Deep Learning Explained

Module 1: Introduction and Overview

Sayan D. Pathak, Ph.D., Principal ML Scientist, Microsoft
Roland Fernandez, Senior Researcher, Microsoft

Course outline

What is deep learning?

Who are the audience?

What to expect from this course?

What to expect

Deep learning basics (6 modules):

- Module 1: Introduction and high-level modeling workflow
- Module 2: Logistic regression for optical character recognition
- Module 3: Multi-layer perceptron
- Module 4: Convolution networks
- Module 5: Recurrence for time-series modeling
- Module 6: Long-short term memory (LSTM) recurrence for text modeling

Assignments:

- 5 hands-on lab and 5 quizzes (required for certification)
- Use Python Jupyter notebooks

Upon completion, students will:

- Have a working knowledge of deep learning concepts and algorithms
- Be able to build deep neural network models
- Solve practical problems in AI involving large data

Who are the audience?

Data Scientists and Engineers who are new to deep learning

Technical Managers who are interested and involved in development of AI based technologies

Machine learning scientists looking to build deep models with agility and ease

Who are the instructors?

Sayan Pathak, PhD.

Industry

- Principal ML Scientist @ Microsoft
- Principal Investigator, National Institutes of Health funded projects in Healthcare
- Allen Institute for Brain Sciences
- Practical ML experience in Computer Vision, On-line Advertisement, Social Networks, Neuro and healthcare informatics

Academic

- Instructor at AI School @ Microsoft
- Affl Faculty @ University of Washington in Bioengineering / EE (from 2001)
- Affl Professor @ IIT (Indian Institute of Technology), Kharagpur in CS (from 2012)
- Courses taught:
  - Image Computing Systems, Information Retrieval, Social Computing, Machine Learning

Who are the instructors?

Roland Fernandez

- Researcher and AI School instructor @ Microsoft

- Research Areas:
  - reinforcement learning, autonomous multitask learning, symbolic representation,
  - information visualization, and Computer Human Interaction

- Worked in the areas of Natural User Interfaces, activity based computing, advanced prototyping, programmer tools, operating systems, and databases.

- Actively involved in AI education efforts within Microsoft

Programming environment

Toolkit:

- Microsoft Cognitive Toolkit (CNTK) Tutorials with Python Notebooks

Environments:

- Install locally (recommended)
- Use pre-installed notebooks in Azure Notebooks (CPU only) for free
- Azure Data Science Virtual Machines (sign up trial account)

References

Deep Learning by I. Goodfellow, Y. Bengio & A. Courville, MIT Press, http://www.deeplearningbook.org, 2016

Stanford cs231n: http://cs231n.github.io/

Neural Networks Tutorials by G. Hinton
https://www.cs.toronto.edu/~hinton/nntut.html

Acknowledgement

Chris Basoglu, Ph.D., & Cha Zhang, Ph.D., @ Microsoft

for several support and suggested improvement.

Surya Pathak, Ph.D. Assoc. Prof. @ University of Washington

for helping with pedagogical approach to teaching and content updates.

Jonathan Sanito, Content Developer @ Microsoft,

for assistance with graphics especially for the convolution networks.

Nikos Karampatziakis, Ph.D.,

for help with course material preparation (especially the Text Classification tutorial).

Machine Learning (ML)

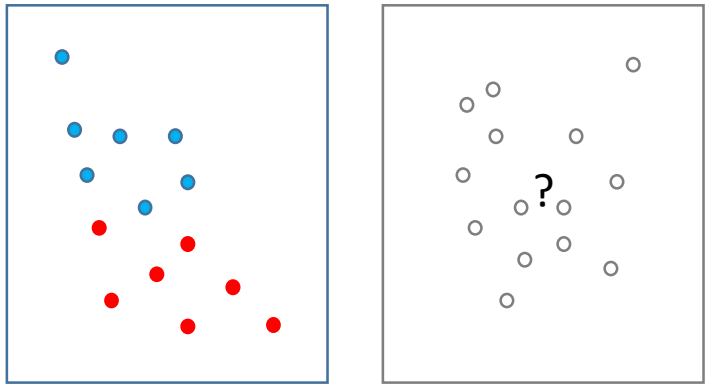

Supervised Learning

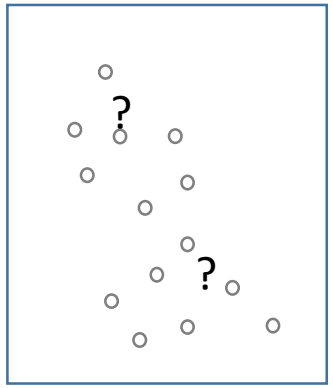

Unsupervised Learning

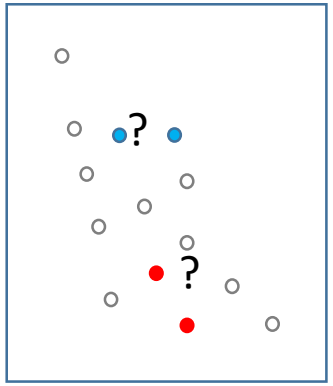

Semi - supervised Learning

Supervised Machine Learning

Input → Model → Predicted

Classification

Data: E-mail
Label: Spam / not Spam

SPAM
vs.
Not SPAM

Categories: Spam / Not Spam

Regression

Data: House features
Label: Price

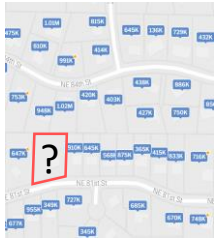

850k

House Price: a real number

Real-world applications

Image

- Autonomous Driving
- Disease detection

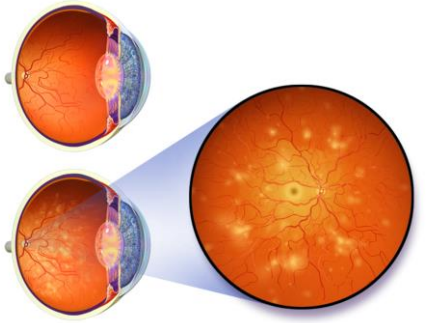

Text

- Machine translation
- Document comprehension

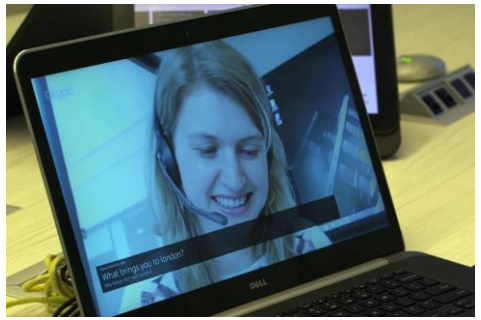

Speech

- Voice recognition
- Speech to text

https://en.wikipedia.org/wiki/Autonomous_car

What is deep learning?

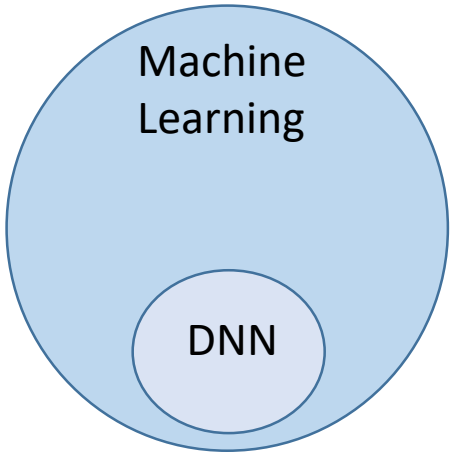

Deep learning = Deep Neural Networks (DNN)

- Mimics several layers in the brain

Deep Neural Networks

- Have multiple layers
- Each layer learns a higher abstraction on the input from the layer before it
- Requires fitting a large number of parameters (100+ Millions)
- Facilitated by (1) large amount of data and (2) computing capabilities

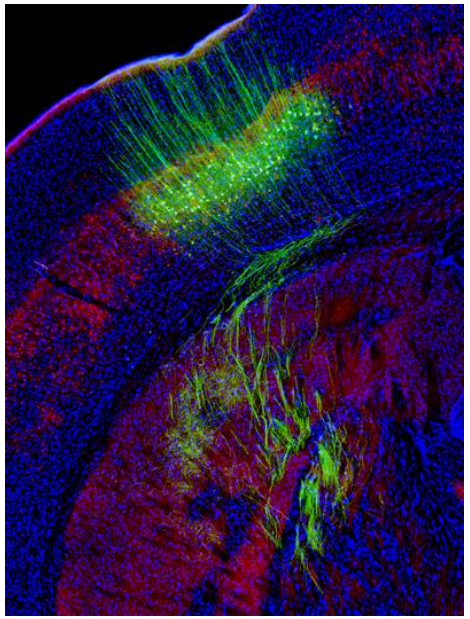

Application domains

- Image / Videos
- Speech
- Text
- Multimodality and IOT data

https://www.nimh.nih.gov/news/science-news/2013/transgenic-mice-lines-aid-in-brain-circuit-imaging.shtml

ML recap

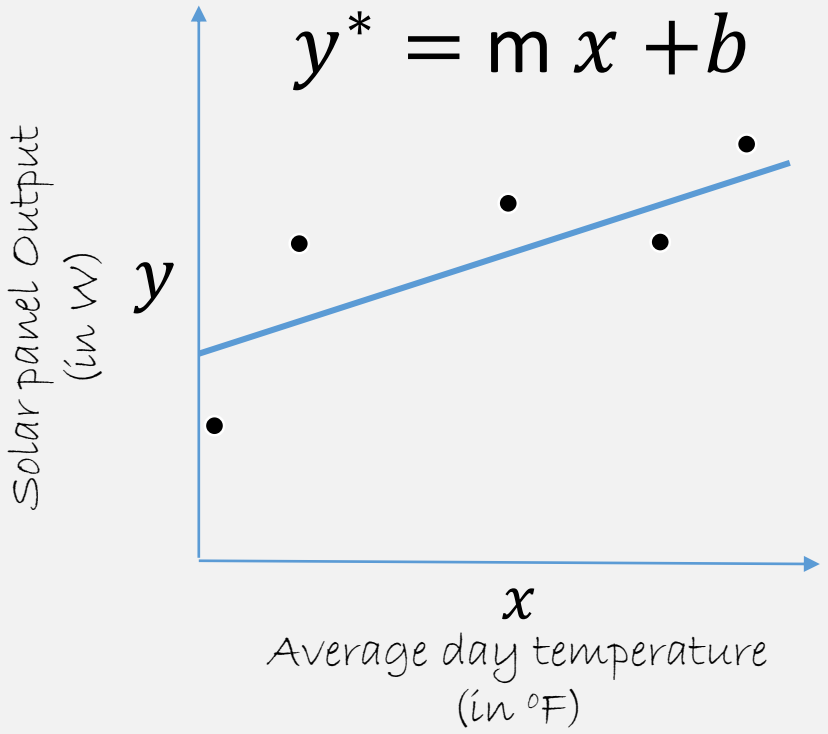

Input

x = Feature

Output

y = observed output (labels)

y^* = predicted output

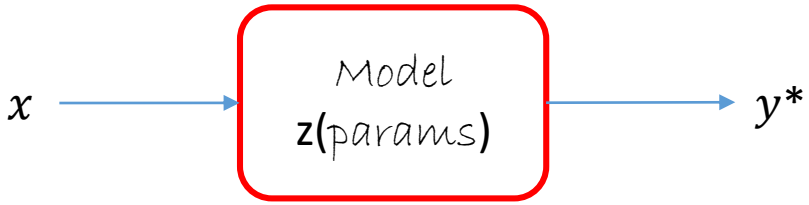

Model function (z)

m: Slope

b: Intercept

are model parameters

ML Recap

Input

Model

Predicted

Classification

Data: E-mail
Label: Spam / not Spam

SPAM
vs.
Not SPAM

Categories: Spam / Not Spam

Regression

Data: House features
Label: Price

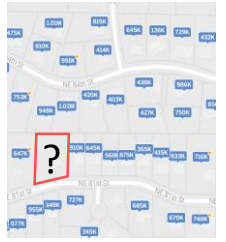

850k

House Price: a real number

Train Workflow

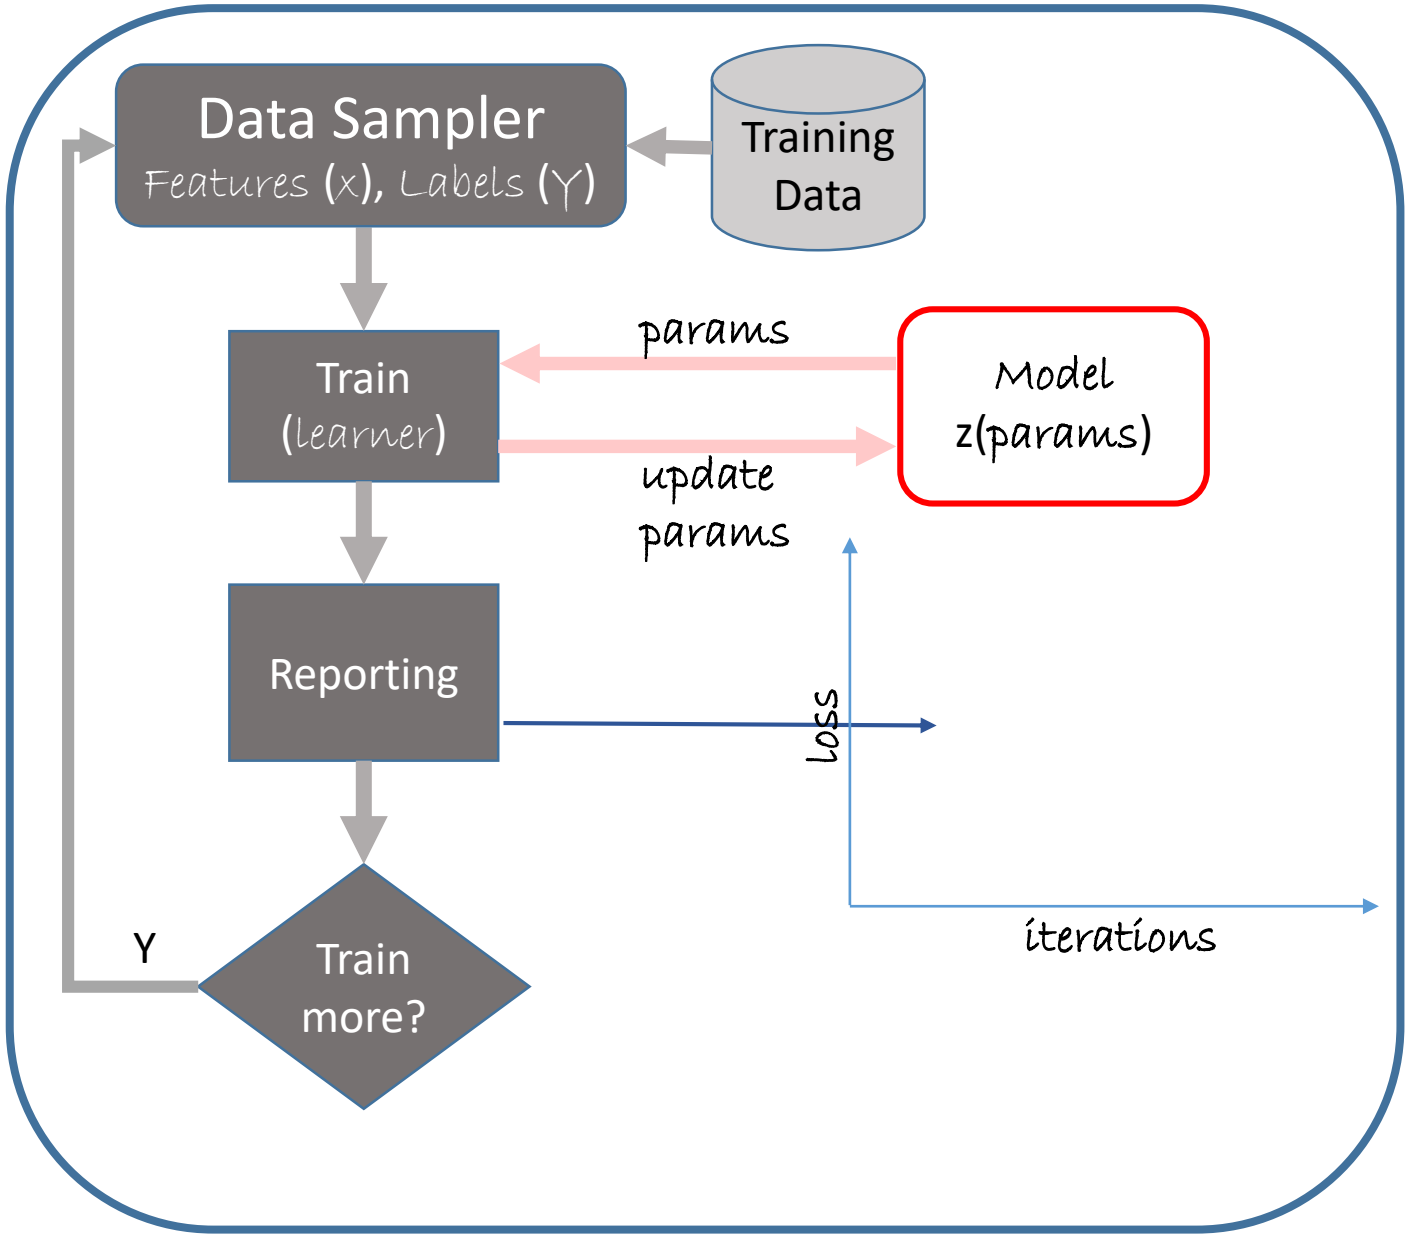

Validation Workflow

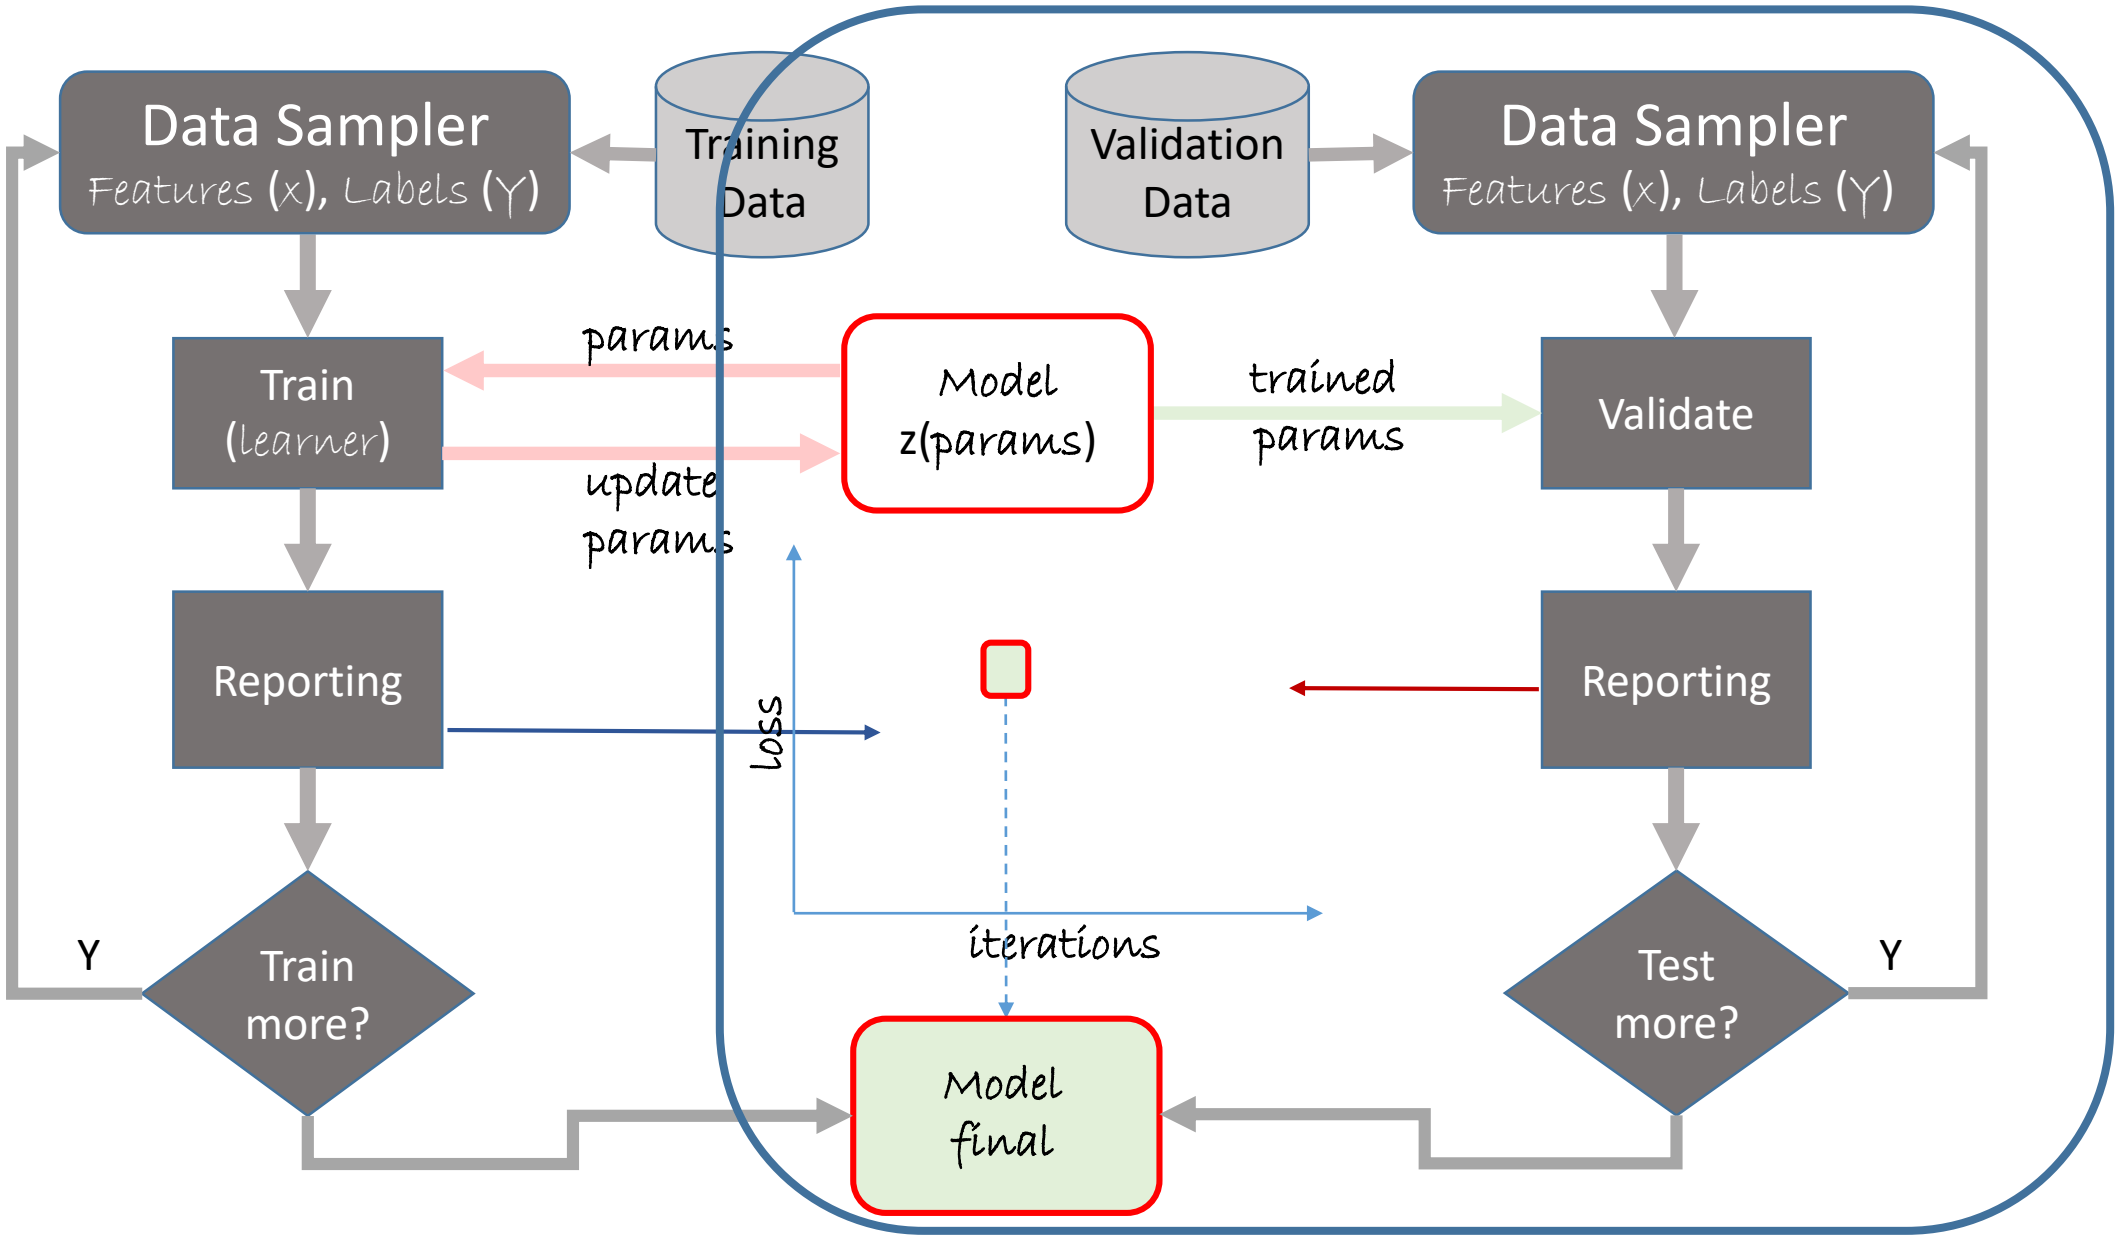

Test Workflow

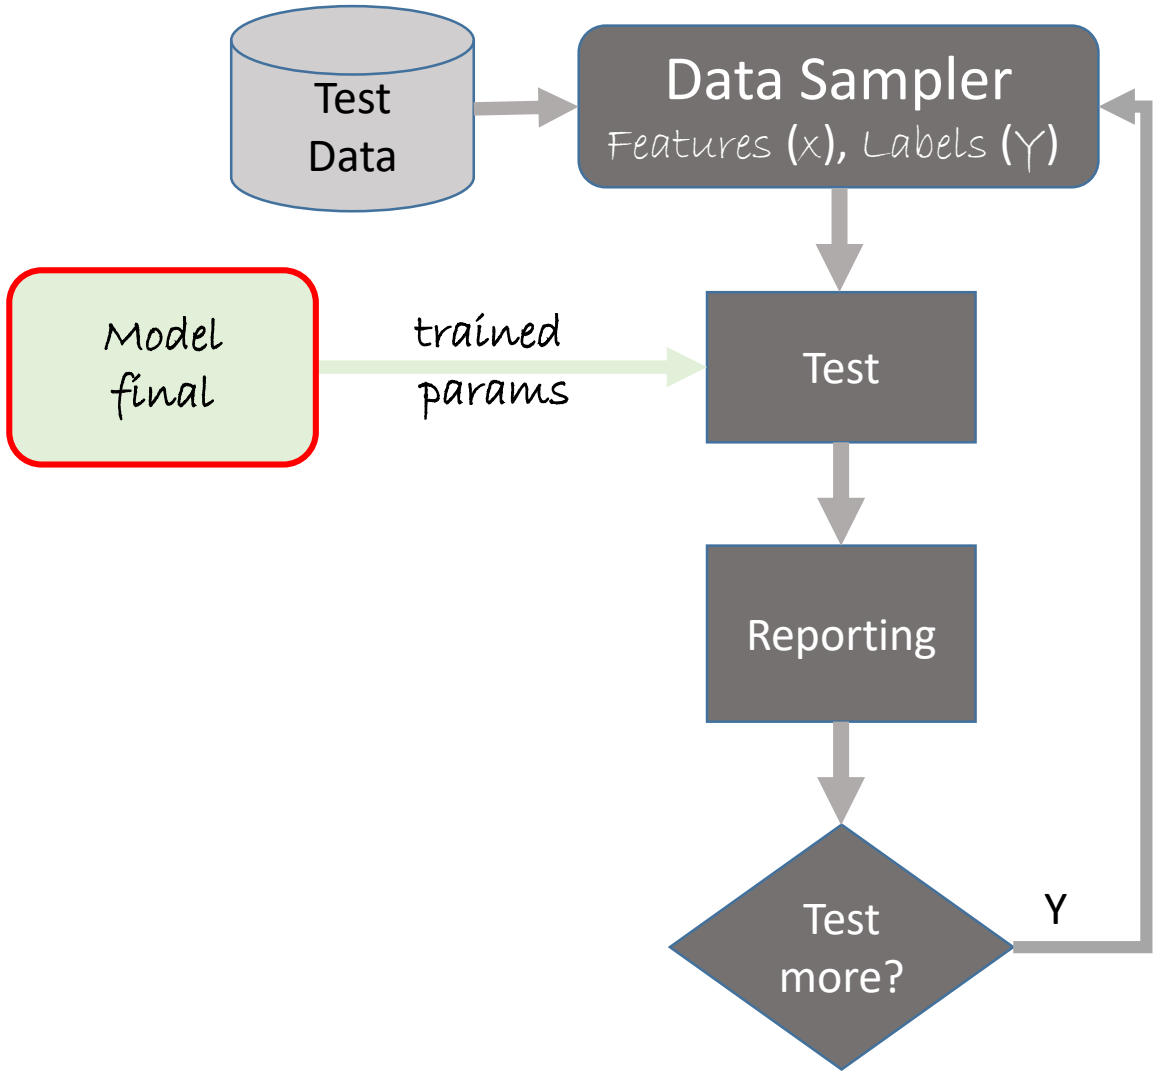

Where to begin with Cognitive Toolkit ?

On GitHub: https://github.com/Microsoft/CNTK/wiki

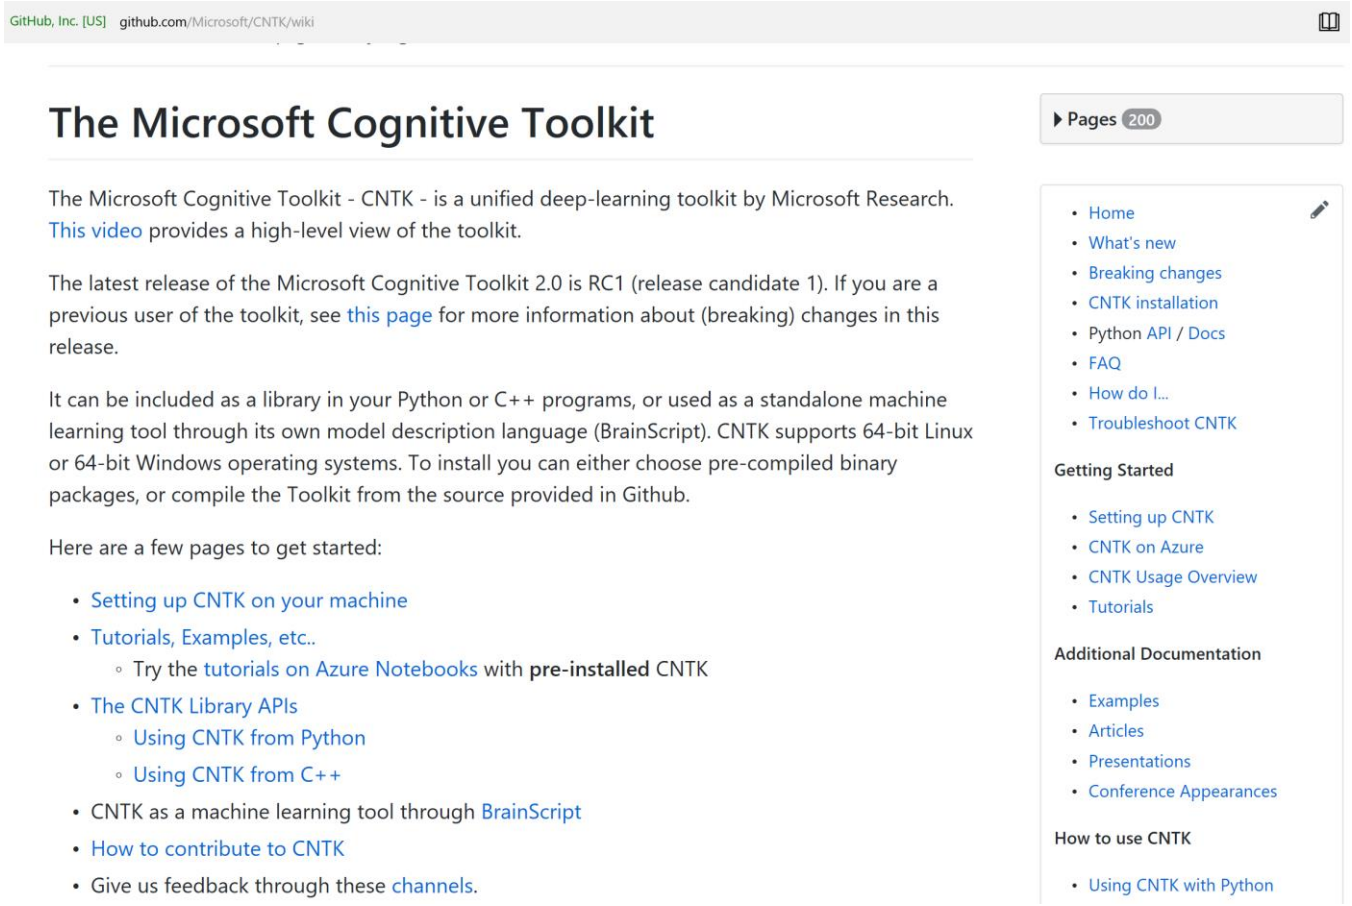

Seek help on Stack Overflow:
http://stackoverflow.com/search?q=cntk (please add *cntk* tag)