CS-E4800 Artificial Intelligence

Jussi Rintanen

Department of Computer Science
Aalto University

March 21, 2019

Impact of A.I. (and Automation in General)

- What will A.I. do in the next 20 years?
  - What did automation and computer technology do before?
  - Where is A.I. needed and why?
  - What impact will A.I. (and automation) have on things?
- What has happened in A.I. in the last 50 years?
- Where are industrial applications of A.I.?

Impact of Automation in the Past

Industrial Revolution 1760-1840 (steam engine, water power)

- textile industry (weaving, sewing, ...)
- metallurgy, mining
- machine tools
- transportation (railways, canals, roads)

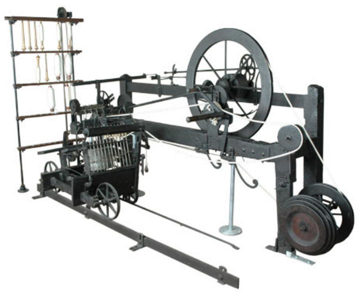

Impact of Automation in the Past

Automation 1880-1950 (electricity, combustion engines)

- transportation (vehicles with combustion engines)
- communication (telegraph 1835-, radio 1897-)
- farming (tractors, etc.)
- all fields of manufacturing
- mining

Impact of Automation in the Past

Automation 1890- (rudimentary mechanical IT)

- punched card systems (1890 US census)
- mechanical calculators 1930- (add, subtract, multiply, divide)

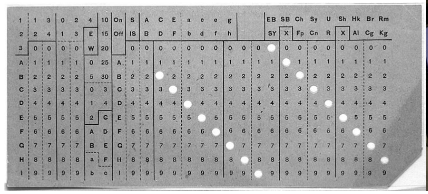


Impact of Automation in the Past

Automation 1960- (rudimentary electronic IT)

- accounting
- banking
- engineering calculations


Impact of Automation in the Past

Automation 1980- (advanced IT, communication)

- office automation 1980-
- Internet banking 1995-
- Internet travel agencies 1995-
- other Internet services

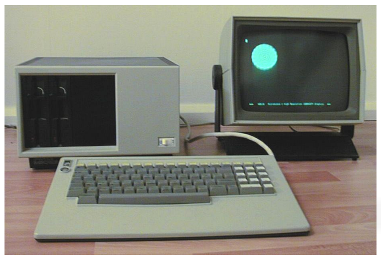


Impact of Automation Now and Future

Automation 2010-2040 (complex IT, “A.I.”, A.I.)

- automation of the physical world accelerates
  - automated manufacturing, warehouses
  - autonomous vehicles (buses, taxis, trucks)
- software production (still need programmers?)
- automation of most office work

Unprecedented flexibility and adaptability of IT

Turing test (1950)

- "If a machine could carry on a conversation (over a teleprinter) that was indistinguishable from a conversation with a human being, then the machine could be called intelligent."
- Suggested by Turing as a thought experiment, not an actual test
- Some programs have now passed "Turing tests"
- Broad indistinguishability from humans far away

First Golden Era 1956–1974

- Term "A.I." coined by McCarthy, Minsky, Shannon, Rochester in 1956
- Lots of research action started
  (with great hype and promises)
- "In from three to eight years we will have a machine with the general intelligence of an average human being." (Minsky, 1970)
- "Computers would never be able to play chess." (Dreyfus, 1972)
- First A.I. winter 1974–1980
  (funding for A.I. research dropped dramatically)

Second Golden Era 1980–1987

- Expert systems (example: medical diagnosis)
- Case-based reasoning (find examples from past matching the current situation, mimic the past human solution)
- Neural networks trained by backpropagation (popularized in 1986)
- Second A.I. winter 1987–1993 (funding again dropped)

New Coming of Machine Learning

- A.I. + statistics ⇒ Machine Learning (Big Data)
- Bringing ad hoc methods in a unified framework
- Neural networks focus of interest, again
- Enabled by scalability due to
  - faster CPUs, more CPUs
  - massive amounts of memory

New Coming of Symbolic Model-Based A.I.

- More detailed models, scalable search and reasoning
- Autonomous vehicles, aircraft and robots
  - Detailed physical models
  - Fast model-based decision-making
- A.I. embedded in all form of software
- Similarly to ML, enabled by scalability due to
  - faster CPUs
  - massive amounts of memory

Latest Hype Cycle 2015-

- Lots of publicity for A.I.
  - self-driving cars
  - voice recognition, machine translation (deep learning)
- Most of this not really fully works (yet)!
- No actual breakthroughs (arguably)
- Best applications: huge amounts of engineering effort
- A.I. start-ups promising too much

Singularity

Good, 1965; Vinge, 1983 (term); Solomonoff, 1985; Kurzweil, 2005

Story (quite irrelevant at the present time):

- A.I. will be more intelligent than humans, sometime in the future.
- This A.I. will accelerate development of more advanced A.I.
- What effect will this have?
  - A.I. will "take control"?
  - Humans still able to use A.I./robots as "slaves"?
- Lots of (philosophical) articles on this + series of conferences

A.I.: Artificial Human Brains or Automation?

Is Artificial Intelligence

- Walking robots
- Talking robots
- Talking computer programs

or is it

- Fully automated primary industries (mining, farming, fishing)
- Fully automated energy generation and distribution
- Fully automated manufacturing
- Fully automated transportation and logistics
- Fully automated services

What is A.I. in Practice Today?

- several different types of technologies
- several different types of applications
- no *single* general-purpose AI technology in sight
- no human-style intelligence in sight

What is A.I. in Practice Today?

- Reproduction of human abilities, human-computer interaction
  - vision
  - speech recognition
  - natural language processing (text-based data-intensive methods)

  Yardstick: How good match with human abilities?
- Industrial A.I. (autonomous systems, autonomous vehicles, autonomous infrastructure)
  - combinatorial search and optimization
  - large-scale constraint solving, optimization (numeric, relational constraints)
  - decision-making

  Yardstick: How good solutions?

  (No need to compare to humans; exceed human capability already)

Vision and Human Computer Interaction

- Speech recognition: becoming practically usable
- Image classification: good medical and life sciences applications
- Machine translation: becoming usable (but not yet high quality)
- Natural language processing: applications emerging
- Big open problems:
  - Natural language understanding
  - Human level vision
  - Understanding of human behavior (gestures, movements, ...)

Industrial Applications

Examples of existing and emerging applications in:

1. Autonomous Vehicles
2. Distributed Systems (Power, Telecom)
3. Software Industry
4. Commerce

Industrial Automation and Control Rooms

Control Rooms

Many industrial facilities employ few humans only, in a control room

- Connection often through a SCADA system (Supervisory Control and Data Acquisition)
- Controls operated through SCADA
- Alarms and other messages come through SCADA
- Earlier: control panels with buttons, levers, etc. & lamps and other electrical and mechanical displays
- Now: computer displays, keyboards (sometimes joysticks, levers)

Why Control Room?

Automation already covers much of system functionality (Control Theory, basic computer automation)

What remains is complex cognitive tasks, requiring deeper and broader expertise (and intelligence).

- Situational awareness (What is going on?)
- Detect (non-routine) fault situations
- Diagnose fault situations (What is wrong?)
- Recover from fault situations (How to fix it?):
  - Take control actions (through computer interface)
  - Deploy humans to fix the problem

Control Rooms

Manufacturing:

- industrial plants, manufacturing
- process industries (metal, oil, chemical)

Energy, communications and transportation:

- power stations (nuclear, coal)
- electricity networks (distribution, transmission)
- cellular communication networks
- rail networks (local, underground, long-distance)
- metropolitan road-traffic (traffic light control)

Military, Space, ...

Control Rooms

One could also view the following as control rooms:

- command deck of a ship
- cockpit of an airplane
- driver in a motor vehicle

General Approach: Models

- Model each component type
  - discrete + continuous behavior
  - Also faulty behavior must be modeled!
- Model interaction between components of different types
- Model how components connected (graph)
- Languages:
  - notion of component (or module)
  - time (real, rational or integer)
  - concurrency of events and actions
  - discrete and continuous change (hybrid systems)

General Approach: Reason about Models

Achieve model-based intelligence by

- state-space search (BFS, Greedy Best First, A*, ..)
- constraint programming (IP, MILP, SAT, SMT, ...)

in order to

- Analyze system: What behaviors are possible?
- State estimation: Explain observations
- Control: Which control actions to take?

The Smart Grid (Electricity)

Generation	Transmission & Distribution	Consumers
more renewables	more proactive	distributed generation
less predictable	better network utilization	controllable demand

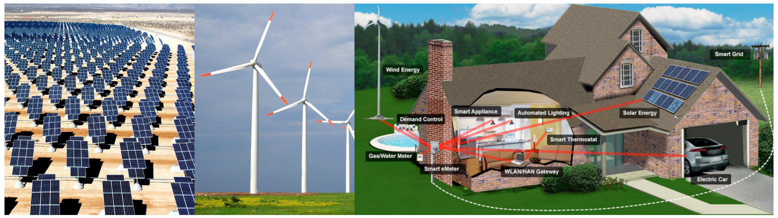

Future Energy Generation and Distribution

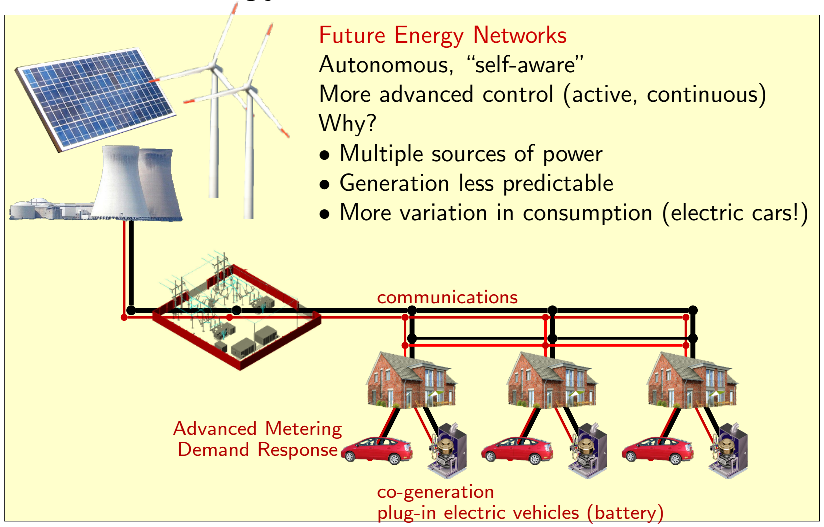

Future Energy Generation and Distribution

Problem:

- More control is needed, in far smaller scale
- Use of human operators not feasible (costs!)
- Control tasks are complicated:
  - Network topologies and devices very heterogenous
  - Control heavily depends on larger context

Power Outages

Existing automation circuit-breakers:

- Protective device opens switch upon detecting fault current
- Recloser closes switch e.g. after 1 and 10 seconds
- These handle faults without human intervention

This failing, control room operators take over:

- Operate remote-controlled switches (to reduce outage area)
- Dispatch human crews to fix

Distribution Network Reconfiguration

- Open and close (remotely-operated) switches to change active network topology
- Outage recovery: re-supply part of outage area
- Loss reduction and load balancing
- Objective:
  - Radial/tree-like configuration maintained (no cycles!)
  - Objective: maximize area supplied with power
  - Objective: minimize losses, balance loads
  - Objective: minimize switching-actions (wear and tear)

Distribution Network Reconfiguration

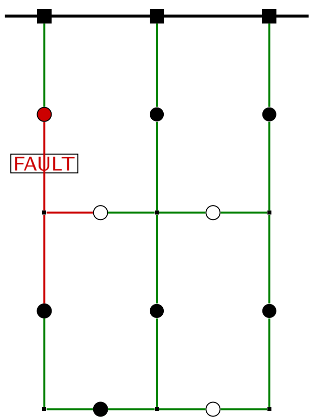

Default configuration
Fault: Switch opened by protective device
Outage!
Re-supply part of network

Real-world problem complicated:

- Complex urban network topologies
- No-shed loads (hospitals!)
- Capacity limits → Transfer loads between substations after initial reconfiguration

Distribution Network Reconfiguration

- state-space search (A$*$ etc.)
- constraint programming, including
  - Integer Programming
  - Mixed-Integer Linear Programming
  - SAT (propositional logic)
- genetic algorithms
- ant colony optimization
- etc etc

Situational Awareness

- Maintain an understanding of the state of the system, at an abstract level
- Control room views and interprets event logs: all alarms and other messages from the devices with communications to the control room
- Manufacturing plants, power stations, electricity networks, traffic networks, railway networks, airline operations, ...

Situational Awareness

Accidents caused/worsened by poor situational awareness

- airplane crashes (e.g. Air France AF447 in 2009)
- nuclear power stations (Three Mile Island 1979)

A.I. to create, or assist in creating better situational awareness

Situational Awareness

```
00:00:00 CB 1B A-B OPEN
00:00:00 CB 2B A-B OPEN
00:00:00 CB 2A A-B OPEN
00:00:00 CB 1A A-B OPEN
00:00:01 Line A-B  KV LIMIT LOW
00:00:04 Line C-D  KV LIMIT NORMAL
00:00:15 CB 1A A-B CLOSED
00:00:17 Line A-B  KV LIMIT NORMAL
00:00:17 CB 1B A-B CLOSED
00:00:17 CB 2B A-B CLOSED
00:00:20 CB 2A A-B CLOSED
00:00:20 Line C-D  KV LIMIT HIGH
```

- Whole event log potentially useful
- Many log entries can arise in short time: difficult to understand
- Need to focus on important parts
- Causal connections between log entries
- Model-based approach: interpret event logs against a system model, generate concise explanation of what is going on

Adaptive Management of Mobile Networks

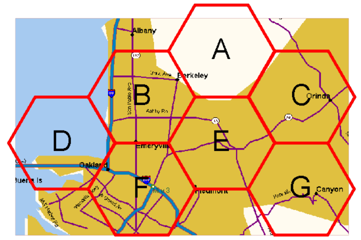

- Location and activity of mobile terminals (phones) change hour to hour, month to month
- Need to adjust antenna: power, tilt (electric)
- Future 5G networks have too many base stations $\rightarrow$ need to automate

Adaptive Management of Mobile Networks

Solution:

- Reinforcement learning
- Learn to adapt to traffic situation
  - Collect data on Quality of Service
  - Explore antenna parameter values
  - Choose best ones based on time-of-day and variables

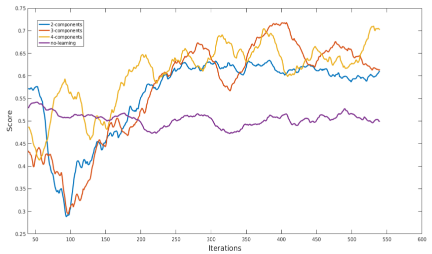

Need for Autonomy: Space

- Cost (weight, room) of carrying crew is high
  - food, water
  - air
  - living quarters
- Crew limits mission duration, reach
- Ground-control limiting for distant spacecraft
  - light to Moon 1 sec
  - light to Mars 4 to 20 min

Need for Autonomy: Military

- Disadvantages of pilot (cost, functionality):
  - Pilot is not payload
  - Limits aircraft design (cockpit, catapult seat, oxygen)
  - Maneuverability (acceleration limits)
- Risk to pilot life eliminated by eliminating pilot
- Remote control susceptible to jamming

"Autonomy is the biggest thing in military technology since nuclear weapons"
"Some experts say autonomous weapons are potentially weapons of mass destruction"
"The availability on the black market of mass quantities of low-cost, anti-personnel micro-robots that can be deployed by one person to anonymously kill thousands or millions of people who meet the users targeting criteria. Autonomous weapons are potentially weapons of mass destruction."

Need for Autonomy: Ships

- Crew salary costs
- "75 to 96 per cent of marine accidents result from human error, often because of fatigue"
- piracy: no hostages to take!
- energy efficiency and construction costs:
  - wind resistance (command deck, living quarters)
  - ventilation, heating, water, sewage
- Remote-control feasible, but mostly unnecessary

Autonomous Systems: Implementation

Vehicles (air, space, land, water), robots, everything in between, have similar requirements for control:

1. Short-term motion planning and execution
   (Robotics courses)
2. Medium-term and long-term task planning
   (This course: planning, scheduling, constraint solving)

Motion planning and execution

- intersection of Computer Science, Control Theory
- Physical models of the device + environment
- Search for best motion plan, based on the model
- Managing uncertainty: discrepancy between
  - intended (predicted) behavior, and
  - observations.

  Both include imprecision and uncertainty!
- Fusion of sense data and predicted behavior:
  - Particle filter
  - Kalman filter
  - Recursive Bayesian estimation (Bayes filter)

  Covered in robotics courses

Autonomous Vehicles

Waymo / Google Self-Driving Car

- Control based on detailed 3D models created from
  - detailed 3D model of all roads (acquired earlier!)
  - LIDAR: Velodyne 64 beam laser HDL-64E (range: 120 m; resolution: 2 cm)
  - radars (range: 200 m; resolution: low)
  - cameras (for traffic lights, signs)
  - GPS
- likely behaviors of pedestrians, bicyclists, other vehicles is learned

Humans responsible for 13 out of 14 crashes with Google Car; 14th was a minor collision with a bus

Autonomous Vehicles

Tesla's Autopilot (semi-autonomous driving)

- Cameras + radar for sensing
- No detailed 3D model
- \$50000+ LIDAR too expensive; Tesla: "overkill"
- Requires constant driver supervision
- Fatal Autopilot accident in 2016:
  1. Large white truck+trailer turns, blocking the lane
  2. Tesla does not "see" it
  3. Driver distracted (watching DVD?), doesn't intervene
  4. Car crashes under the truck's trailer

Autonomous Spacecraft

- Deep Space 1 probe 1997
  - EUROPA planner (constraint-based temporal planner)
  - EXEC plan-execution system
  - Livingstone diagnostics system
  - full autonomy demonstrated
  - Deep Space 2 crashed before deployment (mission: to drill to Mars soil, 1999)
- Mars rovers
  - Early rovers remote-control only (long comms delay)
  - Curiosity rover has (partial) autonomy (2013-)

Autonomous Military Aircraft

- Predator MQ-1 drone (1994-): remote-controlled
- Reaper MQ-9 drone (2001-): can fly pre-programmed routes; no combat autonomy
- Perdix drone (experimental 2014-):
  - purpose: surveillance
  - swarms of dozens of drones
  - wingspan 30 cm, weight 290 g, speed 70 mph, flighttime 20 min
  - developed by MIT students
  - less vulnerable than large drones (Reaper costs \$18M)
  - too expensive to control by pilot: full autonomy
- Very active topic in military

Applications

- Autonomous car: thousands of applications
- Autonomous aircraft: thousands of applications
- Human-shaped robot: thousands of applications
- Compare: microprocessor
- Given autonomous X,
  applications are at abstract level
  → discrete problems
  → state-space search, scheduling, constraints, ...

Model-Based Software Development

1. Develop modeling language for product category
2. Automate the use of the modeling language
3. Software not *programmed*, but modeled

Advantages:

	programming	modeling
cost	expensive	cheap
time	slow	fast
errors	lots	fewer
modifiability	bad	good
extendibility	bad	good

Systems that are Fleets of Vehicles

- Fleets of
  - robots
  - autonomous cars
  - autonomous trucks
  - ships
  - aircraft
  - spacecraft
- Needed in
  - transportation and shipping
  - mining, construction, agriculture
  - military
- Large-scale planning, scheduling (state-space search, constraint programming)

Model-Based Software Development

Potential software categories:

- Control software for physical systems
- Computer games
  - 3D physical world (single-player, multi-agent)
  - Strategy

  Model agents' etc. behaviors, goals, ...
- Information systems (my current project)
  - logic-based representation of complex data
  - system's objectives
  - system's and users' actions
  - high-level software synthesis

Economic Impact of A.I. and Intelligent Robotics

- Cheap A.I. and intelligent robots = Flux of very competent and cheap labor on the labor market
- Earlier revolutions of automation have created new jobs
- Will the ongoing revolution do the same? How?

Legal Implications of A.I.

- Until now: Only humans able to do things.
- Future: Robots do almost everything humans can.
- Separation of agency and responsibility
- System's designer/owner relieved of responsibility, because system acquired a behavior (through learning) that could not have been anticipated? No.

Autonomous Vehicles: Liability in Accidents

- Assume traffic deaths reduced 50 per cent by full automation of road traffic (40000 annually in US)
  - 20000 "thank you" letters, vs.
  - 20000 lawsuits?
- All deaths caused by technology, not human error

Autonomous Vehicles: Ethics of Design

- Autonomy in vehicles unlikely to eliminate all accidents
- Design decisions have impact on who gets hurt (passengers vs. others)
- How does the manufacturer decide how the car behaves when an accident is imminent?

Autonomous Systems: Crime

- High-risk crimes (physical risk, prosecution) attractive for automation
  - drug-trafficking (UAV, submarines)
  - murder
  - spying
- Automation will reduce costs and risks
- As a result, crime will become more attractive
- What can be done about that?

A.I. Crime: Anonymity

- Hijack Internet-connected robots, cars, UAV; commit crime Existing Internet security issues, but with implications on the physical world
- Assassinations by autonomous or tele-operated robots, drones, quadcopters
- All autonomous systems must be registered, traceable; must develop infrastructure to detect and prevent the use of unregistered systems?

Military A.I.

- A.I. technologies, like IT in general, are dual use: anything useful for civilian purposes, also applicable to warfare
- Autonomous weapons reduce/eliminate need of military personnel on the field
- Loss of life major disincentive for use of military force (especially in industrialized democracies)
- Will automation increase military aggression?

Military A.I.

Warfare will change.

- Soldiers now for carrying and operating weapons
- One day both sides of a conflict use autonomous weapons only
- Enables more aggressive warfare
- Advantage in technology means direct and dramatic advantage in warfare

Exam April 8, 2017 at 1pm

- Course exam on April 8
- Later 1 or 2 exams (autumn)

What to Read for the Exam?

- Material in MyCourses
- Course presentation slides (MyCourses frontpage Schedule)
- additional material: Russell and Norvig Artificial Intelligence: A Modern Approach, topics covered in the lectures:
  - 1 Introduction, 2 Intelligent Agents, 3 Search
  - 5 Adversarial Search: 5.1-5.5
  - 13 Quantifying Uncertainty: 13.5 *Bayes rule* only
  - 17 Making Complex Decisions: 17.1-17.3, 17.5
  - 18 Learning from Examples: 18.2, 18.6, 18.7
  - 21 Reinforcement Learning: 21.3

  (Section numbering in R&N varies between editions. Above is for the 3rd.)
  R&N is optional: Everything is explained in material in the Mycourses.)

Missing Exercises?

- Late deadline for all exercises is April 3 at 23:59.
- Exercises submitted late yield only 50 per cent of the points

Grading

- lectures + exercises + exam
- Every week, max 100 points given from completed exercises
- Max 200 points from exam
- Minimums to pass the course with lowest grade:
  - at least 300 points total
  - at least 15 points from exercises every week
  - at least 50 per cent of exam points
- A passing grade is determined by:
  - 300 ... 449 points $\longrightarrow$ grade 1
  - 450 ... 599 points $\longrightarrow$ grade 2
  - 600 ... 749 points $\longrightarrow$ grade 3
  - 750 ... 899 points $\longrightarrow$ grade 4
  - at least 900 points $\longrightarrow$ grade 5

Need more A.I.?

- B.Sc. thesis, M.Sc. thesis, Ph.D. thesis
  - Come and discuss or email Jussi.Rintanen@aalto.fi (or other faculty)
- Aalto courses
  - CS-E3220 Declarative Programming (Autumn 2019, revised content)
    - advanced software engineering (SW synthesis, validation)
    - A.I. base technologies (industrial A.I.)
  - CS-C3160 Data Science
  - CS-E3210 Machine Learning: Basic Principles
  - Many other courses (especially in ML)
- Be a teaching assistant for CS-E4800 in Spring 2020!