ARE MACHINE LEARNING AND AI THE MAGIC TOOLS IN INDUSTRY 4.0?

Jan Larsen, Professor PhD
DTU Compute
Department of Applied Mathematics and Computer Science

DTU

A copy of the physical world through digitization makes it possible for cyber-physical systems to communicate and cooperate with each other and with humans in real time and perform decentralized decision-making

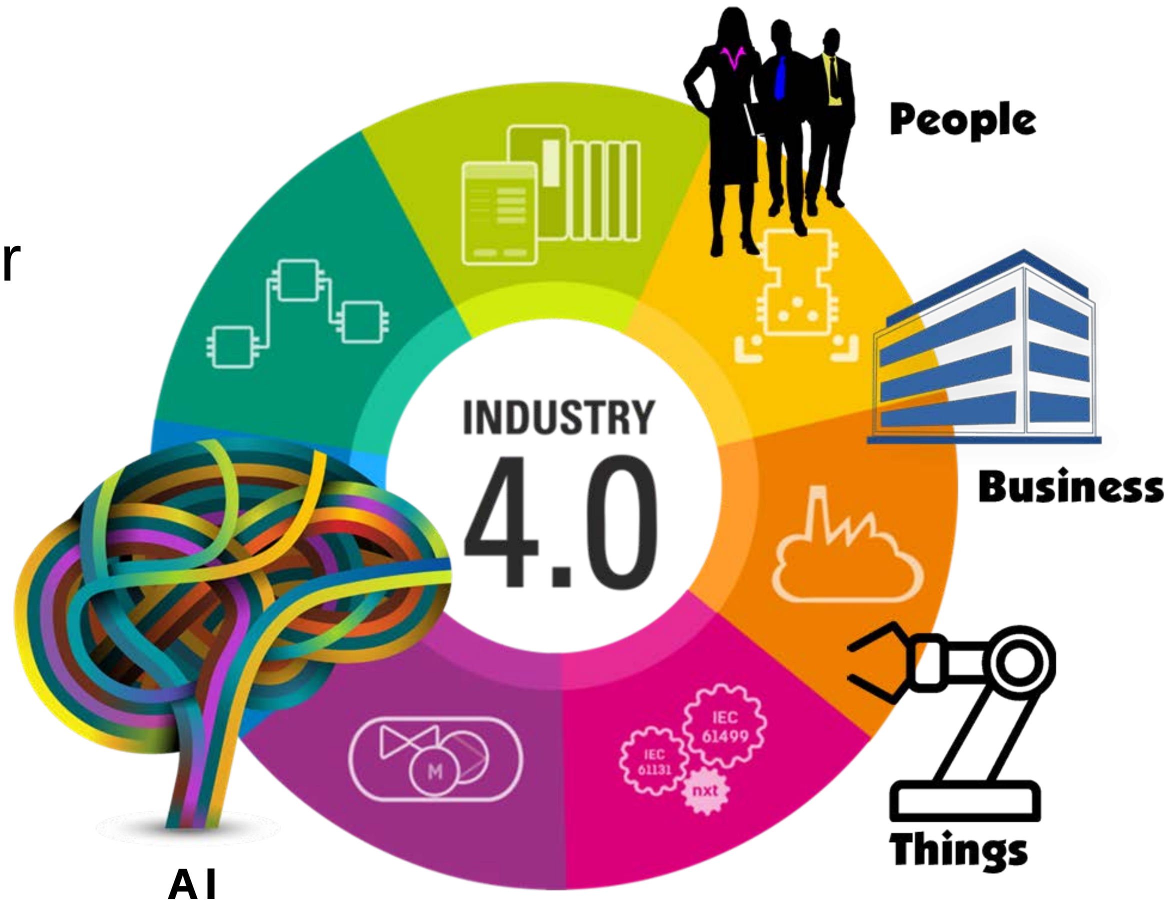

https://en.wikipedia.org/wiki/Industry_4.0

B. Marr: Forbes, June 20, 2016, http://www.forbes.com/sites/bernardmarr/2016/06/20/what-everyone-must-know-about-industry-4-0/#4c979f804e3b

http://www.enterrasolutions.com/2015/10/industry-4-0-facing-the-coming-revolution.html

Brief history of AI

Late 40's Allan Touring: theory of computation

1948 Claude Shannon: A Mathematical Theory of Communication

1948 Norbert Wiener: Cybernetics - *Control and Communication in the Animal and the Machine*

1950 The Touring test

1951 Marvin Minsky's analog neural networks (1st revolution)

1956 Dartmouth conference: Artificial intelligence with aim of human like intelligence

1956-1974 Many small scale "toy" projects in robotics, control and game solving

1974 Failure of success and Minsky's criticism of perceptron, lack of computational power, combinatorial explosion, Moravec's paradox: simple tasks are not easy to solve

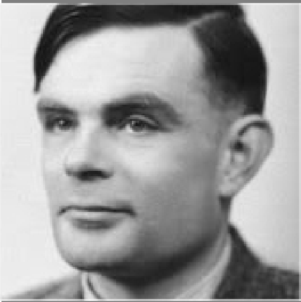

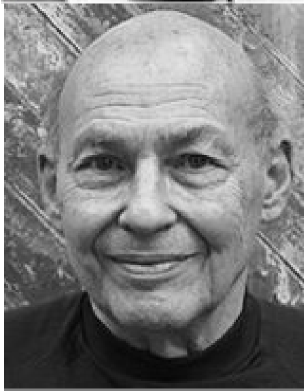

1980's Expert systems useful in restricted domains

1980's Knowledge based systems – integration of diverse information sources

1980's The 2nd neural network revolution starts

Late 1980's Robotics and the role of embodiment to achieve intelligence

1990's and onward AI and cybernetics research under new names such as machine learning, computational intelligence, evolutionary computing, neural networks, Bayesian networks, complex systems, game theory, deep neural networks (3rd generation) cognitive systems

http://en.wikipedia.org/wiki/Timeline_of_artificial_intelligence

http://en.wikipedia.org/wiki/History_of_artificial_intelligence

Gartner®

Emerging Technology Hype Cycle

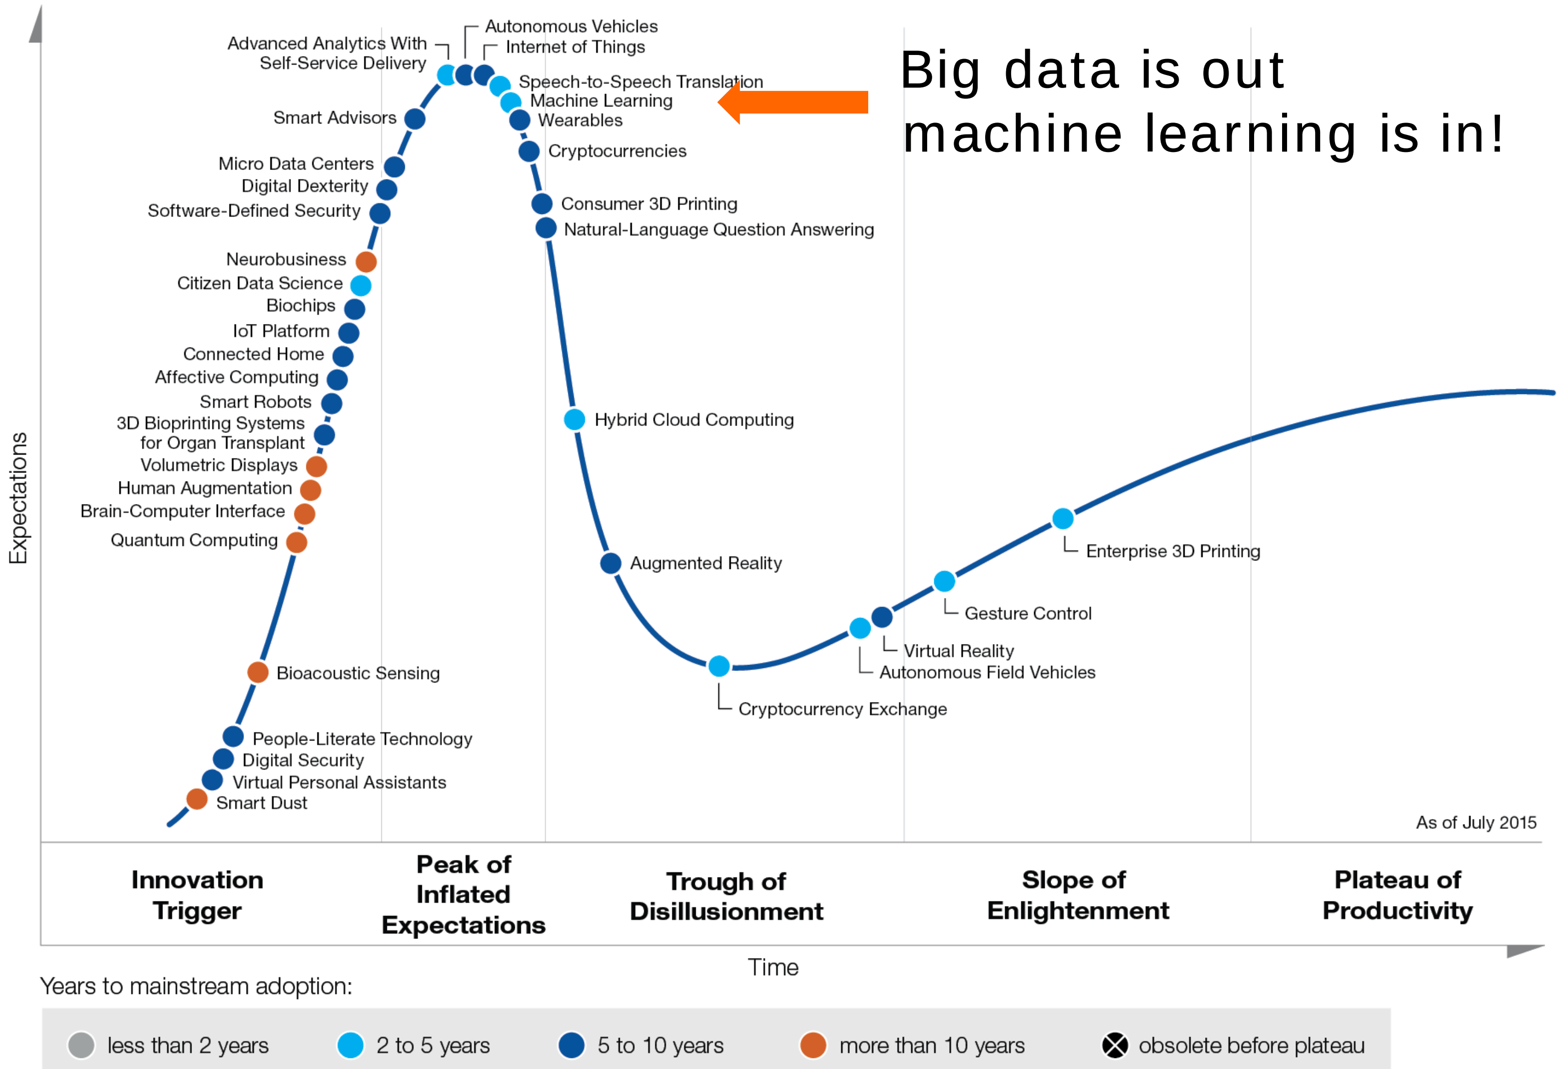

Big data is out
machine learning is in!

The digital revolution makes data science and AI increasingly relevant and important and will eventually disrupt most procedures and aspects of human life

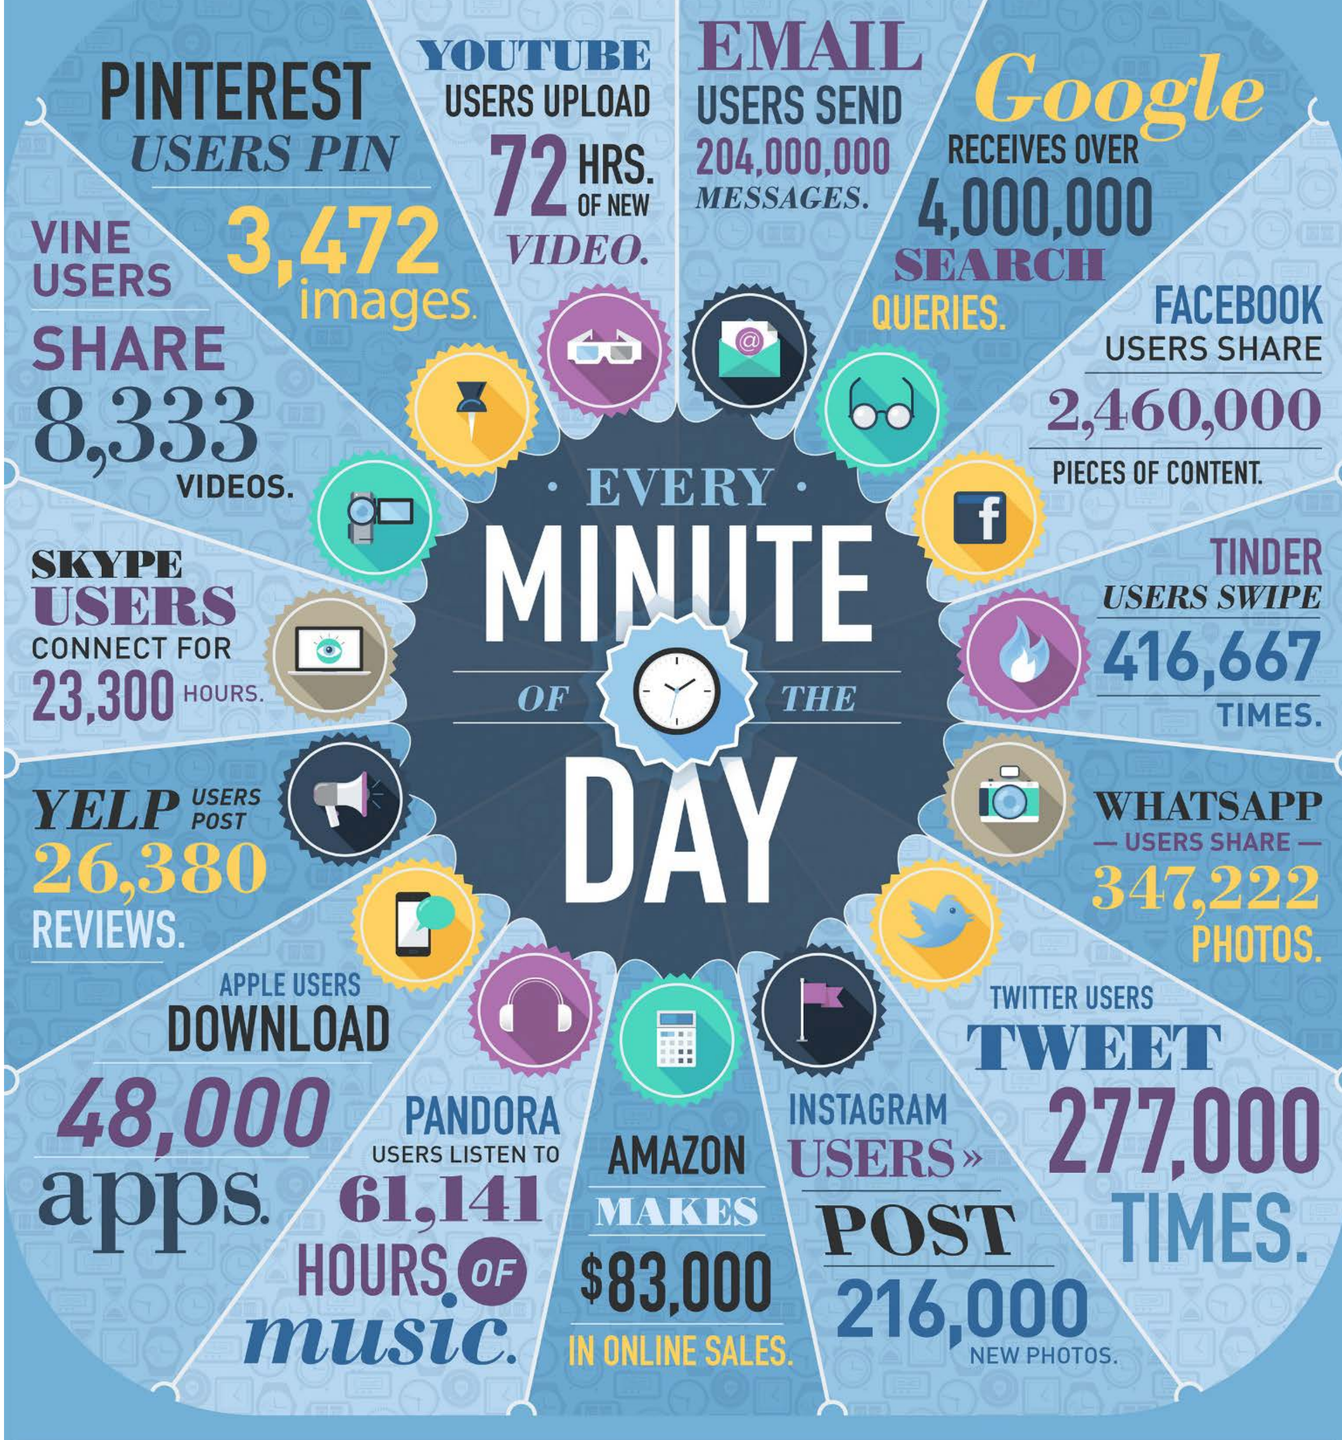

Social metadata according to domo.com

Big data drives industry 4.0

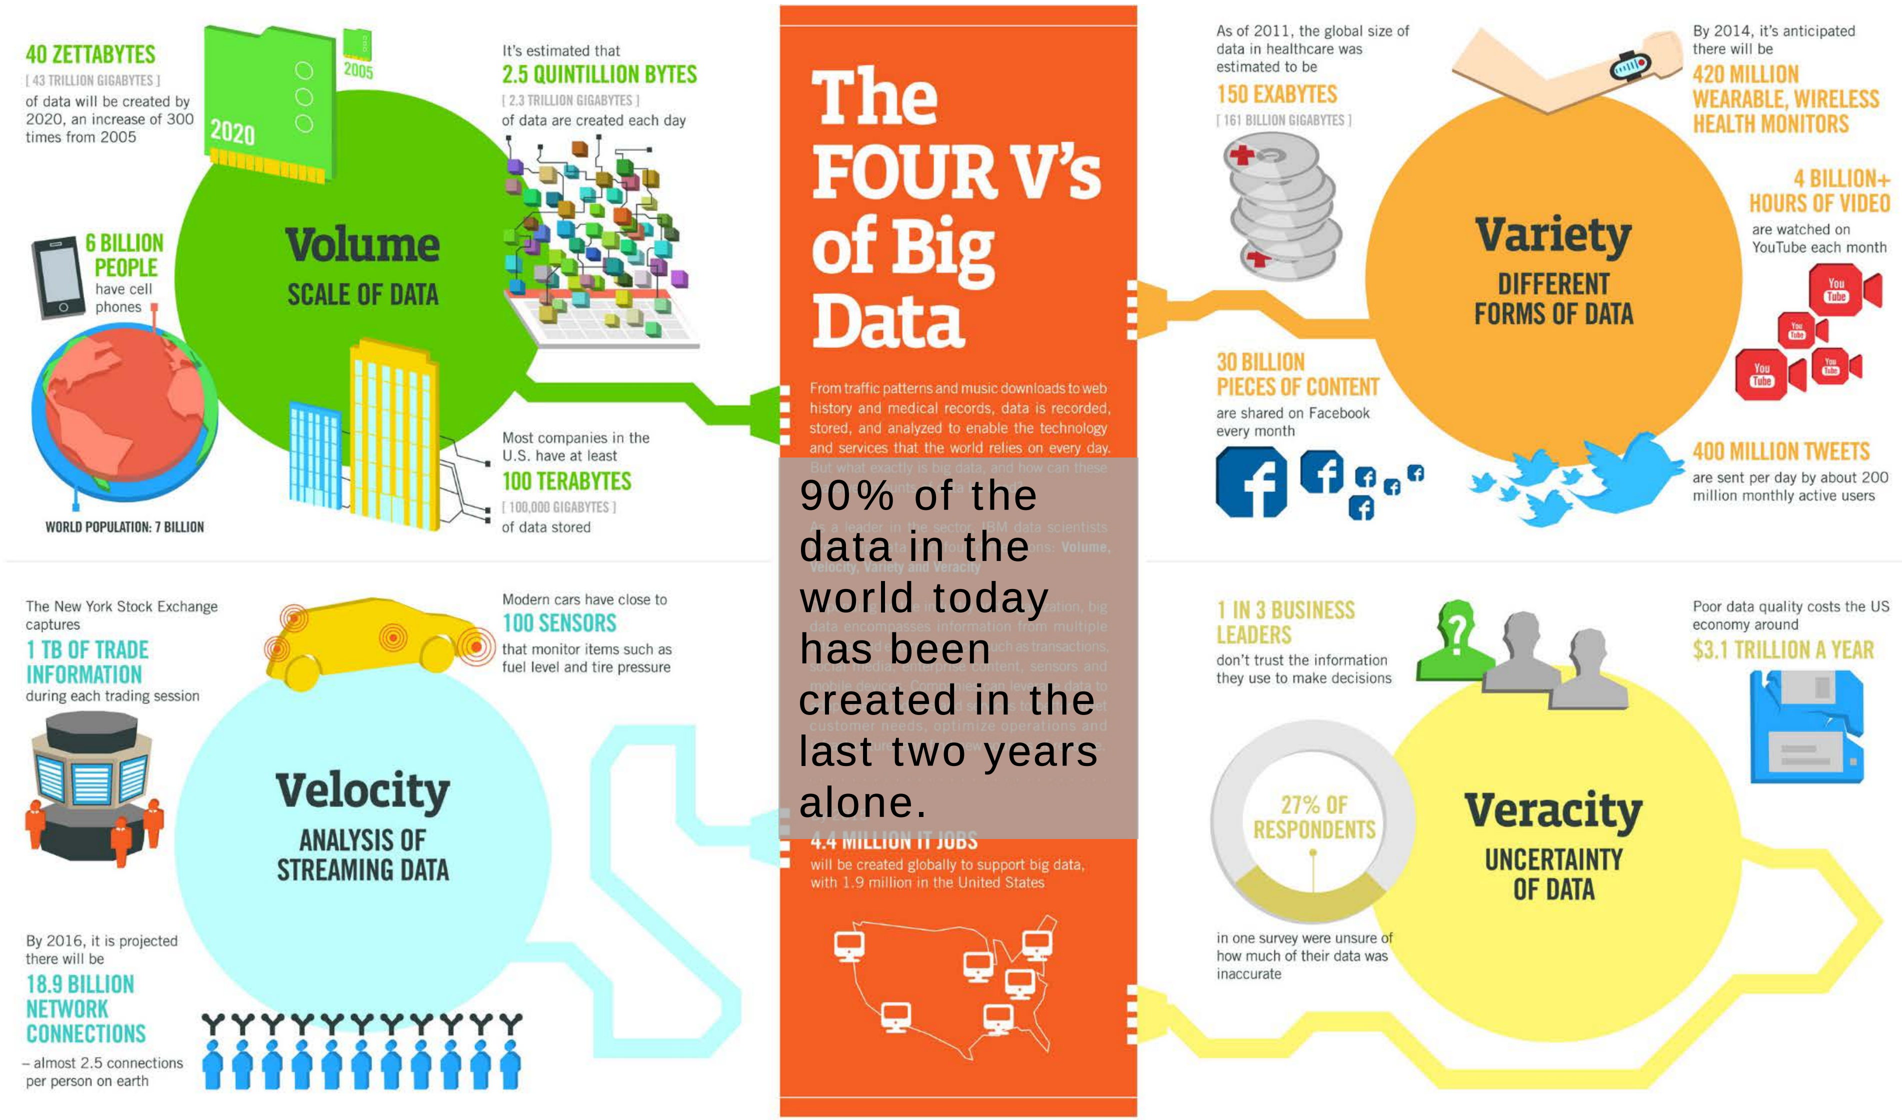

90% of the data in the world today has been created in the last two years alone.

IBM, www.ibmbigdatahub.com

Technological singularity and artificial general intelligence (AGI)

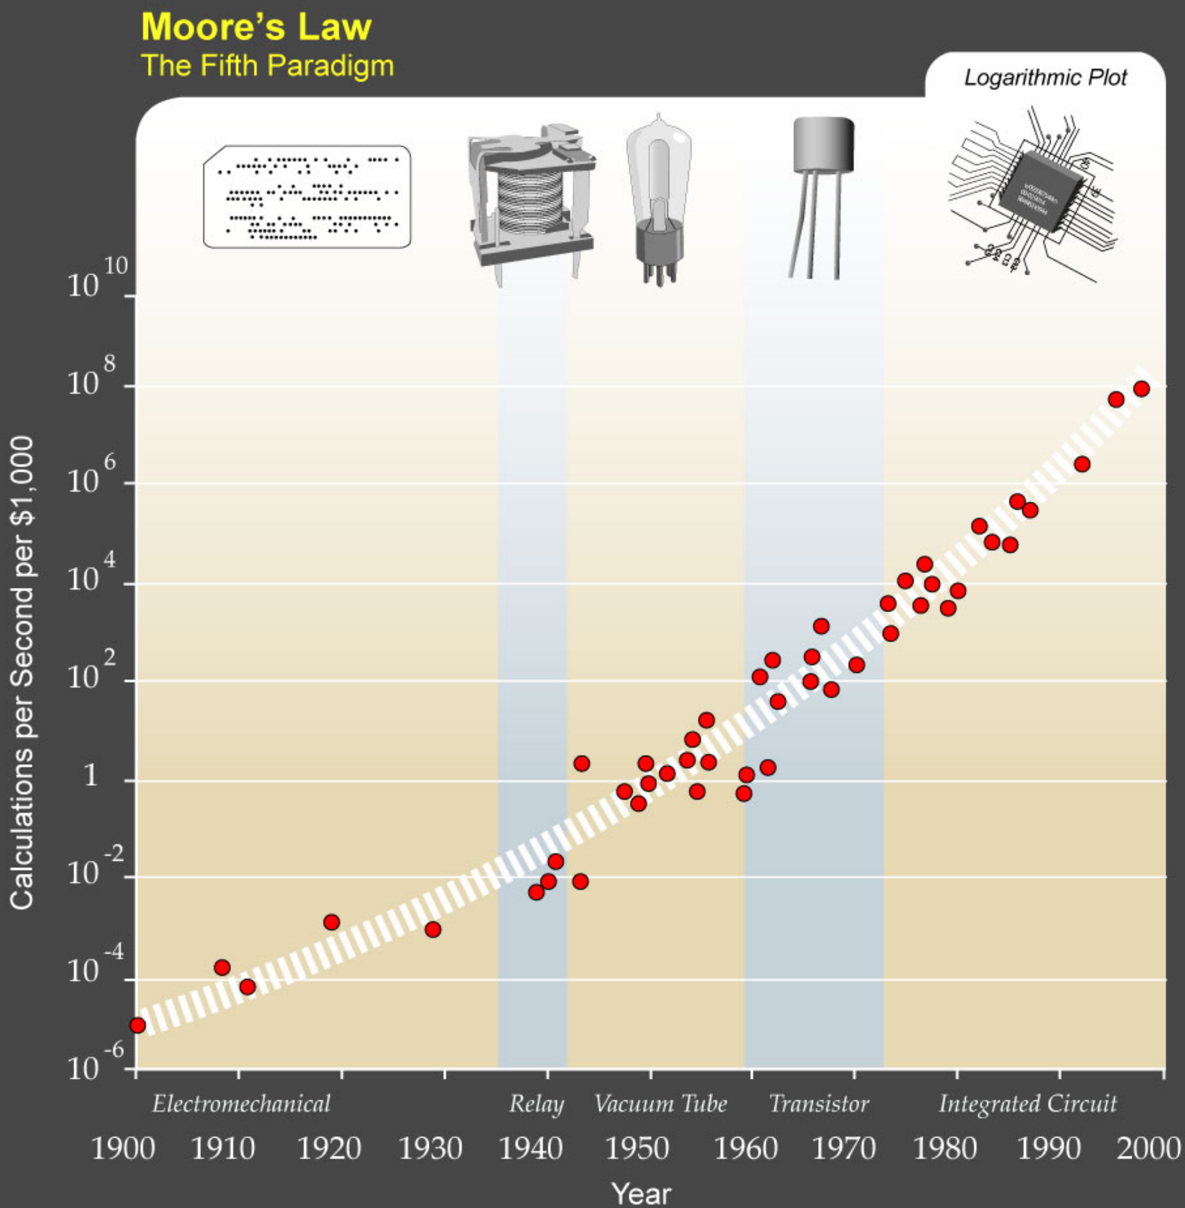

Technological paradigm cause exponential growth extends Moore's law from integrated circuits to earlier transistors, vacuum tubes, relays, and electromechanical computers.

In a few decades the computing power of all computers will exceed that of human brains, with superhuman artificial intelligence appearing around the same time

Ray Kurzweil: The Singularity is Near, Penguin Group, 2005.

IBM's TrueNorth chip and SyNAPSE and Quantum Computing Chips

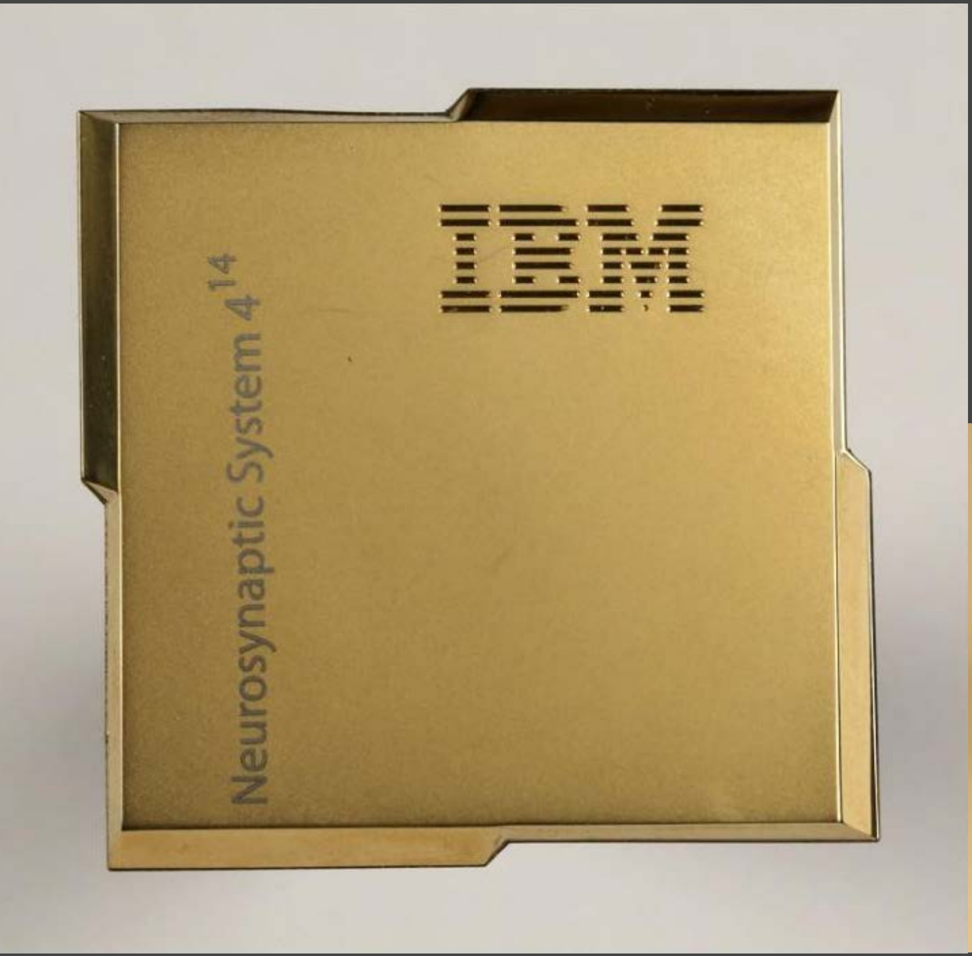

4096 cores in the current chip, each one simulating 256 programmable silicon "neurons" for a total of just over a million neurons

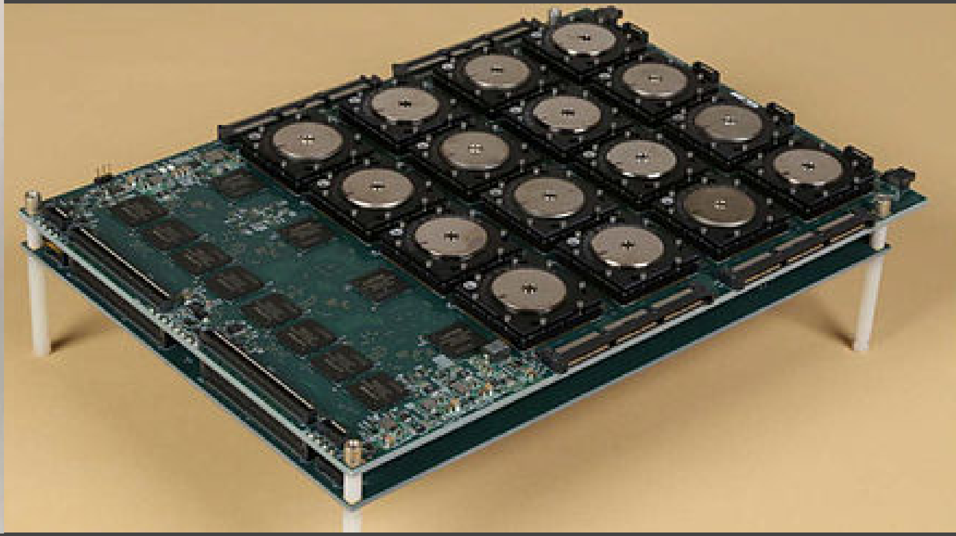

*Merolla, P. A.; Arthur, J. V.; Alvarez-Icaza, R.; Cassidy, A. S.; Sawada, J.; Akopyan, F.; Jackson, B. L.; Imam, N.; Guo, C.; Nakamura, Y.; Brezzo, B.; Vo, I.; Esser, S. K.; Appuswamy, R.; Taba, B.; Amir, A.; Flickner, M. D.; Risk, W. P.; Manohar, R.; Modha, D. S. (2014). "A million spiking-neuron integrated circuit with a scalable communication network and interface". Science. **345** (6197): 668.*

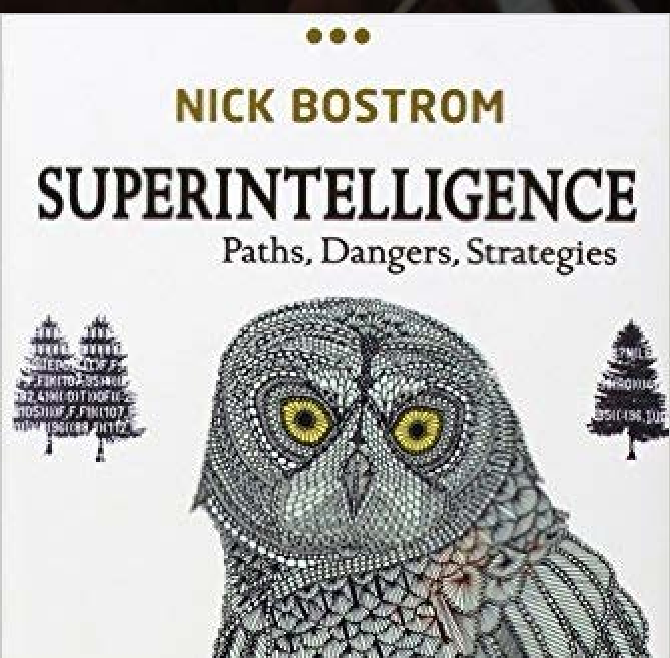

AI run-away?

Argues for the possibility of a fast-leap in intelligence and discusses hypothetical example scenarios where an AI rapidly acquires a dominant position over humanity.

Kaj Sotala, How Feasible Is the Rapid Development of Artificial Superintelligence?, Sept. 2016

Professor Neil Lawrence, University of Sheffield

AI run-away?

fundamental limits on predictability

"We cannot predict with infinite precision and this will render our predictions useless on some particular time horizon."

"This limit on our predictive ability places a fundamental limit on our ability to make intelligent decisions."

Kaj Sotala, How Feasible Is the Rapid Development of Artificial Superintelligence?, Sept. 2016

Algorithms Among Us: The Societal Impacts of Machine Learning, NIPS2015 Symposium.

N. Lawrence: http://inverseprobability.com/blog

Professor Stephen Hawking,
Cambridge University

AI run-away?

AI will be 'either best or worst thing' for humanity.

AI will develop itself and be in conflict with or not understandable by humans.

It challenge what it means to be human, every aspect of live will change, and be the biggest change to civilization maybe also the last.

can remedy damages to the world that industry 3.0 did such as eradicating poverty and cure health problems.

Leverhulme Centre for the Future of Intelligence, Opening, Oct 19, 2016

Industry 4.0 = Civilization 4.0

It is a cognitive revolution that could be even more disruptive than earlier as it concerns not only the industry but the whole way we live our lives

√ Big data through cyber-physical systems and IoT constitute the necessary **resource**/raw material.

√ Low cost, large scale computational platforms constitute the **engine**.

√ Robust high-speed communication **link** resources.

But how do we process and convert data into actionable results **?**

Machine learning has shown to be very a promising methodology!

Big players provide open source and premium storage, computing, and analytics tools

Trifacta, Alteryx, Paxata and Informatica Rev are making data preparation easier (now 80% time data prep, 20% analysis

Amazon Redshift: fast, fully managed, petabyte-scale data warehouse

Machine Learning APIs: IBM Watson, Microsoft Azure Machine Learning, Google Prediction API, Amazon Machine Learning API, and BigML.

Amazon Web Services

Google Deep Mind: methods and technology

Apache Hadoop, Apache Spark are open-source software framework for distributed storage and processing of very large data sets

ML Software platforms: Google Tensor flow, MS CNTK, Apache Mahout, Facebook Learner Flow

IBM Blue Mix cloud based platform

Top 7 Trends in Big Data for 2015, Tableau Software. 5 Best Machine Learning APIs for Data Science, blog.

What is machine learning?

Learning structures and patterns form from historical data to reliably predict outcome for new data

Computers only do what they are programmed to do. ML infers new relations and patterns, which were not programmed they learn and adapt to changing environment.

M. I. Jordan and T. M. Mitchell. *Machine learning: Trends, perspectives, and prospects*. Science, July 2015.

Samuel J. Gershman, Eric J. Horvitz, Joshua B. Tenenbaum. *Computational rationality: A converging paradigm for intelligence in brains, minds, and machines.* Science, July 2015.

Deep Learning: Automating Feature Discovery

Geoff Hinton, Yoshua Bengio, Yann LeCun, Deep Learning Tutorial, NIPS 2015, Montreal.

Deep learning is a disruptive technology

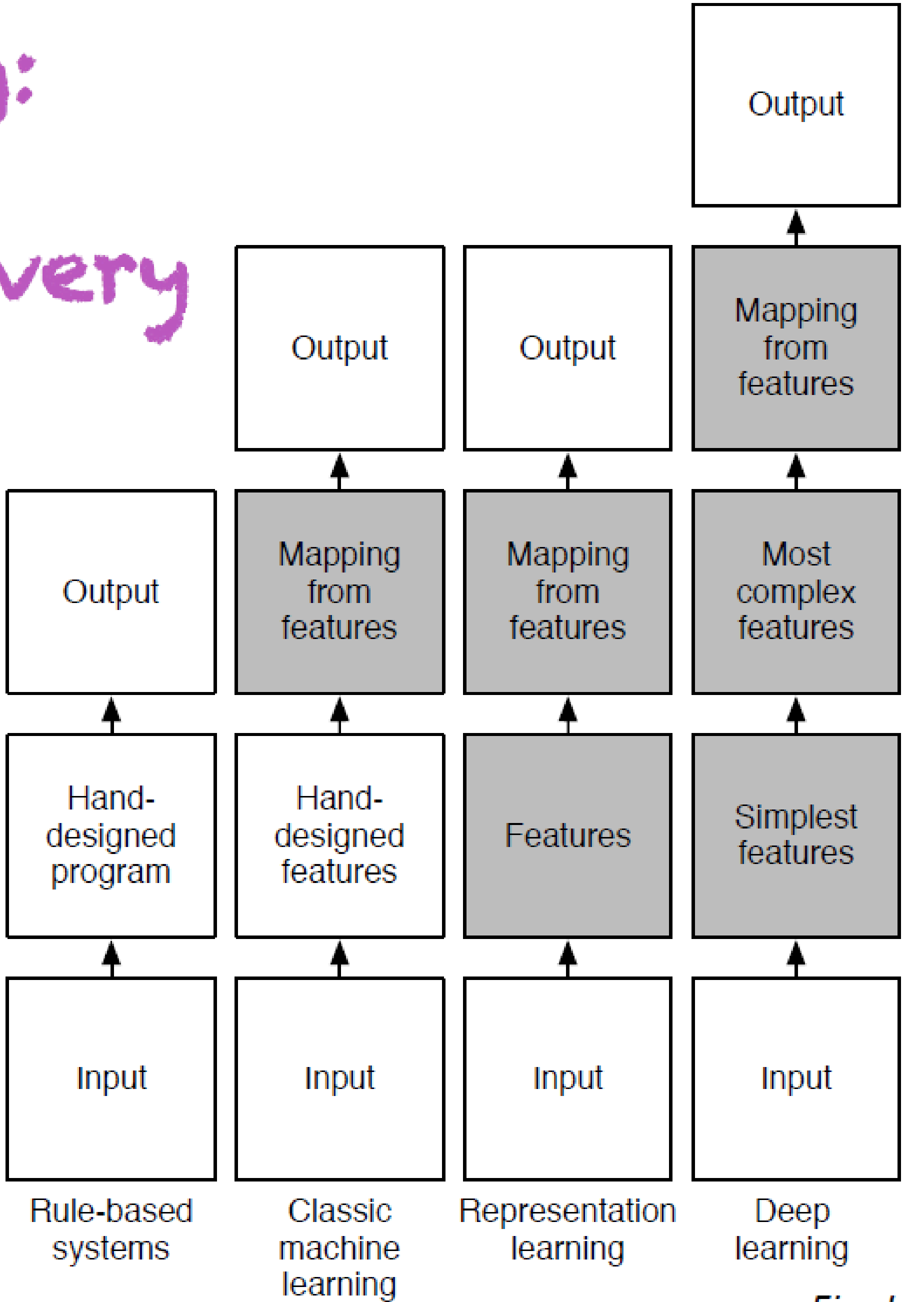

Fig: I. Goodfellow

The unreasonable effectiveness of

Mathematics E. Wigner, 1960

Data Halevy, Norvig, Pereira, 2009

RNNs Karpathy, 2015

Machine learning is very successful: playing GO

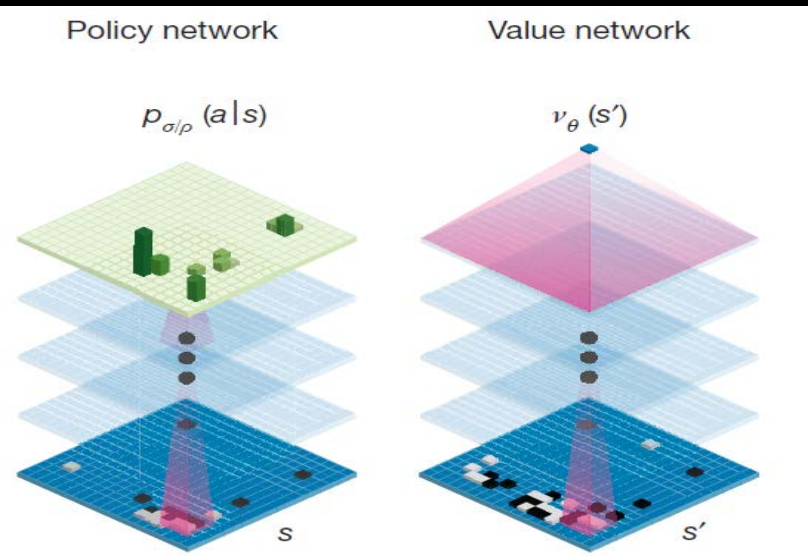

Deep neural ‘value networks’ evaluate board positions and other ‘policy networks’ select moves.

Networks are trained by a novel combination of supervised learning from human expert games, and reinforcement learning from games of self-play.

Silver, David; Huang, Aja; Maddison, Chris J.; Guez, Arthur; Sifre, Laurent; Driessche, George van den; Schrittwieser, Julian; Antonoglou, Ioannis; Panneershelvam, Veda. *Mastering the game of Go with deep neural networks and tree search. Nature 529(7587): 484–489, 2016*

Machine learning is very successful: computer vision

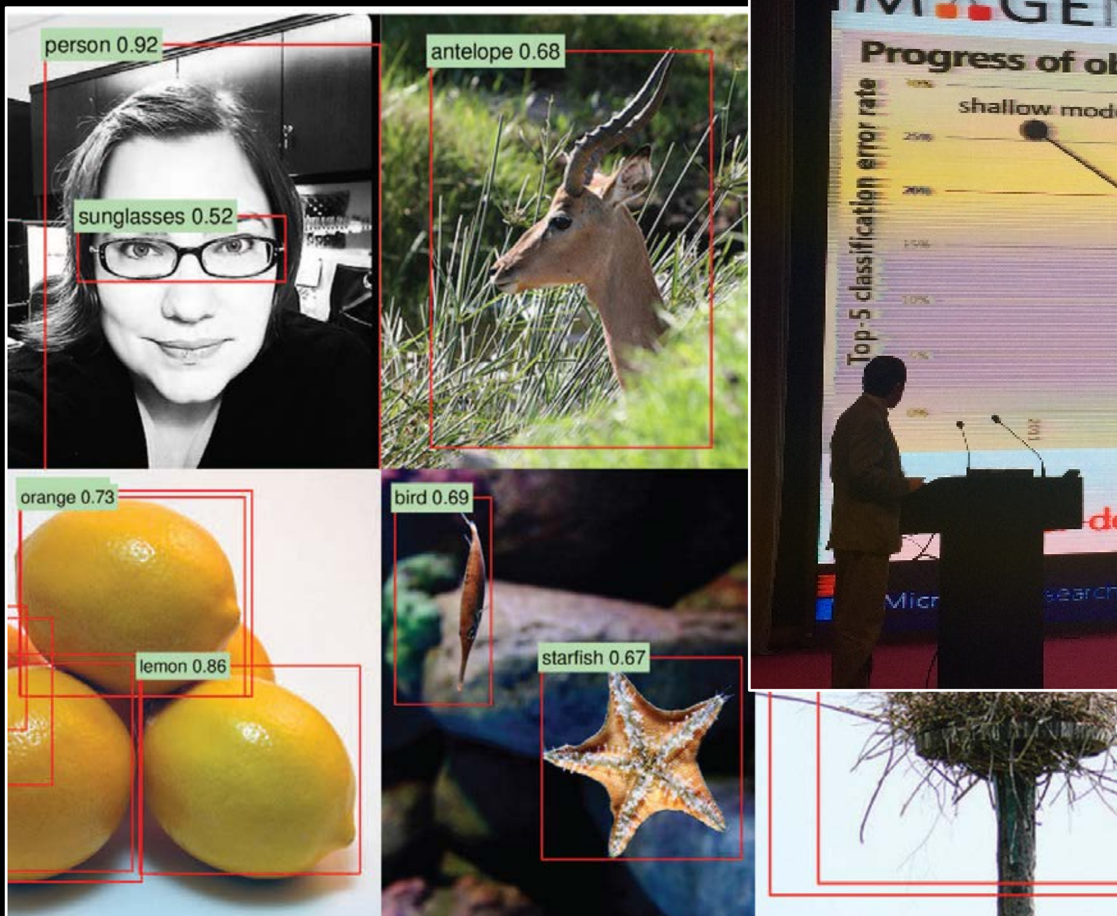

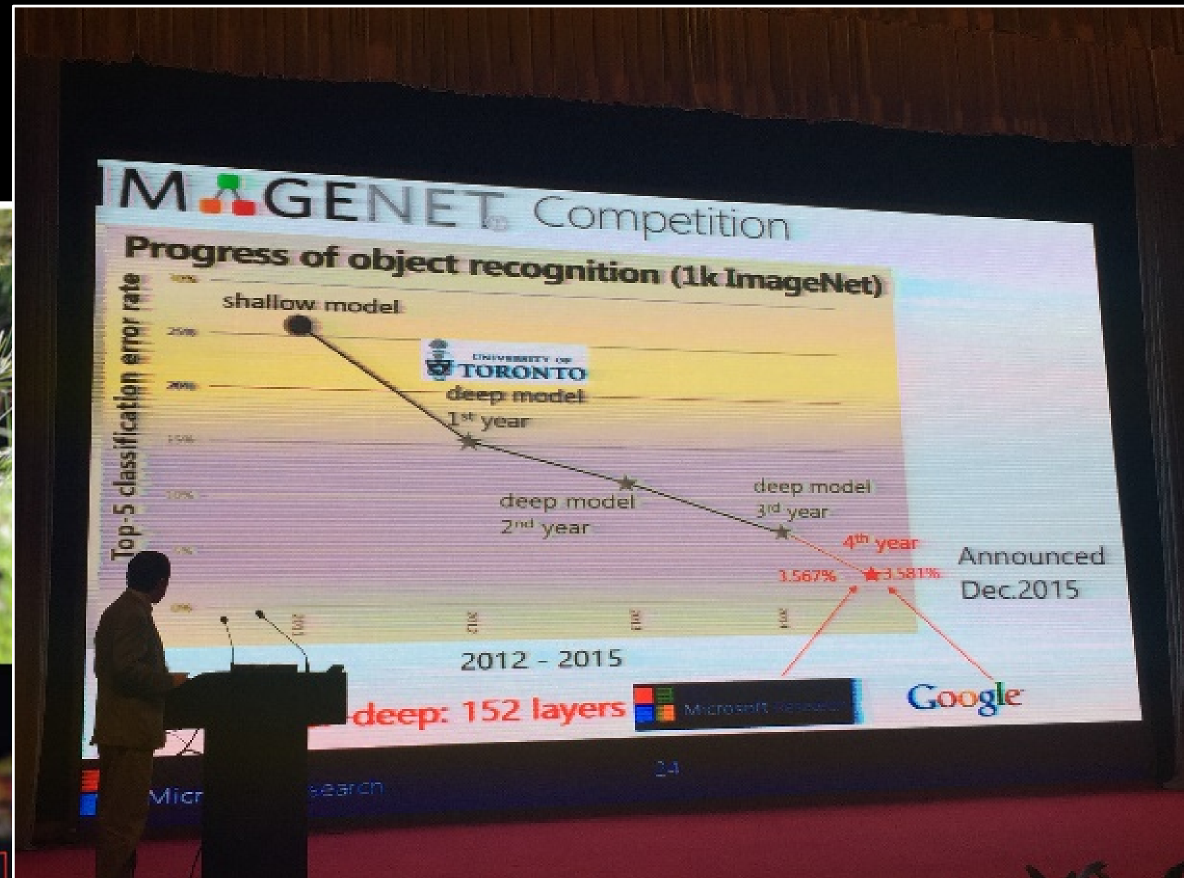

M. I. Jordan and T. M. Mitchell. *Machine learning: Trends, perspectives, and prospects*. Science, July 2015.

Machine learning is very successful: speech recognition and chat bots

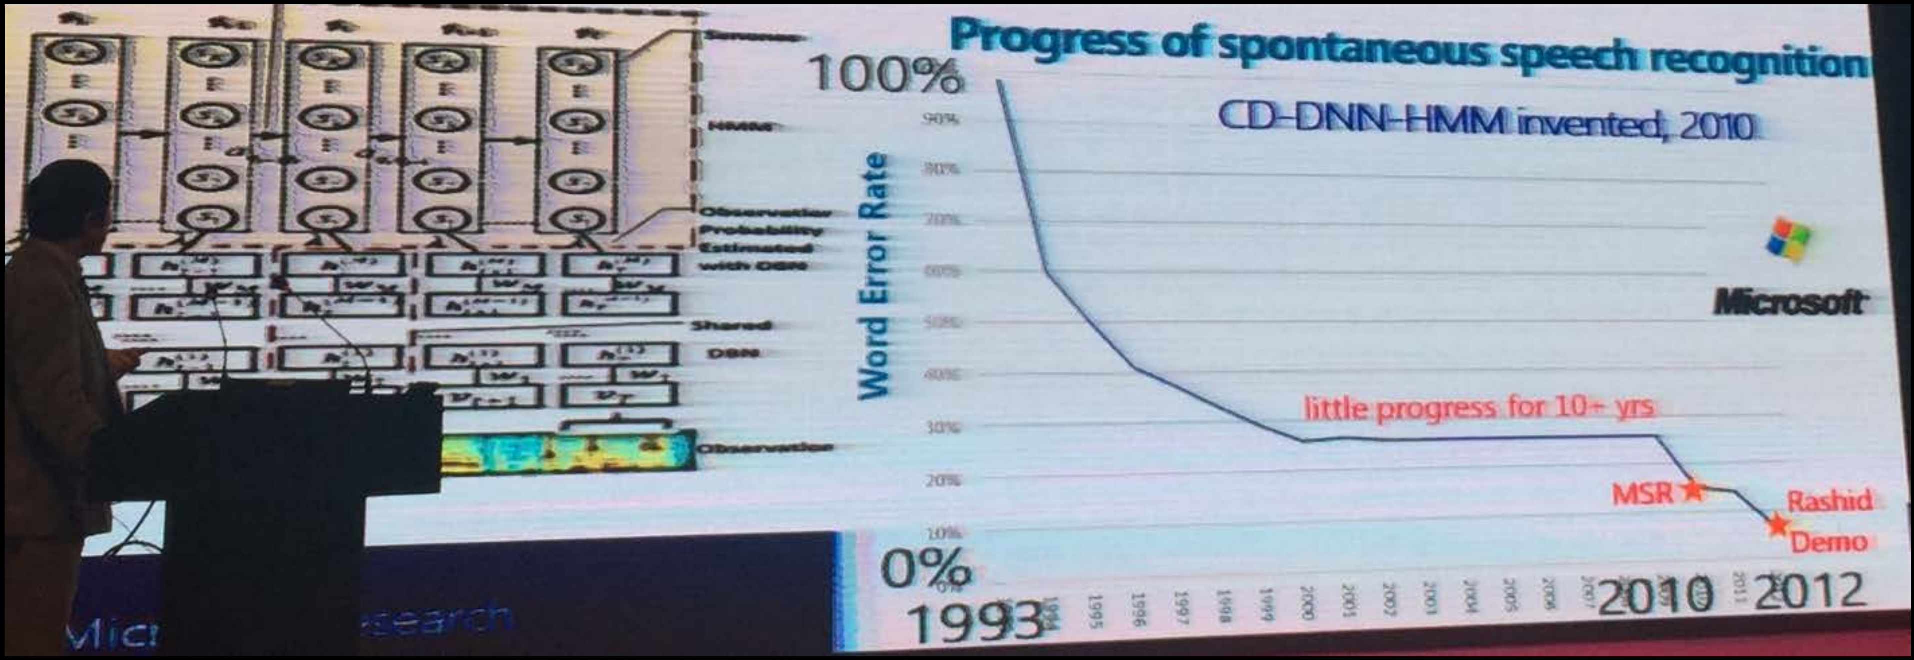

Geoffrey Hinton, Li Deng, Dong Yu, George Dahl, Abdel-rahman Mohamed, Navdeep Jaitly, Andrew Senior, Vincent Vanhoucke, Patrick Nguyen, Tara Sainath, and Brian Kingsbury*. Deep Neural Networks for Acoustic Modeling in Speech Recognition.* IEEE Signal Processing Magazine, 82, Nov. 2012.

Machine learning is sucessful: preditive and personalized medicine

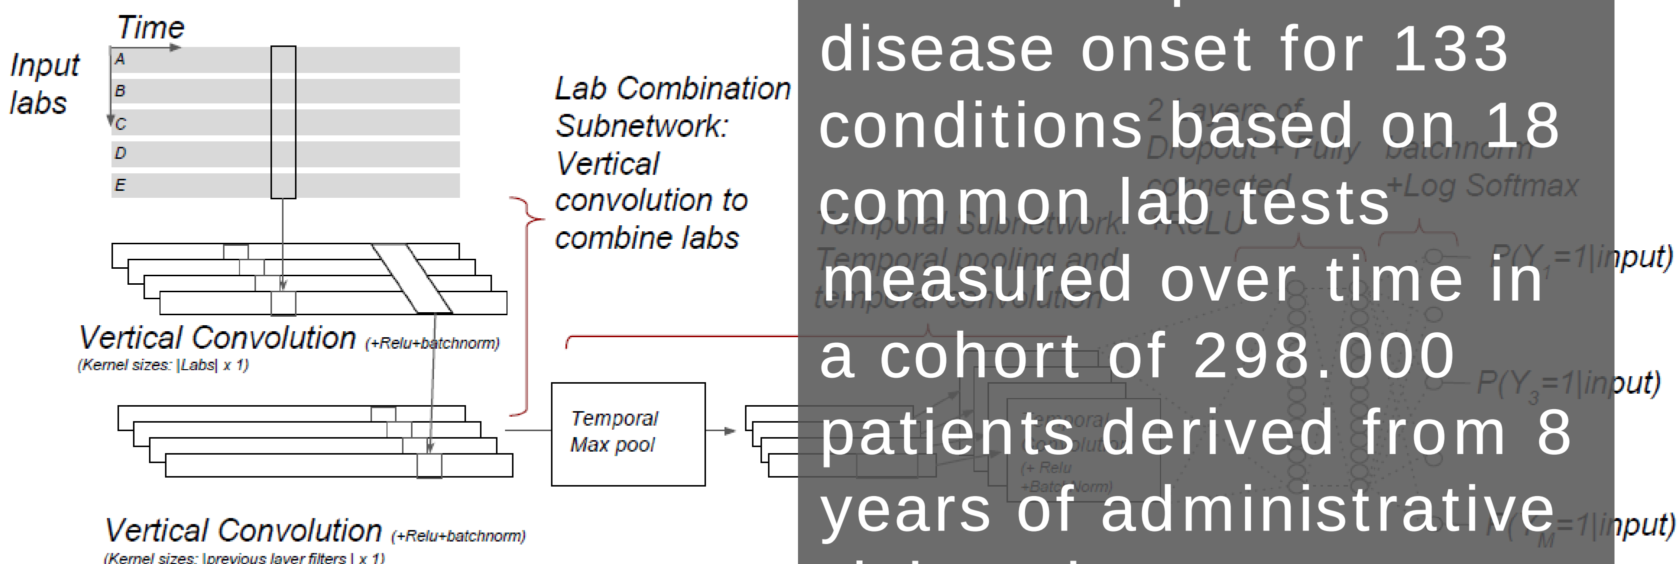

multi-task prediction of disease onset for 133 conditions based on 18 common lab tests measured over time in a cohort of 298.000 patients derived from 8 years of administrative claims data

Figure 3: Architecture for Convolutional Neural Network over Time and Input dimensions (CNN2)

N. Razavian, J. Marcus, D. Sontag: Multi-task Prediction of Disease Onsets from Longitudinal Lab Tests, NYU, ArXiv, 2016

BLACK BOX OF AI

Objectives:

Trust

Causality

Transferability

Decomposability

Informativeness

Legal issues: European Union regulations on algorithmic decision-making and a “right to explanation”

Davide Castelvecchi, http://www.nature.com/polopoly_fs/1.20731!/menu/main/topColumns/topLeftColumn/pdf/538020a.pdf, Nature, Vol. 538, 6 Oct. 2016

Lipton, Z.C. The mythos of model interpretability. arXiv:1606.03490 (2016).

European Union, https://arxiv.org/pdf/1606.08813v3.pdf

Computational creativity using deep nets

Representations of content and style in the Convolutional Neural Network are separable hence can be manipulated independently to produce new, perceptually meaningful images

L.A. Gatys, A. S. Ecker, M. Bethge:A Neural Algorithm of Artistic Style, arXiv:1508.06576v1, 26 Aug. 2015

WaveNet: A Generative Model for Raw Audio

Wave
model
from D

It is sh
able to
mimics
which
than th
Speech
gap wi
by over 50%.

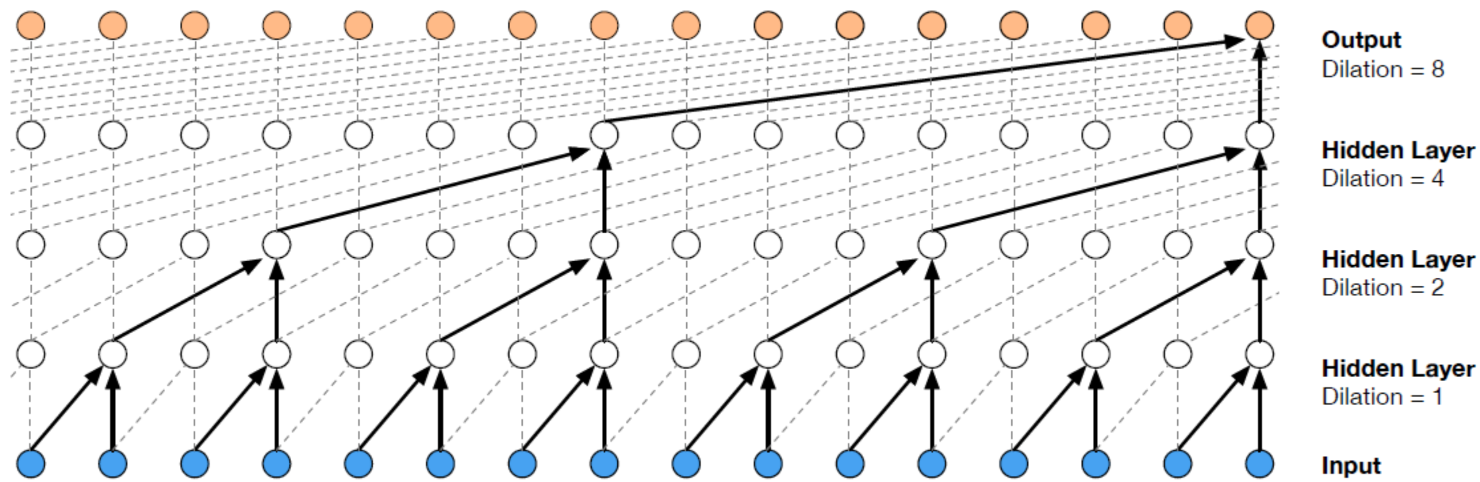

Figure 3: Visualization of a stack of *dilated* causal convolutional layers.

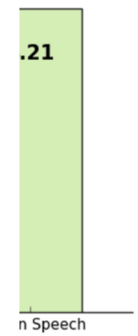

References:

https://deepmind.com/blog/wavenet-generative-model-raw-audio/ https://arxiv.org/pdf/1609.03499.pdf

WaveNet: A Generative Model for Raw Audio

The network generated and out sequences not condition an input sequence telling it what to play (such as a musical score)

Trained it on a dataset of classical piano music

References:

https://deepmind.com/blog/wavenet-generative-model-raw-audio/ https://arxiv.org/pdf/1609.03499.pdf

Todd, P.M. (1989). "A connectionist approach to algorithmic composition". Computer Music Journal. ***13*** *(4): 27–43.*

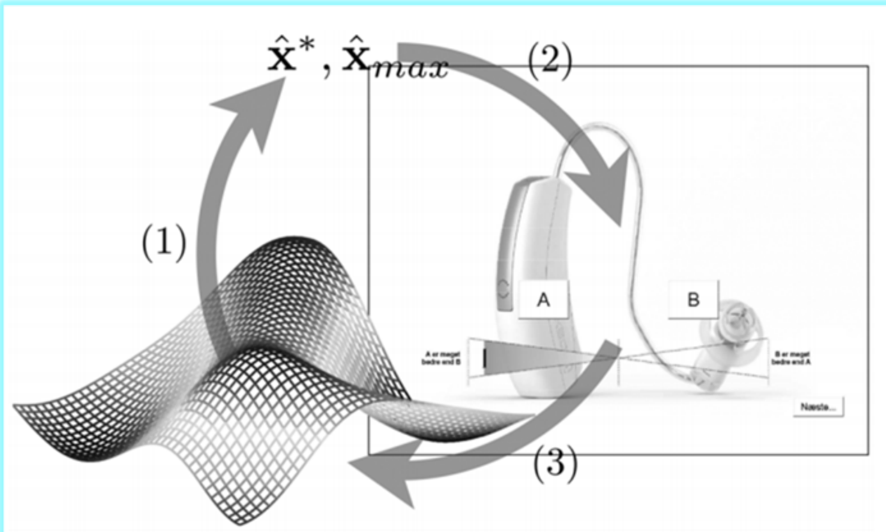

Humans-in-the-loop: Optimization of hearing aids using Bayesian optimization

- Highly personalization needs.
- Dynamic environment and use with different needs.
- Latent, convoluted object functions which are difficult to express though verbal and motor actions.
- Users with disabilities – and often elderly people - provide inconsistent and noisy interactions.

WIDEX®
HIGH DEFINITION HEARING

J.B.B. Nielsen, J. Nielsen, J. Larsen: Perception-based Personalization of Hearing Aids using Gaussian Processes and Active Learning, IEEE Transactions on Audio, Speech, and Language Processing, IEEE, vol. 23, no. 1, pp. 162–173, 2015.

How do we move ahead?

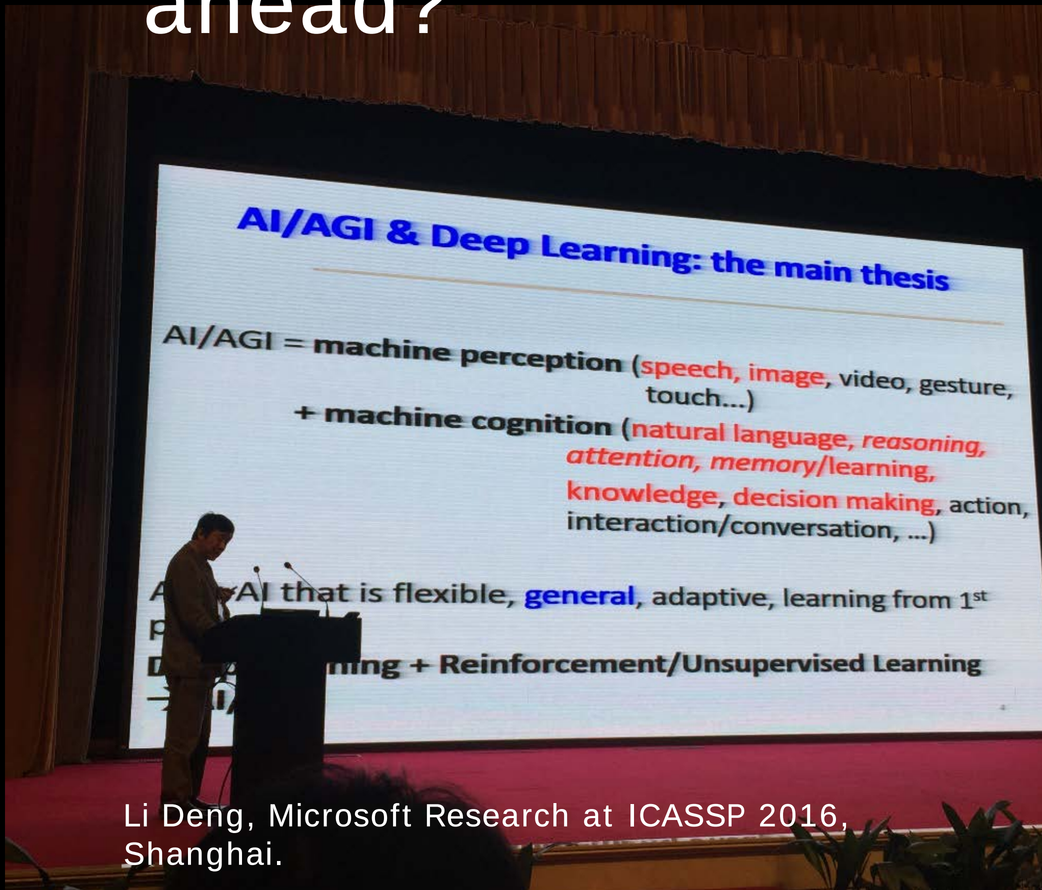

Li Deng, Microsoft Research at ICASSP 2016, Shanghai.

Machine Learning, AI & No Free Lunch

- Four key ingredients for ML towards AI
  1. Lots & lots of data
  2. Very flexible models
  3. Enough computing power
  4. Powerful priors that can defeat the curse of dimensionality

Geoff Hinton, Yoshua Bengio & Yann LeCun, Deep Learning Tutorial, NIPS 2015, Montreal.

Computational rationality: A converging paradigm for intelligence in brains, minds, and machines

Samuel J. Gershman,[1*] Eric J. Horvitz,[2*] Joshua B. Tenenbaum[3*]

After growing up together, and mostly growing apart in the second half of the 20th century, the fields of artificial intelligence (AI), cognitive science, and neuroscience are reconverging on a shared view of the computational foundations of intelligence that promotes valuable cross-disciplinary exchanges on questions, methods, and results. We chart advances over the past several decades that address challenges of perception and action under uncertainty through the lens of computation. Advances include the development of representations and inferential procedures for large-scale probabilistic inference and machinery for enabling reflection and decisions about tradeoffs in effort, precision, and timeliness of computations. These tools are deployed toward the goal of computational rationality: identifying decisions with highest expected utility, while taking into consideration the costs of computation in complex real-world problems in which most relevant calculations can only be approximated. We highlight key concepts with examples that show the potential for interchange between computer science, cognitive science, and neuroscience.

What defines simple and complex problems and how do we solve them them?

passive

active and autonomous

exploration and summarization

prediction

continuous learning
reflection
pro-activeness
engagement
experimentation
creativity

Cognitive systems - a vision for the future: beyond human capabilities

An artificial cognitive system is the ***ultimate learning*** and thinking machine with ability to operate in ***open-ended environments*** with ***natural interaction*** with humans and other artificial cognitive systems and plays key role in the transformational society in order to achieve augmented ***capabilities*** *beyond* human and existing machines.

Jan Larsen, Cognitive Systems Tutorial, MLSP2008, Cancun, Mexico, Oct. 2008.