BIGDATA@CHALMERS
Hosted by INFORMATION AND COMMUNICATION TECHNOLOGY A CHALMERS AREA OF ADVANCE

AI: The New Electricity

Devdatt Dubhashi
Computer Science and Engineering
Chalmers
Machine Intelligence Sweden AB

AI: the New Electricity

"AI is the new electricity.
Just as electricity transformed industry
after industry 100 years ago,
I think AI will do the same."

Andrew Ng, Stanford, Baidu, Coursera

AI

- “I believe that at the end of the century the use of words and general educated opinion will have altered so much that one will be able to speak of machines thinking without expecting to be contradicted.”
  — Alan Turing, *Computing Machinery and Intelligence* (1950)

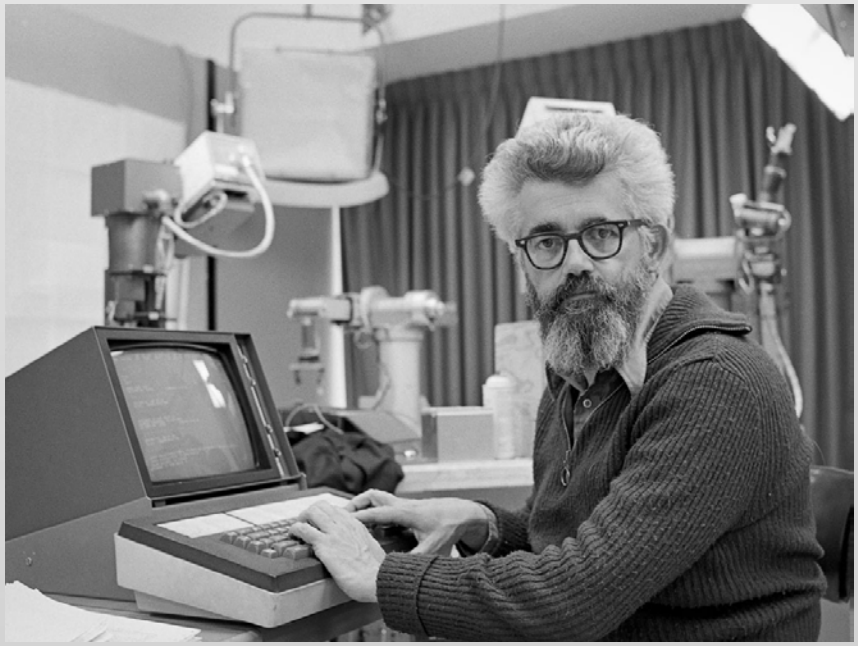

Every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it. An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves. We think that a significant advance can be made in one or more of these problems if a carefully selected group of scientists work on it together for a summer.

John McCarthy,
Dartmouth Workshop 1956

GOFAI (1960-90)

- Knowledge Representation
- First Order Logic and Theorem Proving

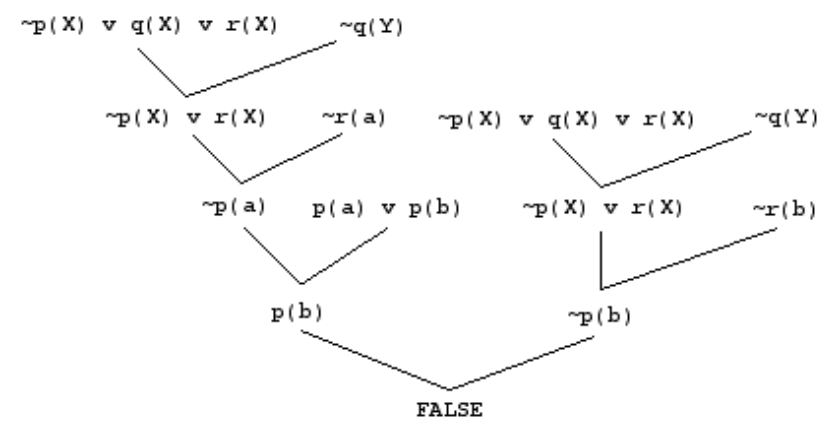

- at(restaurant,Alice)
- at(restaurant,Bob)
- at(restaurant,Carol)
- works_at(restaurant,Carol)
- has_job(restaurant,waitress,Carol)
- orders(Bob,pizza)
- orders(Alice,sushi)
- forall X,Y,Z. orders(X,Y) and has_job(restaurant,waitress,Z) -> serves(Z,X,Y)
- serves(Carol,Bob,pizza) serves(Carol,Alice,sushi)

Statistical Machine Learning 1990-

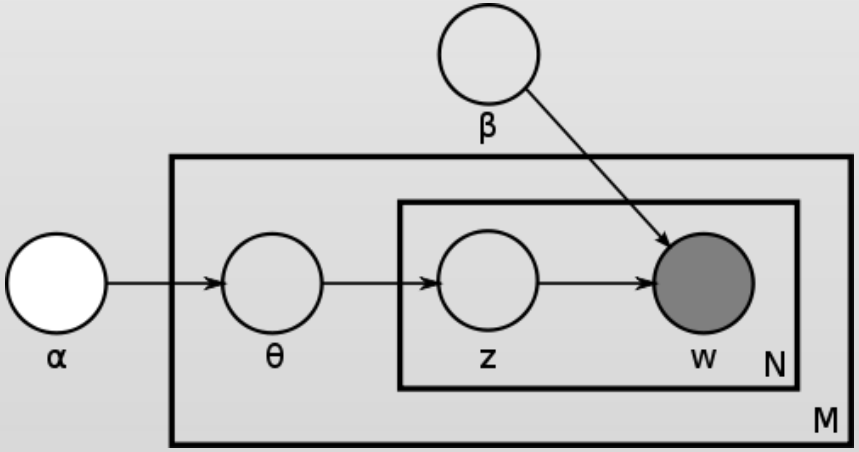

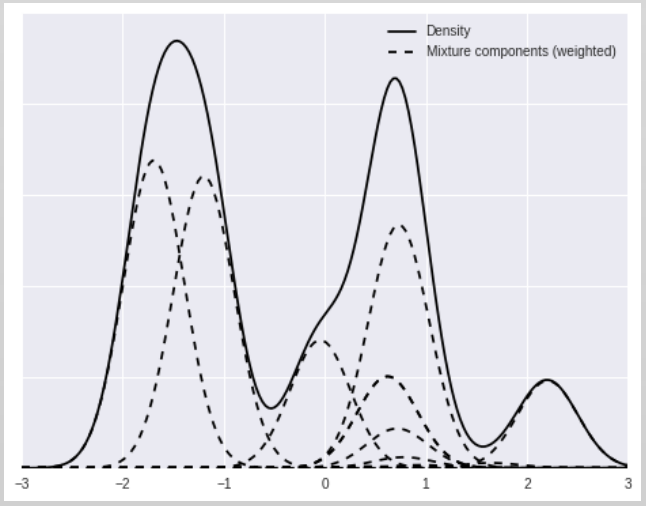

- Probabilistic models
- Statistical learning and inference
- Data driven (no hardcoded rules)

Deep Learning 2005-

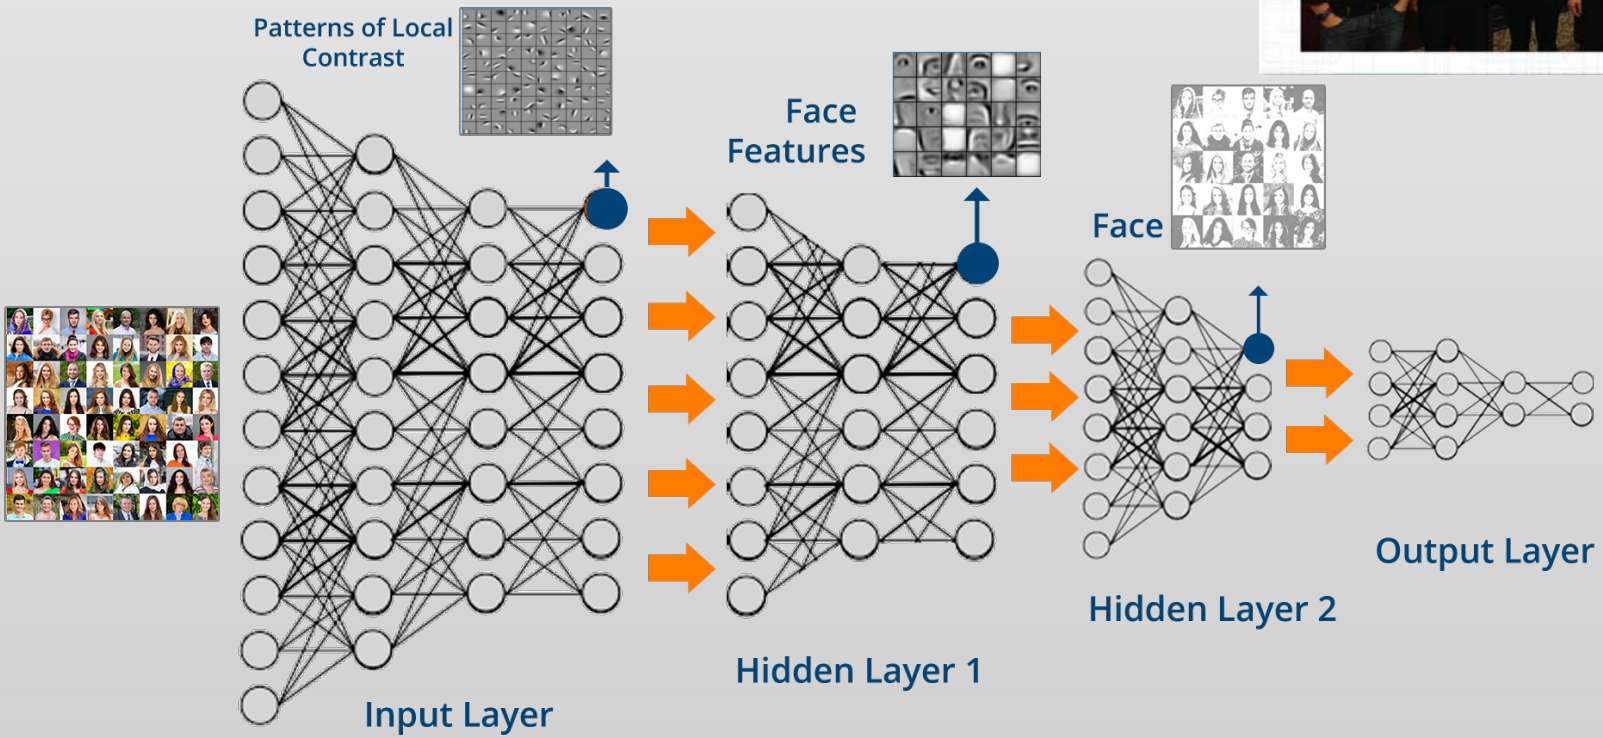

Why Now?
Convergence of Technologies

- Data sensing, acquisition revolution
- Rapid increase in computing power
- Novel algorithms
- Software frameworks

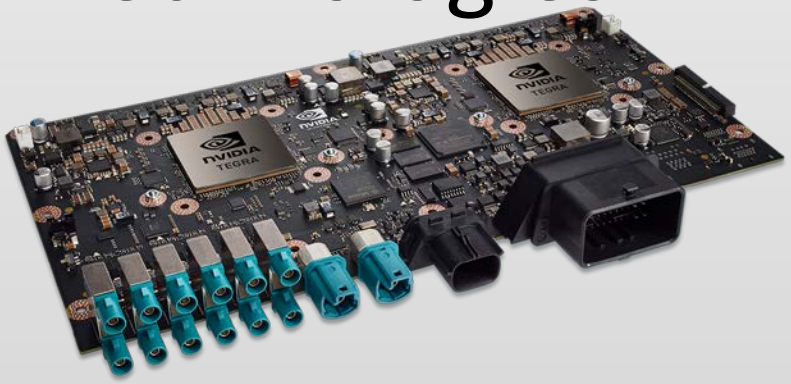

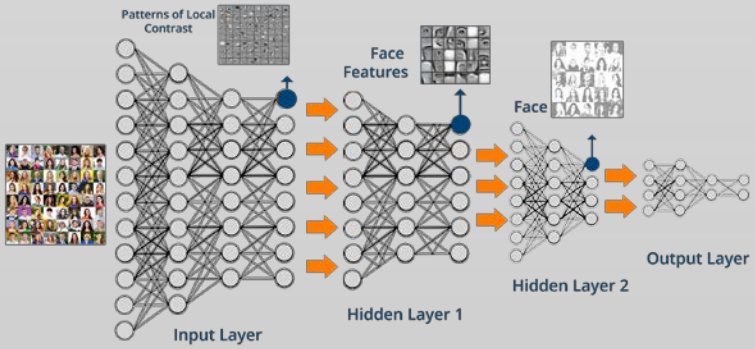

Yann LeCun, NIPS 2016

- "Pure" Reinforcement Learning (cherry)
  - The machine predicts a scalar reward given once in a while.
  - **A few bits for some samples**
- Supervised Learning (icing)
  - The machine predicts a category or a few numbers for each input
  - Predicting human-supplied data
  - **10→10,000 bits per sample**
- Unsupervised/Predictive Learning (cake)
  - The machine predicts any part of its input for any observed part.
  - Predicts future frames in videos
  - **Millions of bits per sample**

- (Yes, I know, this picture is slightly offensive to RL folks. But I'll make it up)


おはようございます。
ジョーンズさん
元気ですか。
GOOD MORNING, MR. JONES. HOW ARE YOU TODAY?

Data-driven Secure Business Intelligence

David Sands, Peter Damaschke, Devdatt Dubhashi, Gerardo Schneider, Andrei Sabelfeld

Recorded Future
CREATING AN INSIGHTFUL WORLD

FINDWISE
SEARCH DRIVEN SOLUTIONS

Inside the Web Intelligence Machine

Recorded Future

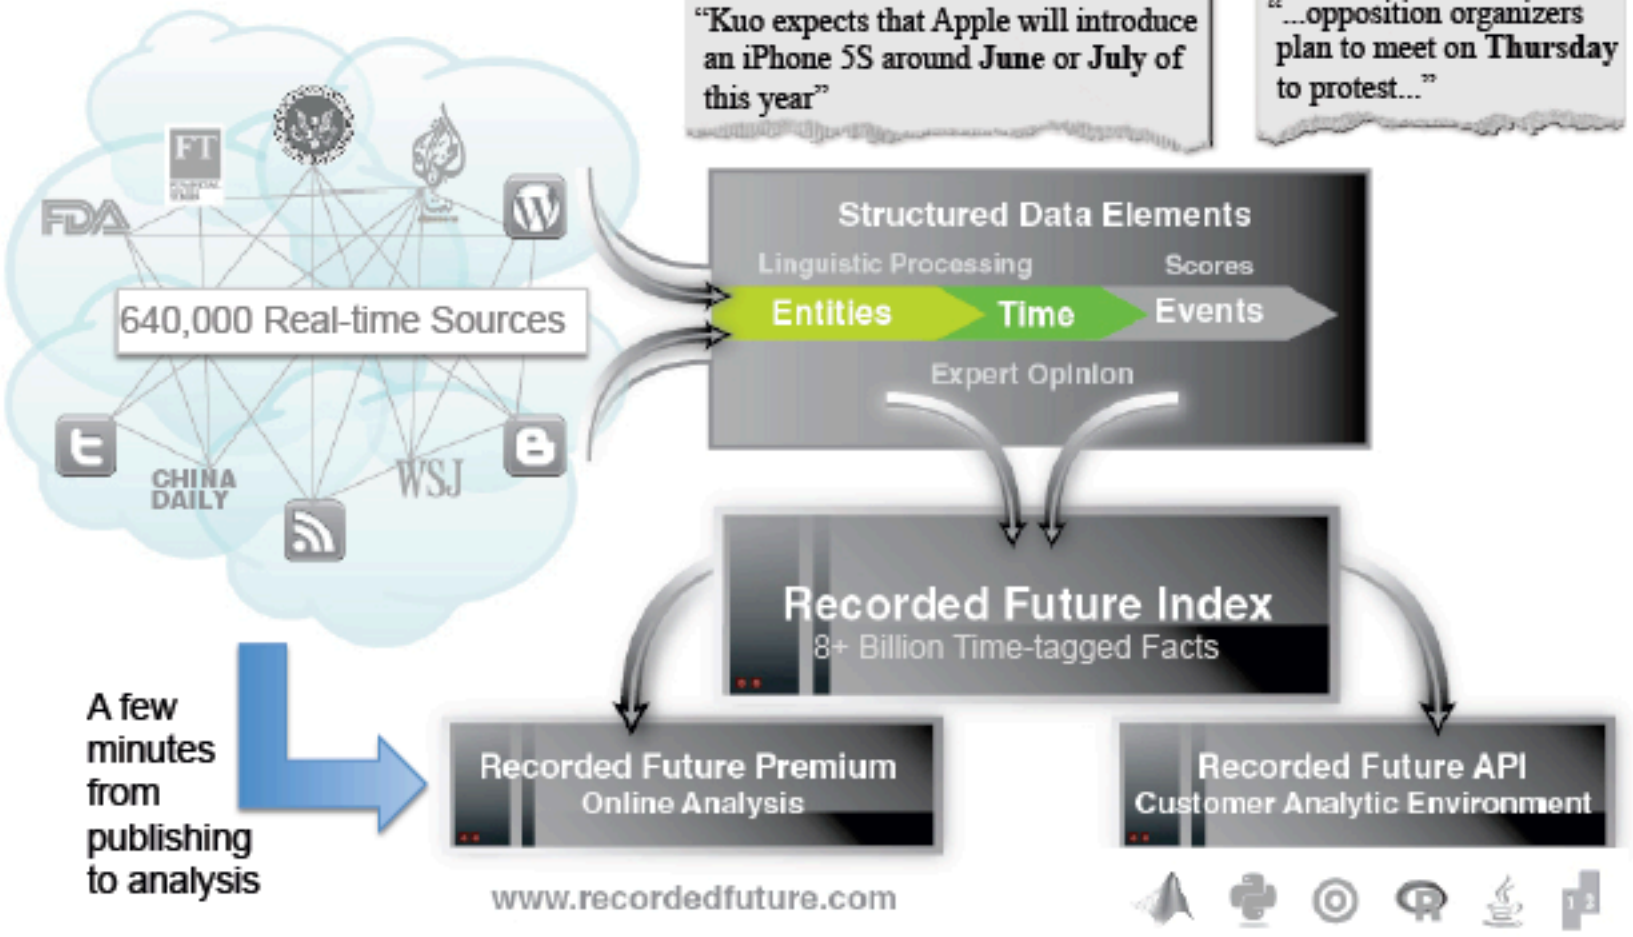

Google Translate

Input sentence:	*Translation (PBMT):*	*Translation (GNMT):*	*Translation (human):*
李克強此行將啟動中加總理年度對話機制，與加拿大總理杜魯多舉行兩國總理首次年度對話。	Li Keqiang premier added this line to start the annual dialogue mechanism with the Canadian Prime Minister Trudeau two prime ministers held its first annual session.	Li Keqiang will start the annual dialogue mechanism with Prime Minister Trudeau of Canada and hold the first annual dialogue between the two premiers.	Li Keqiang will initiate the annual dialogue mechanism between premiers of China and Canada during this visit, and hold the first annual dialogue with Premier Trudeau of Canada.

reduce translation errors across its Google Translate service by between 55 percent and 85 percent

AI Revolution in NLP

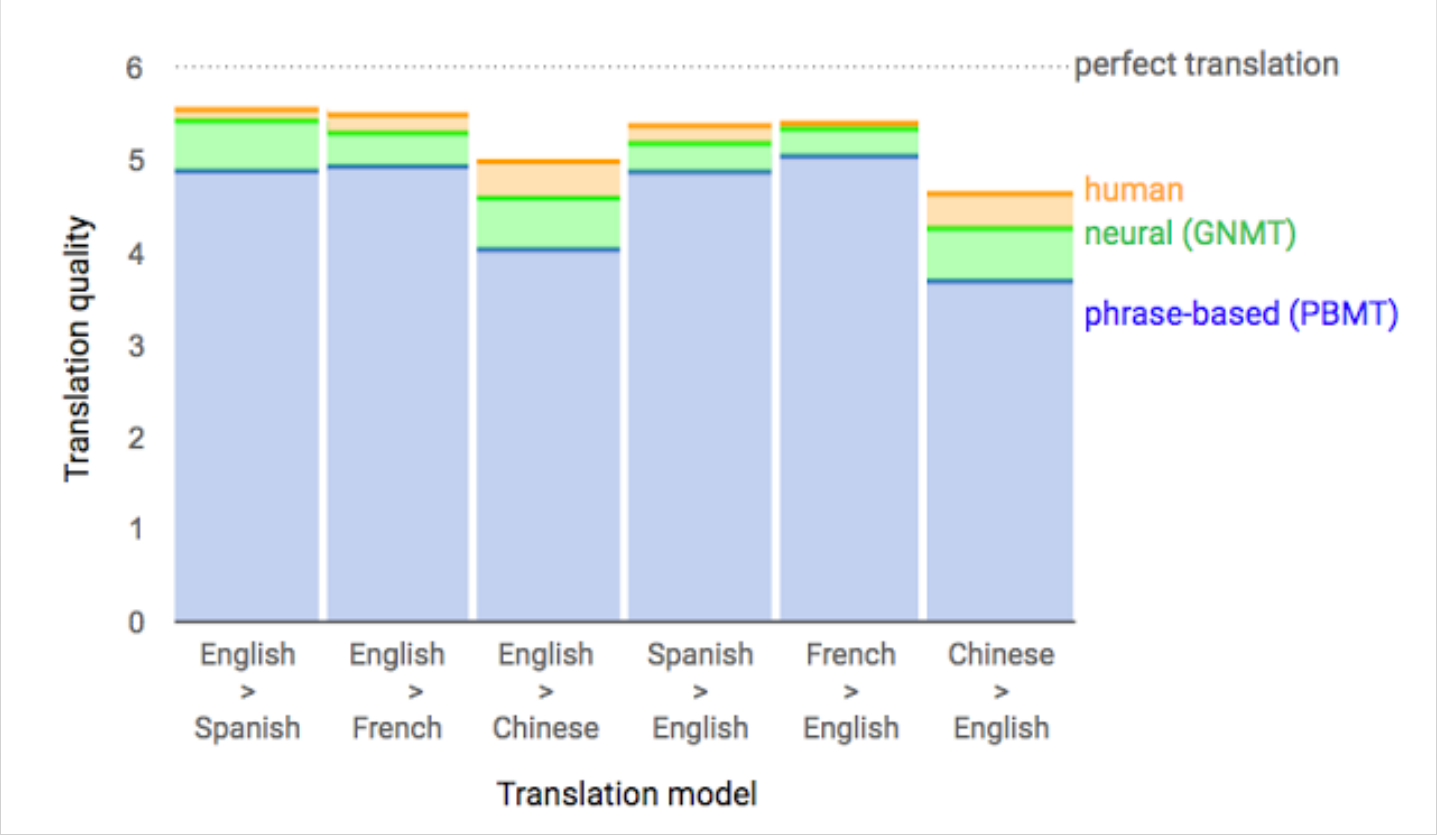

Supervised Learning

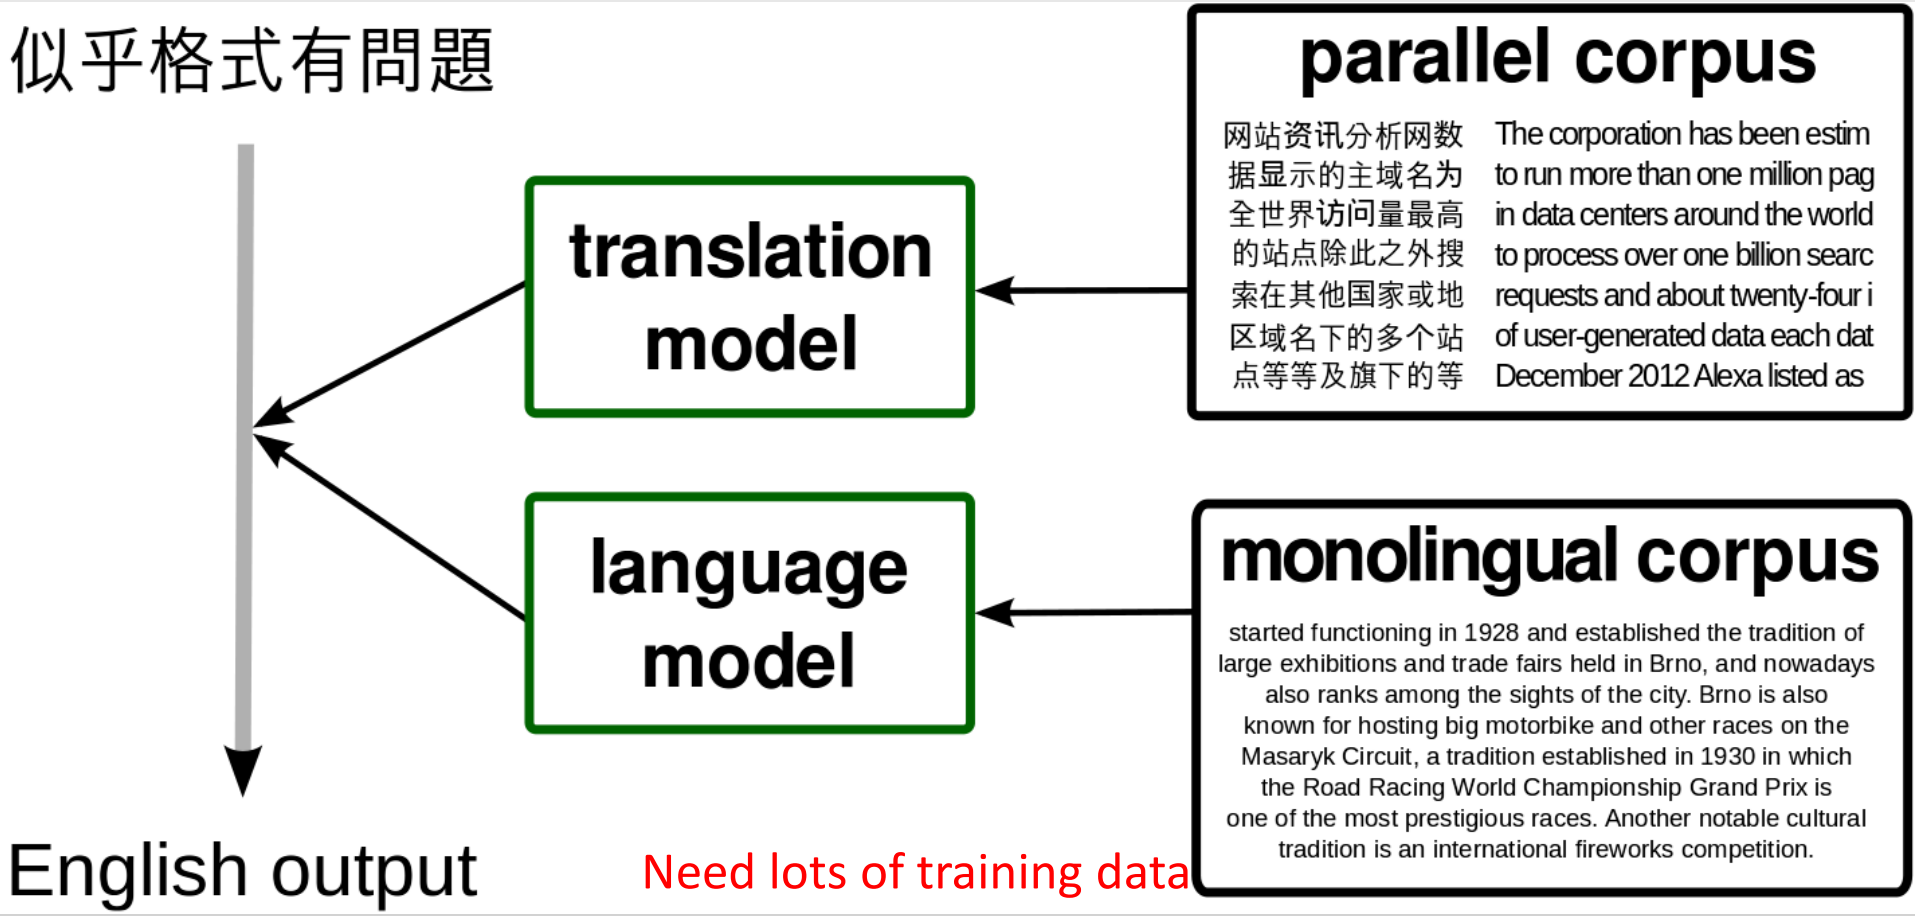

Need lots of training data

- EU Parliament documents in multiple languages
- Bibles in multiple languages

Word Embeddings

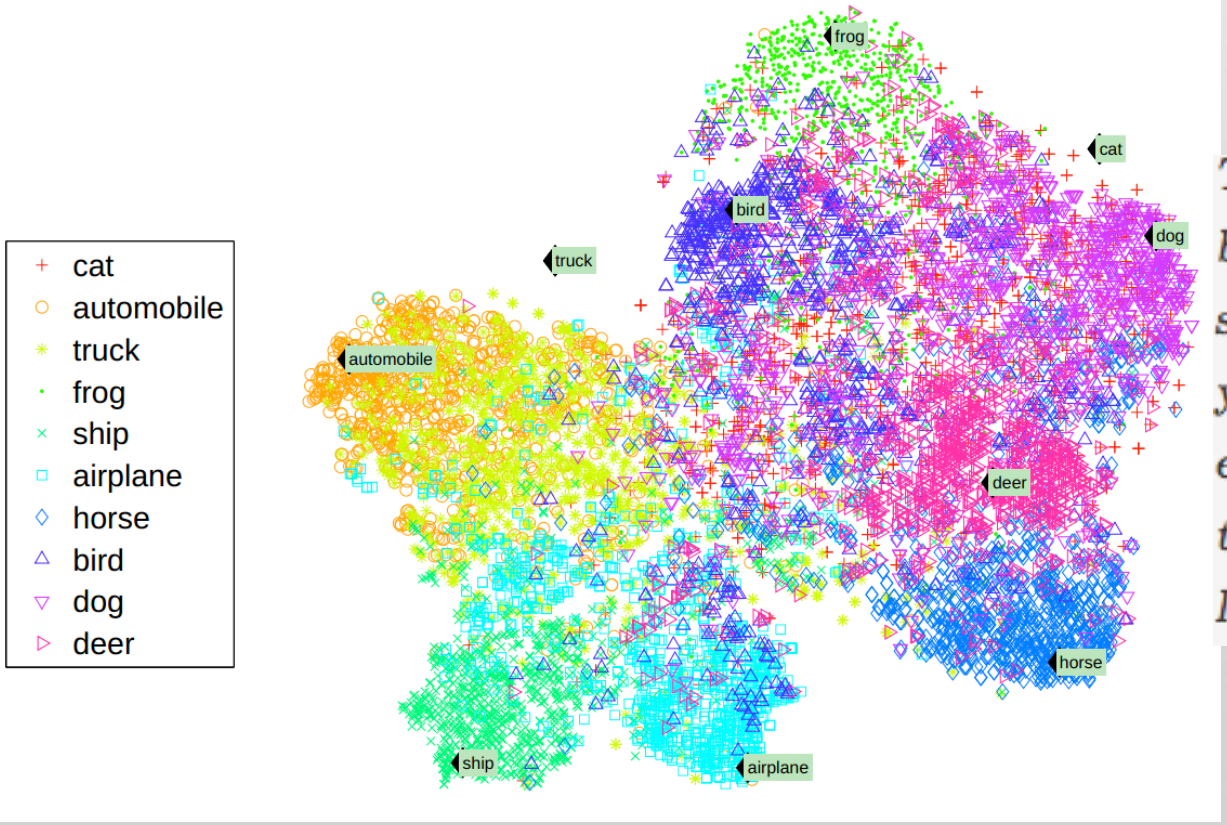

$W : \text{words} \rightarrow \mathbb{R}^n$

The use of word representation
become a key "secret sauce"
success of many NLP systems i
years, across tasks including
entity recognition, part-o
tagging, parsing, and seman
labeling. (Luong et al. (2013))

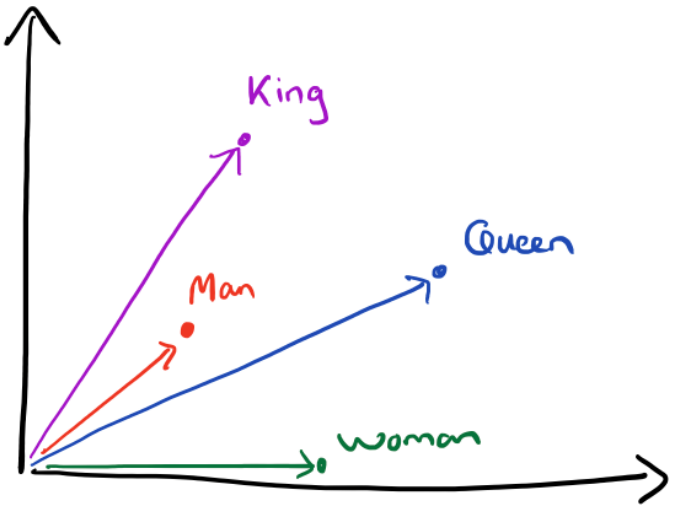

Word Vectors

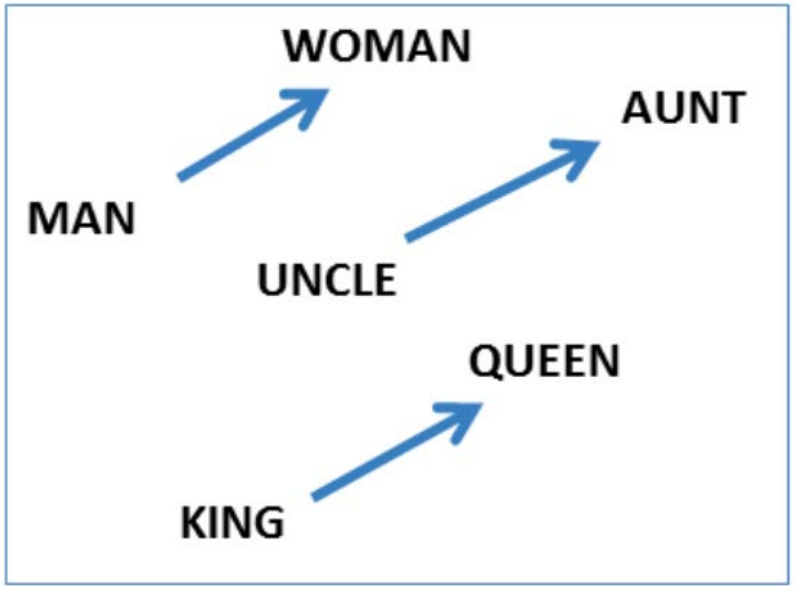

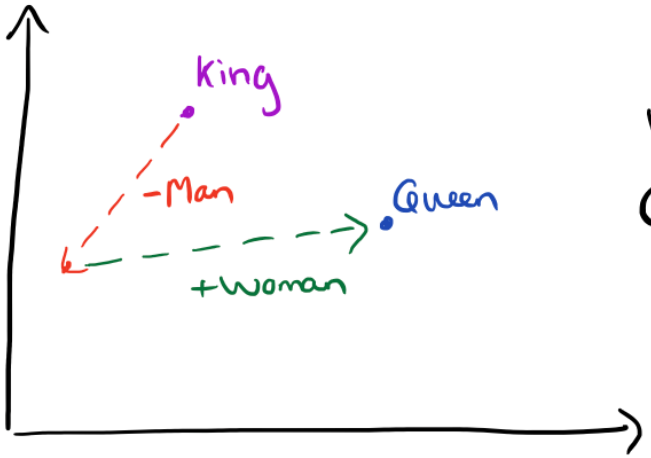

Vector Composition

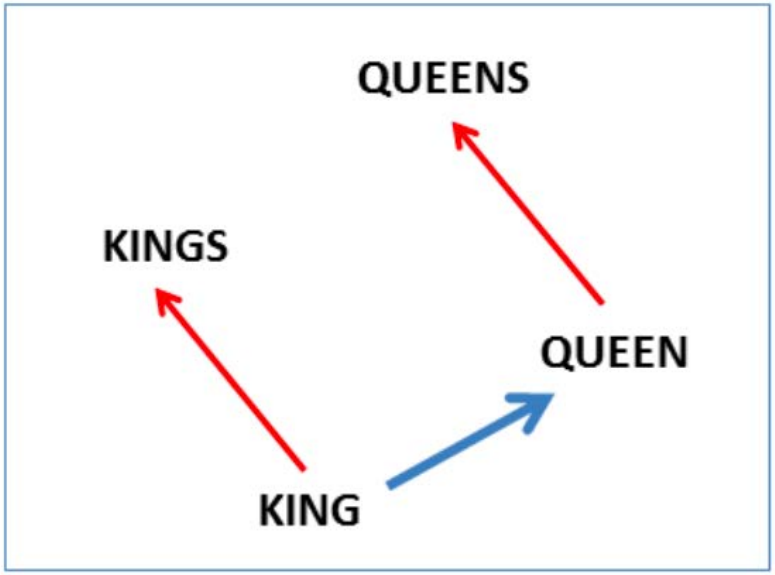

Word senses and Machine Translation

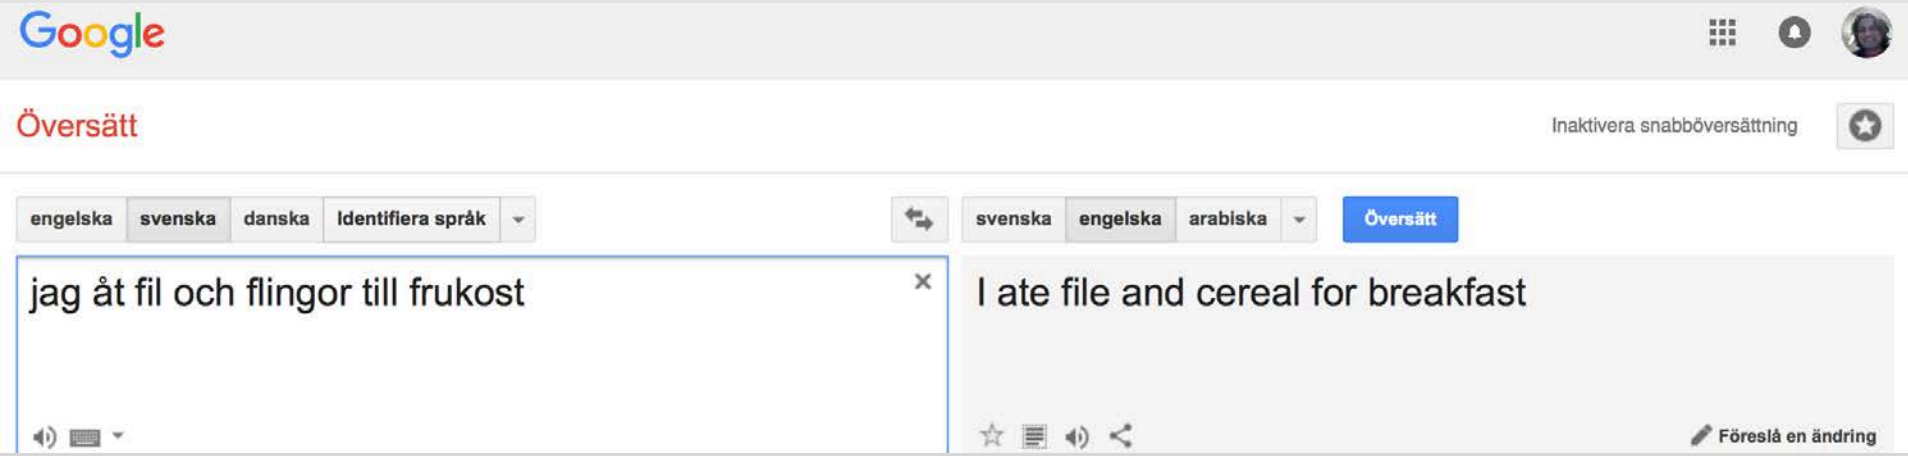


Geology
Name the three types of rock.
1. Classic
2. Punk
3. Hard

Unsupervised Learning of Word senses

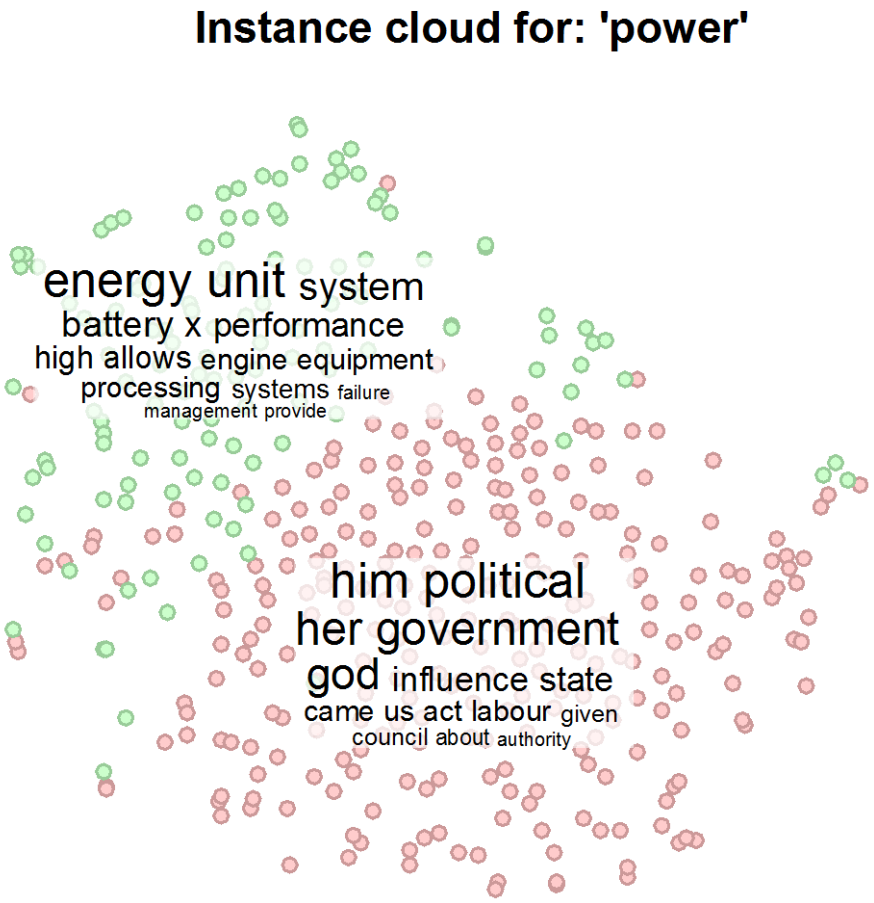

Learn the different senses of a word from raw text without any training data

Document summarization

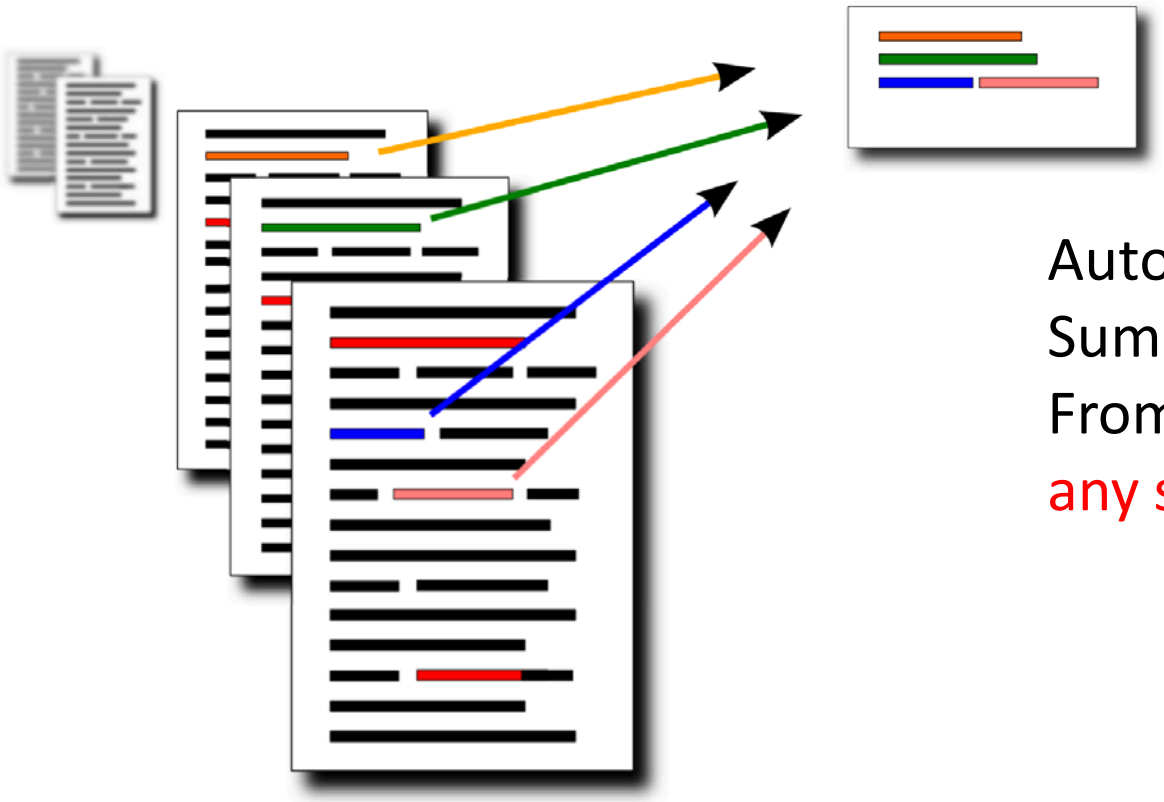

Automatically extract a Summary of documents From raw text without any supervision.

Document Summarization

CHALMERS
UNIVERSITY OF TECHNOLOGY
LAB | RESEARCH GR

FINDWISE
SEARCH DRIVEN SOLUTION

Summaries

(Approx. 40 words)

[Multiple Kernel Learning]

The report said Andreas Lubitz repeatedly set the plane for an unauthorised descent earlier that day. He had locked the flight captain out of the cockpit. Five minutes on the Duesseldorf-Barcelona flight 07:21:10 - Plane told to descend to 21,000ft

[TextRank]

The co-pilot of the Germanwings plane that

Original Text

ness Tech Science Magazine Entertainment & Arts Health Pictures World selected Africa Asia Australia Europe selected Latin America Middle East US & Canada [**Germanwings crash: Co-pilot Lubitz 'practised rapid descent'**] 21 minutes ago From the section Europe [**Germanwings co-pilot Andreas Lubitz is known to have suffered depression in the past**] [**Alps plane crash**] What drives people to murder-suicide? The victims of the Germanwings plane crash Germanwings: Unanswered questions Flight 4U 9525: The final 30 minutes [[**The co-pilot of the Germanwings plane that crashed in the French Alps in March appears to have practised a rapid descent on a previous flight, a report by French investigators says.**]] [**The report said Andreas Lubitz repeatedly set the plane for an unauthorised descent earlier that day.**] Lubitz is suspected of deliberately crashing the Airbus 320, killing all 150 people on board. [[**He had locked the flight captain out of the cockpit.**]] The plane had

Reinforcement Learning

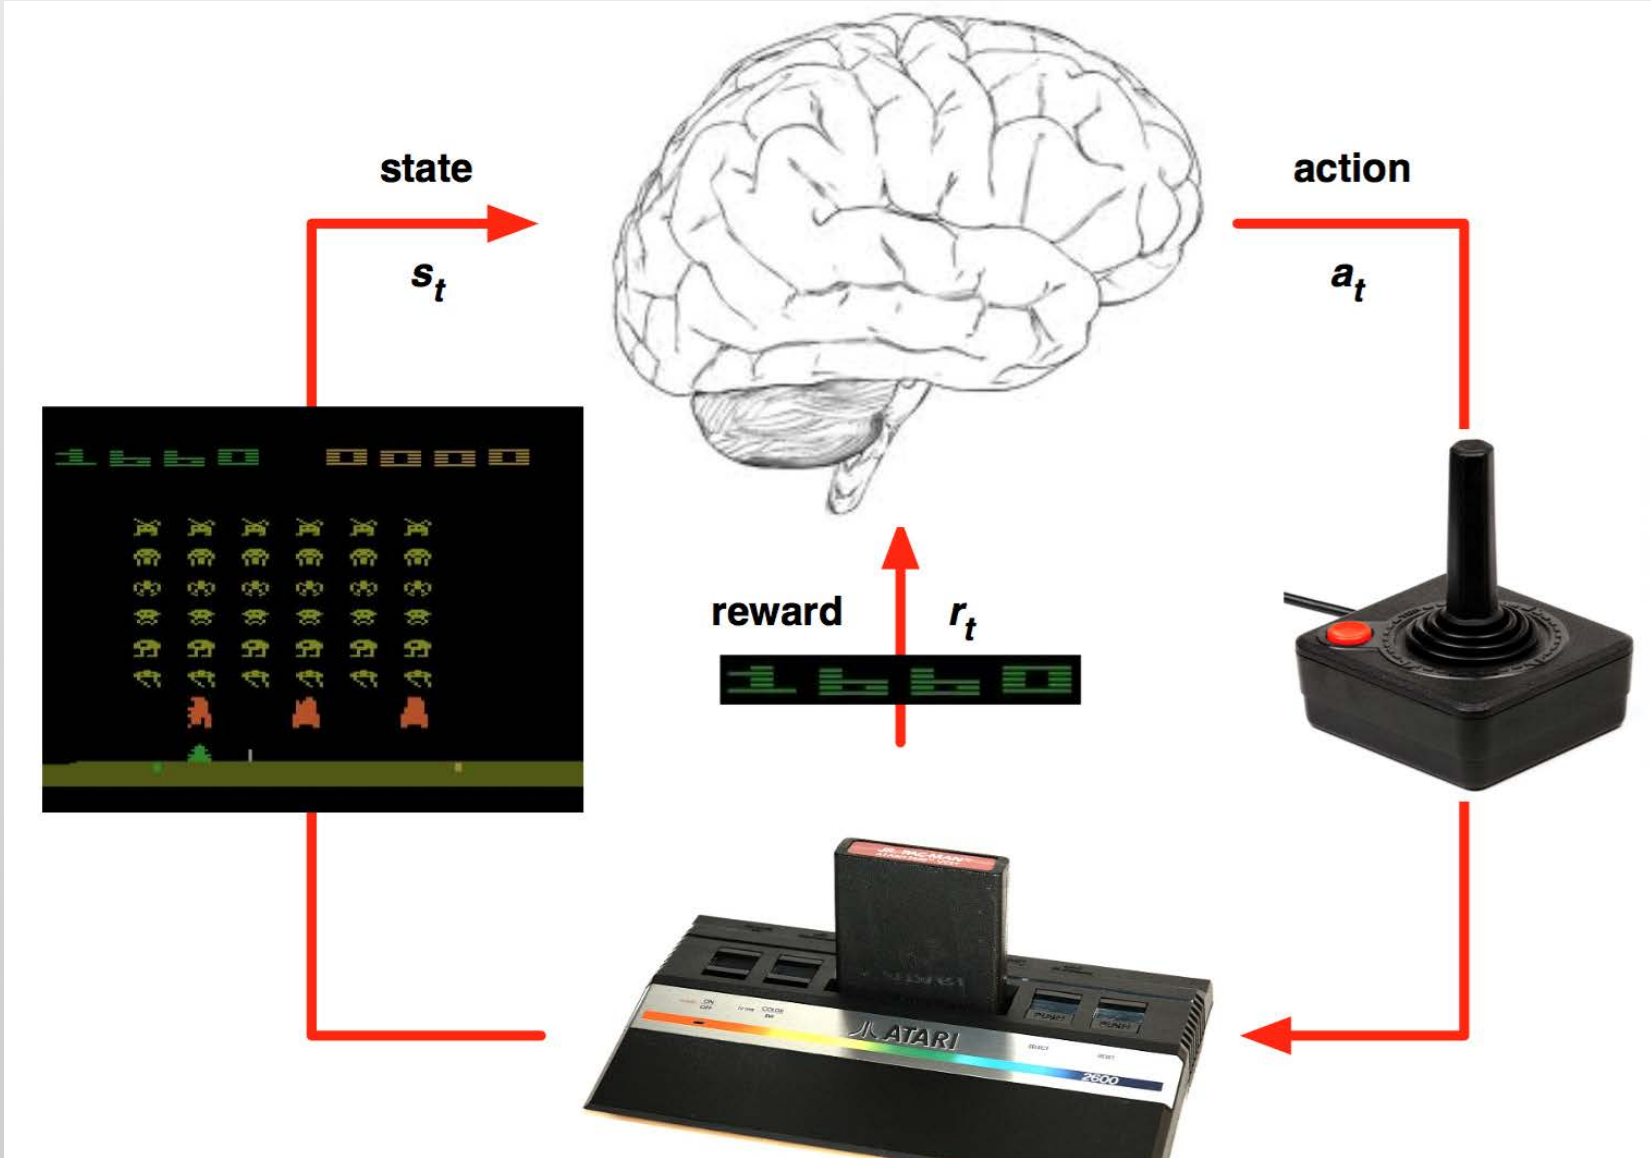

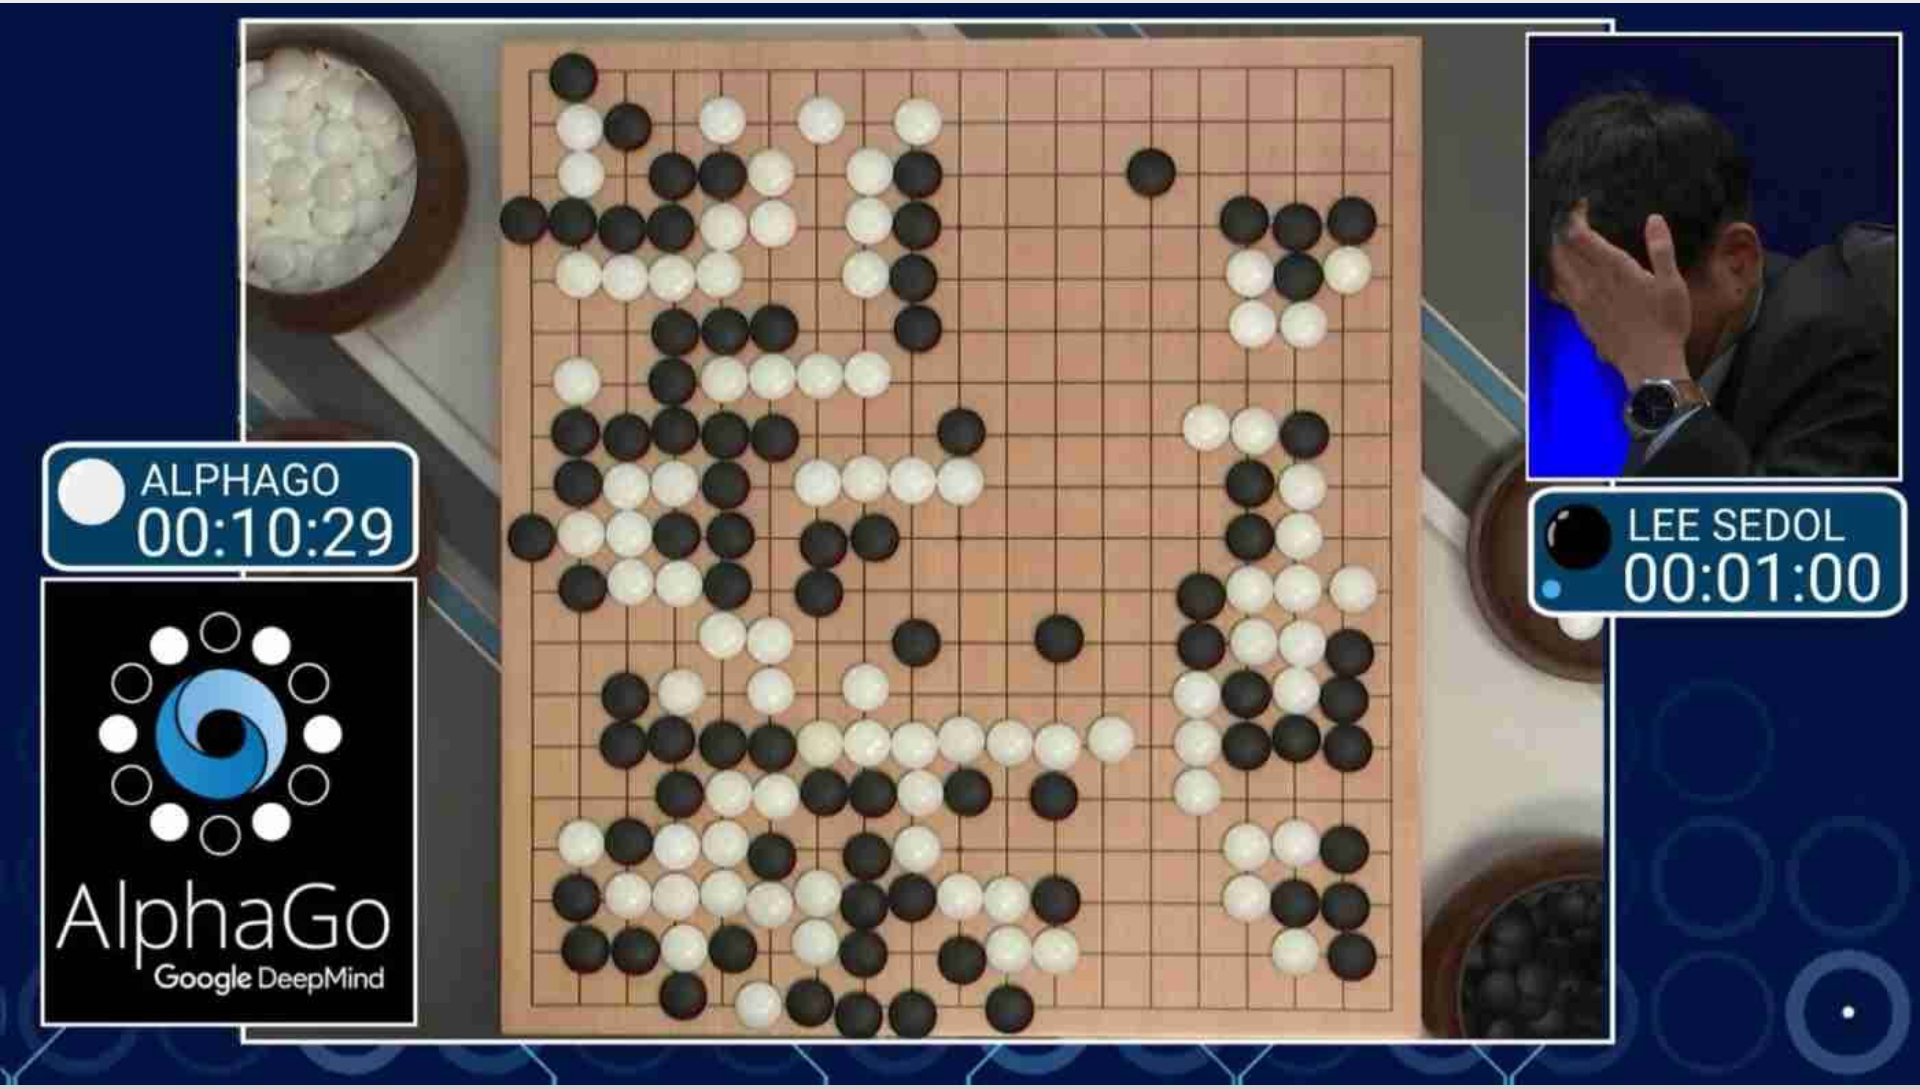
ALPHAGO
00:10:29
AlphaGo
Google DeepMind
LEE SEDOL
00:01:00

AlphaGoZero: AI *Tabula Rasa*

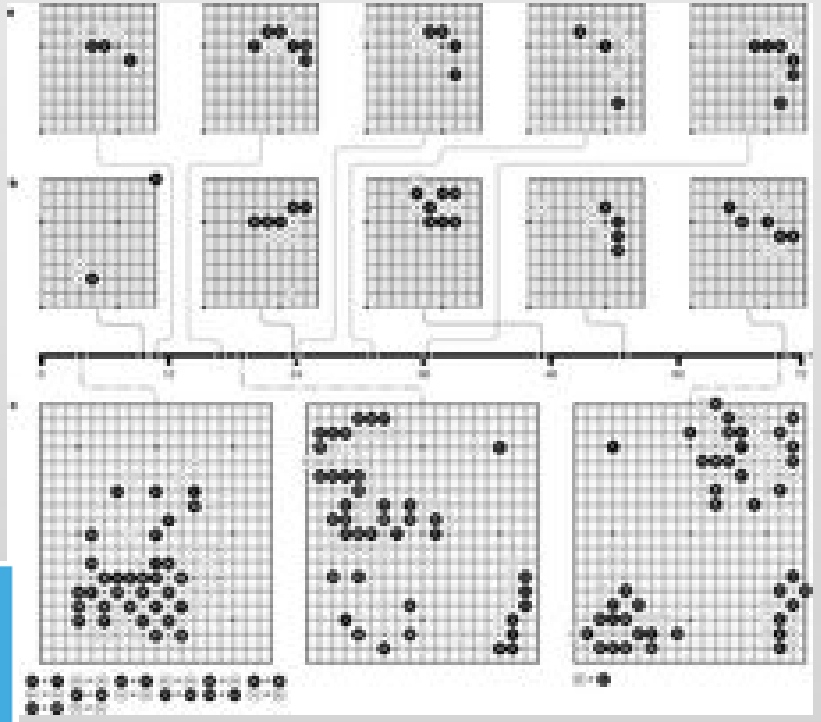

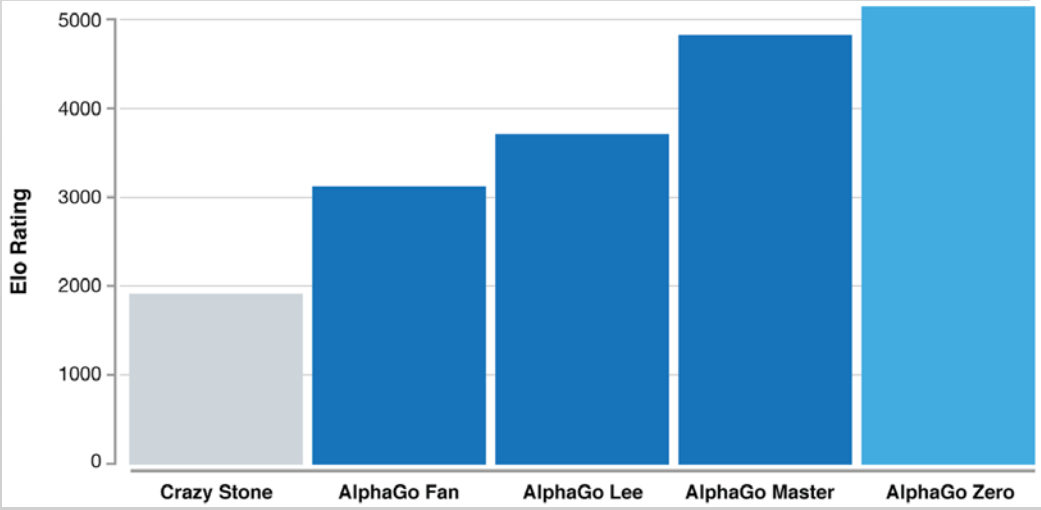

Trained from scratch without any Human input only for 36 hours and beat the previous version 100-0!

Learning to Communicate

- Agents interact sending messages to solve a common task.
- Invent a language grounded in real world objects

E, Jorge, M. Kageback and E. Gustavsson, “Learning to play *Guess Who?* And inventing a Grounded Language as a Consequence”, NIPS Deep RL Workshop (2016)

AI: NEW ELECTRICITY

Electricity and AI as General Purpose Technologies

- Wide scope for improvement and elaboration
- Application across a wide range of uses
- Potential for use in a wide variety of products and processes
- Strong complementarities with existing or potential new technologies

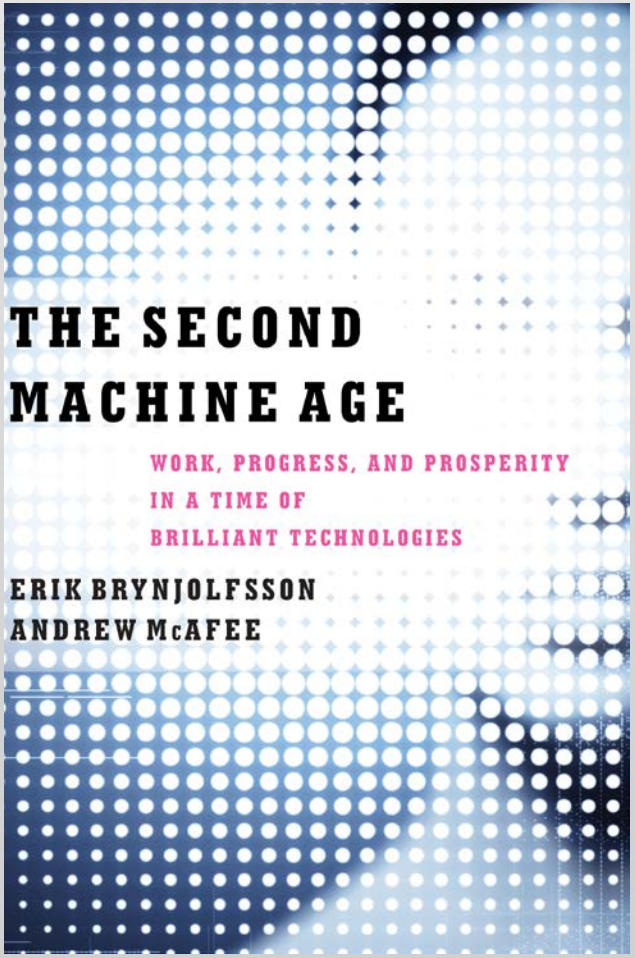

Easy to Use and Improve


“Electricity , communication, manufacturing. I think we are now in that phase where AI technology has advanced to the point where we see a clear path for it to transform multiple industries.”
AI

Data is the New Oil!

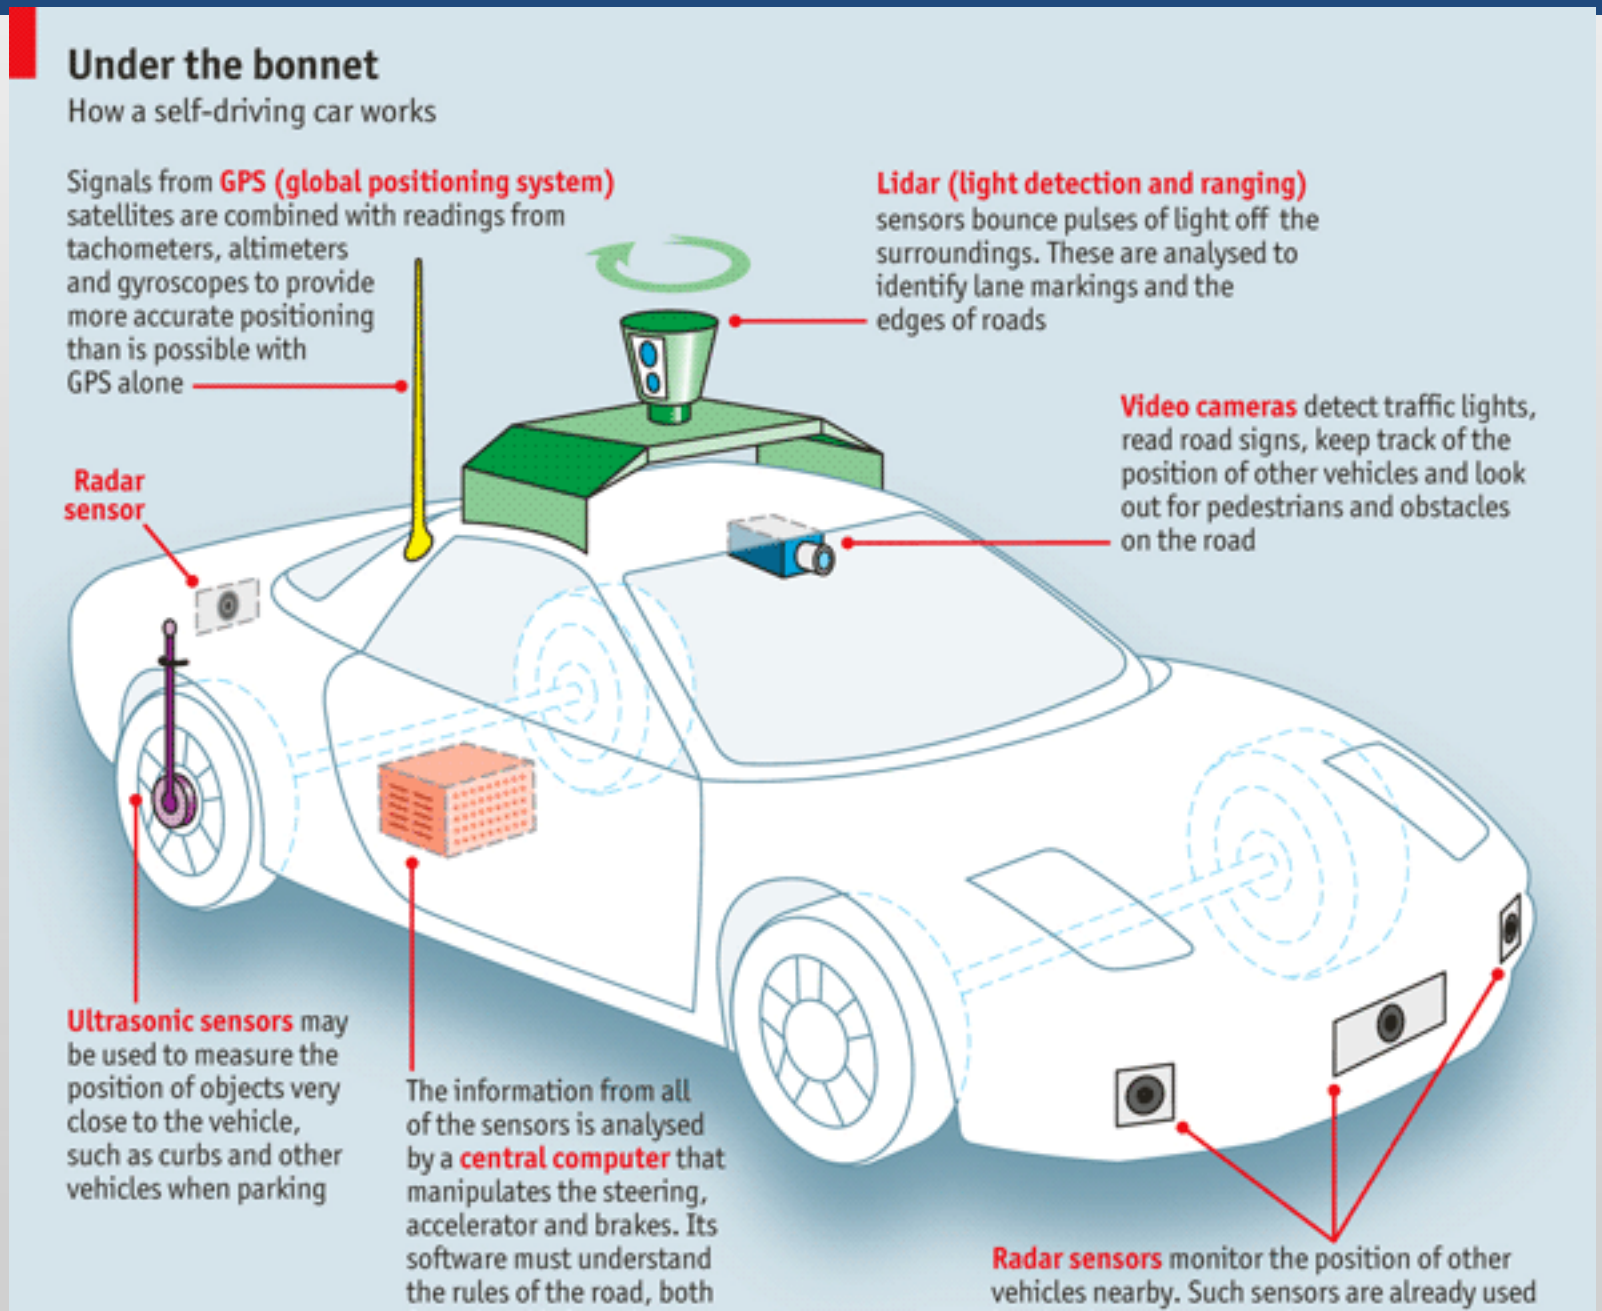

MACHINE INTELLIGENCE 3.0

ENTERPRISE INTELLIGENCE

VISUAL
Orbital Insight · planet · clarifai · DEEPVISION · cortica · Sagacion · SPACE_KNOW · Captricity · netra · deepomatic

AUDIO
Gridspace · TalkIQ · nexidia · twilio · CAPIO · Expect Labs · Clover · Mobvoi · Qurious.AI · popUP archive

SENSOR
PREDIX · C3IoT · MAANA · Sentenai · PLANET OS · UPTAKE · IMUBIT · Preferred Networks · thingworx · KONUX · Alluvium

INTERNAL DATA
PRIMER · IBM WATSON · Cycorp · Palantir · ARIMO · Alation · sapho · Outlier · Digital Reasoning

MARKET
mattermark · Quid · DataFox · PREMISE · Bottlenose · MOTIVA · enigma · CB INSIGHTS · Tracxn · predata

ENTERPRISE FUNCTIONS

CUSTOMER SUPPORT
DigitalGenius · Kasisto · ELOQUENT · Wise.io · ACTIONIQ · zendesk · Preact · CLARABRIDGE

SALES
collective[i] · 6sense · fusemachines · salesforce · INSIDE SALES.COM · Zensight

MARKETING
MINTIGO · Lattice · RADIUS

SECURITY
CYLANCE · DARKTRACE

RECRUITING
textio · entelo · Wade & Wendy · hi · unitive · SpringRole · GIGSTER · HireVue

X+AI

AUTONOMOUS S… … AGENTS

GROUND NAVIGATION
drive.ai · AdasWorks · ZOOX · MOBILEYE · UBER · Google · TESLA · nuTonomy · Auro Robotics

AERIAL
SKYDIO · SHI · Airware · dji · LILY · DroneDeploy · pilot.ai · SKYCATCH

CLEARPATH · fetch · KINDRED · HARVEST · rethink robotics

Cortana · Allo · facebook · Siri · Replika

PROFESSIONAL
butter.ai · pogo · SKIPFLAG · clara · x.ai · slack · talla · Zoom · sudo

INDUSTRIES

AGRICULTURE
BLUE RIVER · mavrx · tule · TRACE GENOMICS · Pivot Bio · Terravion · AGRI-DATA · Descartes Labs · ubiquity · abundant robotics

EDUCATION
KNEWTON · volley · gradescope · CTI · coursera · UDACITY · altschool

INVESTMENT
Bloomberg · sentient · SENTIUM · KENSHO · alphasense · Dataminr · CEREBELLUM CAPITAL · Quandl

LEGAL
blueJ · BEAGLE · Everlaw · RAVEL · Sseal · ROSS · LEGAL ROBOT

LOGISTICS
NAUTO · Acerta · PRETECKT · Routific · clearmetal · MARBLE · PITSTOP

INDUSTRIES CONT'D

MATERIALS
zymergen · Citrine · Eigen Innovations · SIGHT MACHINE · GINKGO BIOWORKS · nanotronics · CALCULARIO

RETAIL FINANCE
TALA · zestfinance · Lendo · earnest · Affirm · MIRADOR · wealthfront · Betterment

HEALTHCARE

PATIENT
PULSE · CareSkore · ZEPHYR HEALTH · IBM Watson Health · Onc Data · SENTRIAN · Atomwise · Numerate

IMAGE
BUTTERFLY · 3SCAN · ARTERYS · enlitic · BAYLABS · imagia · Google DeepMind

BIOLOGICAL
iCarbonX · color · GRAIL · deep genomics · RECURSION · LUMINIST · Numerate · Atomwise · verily · WHOLE BIOME

TECHNOLOGY STACK

AGENT ENABLERS
OCTANE.AI · howdy · Maluuba · KITT.AI · OpenAI Gym · Kasisto · AUTOMAT · semanticmachines

DATA SCIENCE
DOMINO · SPARKBEYOND · rapidminer · kaggle · DataRobot · ŷhat · AYASDI · seldon · yseop · bigml

MACHINE LEARNING
CognitiveScale · GoogleML · context relevant · Cycorp · HyperScience · narrologics · minds.ai · H2O.ai · SCALED INFERENCE · sparkcognition · loop · GEOMETRIC INTELLIGENCE · deepsense.io · reactive · skymind · bonsai

NATURAL LANGUAGE
LYLIEN · LEXALYTICS · Narrative Science · spaCy · LUMINOSO · cortical.io · MonkeyLearn

DEVELOPMENT
SIGOPT · HyperOpt · fuzzyio · pkite · rainforest · lobe · Anodot · Sightful · LAYER 6 · bonsai

DATA CAPTURE
CrowdFlower · diffbot · CrowdAI · import.io · Paxata · DATASIFT · amazon mechanical turk · enigma · WorkFusion · DATALOGUE · TRIFACTA · parsehub

OPEN SOURCE LIBRARIES
Keras · Chainer · CNTK · TensorFlow · Caffe · theano · torch · H2O · DEEPLEARNING4J · DSSTNE · Scikit-learn · AzureML · neon · MXNet · DMTK · Spark · PaddlePaddle · WEKA

HARDWARE
KNUPATH · TENSTORRENT · Cirrascale · NVIDIA · intel nervana · Movidius · tensilica · GoogleTPU · 10^26 Labs · Qualcomm · Cerebras · Isosemi

RESEARCH
OpenAI · maluuba · ELEMENT AI · vicarious · KNOGGIN · Numenta · Kimera Systems · Cogitai

shivonzilis.com/MACHINEINTELLIGENCE · Bloomberg BETA

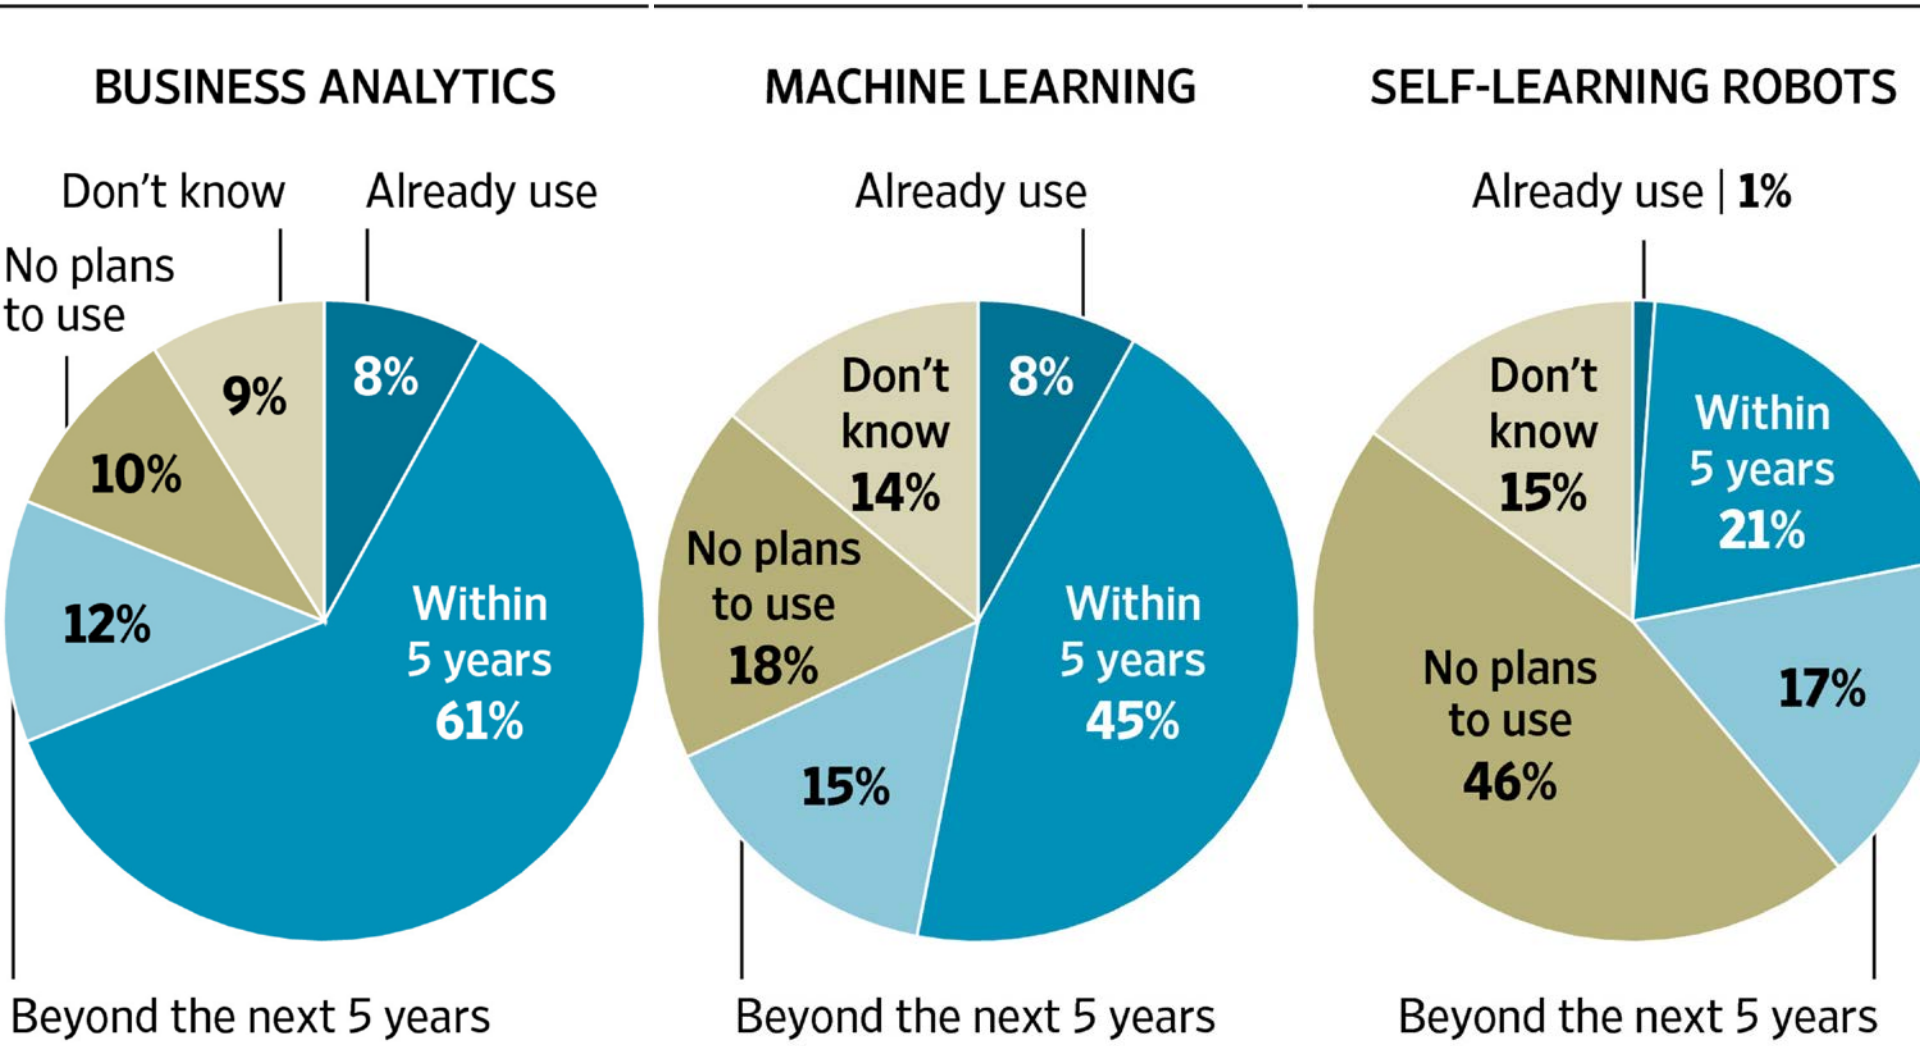
BUSINESS ANALYTICS
Don't know
Already use
No plans to use
8%
9%
10%
12%
Within 5 years 61%
Beyond the next 5 years
MACHINE LEARNING
Already use
Don't know 14%
8%
No plans to use 18%
Within 5 years 45%
15%
Beyond the next 5 years
SELF-LEARNING ROBOTS
Already use | 1%
Don't know 15%
Within 5 years 21%
No plans to use 46%
17%
Beyond the next 5 years

Complementary Technologies

CRISPR/CAS9 Gene Editing for Life Sciences

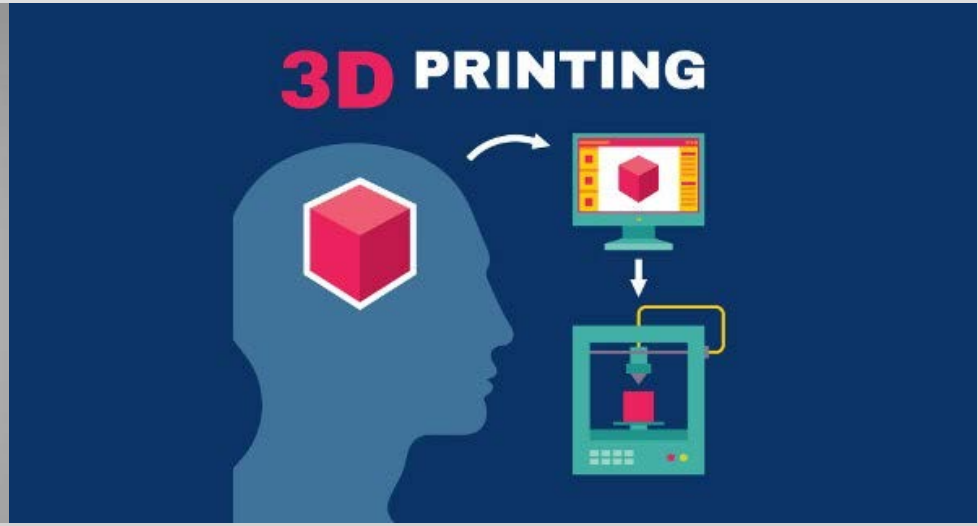

Additive Manufacturing

Music itself is going to become like running water or electricity.

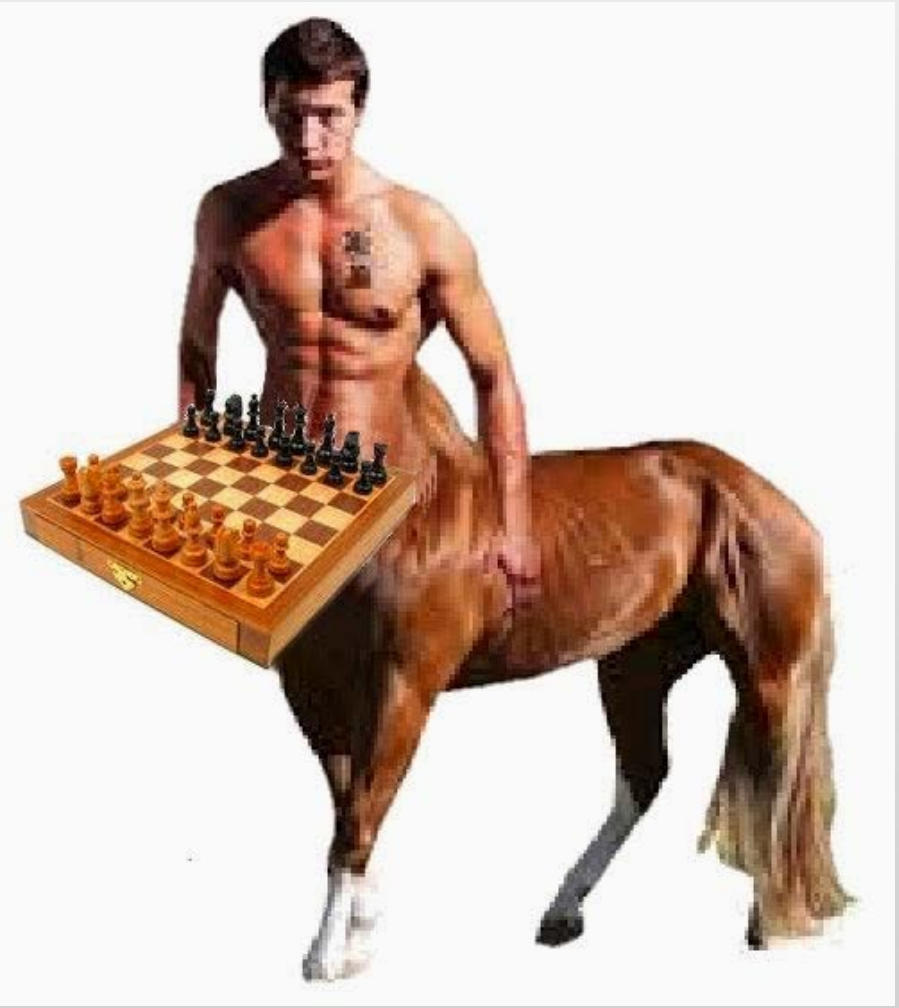

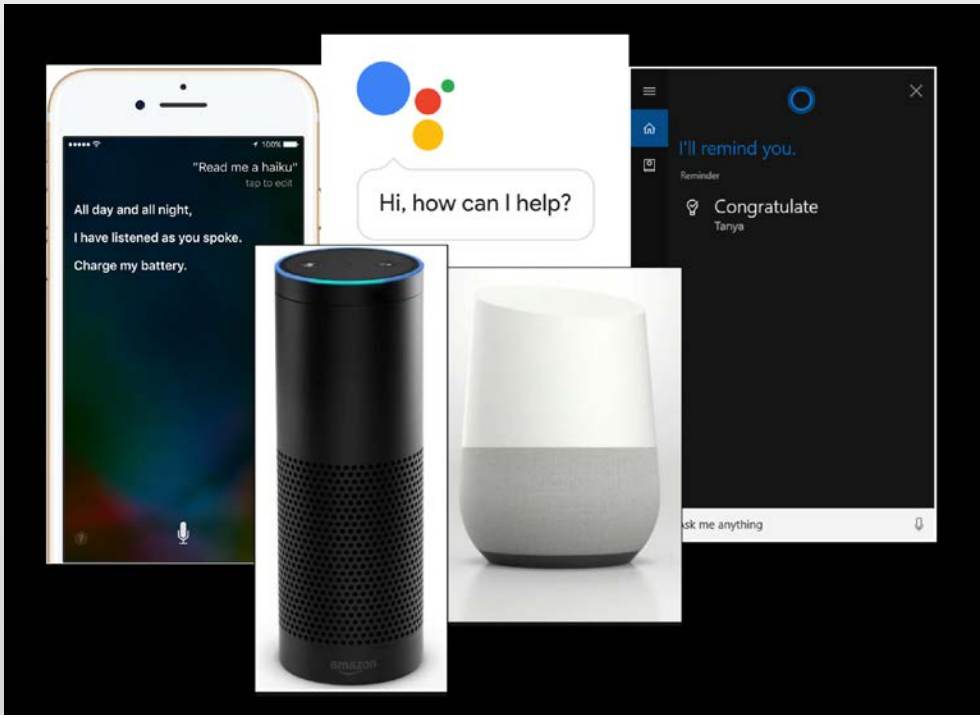
"Read me a haiku"
tap to edit
All day and all night,
I have listened as you spoke.
Charge my battery.
Hi, how can I help?
I'll remind you.
Reminder
Congratulate
Tanya
Ask me anything

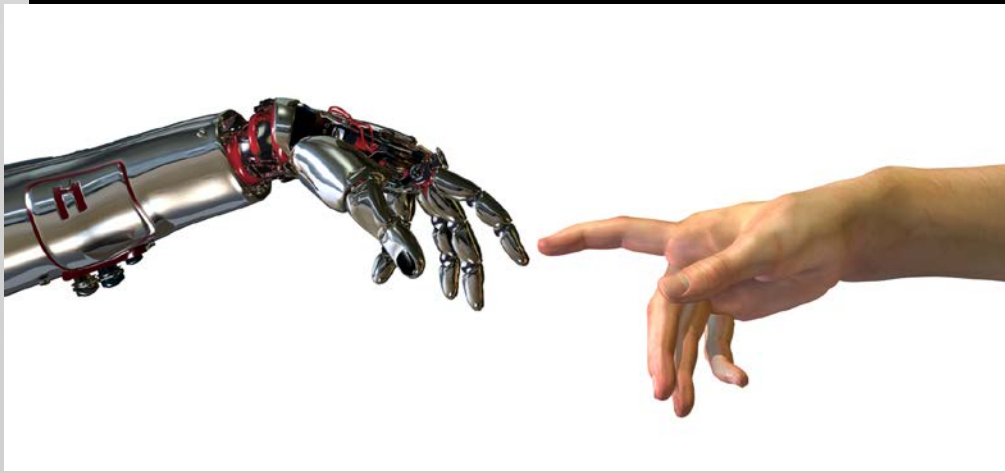

- AI will contribute as much as $15.7 trillion to the world economy by 2030 (PwC)
- $6.6 trillion from increased productivity as businesses automate processes and augment with new AI technology, and $9.1 trillion from consumption side-effects as shoppers snap up personalized and higher-quality goods

BigData@Chalmers Seminars

10 November 2017 - BigData@Chalmers seminar

Artificial Intelligence Starting From a Blank Slate

Speakers: Devdatt Dubhashi, Mikael Kågebäck and others, Department of Computer Science and Engineering, Chalmers

More information >>

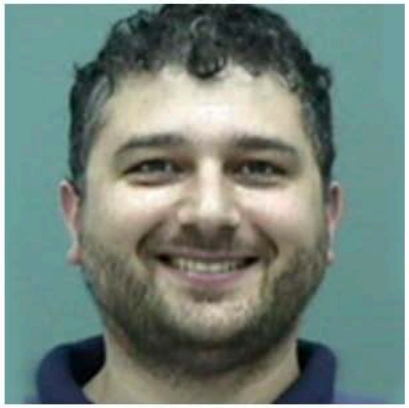

27 October 2017 - BigData@Chalmers seminar

Privacy in the hands of the data analyzer

Speaker: Kobbi Nissim, Georgetown University

More information >>

Data Science: Algorithms

LP 1	LP 2	LP 3	LP 4
Introduction to Data Science	Stochastic processes and Bayesian statistics	DAT340 - Applied machine learning	**DAT345 - Techniques for large-scale data**
Nonlinear optimization	**TDA507 Computational methods in bioinformatics** ***(eller)*** **DAT246 Empirical software engineering**	**TDA206 Discrete optimization** ***(eller)*** **CIU187 Information visualization** ***(eller)*** **TIN093 Algorithms**	**TDA231 Algorithms for Machine Learning and Inference**

LP 1	LP 2	LP 3	LP 4
TIN093 Algorithms	Data science project	Master's thesis	
SSY340 Deep Machine Learning	**TDA251 Algorithms,**		

Data Science: Large Scale Systems

LP 1	LP 2	LP 3	LP 4
Introduction to Data Science	Stochastic processes and Bayesian statistics	DAT340 - Applied machine learning	**DAT345 - Techniques for large-scale data**
Nonlinear optimization	**TDA596 Distributed systems**	**EDA263 Computer Security** ***(eller)*** **TDA297 Distributed systems II**	**TDA231 Algorithms for Machine Learning and Inference**

LP 1	LP 2	LP 3	LP 4
DAT300 Data-driven support for cyber-physical systems	Data science project	Master's thesis	
SSY340 Deep Machine Learning	**DAT295 Autonomous and Cooperative Vehicular Systems**		

Data Science: Optimization

LP 1	LP 2	LP 3	LP 4
Introduction to Data Science	Stochastic processes and Bayesian statistics	DAT340 - Applied machine learning	**DAT345 - Techniques for large-scale data**
Nonlinear optimization	**MVE190 Linear statistical models [MV]**	**TDA206 Discrete optimization** ***(eller)*** **SSY097 - Image analysis [EE]** ***(eller)*** **TMA521 Large-scale optimization [MV]**	**TDA231 Algorithms for Machine Learning and Inference**

LP 1	LP 2	LP 3	LP 4
TMA265 Numerical linear algebra [MV] ***(eller)*** **TMA881 High performance computing [MV]** ***(eller)*** **RRY025 Image processing [SEE]**	Data science project	Master's thesis	
SSY340 Deep Machine Learning	**TDA507 Computational methods in bioinformatics**		

Tillvalskurser

Nedan följer en lista med blandade tillvalskurser på olika institutioner utöver de som nämnts ovan.

LP1

- TMA265 - Numerical linear algebra (Math)
- MVE187 - Computational methods for Bayesian statistics (Math)
- FFR105 - Stochastic optimization algorithms (Physics)
- TIF160 - Humanoid robotics (Physics)
- RRY025 - Image processing (Earth and Space Sciences)
- FFR135 - Artificial neural networks (Physics)

LP2

- DAT246 - Empirical software engineering (CSE)
- TDA507 - Computational methods in bioinformatics (CSE)
- TIF155 - Dynamical systems (Physics)
- FFR120 - Simulation of complex systems (Physics)
- SSY130 - Applied signal processing (E2)
- TDA357 - Databases (CSE) (** First cycle course)

LP3

- TMA521 - Large-scale optimization (Math) (** Given every other year)
- TIF150 - Information theory for complex systems (Physics)
- FFR110 - Computational biology A (Physics)
- SSY097 - Image analysis(S2/Elec)
- VVT105 - Geographical Information Systems (Architecture)

LP4

- TMS016 - Statistical image analysis (Math)
- MVE220 - Financial risk (Math)
- TMS087 - Financial Time Series (Math)
- MVE440 - Statistical Learning for Big Data (Math)
- SSY115 - eHealth (E2)
- SSY315 - Bayesian statistics (E2)


TIME
SPECIAL EDITION
Artificial Intelligence
The Future of Humankind