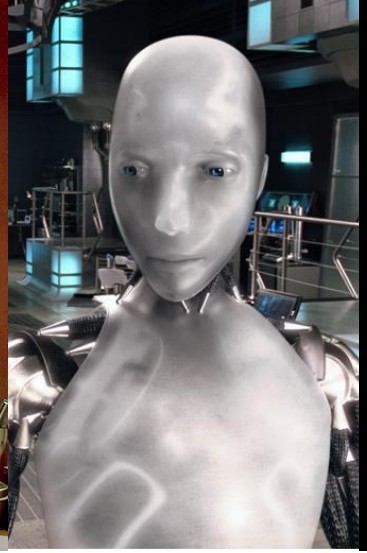

Autonomic Systems that Learn

Botond Virginas

Principal Researcher for Data Science
BT, Research & Innovation

Workshop structure

- Introduction
- From Algorithms to Business Autonomics – Kjeld Jensen
- Business Autonomics – Dave Rohlfing

© British Telecommunications plc

Signposting

I : Free-standing AI – ML and examples (and cautions about them)

II: Autonomics – Closing the loop

III: The need for autonomics in business

© British Telecommunications plc

- AI and Machine Learning

© British Telecommunications plc

What is Artificial Intelligence (AI)?

- Search, Planning & Optimisation
- Language Understanding
- Machine Learning
- Context Understanding, Concept Forming & Abstraction
- Perception (Speech & Vision)
- Reasoning

Weak or specialised AI – intelligently solving specific problems: *many applications exist*
Strong or generalised AI – real human-like machine intelligence: *does not exist yet*

Extended Intelligence – using machine learning to extend the abilities of human intelligence

© British Telecommunications plc

The Three Waves of AI

- Handcrafted Knowledge

  Engineers create sets of rules to represent knowledge in well-defined domains; the computer explores the specifics
  (e.g. chess programs, logistics planning)

- Statistical Learning (now)

  Engineers create statistical models for specific problem domains and train them on big data
  (e.g. Alexa & Siri, Google's image recognition, IBM Watson for Jeopardy, chat bots)

- Contextual Adaptation (future)

  Systems construct contextual explanatory models for classes of real world phenomena

© British Telecommunications plc

Why the Hype?

- Recent advances through “deep learning” or “deep networks” – a flavour of artificial neural networks – that became possible through cheaper hardware, big data and improvements in structuring and training them.
- Impressive applications in image recognition and image labelling, speech recognition, and automatic translation (e.g. Google Brain).
- Deep Mind (British company acquired by Google) won against Go champion
- Applications are pushing into consumer electronics through virtual assistants like Amazon’s Alexa, Google Now, Apple’s Siri.
- New players heavily invest in AI and ML: Google, Baidu, Facebook, Amazon, Uber, AirBnB, …
- Traditional players like IBM and Microsoft are pushing ML in their offerings.
- Vendors have declared ML & AI to be the next big thing after Data Science.
- Increase the temperature of the hype: fears about the jobs and fears about ethics

© British Telecommunications plc

Gartner: Statistics

- By 2018, 20% of all business content will be authored by machines.
- By 2018, 6 billion connected things will be requesting support.
- **By 2020, autonomous software agents outside of human control will participate in 5% of all economic transactions**.
- By 2018, more than 3 million workers globally will be supervised by a "roboboss."
- By YE18, 20% of smart buildings will have suffered from digital vandalism.
- **By 2018, 50% of the fastest-growing companies will have fewer employees than instances of smart machines.**
- By YE18, customer digital assistants will recognize individuals by face and voice across channels and partners.
- By 2018, 2 million employees will be required to wear health and fitness tracking devices as a condition of employment.
- **By 2020, smart agents will facilitate 40% of mobile interactions, and the post-app era will begin to dominate.**
- Through 2020, 95% of cloud security failures will be the customer's fault.

© British Telecommunications plc

What is Machine Learning (ML)?

- Field in Computer Science that looks for algorithms that can learn from data to make predictions without explicit programming.
- Also called statistical learning and closely related to computational statistics and mathematical optimisation
- Part of AI (Artificial Intelligence)
- Used in data analytics (data science, data mining, …)

© British Telecommunications plc

State of the art in ML: Deep Networks

From:

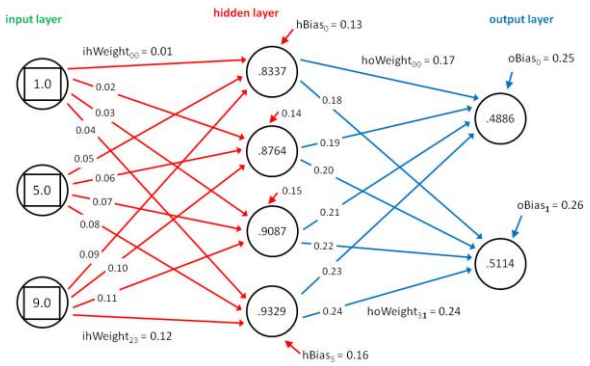

Multilayer Perceptron

- Trained through Backpropagation (invented in 1980s)
- For large number of hidden layers training is slow
- Real world problems need many hidden units and a lot of labelled data
- Advances through structured hidden layers, improved algorithms, specialised hardware (GPUs) and Big Data.

To:

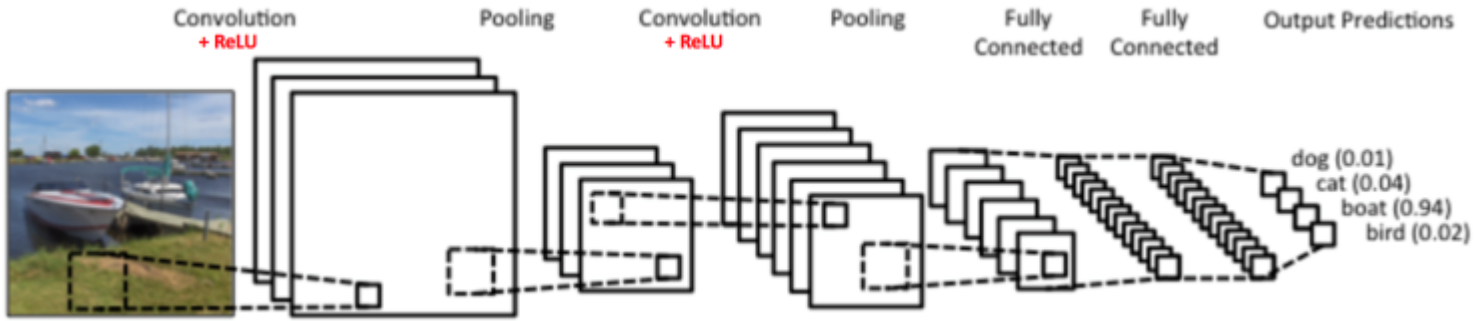

Machine-learning "programmers" design the network structure with experience and by trial and error

© British Telecommunications plc

Complex Example of ML: Combining Deep Networks

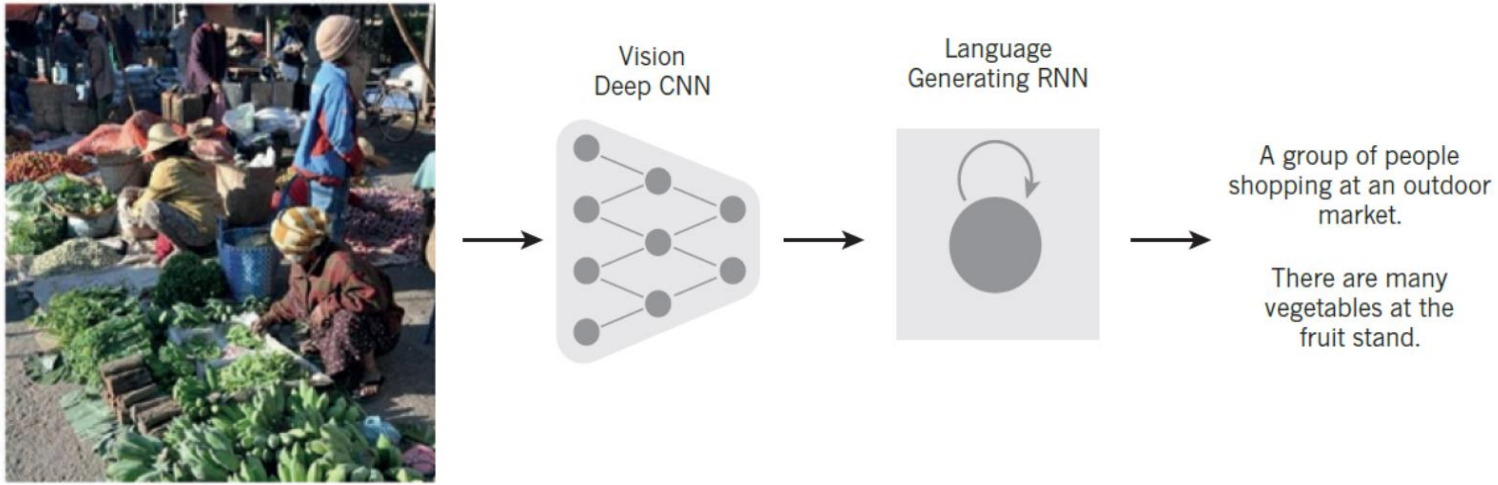

Yann LeCun, Yoshua Bengio, & Geoffrey Hinton (2015). Deep Learning, Nature, Vol. 521, (pp. 436-444)

A deep convolution neural net (CNN) produces a set of outputs (abstract "words")

A language-generating recurrent neural net (RNN) "translates" the abstract "words" into captions

© British Telecommunications plc

Challenges with Deep Networks

Statistically impressive,
but individually unreliable

© British Telecommunications plc

Challenges with Deep Networks

“Panda”

+

<1%
targeted distortion

=

“Gibbon”
(99.3% confidence)

Inherent flaws can be exploited:
vulnerability to adversial perturbatrions

© British Telecommunications plc

Challenges with ML – New Application Areas

Example: Algorithms to Predict Likelihood of Criminal Conduct

DJIA Futures ▲ 18854 0.11% S&P 500 F ▲ 2176.00 0.15% Stoxx 600 ▲ 338.92 0.13% U.S. 10 Yr ▲ 6/32 Yield 2.200% Crude Oil ▲ 46.09 1.14% Euro ▲ 1.0735 0.41%

THE WALL STREET JOURNAL.

Subscribe Now | Sign In

SPECIAL OFFER: JOIN NOW

Home World U.S. Politics Economy Business Tech Markets Opinion Arts Life Real Estate

Search

Minnesota Officer Faces Manslaughter Charges in Shooting of Black Man

In Maine's Tight Job Market, Businesses Look to Immigrants

Sign-Ups Under Affordable Care Act So Far Seem Not Hurt by Trump Win

Trump's Made-in-America Hurdle: Asia

U.S.

Wisconsin Supreme Court to Rule on Predictive Algorithms Used in Sentencing

Ruling would be among first to speak to legality of risk assessments as aid in meting out punishments

The Wisconsin Supreme Court, located in the State Capitol, is set to rule on whether algorithms used to predict criminality can be used in sentencing. PHOTO: GETTY IMAGES/SCIENCE SOURCE

By JOE NAVARRO

June 5, 2016 5:30 p.m. ET

NO COMMENTS

http://www.wsj.com/articles/wisconsin-supreme-court-to-rule-on-predictive-algorithms-used-in-sentencing-1465119008

Models assessing individuals are highly problematic

- Independent tests show that these algorithms only about 60%-70% correct.
- False positives can be devastating for individuals
- High risk for machine bias if features like race are used.
- Models are opaque and cannot be challenged.
- It is illegal in the UK to use algorithms for assigning risk to individuals (exception: credit score).

© British Telecommunications plc

Challenges with ML

- Data Provenance
  Where does the data come from, how was it collected and can I trust it?
- Data Quality
  Is the data error free?
- Data Bias
  Does the data accurately reflect the population/situation I want to model? What is missing?
- Model Bias
  Driven by Data Bias, or bad sampling – to what extent is my model biased by the training data?
- Model Comprehension
  Why are the outputs of my (black box) models the way they are?

© British Telecommunications plc

AI: The next frontier

- The next frontier of computing involves sophisticated systems that perform complex "human-like" tasks, involving complex inferences and predictions.
- Examples of such systems include:
  - Personalized health care
  - Safe driving and autonomous vehicles
  - Automated industrial machines and household gadgets
  - Conversational engines
  - Document summarizations
  - Real-time objects recognition from video
  - Pattern and anomaly detection
  - Inferences over large graphs and social structures
  - Large-scale distributed systems and network architectures

© British Telecommunications plc

Examples: IBM Watson

"more than 100 different techniques are used to analyze natural language, identify sources, find and generate hypotheses, find and score evidence, and merge and rank hypotheses (e.g. advanced natural language processing, information retrieval, knowledge representation, automated reasoning, and machine learning technologies to the field of open domain question answering)"

Application domains: health care, chatterbot, weather forecasting, teaching assistant,fashion, tax preparation

© British Telecommunications plc

Google Self driven car – urban mobility

- Waymo stands for a new way forward in mobility. We are a self-driving technology company with a mission to make it safe and easy for people and things to move around.
- Machine learning algorithms are used to create models of other people on the road. Every single mile of driving is logged, and that data fed into computers that classify how different types of objects act in all these different situations. While some driver behaviour could be hardcoded in ("When the lights turn green, cars go"), they don't exclusively program that logic, but learn it from actual driver behaviour.

© British Telecommunications plc

Self driven car: Tesla car learning

- “The whole Tesla fleet operates as a network. When one car learns something, they all learn it. That is beyond what other car companies are doing”. When it comes to the autopilot software, each driver using the autopilot system essentially becomes an “expert trainer for how the autopilot should work.
- The Tesla fleet does machine learning by reporting back to the mothership any time the driver needs to correct the autopilot.

  For example, if the autopilot thinks the lane is getting wider and stays in the center, but the real situation is that the right-hand lane marking is diverging to create an exit ramp, the driver will tug the wheel to correct it. This event goes back to the central database tagged with GPS coordinates. The next time the car passes this spot, it knows to follow the left lane marker for a while. But, so does every other Tesla. It's a continuously expanding collection of driving tidbits available to the entire fleet.

© British Telecommunications plc

Universities: Systems that Learn initiative - MIT

- Goal: To accelerate the development, deployment, and evolution of large-scale software systems that incorporate machine learning and artificial intelligence.

© British Telecommunications plc

Heterogeneous architectures

- The data and features that drive learning in these systems and applications increasingly come from diverse distributed infrastructure, including phones, sensors, or other bandwidth and power impoverished endpoints. Thus even acquiring data for learning may require adaptive allocation of computation over heterogeneous infrastructure. Furthermore, the rise in heterogeneous hardware, such as GPUs and many-core processors, which excel at certain aspects of the learning pipeline, suggests a diversity of computational resources will be brought to bear.

© British Telecommunications plc

Predictable composition

- Successfully designing and training machine learning methods for the desired task once data is available, i.e., programming at the level of learning components, and reasoning about the behaviour of the composition of such components, calls for skill and expertise which is not yet well-supported or automated

© British Telecommunications plc

Distributed execution

- In terms of the underlying infrastructure, complex machine learning methods also demand considerable parallel resources to train effectively. Once trained, models may be deployed either on massive parallel infrastructures (e.g., data centers) or may have to be reduced and distributed back to the heterogeneous components to be utilized where needed (e.g., mobile devices), requiring new distributed algorithms and execution frameworks.

© British Telecommunications plc

Seamless integration of training and deployment.

- Many machine learning solutions today are trained and deployed in well-separated phases of training and testing (deployment), but this will change. Learning will increasingly become an ongoing, integrated process. The tighter integration of learning and computer systems offer exciting possibilities in terms of new capabilities, but requires us to overcome challenging hurdles pertaining to programming abstractions,

© British Telecommunications plc

Some projects – CSAIL MIT

- Deep Learning Architectures that enable Interpretable and Informative Classification
- Robust Automated Intelligence for Health Data via Bayesian Model Discovery
- Convex and submodular optimization for machine learning
- Data discovery and meta data flows
- Competing genetic algorithms for network security

© British Telecommunications plc

II

Autonomics: closing the loop

© British Telecommunications plc

Autonomics: IBM whitepaper 2005

- Autonomic systems are systems that could manage themselves in conformance to high-level objectives set out by administrators (IBM 2005)
- It is analogous to how the autonomic nervous systems automatically manages the basic body functions
- IBM defined An architectural blueprint for autonomic computing:
  - Architectural building blocks
  - The Autonomic Computing Adoption Model
  - The role of the human in autonomic systems, including delegation of tasks
  - New developments in standardization

© British Telecommunications plc

MAPE cycle

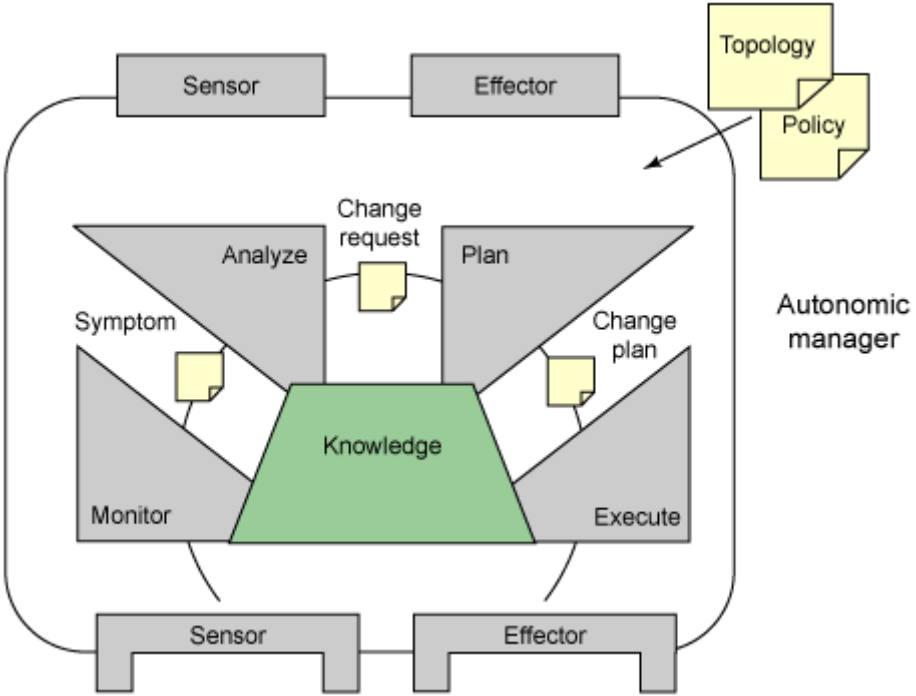

© British Telecommunications plc

Main characteristics of autonomic systems (CHOP)

- Self-configuration: Automatic configuration of components;
- Self-healing: Automatic discovery, and correction of faults;
- Self-optimization: Automatic monitoring and control of resources to ensure the optimal functioning with respect to the defined requirements;
- Self-protection: Proactive identification and protection from arbitrary attacks.

© British Telecommunications plc

Alternative perspectives of autonomic systems

- Self-regulation: A system that operates to maintain some parameter, e.g., Quality of service, within a reset range without external control;
- Self-learning: Systems use machine learning techniques such as unsupervised learning which does not require external control;
- Self-awareness (also called Self-inspection and Self-decision): System must know itself. It must know the extent of its own resources and the resources it links to. A system must be aware of its internal components and external links in order to control and manage them;
- Self-organization: System structure driven by physics-type models without explicit pressure or involvement from outside the system;

© British Telecommunications plc

Minimum set of properties

- **Automatic**: being able to self-control its internal functions and operations.
- **Adaptive**: An autonomic system must be able to change its operation (i.e., its configuration, state and functions). This will allow the system to cope with temporal and spatial changes in its operational context either long term (environment customisation/optimisation) or short term (exceptional conditions such as malicious attacks, faults, etc.).
- **Aware**: An autonomic system must be able to monitor (sense) its operational context as well as its internal state in order to be able to assess if its current operation serves its purpose. Awareness will control adaptation of its operational behaviour in response to context or state changes.

© British Telecommunications plc

LHS terms from “thoutonomy”

IBM terms (broad)

Levels of autonomic capabilities

Aware

Machine learning

AI-based

Rules-based, heuristics

Adaptive

Active and passive autonomics

Automatic

Non-autonomic

	Single domain	Interacting elements	Multi-agent shared learning	Multi-agent state detection	Achievement of global goals (state)
Introspective learning from state history: agility					
Self-regulation of learning strategies					
Adjustable learning rates and risk levels					
Ability to balance multiple business goals					
Self-learning maintenance of goal					
Automated action based on repertoire					
Identify actions from repertoire of states					
Interactive visualisations					
Automation					

© British Telecomm

Survey highlights benefits of robotic process automation

For its research, Mindfields evaluated 12 leading IT and business process-outsourcing providers on a wide range of parameters, including automation maturity, experience, expertise, and scale of robotic automation initiatives. Out of the twelve providers in the study, Cognizant was the only firm to receive top scores on four key evaluation parameters.

In surveying the market of end clients and service providers, Mindfields noted several positive and emerging trends as robotic-process-automation adoption rates increase globally. Companies and providers are now achieving benefits from this automation category, including the following:

- Significant cost and operational benefits for client organizations
- Enabling a new business model for service providers
- Establishing a new, nonlinear version of outcome-based pricing
- Redefining of hiring strategies

© British Telecommunications plc

III

The need for autonomics in business

© British Telecommunications plc

Enterprises are unprepared for autonomics

FIGURE 1 Enterprises Are Unprepared For Smart Machines From Both Human And Technical Standpoints

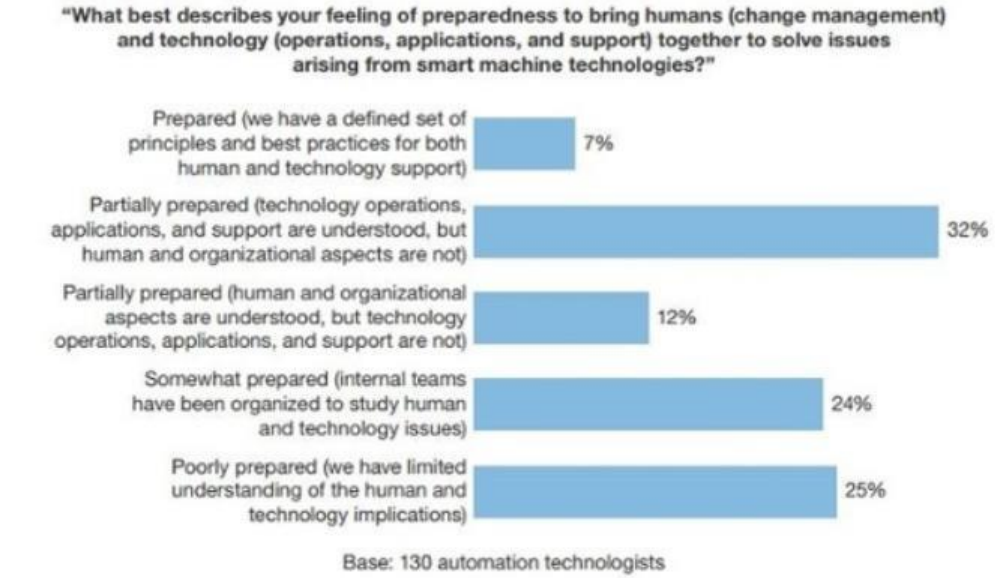

Source: Forrester's Q2 2016 Global Automation Technology Online Survey

© British Telecommunications plc

Business Autonomics

- Motivation: closing the loop of a self learning business system
- It is very different thinking
- Never a certain algorithm is fit for purpose
- Wouldn’t be nice if an algorithm would change overtime as required by changing conditions
- Particular example in the second part – Kjeld Jensen
- Abstraction, concepts, etc : Dave Rohlfing

© British Telecommunications plc