EECS 510: SOCIAL MEDIA MINING
SPRING 2015

Introduction and Social Application of Deep Learning

Rosanne Liu
rosanne.liu@northwestern.edu

NORTHWESTERN
UNIVERSITY

Outline

- Deep Learning: What and Why
- Social Applications
- Theories

NORTHWESTERN
UNIVERSITY

Deep Learning: What is it?

NORTHWESTERN
UNIVERSITY

What is Deep Learning

- First you need to know ML + AI
  - the study of how computer systems emulate human intelligence
- Deep Learning: the most cutting edge ML/AI research
- "A method which makes predictions by using **a sequence of non-linear processing stages**. The resulting intermediate representations can be interpreted as feature hierarchies and the whole system is jointly learned from data."

-- Facebook Research

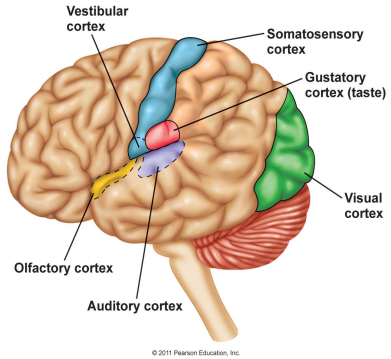

NORTHWESTERN UNIVERSITY

The Space of Machine Learning Methods

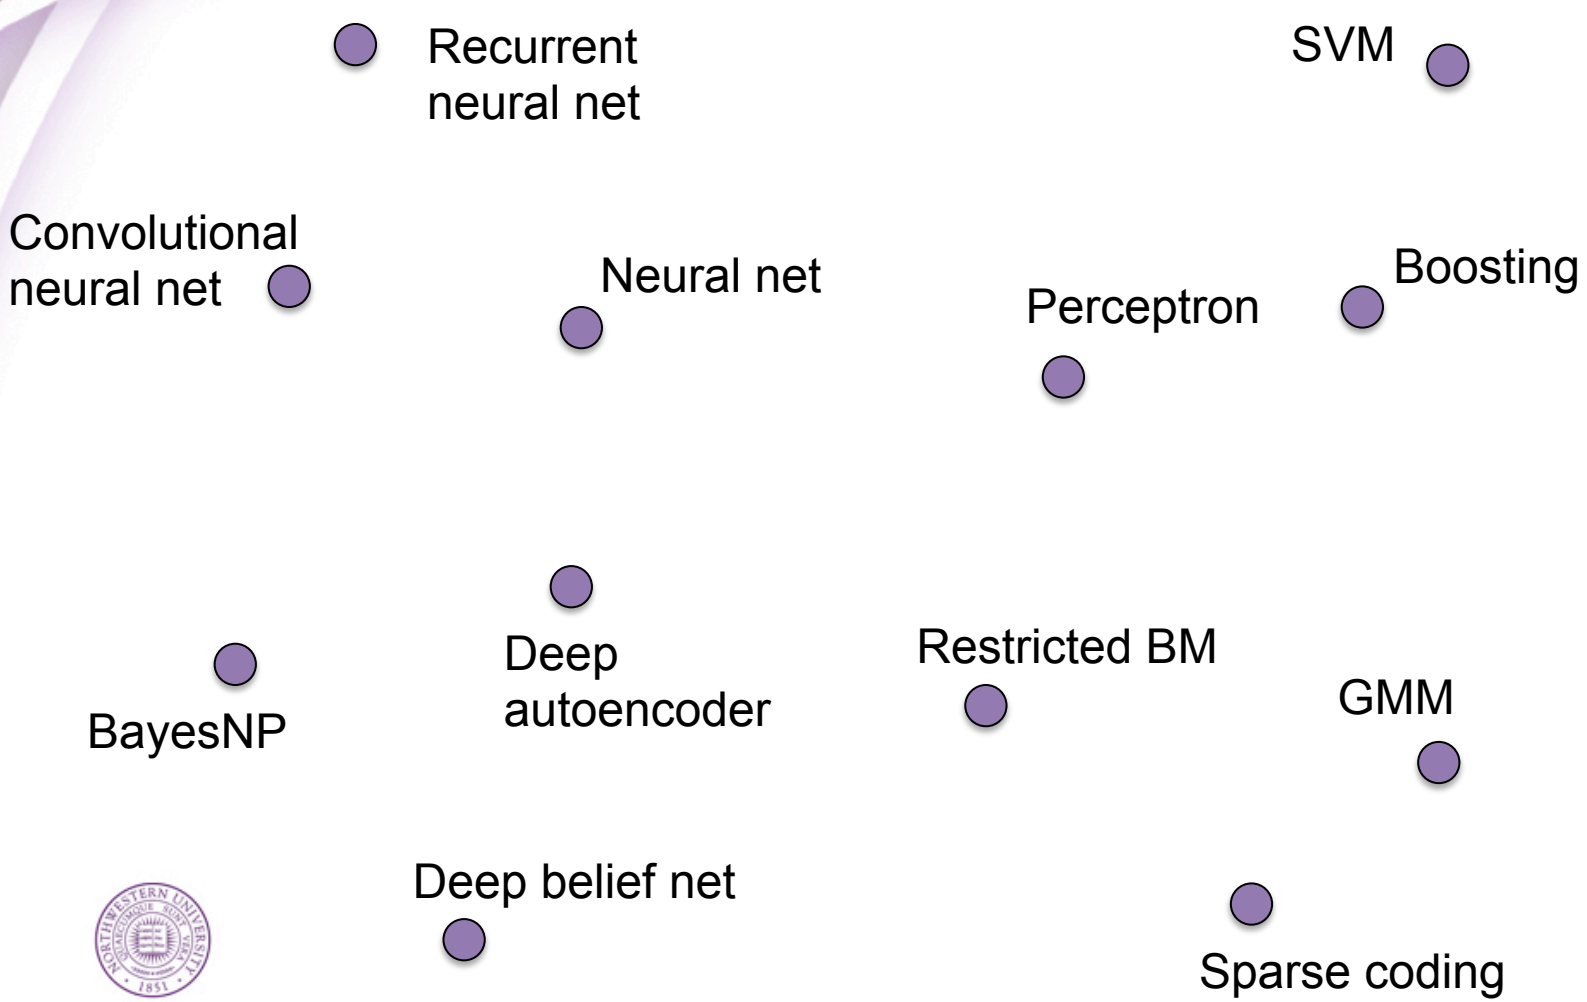

The Space of Machine Learning Methods

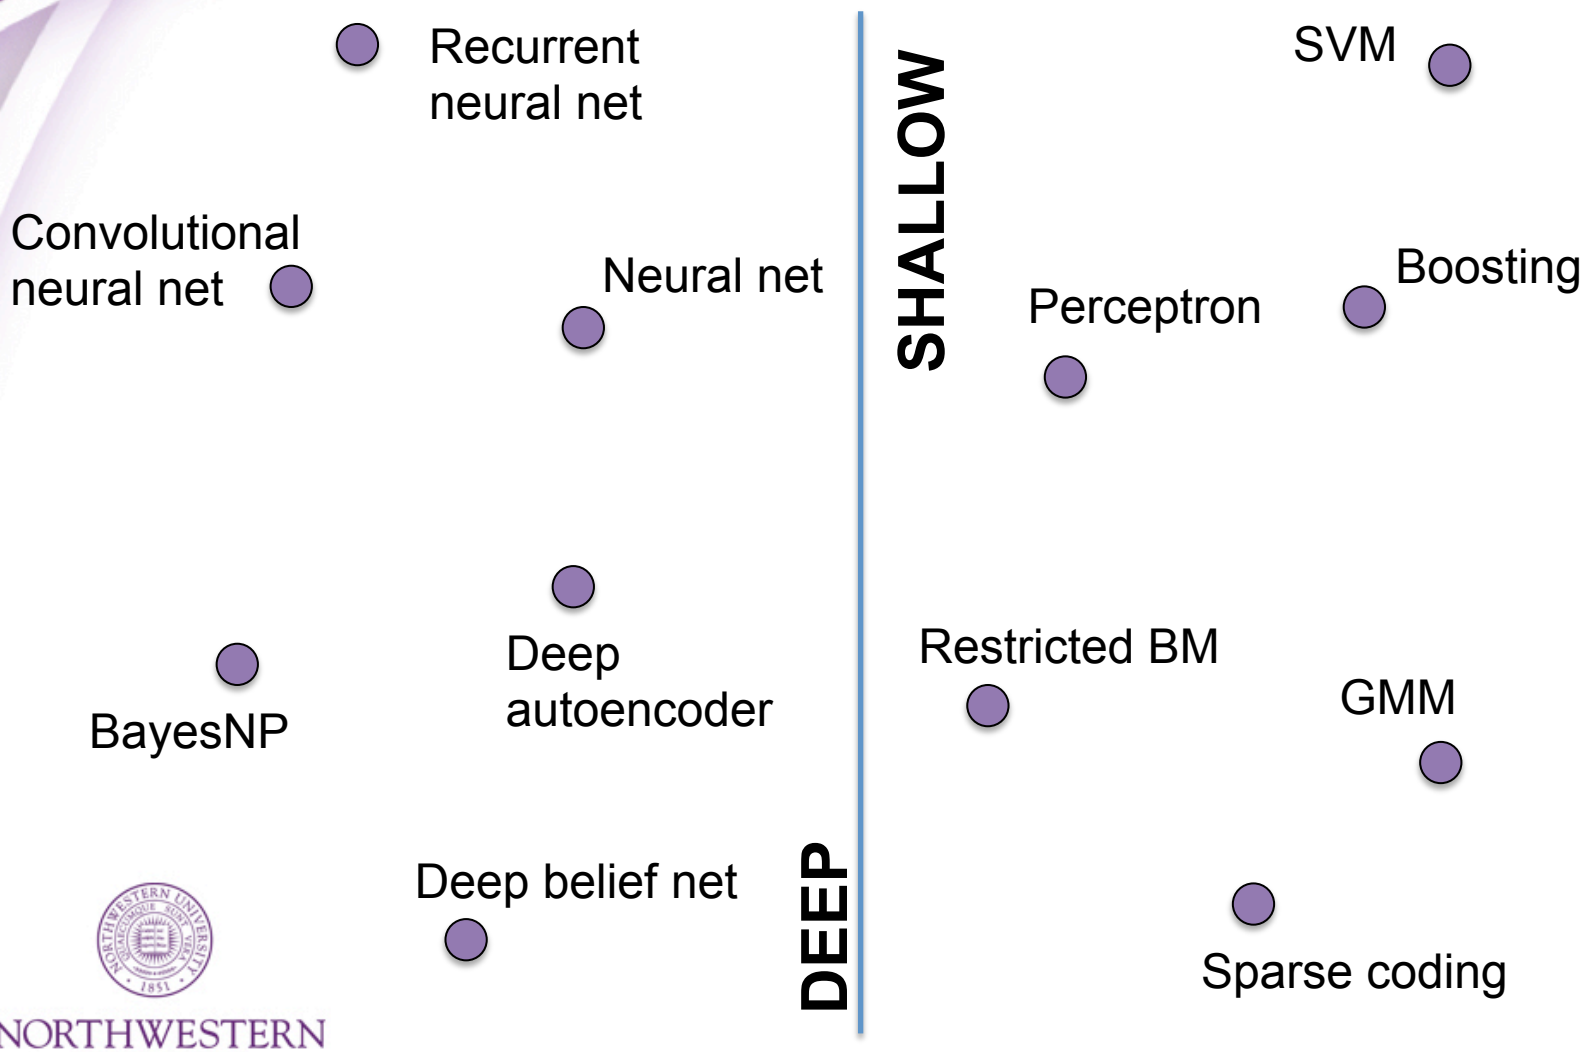

NORTHWESTERN UNIVERSITY

The Space of Machine Learning Methods

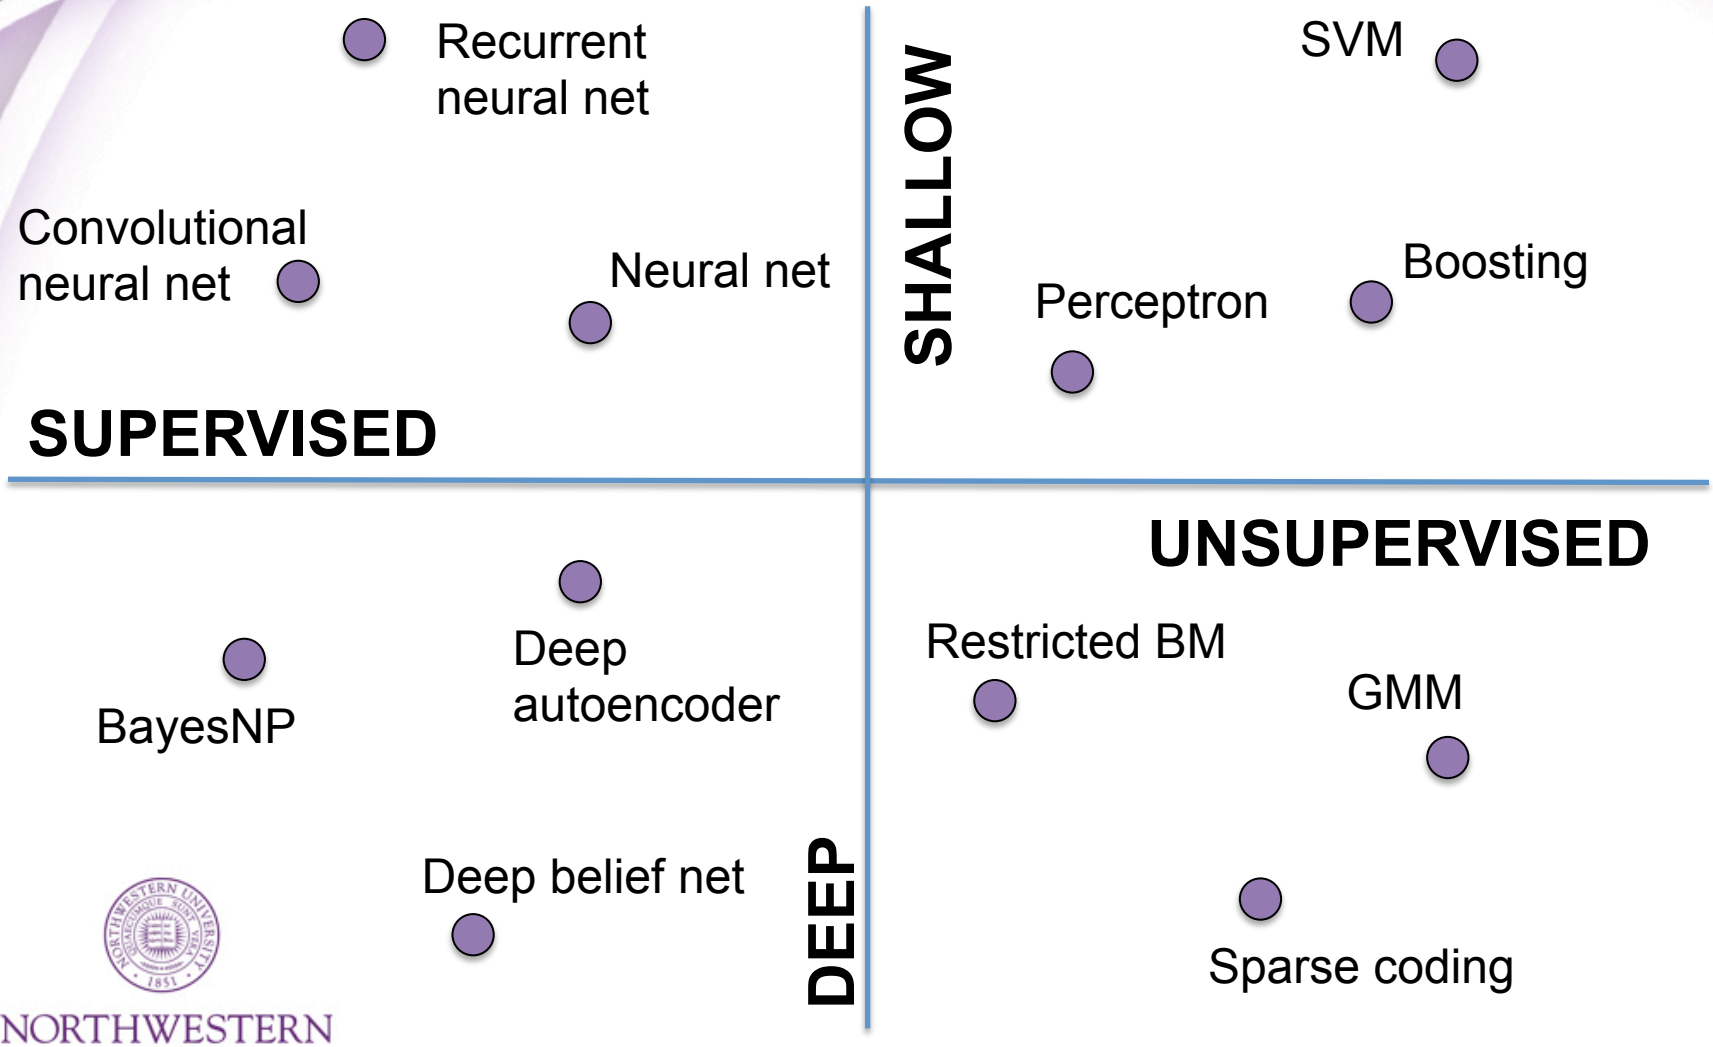

NORTHWESTERN UNIVERSITY

What is Deep Learning

- First you need to know ML + AI
- Second you need to know neural nets
  - Supervised, non-linear function, back-propagation

A "neuron"

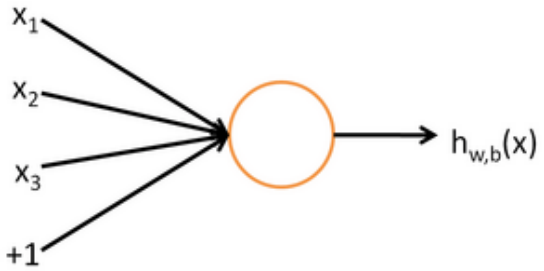

A small neural net

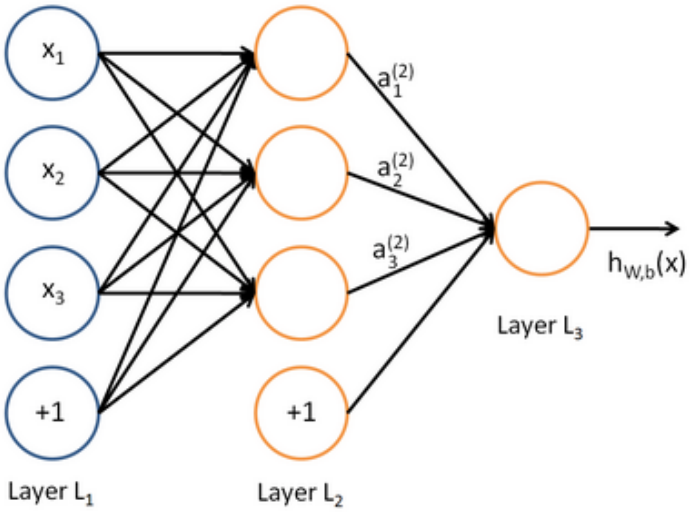

Training of a NN

Loop until tired:

1. **Sample** a batch of data.
2. **Forward** it through the network to get predictions.
3. **Backprop** the errors.
4. **Update** the weights.

NORTHWESTERN UNIVERSITY

Deep Learning: Why is it (so popular)?

(a game changer)?

NORTHWESTERN
UNIVERSITY

First You Need to Know that It Is Indeed Popular

Facebook AI Research

Overview Events Blog Publications

What's New

BLOG POST

FAIR open sources deep-learning for Torch

PUBLICATION

Question Answering with Subgraph Embeddings

About Facebook AI Researc

We're committed to advancing the field of machine intelli
people better ways to communicate. In the long term, we
intelligent machines. How will we accomplish all this? By

Research at the lab covers the full spectrum of topics relat

Nov 21, 2014 ASIA

Baidu's Andrew Ng on Deep L
and Innovation in Silicon Valle

ARTICLE COMMENTS (2)

APPS ARTIFICIAL INTELLIGENCE ASIA BAIDU CHINA DEEP LEARNI

Email Print

By GILLIAN WONG CONNECT

Baidu Chief Scientist Andrew Ng — Baidu

Six months ago, Chinese Internet-search giant Baidu signaled its ambitions to innovate by opening an artificial-intelligence center in Silicon Valley, in Google GOOGL -0.50%'s backyard. To drive home the point, Baidu hired Stanford researcher Andrew Ng, the

Europe acquisition Google DeepMind Technologies DeepMind

Google Acquires Artificial Intelligence Startup DeepMind For More Than $500M

Posted Jan 26, 2014 by Catherine Shu (@catherineshu)

12.7k SHARES

Next Story

Google will buy London-based artificial intelli
that the acquisition price was more than $50
the startup late last year. DeepMind confirme
terms.

LinkedIn Snatches Up Data Savvy Job Search Startup Bright.com For $120M, In Its Largest Acquisition To Date

Posted Feb 6, 2014 by Rip Empson (@ripemp)

1,607 SHARES

Next Stor

Today, alongside a fourth quarter earnings report in which it beat Wall Street estimates yet again, LinkedIn announced its intentions to acquire data-savvy job search startup, Bright.com, for $120 million. The deal, which was 70 percent stock and 30 percent cash the company said, will be completed during the firs quarter of this year.

In a statement today, LinkedIn said that "severa members of Bright's team," which now number over 50 –particularly those on its engineering and product teams — will be joining LinkedIn ir the coming weeks. However, one notices that the announcement conspicuously leaves out any mention of Bright's founders and whether or not they will be joining LinkedIn's team in Mountain View.

Either way, what is clear is that, unfortunately for Bright.com users and loyalists, as a result of the acquisition, access to the startup's job search products will continue until February 28th, at which point it LinkedIn will pull the plug.

Why did LinkedIn buy Bright, you ask, and whatever happened to that Monster.com fella? While we can't answer for the latter, we do know that the Bright.com purchase is the latest in a fairly short string of acquisitions LinkedIn has made over the last two years. No Yahoo by any means, LinkedIn has been methodically and strategically picking off startups that will either help expand its growing professional content network or its talent solutions products.

Big Acquisitions

- **Facebook**
  - Launched a new research lab dedicated entirely to advanced AI.
  - Yann LeCun
- **Google** acquired DeepMind for $500 million
  - Geoff Hinton
- **LinkedIn** acquired Bright for $120 million
- **Pinterest** acquired VisualGraph
- Chinese Web giant **Baidu** also recently established a Silicon Valley research lab to work on deep learning.
  - Andrew Ng

Top Academic Conferences

- **NIPS** Deep Learning and Unsupervised Feature Learning Workshop (2007 – 2014)
- **ICML** 2014 Workshop on Deep Learning Models for Emerging Big Data Applications
- **KDD** 2014 Tutorial on Scaling Up Deep Learning
- **CVPR** 2014 TUTORIAL ON DEEP LEARNING FOR VISION
- **ICML** 2013 Workshop on Deep Learning for Audio, Speech and Language Processing
- **AAAI** 2013 Deep Learning of Representations
- …

NORTHWESTERN UNIVERSITY

First You Need to Know that It Is Indeed Popular
Second You Have to See It Is Indeed Magical

- Deep Mind

  http://www.nature.com/nature/journal/v518/n7540/full/nature14236.html#videos

  https://www.youtube.com/watch?v=Zt-7MI9eKEo

- ConvNet

  http://cs.stanford.edu/people/karpathy/convnetjs/index.html

NORTHWESTERN
UNIVERSITY

Before Deep Learning…

Much of machine learning is about "feature engineering".

Most learning models are shallow.

Gradient-based training of multi-layer neural networks is difficult to obtain good generalization.

With Deep Learning…

Self-taught Feature Learning	A Multi-layer Structure ("Deep")	An Effective Training Algorithm
Unsupervised, automatic feature discovery from raw inputs.	Many layers, forming a hierarchy of representation.	Unsupervised layer-by-layer pre-training + supervised fine tuning.

NORTHWESTERN UNIVERSITY

Feature Hierarchies

- Learning feature hierarchies

  e.g. pixels ➔edges ➔combinations of edges ➔…

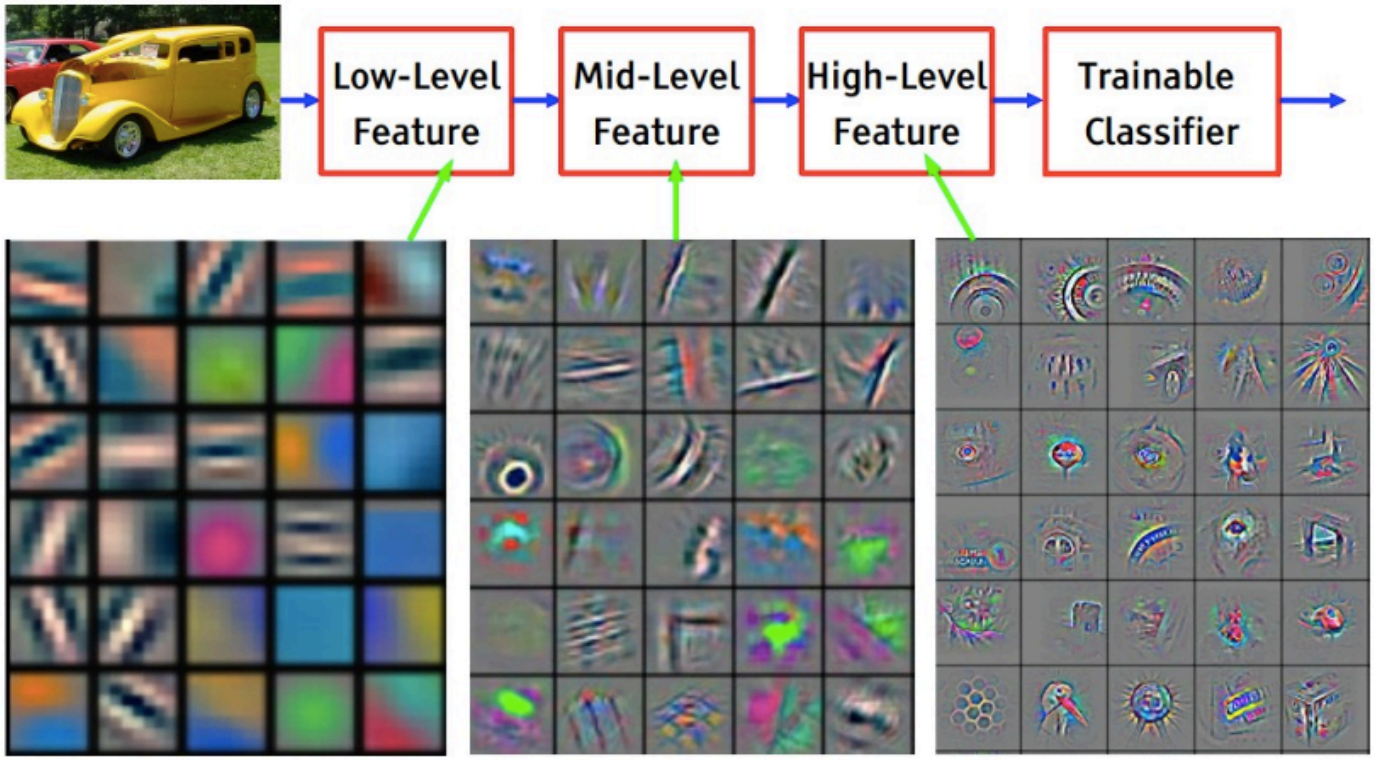

Feature visualization of convolutional net trained on ImageNet from [Zeiler & Fergus 2013]

Social Applications

NORTHWESTERN
UNIVERSITY

Mining for Structure

- Massive increase in the amount of data available from web, video cameras, laboratory measurements…

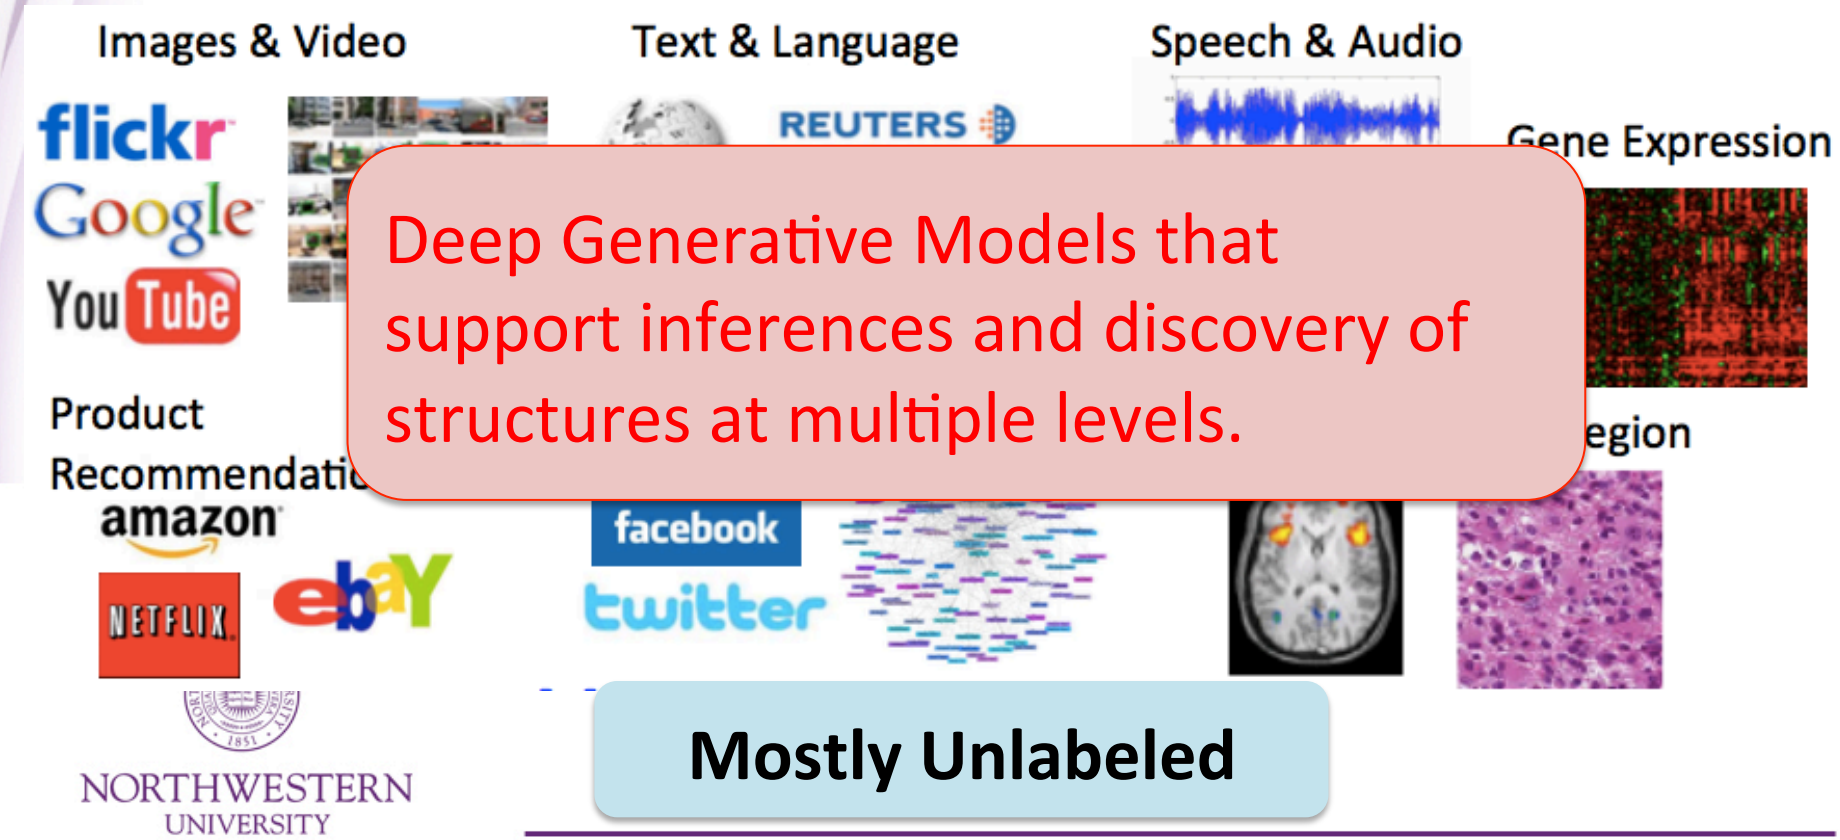

The Era of Big Data

- 71% of CMOs around the globe cited Big Data as their top challenge.
- Social network data grows unboundedly.
  - By 2015, there will be more than **5,300** exabytes of unstructured digital consumer data stored in databases, a large share of that generated by social networks.
  - (1 exabyte = 1 million terabytes, Facebook ingests ~500 terabytes of data each day)
- Facebook > 500 times NYSE
- Twitter > 12 times NYSE

NORTHWESTERN
UNIVERSITY

The Era of Big Data

- Data are mostly “Unstructured”
  - spontaneously user-generated and not easily captured and classified
  - “Structured” data is more akin to data entered into a form
- Data are fragmentary and cryptic

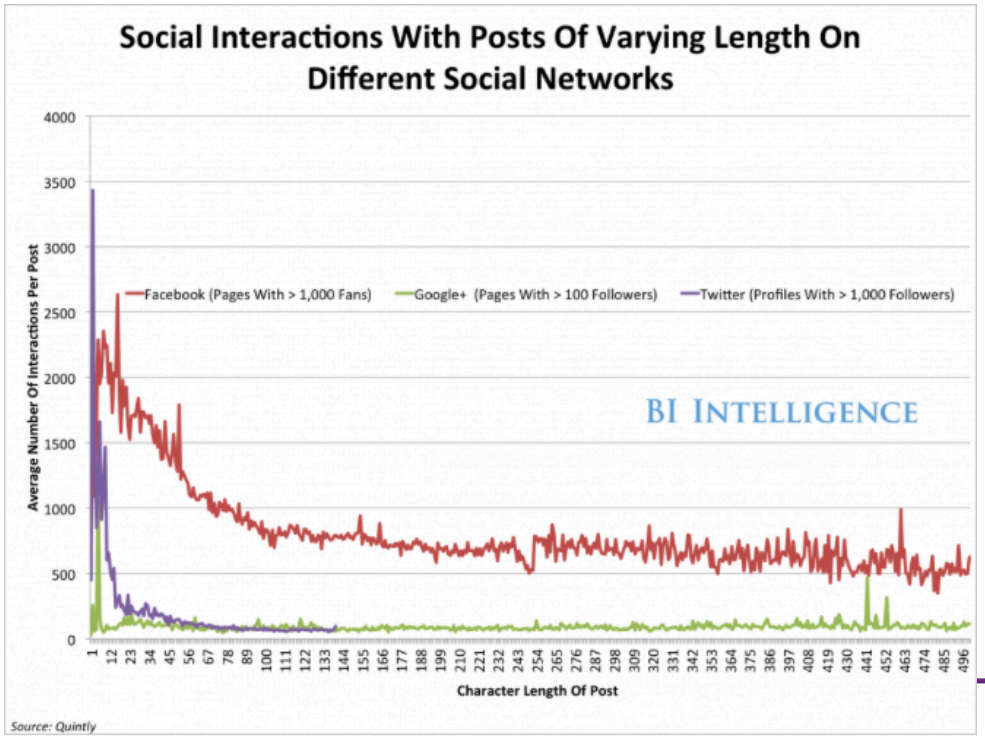

NORTHWESTERN UNIVERSITY

Social Network + ML/AI

- ML and AI are helping marketers and advertisers glean insights from this vast ocean of unstructured consumer data.
  - e.g. Image recognition identifies brand logos in photos.
- Visual content extraction
  - Social networking experiences are becoming increasingly centered around photos and videos.
- Text mining
- New application areas: audience targeting, personalized predictive marketing, brand sentiment analysis…

NORTHWESTERN
UNIVERSITY

Watch the Two Prominent Tasks in AI

- Deep Learning on Image Recognition
- https://vimeo.com/109982315

- Deep Learning on Speech Recognition
- https://www.youtube.com/watch?v=yxxRAHVtafI

NORTHWESTERN
UNIVERSITY

Theories

NORTHWESTERN
UNIVERSITY

Deep Learning is Big and Complex

- Many types of deep architectures

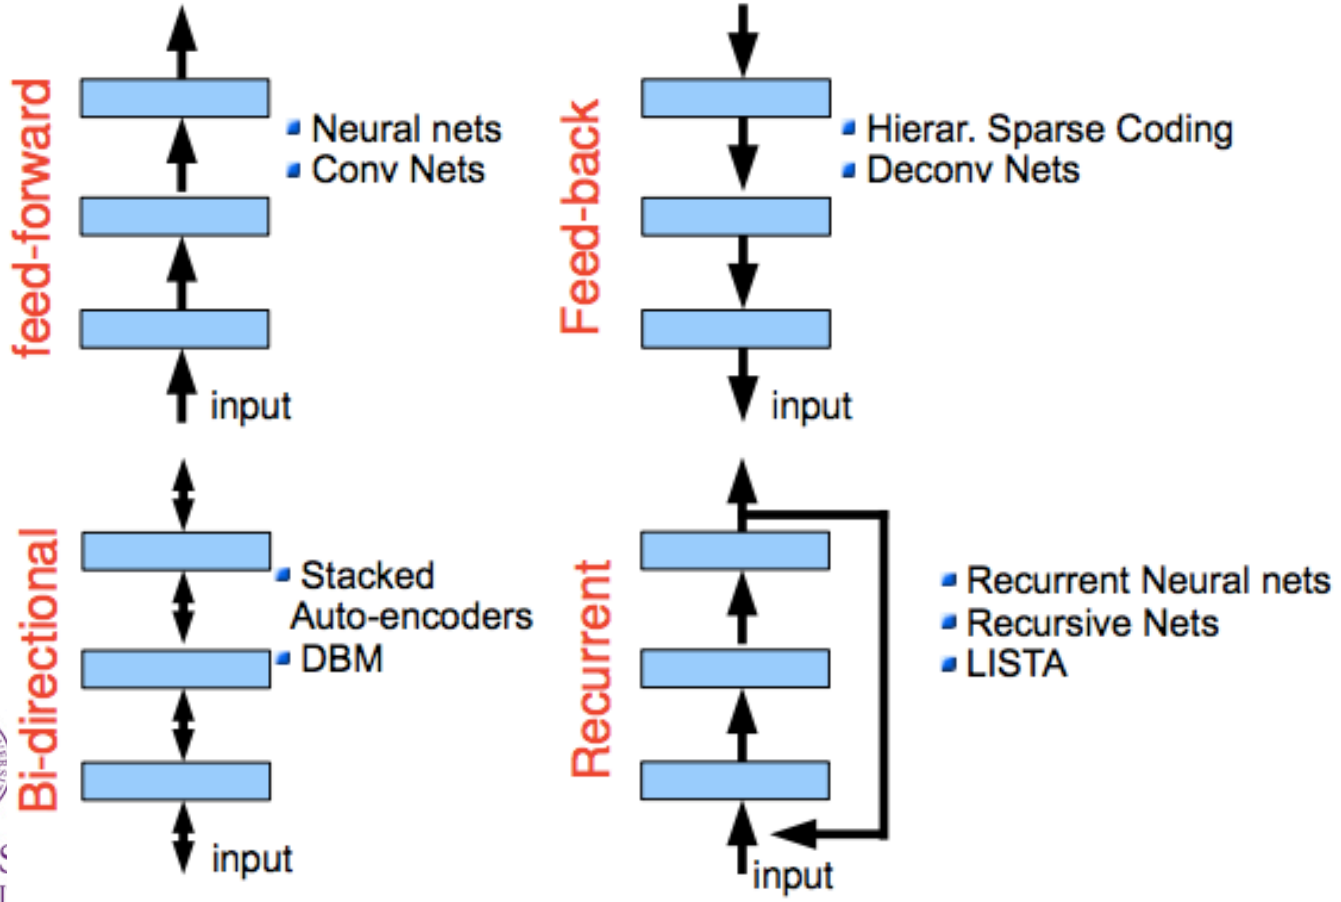

Deep Learning is Big and Complex

- Many types of learning protocols
  - Purely supervised
    - Backprop + SGD
    - Good when there is lots of labeled data
  - Layer-wise supervised + supervised linear classifier
    - Train each layer in sequence
    - Hold fix the feature extractor, train linear classifier on features
    - Good when labeled data is scarce but there is lots of unlabeled data.
  - Layer-wise unsupervised + supervised backprop
    - Train each layer in sequence
    - Backprop through the whole system
    - Good when learning problem is very difficult.

Auto-encoder

- A simple unsupervised learning algorithm
- A one-layer perceptron model

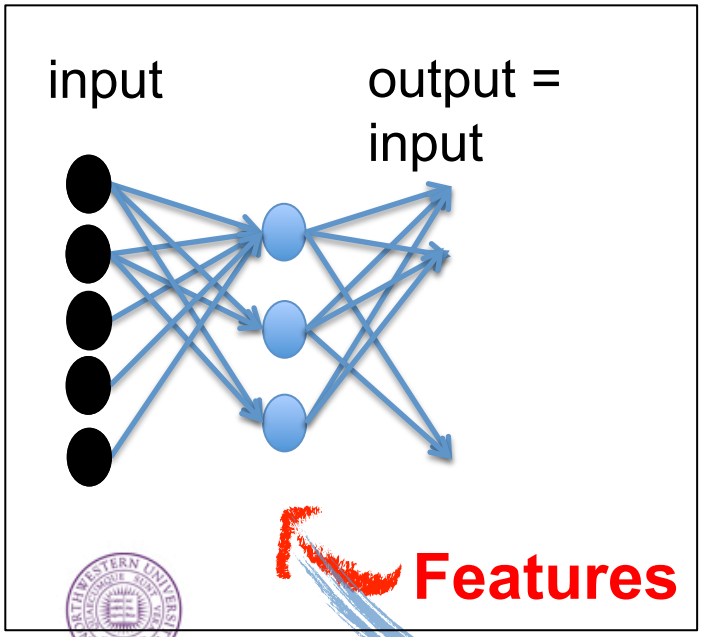

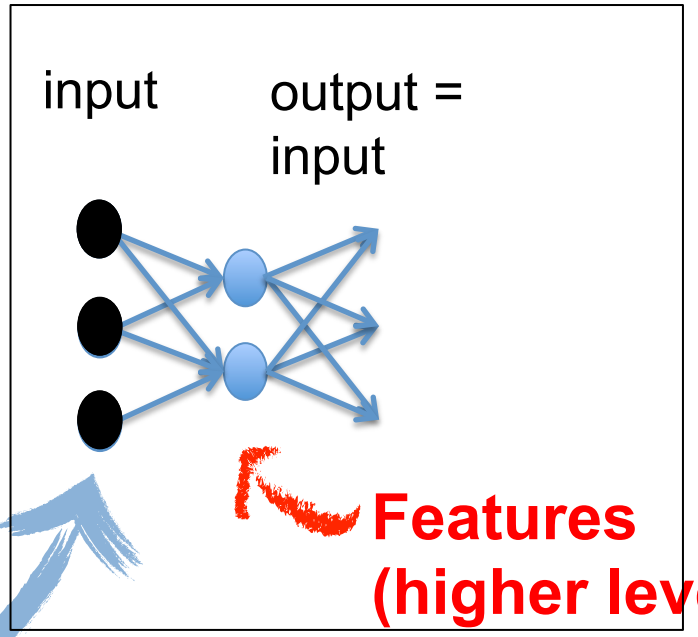

Stacked Auto-encoder

- Stack the feature layers learned separately
- Fine-tune: supervised training as usual

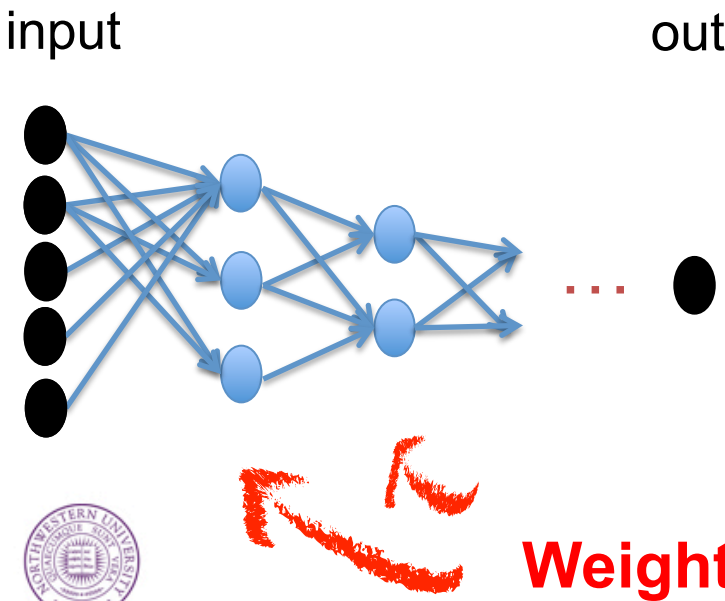

Weights initialized by pre-training

Projects

- Krizhevsky/Google's convnet
  - CNN w/ CUDA
- Berkeley's Caffe
  - Modular CNN package in C++ with both CPU/GPU training
  - Interface in Python & MATLAB
- Bengio's Theano
  - Python project (w/ Numpy & Scipy), works with GPU
- Facebook/NYU's Torch7
  - LuaJIT interface to C

People

- Google
  - Geoff Hinton, Ilya Sutskever, Alex Krizhevsky
- Stanford
  - Chris Manning, Richard Socher, Andrej Karpathy
- Baidu
  - Andrew Ng
- Facebook
  - Yann LeCun, Tomas Mikolov, Jason Weston, Marc'Aurelio Ranzato

Thanks!

NORTHWESTERN
UNIVERSITY