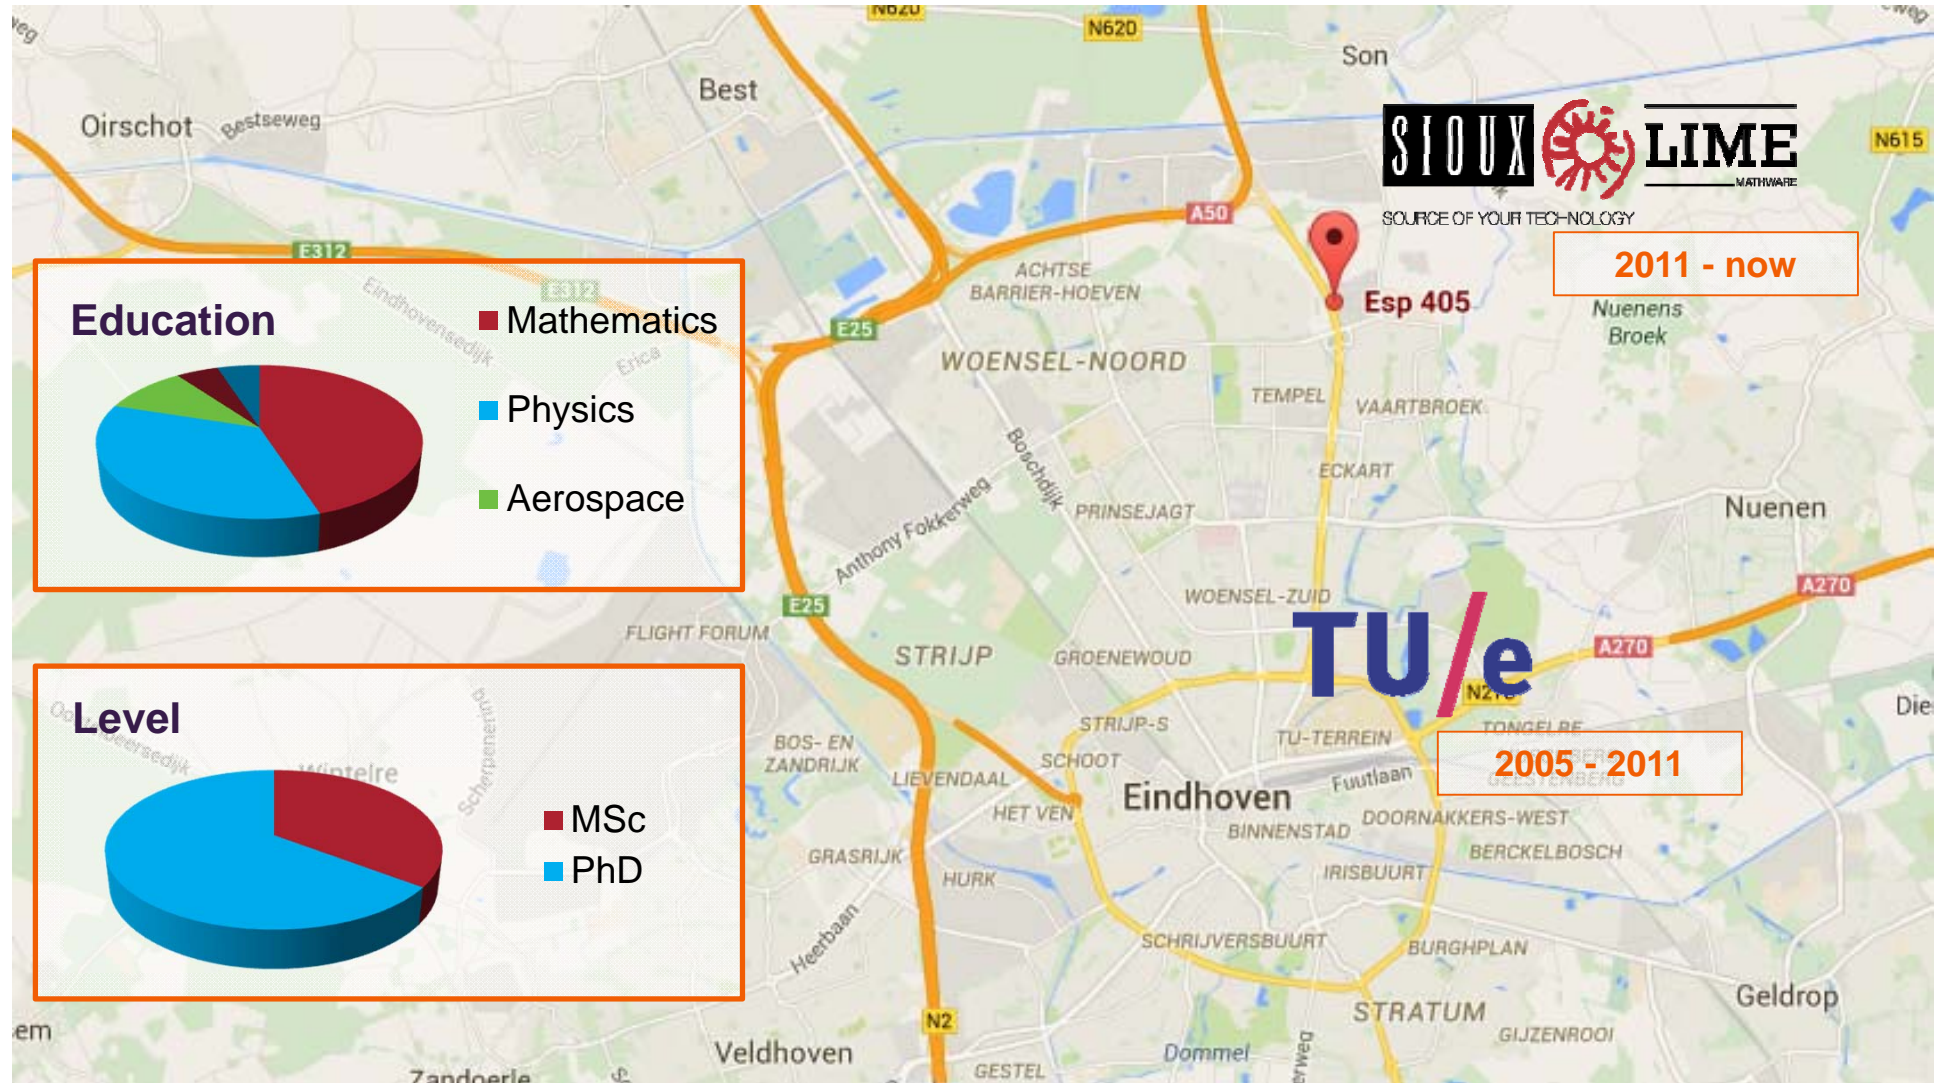
# Big Data, Data Science, Machine Learning...



Source: <http://www.kdnuggets.com/2016/03/data-science-puzzle-explained.html>



# Sioux LIME - Company profile

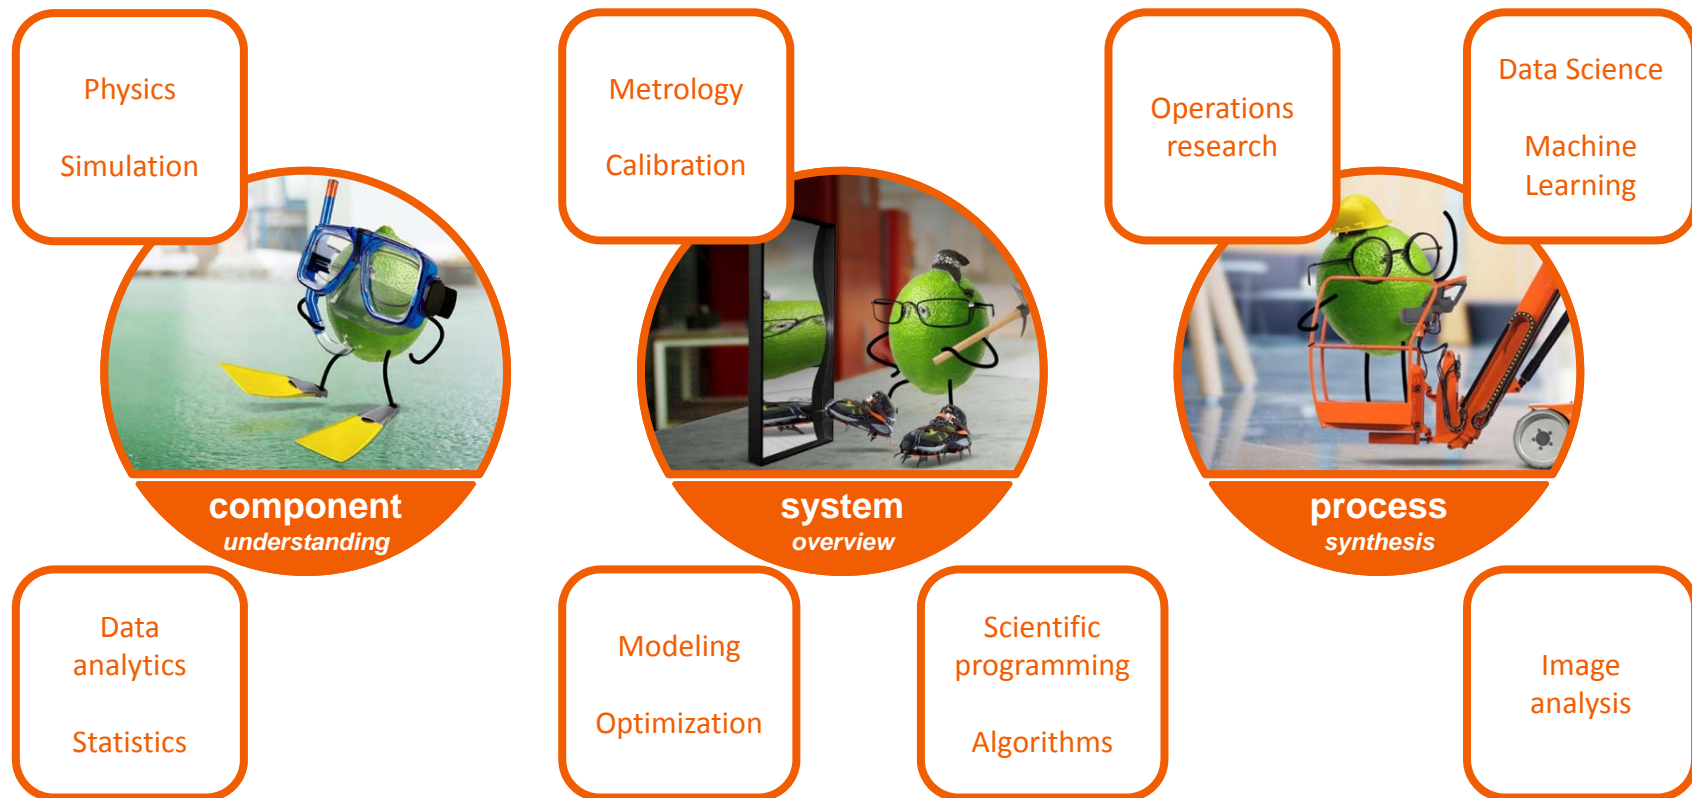


# Mathware



Realize practical solutions for industry through scientific knowledge and expertise.  
(Mathematical consulting)

# Competences



# Clients



component



system



process



# Presentation overview

- ✓ Short introduction of myself and Sioux LIME
- Case: Building a ‘virtual melon expert’ using Deep Learning
- ‘Technical walkthrough’:
  - How to practically tackle a case like this?!
  - Methods and techniques used.
  - Technology and tools used.
- Q&A



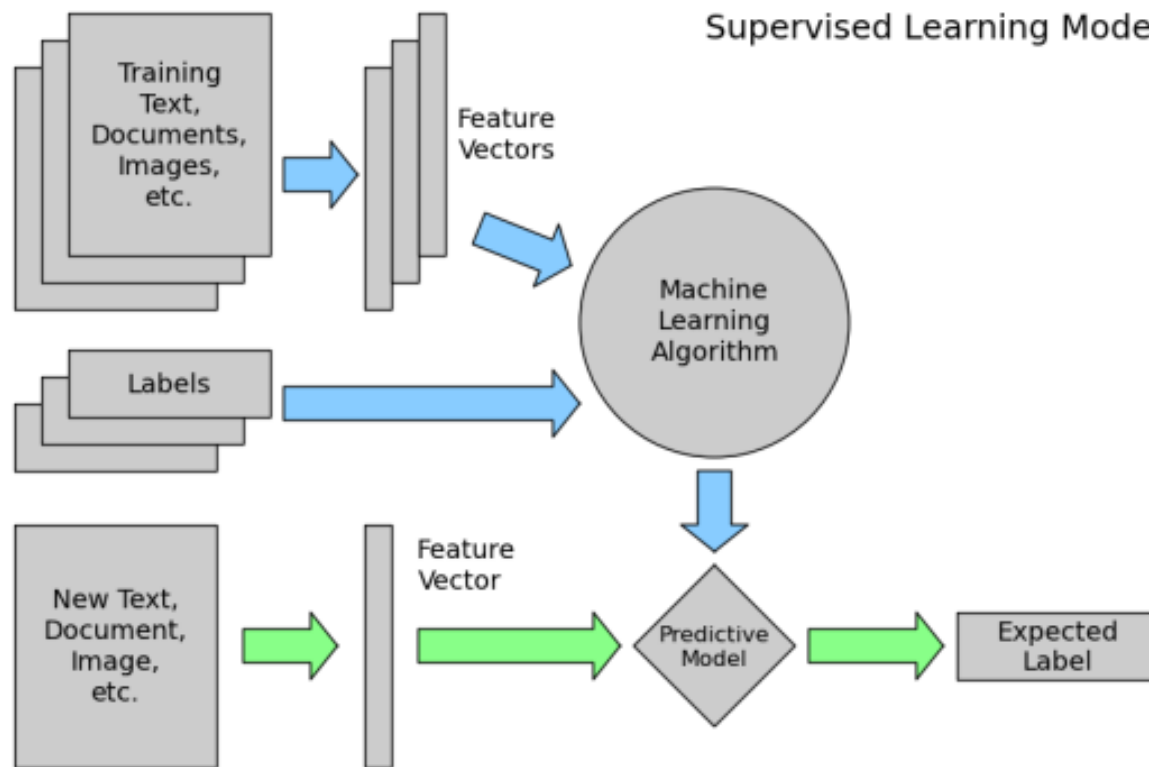
## Case: Build a 'melon expert'




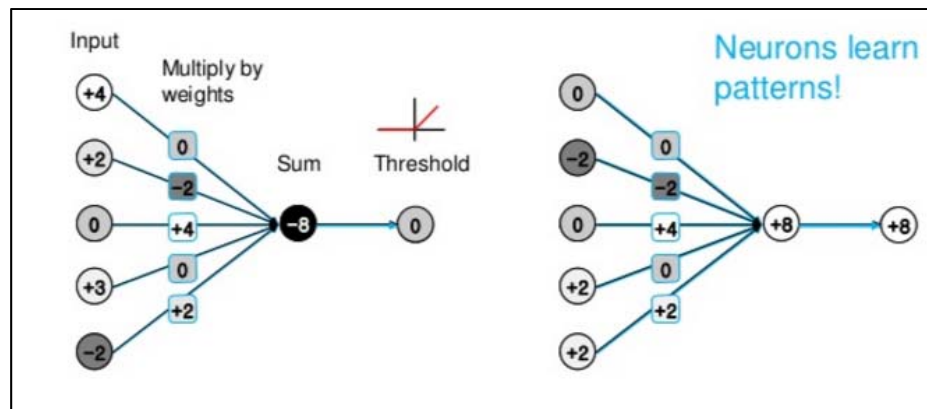


- Build a 'virtual melon expert' using Deep Learning technology.
- Using computer vision (only), assess melons on a set of criteria.
- Can we classify melon images according to their 'net structure'?
- Short project, proof of concept.




# Classification – Supervised Learning



# Neural Networks – key concepts



# Convolutional Neural Networks (CNNs)




# Deep Learning Software

Name	Interface(s)	Remarks
TensorFlow	Python, C++, Java	From Google
<b>Caffe</b>	Python, Matlab	Strong in computer vision
MXNet	Python, R, ...	Choice of Amazon
Microsoft Cognitive Toolkit	Python, C++	
Torch	C, Lua	Coming from FB
Theano	Python	Used a lot in academia
Keras	Python	High-level framework
DeepLearningForJ	Java, Scala	Well-documented
....		

- We have used Caffe in this case.




# What network (architecture) to use?!



\* Image taken from 'AN ANALYSIS OF DEEP NEURAL NETWORK MODELS FOR PRACTICAL APPLICATIONS' by Alfredo Canziani e.a.

- Take an existing network architecture. 'Do not be a hero!'
- Specific 'project' constraints/requirements usually give guidance.
- We have used AlexNet (and also tried SqueezeNet (not in picture)).

# AlexNet-architecture



- AlexNet was developed in 2012 to win the ImageNet challenge.
- It has 5 convolutional layers with 3 intermediate max pooling layers followed by 3 fully-connected layers.
- Number of parameters: around 60 million (!).

# Setting up Caffe for GPU-based learning

- How do we actually use Caffe to train and test a model?
- You can install Caffe on a laptop with Ubuntu, OS X or even Windows...
- Ideally, we would like to be able to do GPU-based training.
- Convenient alternative: “in the cloud” at e.g. Amazon Web Services (EC-2):

## P2 Features

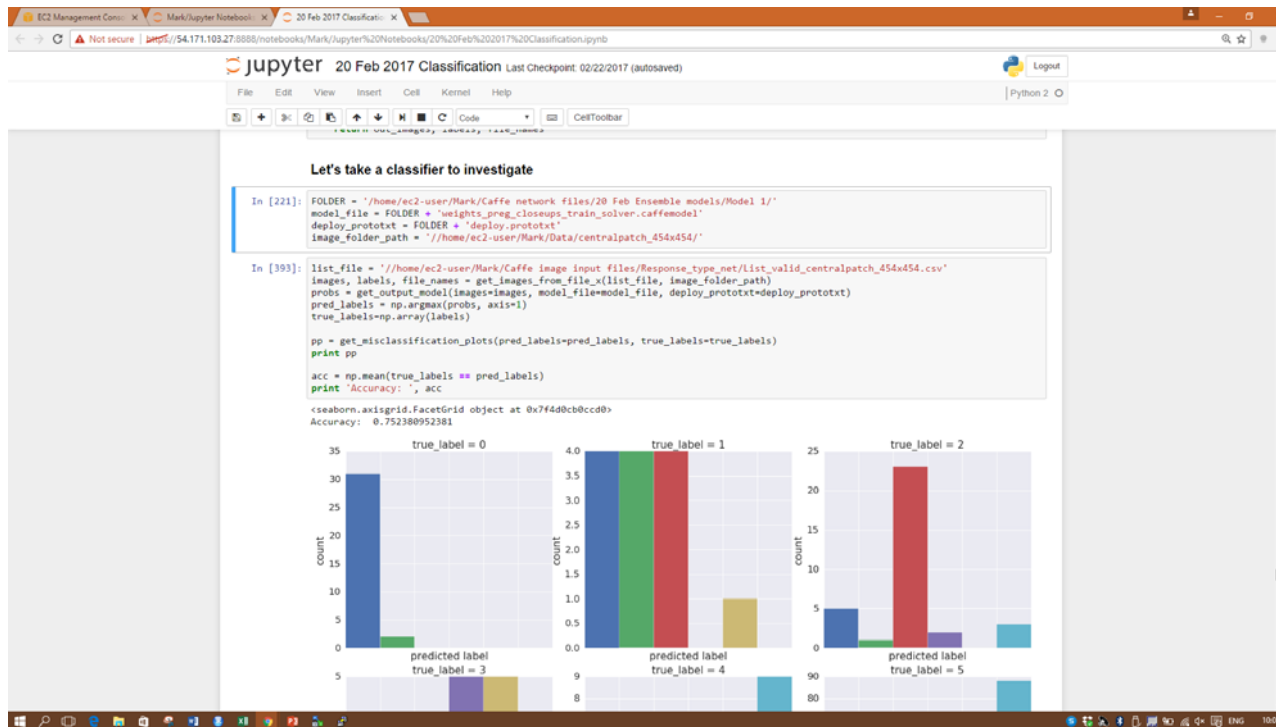


### Powerful Performance

P2 instances provide up to 16 NVIDIA K80 GPUs, 64 vCPUs and 732 GiB of host memory, with a combined 192 GB of GPU memory, 40 thousand parallel processing cores, 70 teraflops of single precision floating point performance, and over 23 teraflops of double precision floating point performance. P2 instances also offer GPUDirect™ (peer-to-peer GPU communication) capabilities for up to 16 GPUs, so that multiple GPUs can work together within a single host.

- A p2.xlarge machine (single GPU) costs around 1 USD/h.
- Machine images with Caffe installed are available, so you can start straight away!

# Running Caffe on AWS via a Jupyter Notebook

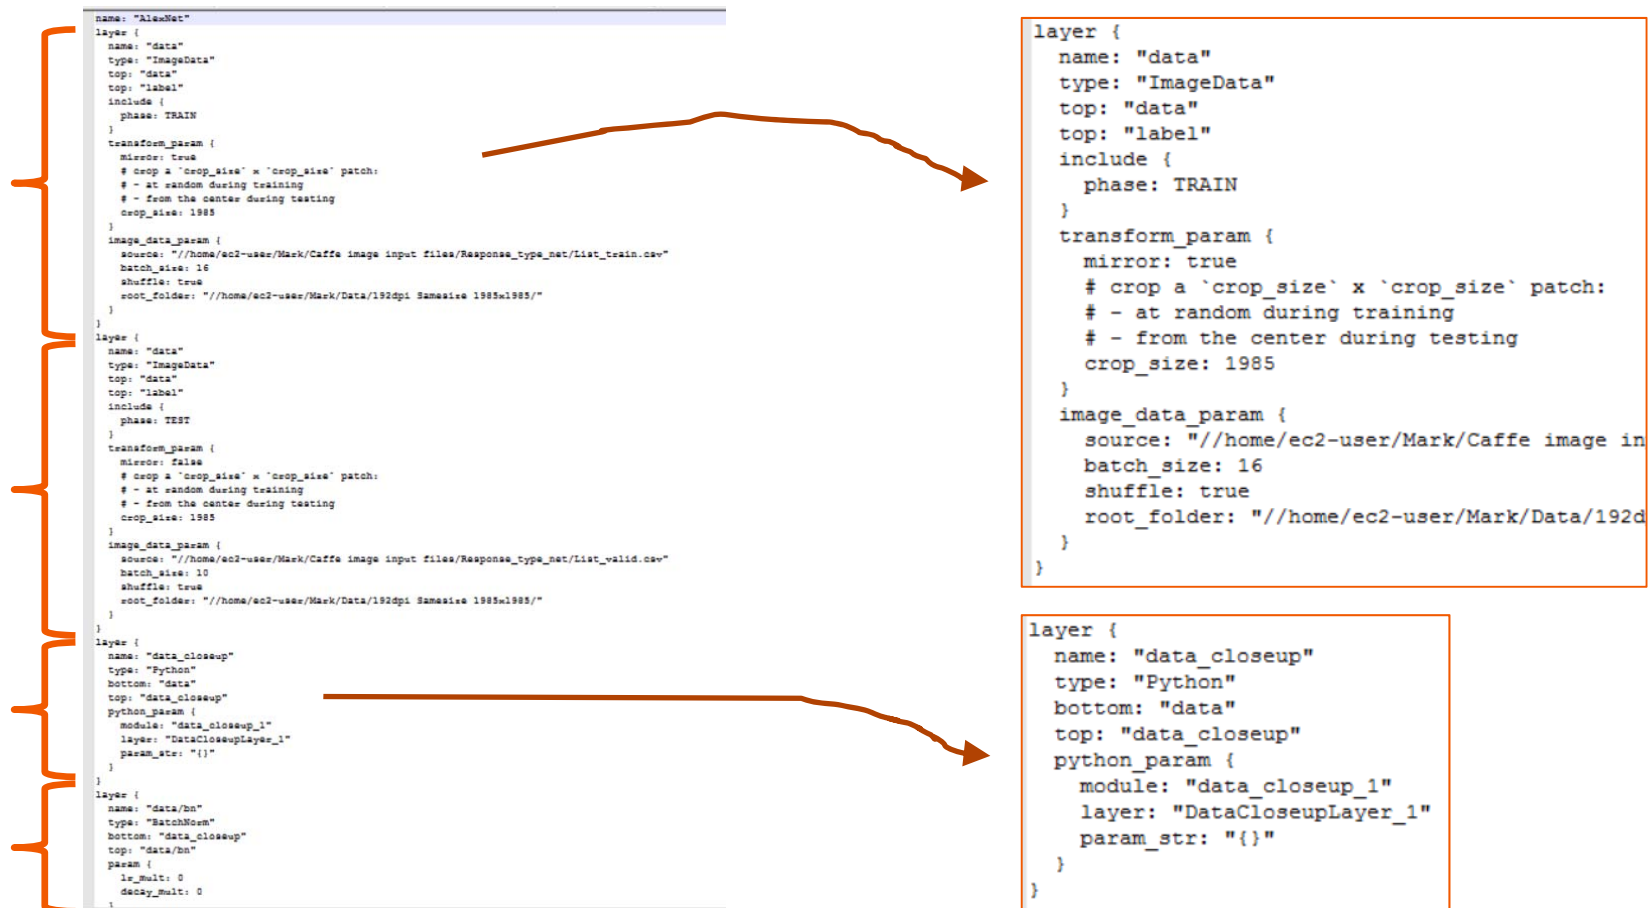


- A convenient way of working with Caffe (PyCaffe) is using a Jupyter Notebook.
  - Start a Jupyter Notebook server from your instance at AWS
  - Connect to this server through your local browser.
  - Open a notebook (located on your AWS instance).
  - Start coding and running notebook cells from your browser.



# Defining a network architecture in Caffe

- In Caffe, a network architecture is defined in a '.prototxt'-file:

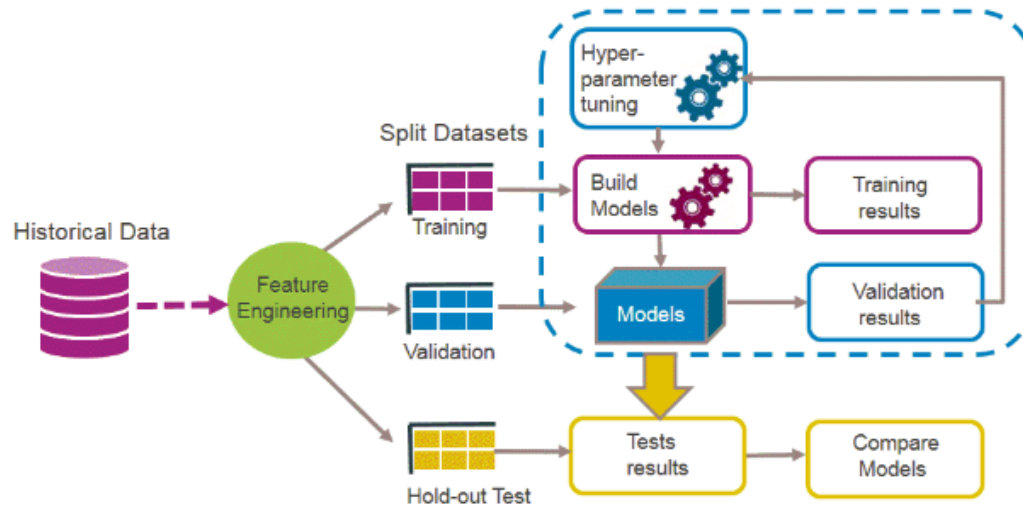


# Setting the solver parameters

```
1 # test_iter specifies how many forward passes the test should carry out.
2 # After every test_interval training iterations, test_iter * batch_size images are fetched for testing.
3 test_iter: 1
4 # Carry out testing every 1 training iterations.
5 test_interval: 1
6 base_lr: 0.01
7 display: 1
8 max_iter: 320000
9 lr_policy: "poly"
10 power: 1
11 momentum: 0.1
12 weight_decay: 0.5
13 snapshot: 5000
14 snapshot_prefix: "alexnet_cvgj"
15 random_seed: 0
16 net: "/home/ec2-user/Mark/Caffe network files/26 januari 2017 Learning rate 1/train.prototxt"
17 test_initialization: true
18 iter_size: 2
```

- We need to tell the solver a few things:
  - Learning rate (scheme), momentum.
  - The network we want to train/solve/optimize.
  - Regularization penalty (weight\_decay).
  - ....

# Preparing our data...




- Generated a train, validation and test set: 60%-20%-20%.

- Resize images to 227x227: AlexNet input dimensions.
- Actually, we have taken central 454x454 patches and downsized these.
- What about data augmentation? → I come back to that quickly.

# How to avoid overfitting?

## What to do when having limited data?

- In general, Supervised Learning is about the 'bias-variance' trade off.



- Underfitting is typically not a problem when using CNNs...but overfitting is...
- We will now look at some methods to avoid overfitting and improve the generalization capabilities of a model.



# Transfer Learning


- In our case (and many other cases), data is limited.
- How to optimize 60 million parameters, based on ca. only 1000 images??
- Transfer Learning offers a potential solution for this problem.
- **Transfer learning:** take a 'pre-trained' model and use this.
- We have taken a pre-trained variant of the AlexNet-architecture.
- See e.g. the Caffe Model Zoo for many (pre-trained) networks.
- How to use a pre-trained model?
  - As a feature extractor (either from top or middle activations).
  - Retrain / fine-tune the weights of the final layer(s).

# Data Augmentation ('on the fly')

- Augmenting the data (images) is a means to avoid overfitting.
- Preferably, this data augmentation is done on-the-fly (and not off-line).
- This means that an input batch during training is augmented/transformed in real-time to a batch containing modified images.
- Possible ways to augment your data, based on original images:
  - Rotation: rotate an image around a random angle.
  - Translation: translate an image randomly in x and y direction.
  - Rescaling: apply a random scaling factor to the image.
  - Contrast/brightness: adjust contrast/brightness with a random amount.
  - Random crops: take a random 'patch' of the (larger) image.
- We implemented this in Caffe using a 'custom Python Layer'.



# Other techniques to improve generalization and/or prevent overfitting

- Drop-out / Drop-connect:
  - Randomly remove neurons/connections from network while training.
- Weights regularization
  - ‘Penalize parameters for being non-zero’.
- Early stopping
  - E.g. via monitoring a validation error.
- Collect more data...
  - So-called ‘learning curve’ gives a clue...



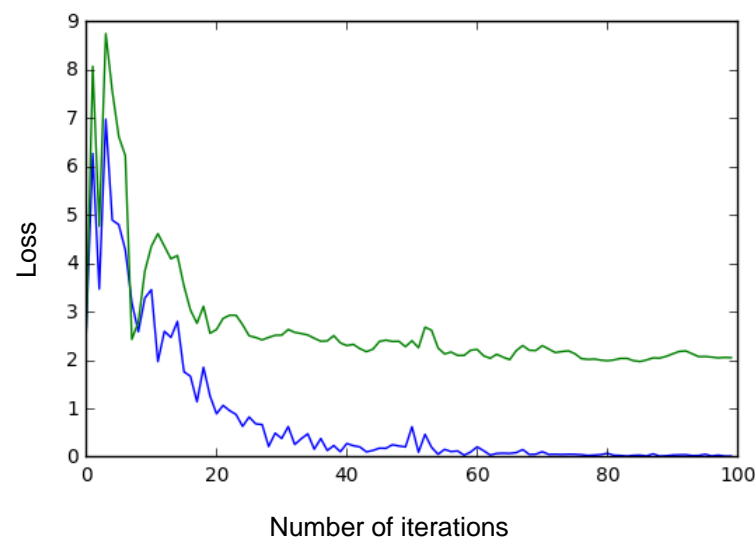

$$\lambda \sum_{j=1}^n \theta_j^2$$

Regularization



# Training our classifier

- Network architecture chosen; modified the last layer to have only 6 outputs.
- Configured the solver file with training (hyper)parameters.
- Loaded the pre-trained model weights / parameters.
- Prepared the data for input: central patch + resize.
- Added (“on the fly”) data augmentation.
- Develop some (basic) code to train a classifier and monitor (validation set) performance.
- Start training and ‘tune’ the hyper-parameters.





# Questions?!

