

# Deep learning from academia to industry

Ole Winther

Dept for Applied Mathematics and Computer Science  
Technical University of Denmark (DTU)

Bioinformatics Centre/BRIC  
University of Copenhagen (KU)



## Take home messages


- Human-level machine perception with deep learning
- Algorithms and frameworks open source
- Short idea-to-product cycles
- Opportunities for companies but
- clear data collection strategy needed!

## Take home messages

- Human-level machine perception with deep learning
- Algorithms and frameworks open source
- Short idea-to-product cycles
- Opportunities for companies but
- clear data collection strategy needed!
- Probabilistic programming is trending!

# Take home messages

- Human-level machine perception with deep learning
- Algorithms and frameworks open source
- Short idea-to-product cycles
- Opportunities for companies but
- clear data collection strategy needed!
- Probabilistic programming is trending!
- A bit about myself



# Computing - why we are here!

## 1 The accelerating pace of change ...



## 2 ... and exponential growth in computing power ...

Computer technology shown here climbing dramatically by powers of 10, is now progressing more each hour than it did in its entire first 90 years

### COMPUTER RANKINGS

By calculations per second per \$1,000

**Analytical engine**  
Never fully built, Charles Babbage's invention was designed to solve computational and logical problems

**Colossus**  
The electronic computer, with 1,500 vacuum tubes, helped the British crack German codes during WW II



**UNIVAC I**  
The first commercially marketed computer, used to tabulate the U.S. Census, occupied 943 cu. ft.




**Apple II**  
At a price of \$1,298, the compact machine was one of the first massively popular personal computers

## 3 ... will lead to the Singularity



**Power Mac G4**  
The first personal computer to deliver more than 1 billion floating point operations per second



# Data - why we will have self-driving cars



The Cityscapes dataset

# Major areas in AI

- Speech recognition
- Image classification
- Machine translation
- Question-answering
- Self-driving vehicles
- Dialogue systems
- General unsupervised learning



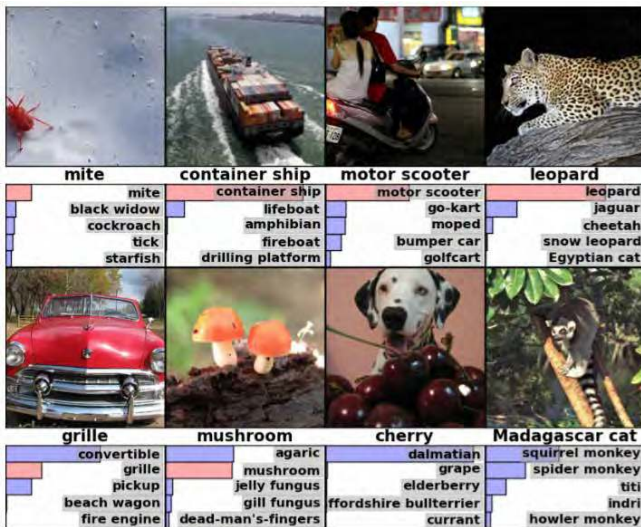
# Major areas in AI

- Speech recognition
- Image classification
- Machine translation
- Question-answering
- Self-driving vehicles
- Dialogue systems
- General unsupervised learning






# ImageNet - image classification



# ImageNet classification challenge

ILSVRC top-5 error on ImageNet



- AlexNet - A Krizhevsky et al. (2012) won with huge margin!
- Soon everyone started using **deep learning** and **GPUs**.


# Feature engineering vs engineered models

## ImageNet Classification with Deep Convolutional Neural Networks

**Alex Krizhevsky**  
University of Toronto  
kriz@cs.utoronto.ca


**Ilya Sutskever**  
University of Toronto  
ilya@cs.utoronto.ca

**Geoffrey E. Hinton**  
University of Toronto  
hinton@cs.utoronto.ca



[www.cs.toronto.edu/~fritz/absps/imagenet.pdf](http://www.cs.toronto.edu/~fritz/absps/imagenet.pdf)

# Convolutional neural networks




$$\begin{bmatrix} 10 & 0 & -10 \\ 0 & 0 & 0 \\ -10 & 0 & 10 \end{bmatrix}$$




# We need bigger brains

- AlexNet (2012): 16.4% error, 8 layers, 1.4 Gflop




# We need bigger brains

- AlexNet (2012): 16.4% error, 8 layers, 1.4 Gflop



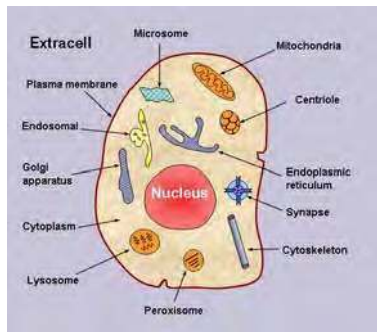
- ResNet (2016): 3.5% error, 152 layers, 22.6 Gflop.



- (This is a so-called DenseNet and not a ResNet.)
- **Source:** [Source Jen-Hsun Huang, CEO NVIDIA, GTC Europe, 2016](#)

# Deep learning – DTU and KU research group


- End-to-end!
- Structured data - sequences+
- Bioinformatics



- Information retrieval - search in `findzebra.com`

# Deep learning – DTU and KU research group

- End-to-end!
- Structured data - sequences+
- Bioinformatics



- Information retrieval - search in [findzebra.com](http://findzebra.com)

- Condition monitoring – [Siemens Wind Power](http://Siemens Wind Power) and [greengoenergy.com](http://greengoenergy.com)
- Document interpretation - [tradeshift.com](http://tradeshift.com)

An invoice from Tyrell Corporation with red arrows pointing to specific fields for document interpretation. The fields are: SENDER, INVOICE NUMBER, DATE, ITEMS, TAX, and TOTAL.

ITEM NO.	ITEM DESCRIPTION	QTY	PRICE	TOTAL
1001	Flow capacitor (Discussion completed)	5	18,000	90,000
1002	Flow capacitor (Discussion completed)	5	18,000	90,000
1003	Flow capacitor (Discussion completed)	5	18,000	90,000
1004	Flow capacitor (Discussion completed)	5	18,000	90,000
1005	Flow capacitor (Discussion completed)	5	18,000	90,000
1006	Flow capacitor (Discussion completed)	5	18,000	90,000
1007	Flow capacitor (Discussion completed)	5	18,000	90,000
1008	Flow capacitor (Discussion completed)	5	18,000	90,000
1009	Flow capacitor (Discussion completed)	5	18,000	90,000
1010	Flow capacitor (Discussion completed)	5	18,000	90,000

NET TOTAL: USD 225,000.00  
TAX (9.000%): USD 20,250.00  
TOTAL: USD 245,250.00

- Variational un- and semi-supervised learning



# Document interpretation with tradeshift.com

**SENDER** → **TYRELL CORPORATION**  
Dr. Eldon Street, 1, Los Angeles, CA 91020, USA

**RECEIVE** → **Stark Industries**  
Howard Stark, 40  
Palmdale,  
CA 93550, USA  
98-7654321

**INVOICE NUMBER** → Invoice number.: inv. 123456

**DATE** → Invoice date: 14/02/2015

ITEM REF.	ITEM DESCRIPTION	QTY	PRICE	TOTAL
RF673	Flux capacitor (DeLorean compatible)	1	95.000,-	95.000,-
AS245	Cell phone Samsung SPH-N270	10	345,-	3450,-
ZS304	Psychomagnotheric Slime	100 l.	17.00	1700,-

**ITEMS** →

**TAX** → **SUBTOTAL:** USD 100.150,00  
**TAX (20% VAT):** USD 20.030,00

**TOTAL** → **TOTAL:** USD 120.170,00

The payment must be done 14 days after the invoice date, while any claim must be done within 30 days. Bunch of other words in very small font that nobody reads. They do not follow any formal unit may repeat that it was previously included in the invoice. For our purposes, there is nothing interesting here.


“When you hear hoofbeats behind you, don’t expect to see a zebra”

**Filters**

**Group by:**  disease  gene

Mediastinal tumor		<b>Mediastinal tumor</b>
Desmoplastic small round cell ... (2)		Retrieved: 28-09-2014 Source: WIKIPEDIA (Original article)
Thymoma (2)		
Adult-onset Still's disease (3)		
Large-cell lymphoma (3)		
Follicular lymphoma		
Periodic fever, familial, autosomal... (2)		
Thymic carcinoma		
TNF receptor associated periodic...		
Japanese encephalitis		
Periodic fever, familial, autosomal ...		

The mediastinum is the cavity that separates the lungs from the rest of the chest. It contains the heart, esophagus, trachea, thymus, and aorta. The mediastinum has three main parts: the **anterior** mediastinum (front), the middle mediastinum, and the posterior mediastinum (back). The most common **mediastinal** masses are neurogenic tumors (20% of **mediastinal** tumors), usually found in the posterior mediastinum, followed by thymoma (15-20%) located in the **anterior** mediastinum. Masses in the **anterior** portion of the mediastinum can include thymoma, lymphoma, pheochromocytoma, germ cell tumors including teratoma, thyroid tissue, and parathyroid lesions. Masses in this area are more likely to be malignant than those in other compartments. Masses in the posterior portion of the mediastinum tend to be neurogenic in origin, and in adults tend to be of neural sheath origin including neurilemmomas and neurofibromas. Lung cancer typically spreads to the lymph nodes in the mediastinum.



**Diagnosis**

In several editions of Physical Diagnosis, concerning **mediastinal** tumors the author writes: According to Christian 1 the **mediastinal** neoplasms which are neither so rare nor so obscure as to make diagnosis practically impossible are:


Navigation icons: back, forward, search, etc.

# Industrial PhD with Siemens Wind Power

- Data is there
  - 10k turbines monitored for 5+ years with
  - detailed vibration and other sensor data
  - 100s of faults of different types
- Organisation already taking a data-driven approach.

# Industrial PhD with Siemens Wind Power

- Data is there
  - 10k turbines monitored for 5+ years with
  - detailed vibration and other sensor data
  - 100s of faults of different types
- Organisation already taking a data-driven approach.



- Work by Martin Bach-Andersen, to appear in Wind Energy

# Convolutional neural network for vibrational data

