AI current research and challenges. Applicability to ATM automation

Jesús García

Departamento de Informática

Universidad Carlos III de Madrid

*Grupo **GIAA** (Applied Artificial Intelligence Research Group)*

Hosted by SESAR JOINT UNDERTAKING

Outline

- Introduction. Current state and waves of research in AI
    - **(Crafted knowledge in expert systems, Statistical Learning and Big Data, Contextual Adaptation)**
- Fast Review of some paradigms, emphasis in deep learning and big data (current wave), limitations and strategies towards generalization
    - **Next wave of AI: Adaptation to context and role of context in AI**
- Other paradigms: Expert Systems with uncertainty, distributed AI, cognitive agents
- Connections with scenarios for ATM automation roadmap

AI today

- Spectacular recent successes in some challenges defeating human competitors
  - **IBM's Deep blue beat world champion (G. Kasparov) in 1997**
  - **IBM's Watson Jeopardy quiz show in 2011**
  - **Google's AlphaGo beat world's top player (Ke Jie) in 2017**
- Huge investments in Artificial Intelligence
  - **HW: Nvidia (GPUs), Intel (Movidius), IBM (AI server), …**
  - **SW: Google, Microsoft, Facebook, Baidu, Alibaba, …**
- Artificial Intelligence in everyday applications
  - **Smartphones, speech recognition, machine translation**
  - **Homes: smart devices, ambient intelligence, robotics…**
  - **Cars: drivers assistants**

Explosive growth of big data/data science

- Data analytics growing exponentially
  - **Large data sets**
  - **Fast computing, cloud resources**
  - **New data streams:**
    - **Internet of Things (IoT)**
    - **Social media (Facebook, Twitter, …)**
- New applications continuously appearing
  - **Business in retail, marketing, manufacturing,…**
  - **Pharmaceutical drug discovery**
  - **Personal medicine**
  - **…**

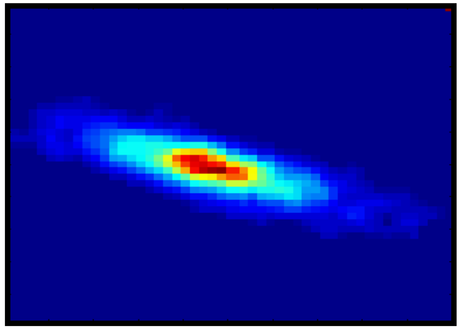
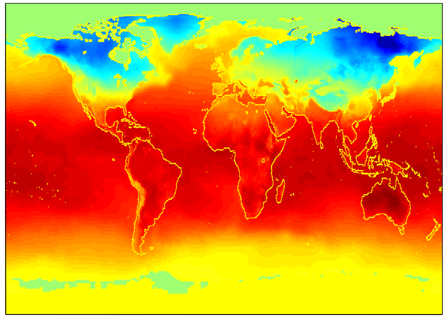
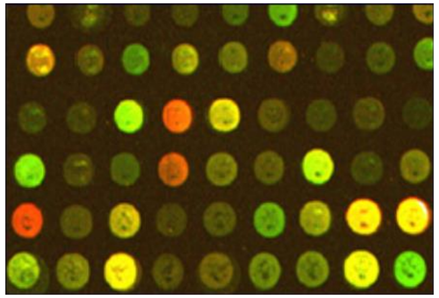

AI and autonomous cars

2005: Stanford. Desert

2007: Carnegie Mellon. Urban

2011: Googles' Self-driving car project

AI Waves

- The goal of AI research is to bestow on machines the ability to solve problems by mimicking human intelligence.
- Artificial Intelligence (AI), popularized in the 1950s, has gone through three major instantiations.
  - **First AI methods trained on simple use cases employing handcrafted knowledge.**
  - **Second wave of AI focused on machine learning methods, first limited by insufficient data for training and development. The extension of second wave is based on statistical-based deep learning which requires many training exemplars ("Big Data" resources), hope for the realization of advanced capabilities.**
  - **Third wave: contextual adaptation, a path towards generalization with explainable processes.**

Source: J. Launchbury. "DARPA perspective of AI", 2017.

AI Waves

	First Wave	Second Wave	Second Wave +
Years	1960s-1980s	1980s-2010s	2010s-
Technology	Expert-Systems	Machine Learning	Deep Learning
Algorithms	Logical rules	Statistical methods	Statistical methods
Expert knowledge	Expert knowledge ↓ Rules	Expert knowledge ↓ Model, Features	Expert knowledge ↓ Model
Learning		Parameters ↑ Data	Parameters ↑ Data
Algorithm application	Rules Data	Model ↓ Data	Model ↓ Data
Uncertainty handle	NO	YES	YES
Abstraction	NO	NO	YES
Interpretable	YES	NO	NO

1st AI Wave

- First wave: *Expert Systems* simulate human reasoning for specific problems in narrow domains (e.g., medical diagnoses)
  - **Intelligence encoded in logic rules - *rule base* - process queries to known facts in a *knowledge base*.**
  - **Unable to learn from data,**
  - **Brittle at handling uncertainty, no ability to abstract new concepts (i.e., develop more general rules) from data.**
  - **Major advantage is interpretability, possibility to trace exactly the chain of reasoning that led to certain conclusions.**
  - **Expert systems are still used today, enhanced for uncertainty management via fuzzy logic or Markov networks**

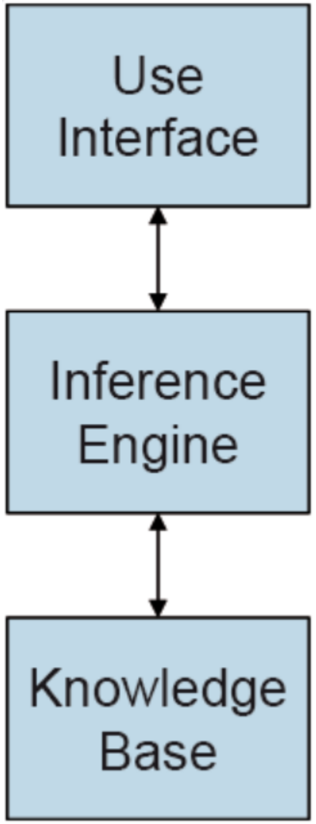

2nd AI Wave

- Advent of *Machine Learning*, methods able to solve certain problems (e.g., classification problem) by learning from examples and corresponding expected output.
- No explicit rules, focused in the most convenient model and parameters
- These ML systems are able to handle never-seen-before and noisy inputs with excellent performance in low-level tasks.
- Most methods, especially neural networks, operate like black boxes, no attribute semantics to the parameters learned during the training phase.
- Performance is sensitive to the training data

Deep learning, 2nd wave extension

- Interest in AI reemerges considerably thanks to the advent of the deep-structured machine learning or *deep learning paradigm*
  - **Deep neural networks comprised of many hidden layers. Training large datasets with sufficient computing power (specific architectures, such as GPUs)**
- Success of deep neural networks (DNNS) mostly relies in their ability to create abstractions from the observed training data.
- Abstraction ability gives them great generalization capabilities that allow them to perform better in the presence of noisy input data.
  - **Multiple levels of representation/abstraction**
  - **Speech and image recognition show spectacular success of DL, overpassing human-level performance. Extension to natural language processing (NLP)**

Deep learning

- **Automatic Feature Discovery**
- **[Yoshua Bengio]**

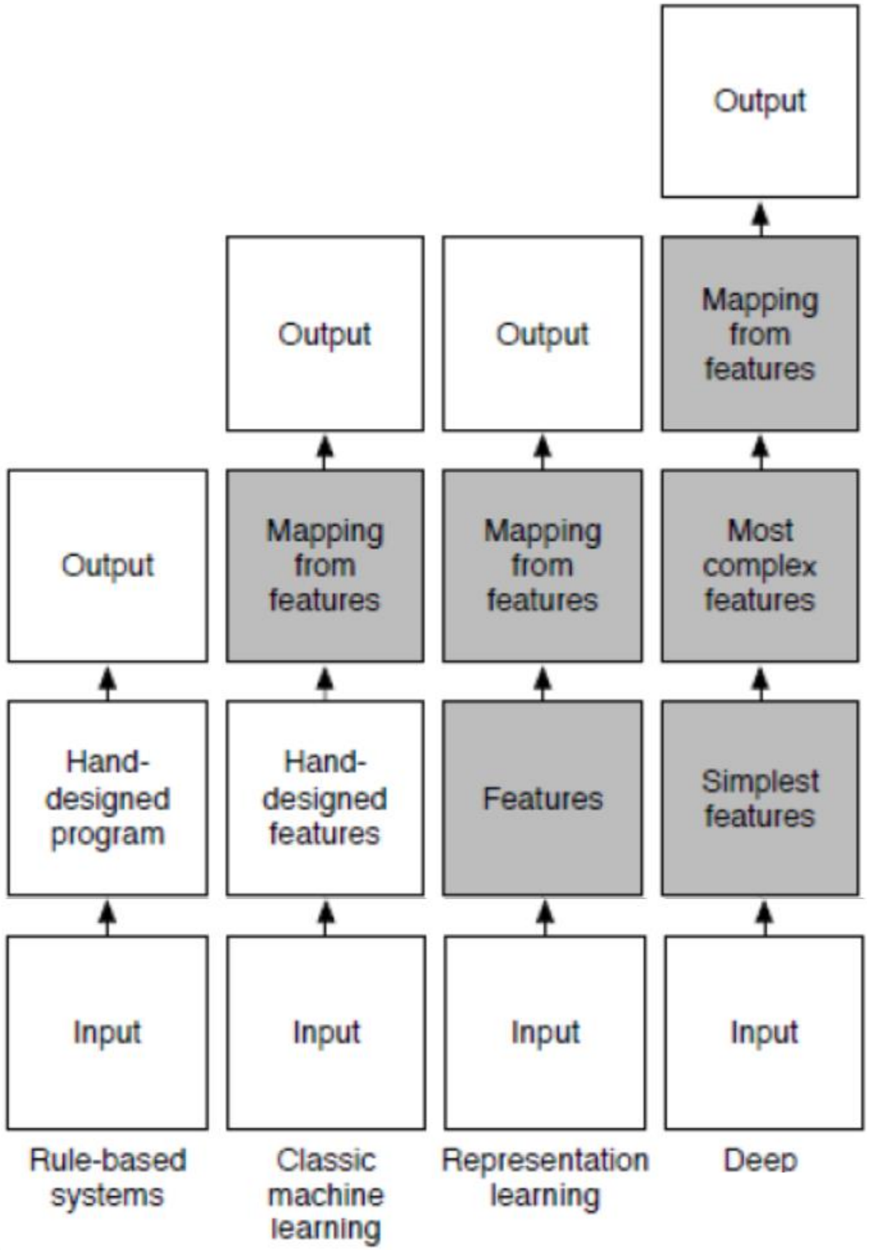

Deep learning history

- **Google Brain is a deep learning research project at Google**
  - **In 2013, Google acquired DNNresearch Inc created by Geoffrey**
  - **Hinton.**
  - **In 2014 bought 'startup' Deepmind Technologies at London**
- **Deep Mind: Start up-2011**
  - **Demis Hassabis Shane Legg, Mustafa Suleyman**

Deep learning success: Images

- **Challenge ImageNet Classification with Deep Convolutional Neural Networks ILSVRC-2012 competition**
  - **over 15 million labeled high-resolution images belonging to roughly 22,000 categories**

- **Playing Atari with Deep Reinforcement Learning**
  - **NIPS Deep learning workshop 2013**
  - **Using only pixels and game points as input can learn to play highly competitively**

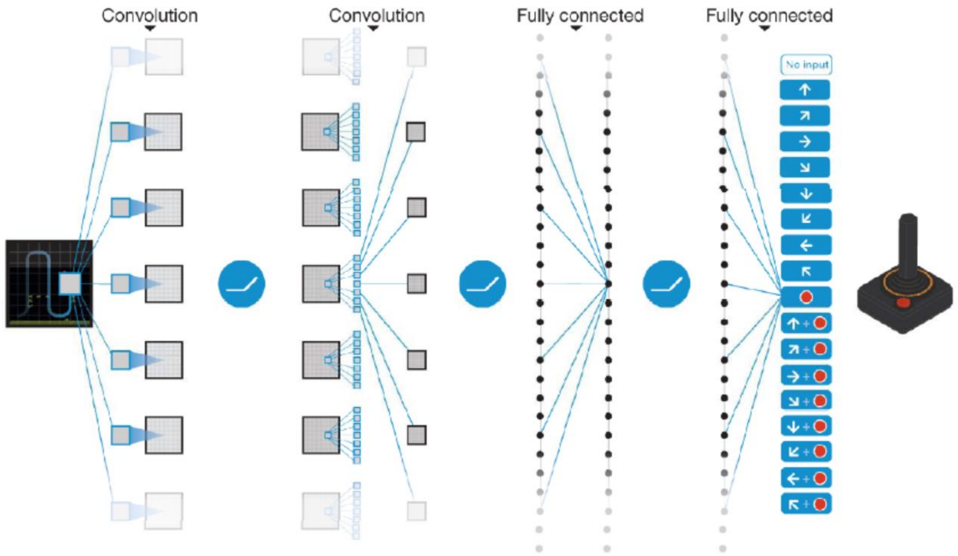

Deep learning success: Games

nature
International journal of science

Article | Published: 27 January 2016

Mastering the game of Go with deep neural networks and tree search

David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel & Demis Hassabis

Nature **529**, 484–489 (28 January 2016) | Download Citation

Emerging Technology From the arXiv
September 14, 2015

Deep Learning Machine Teaches Itself Chess in 72 Hours, Plays at International Master Level

In a world first, a machine plays chess by evaluating the board rather than using brute force to work out every possible move.

The computer that mastered Go

Deep learning issues

- Lack of transparency/explanaiability
- Performance is limited to availability of data
  - **Large amounts of data required**
  - **Training data for rare events sparse**
- Resuts hard to explain
  - **Black box has no visibility**
  - **Research on AI continues in this sense**

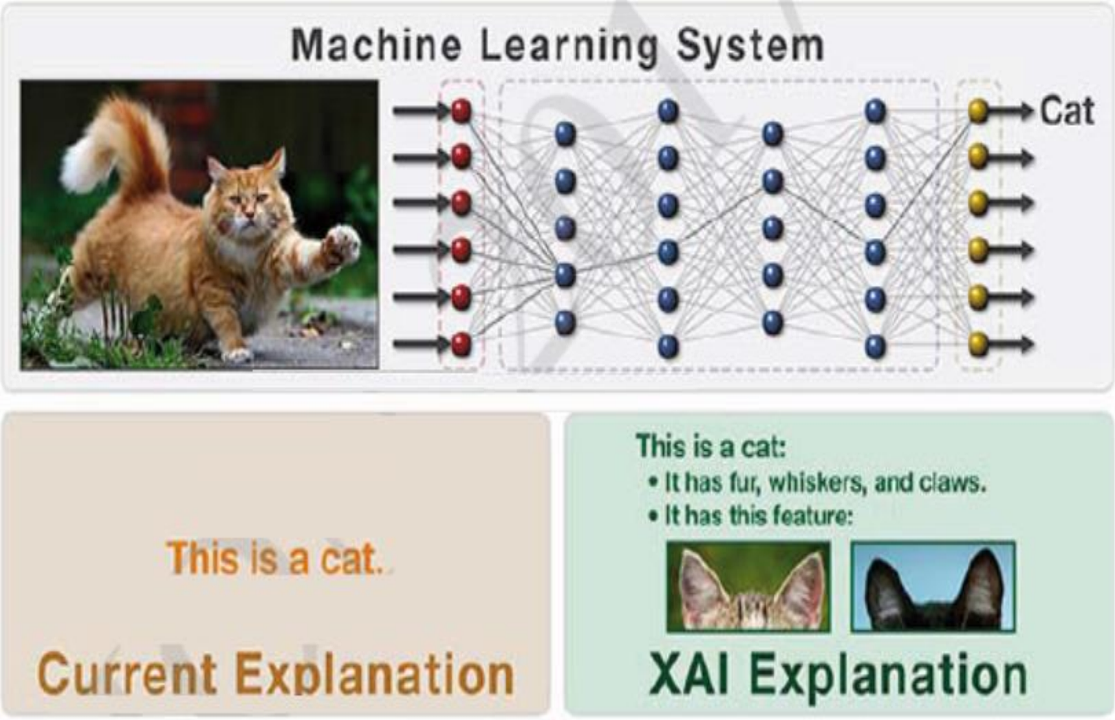

- Key for AI Applied for critical missions or making life/death decisions

Third wave of AI?

- Explaining the AI

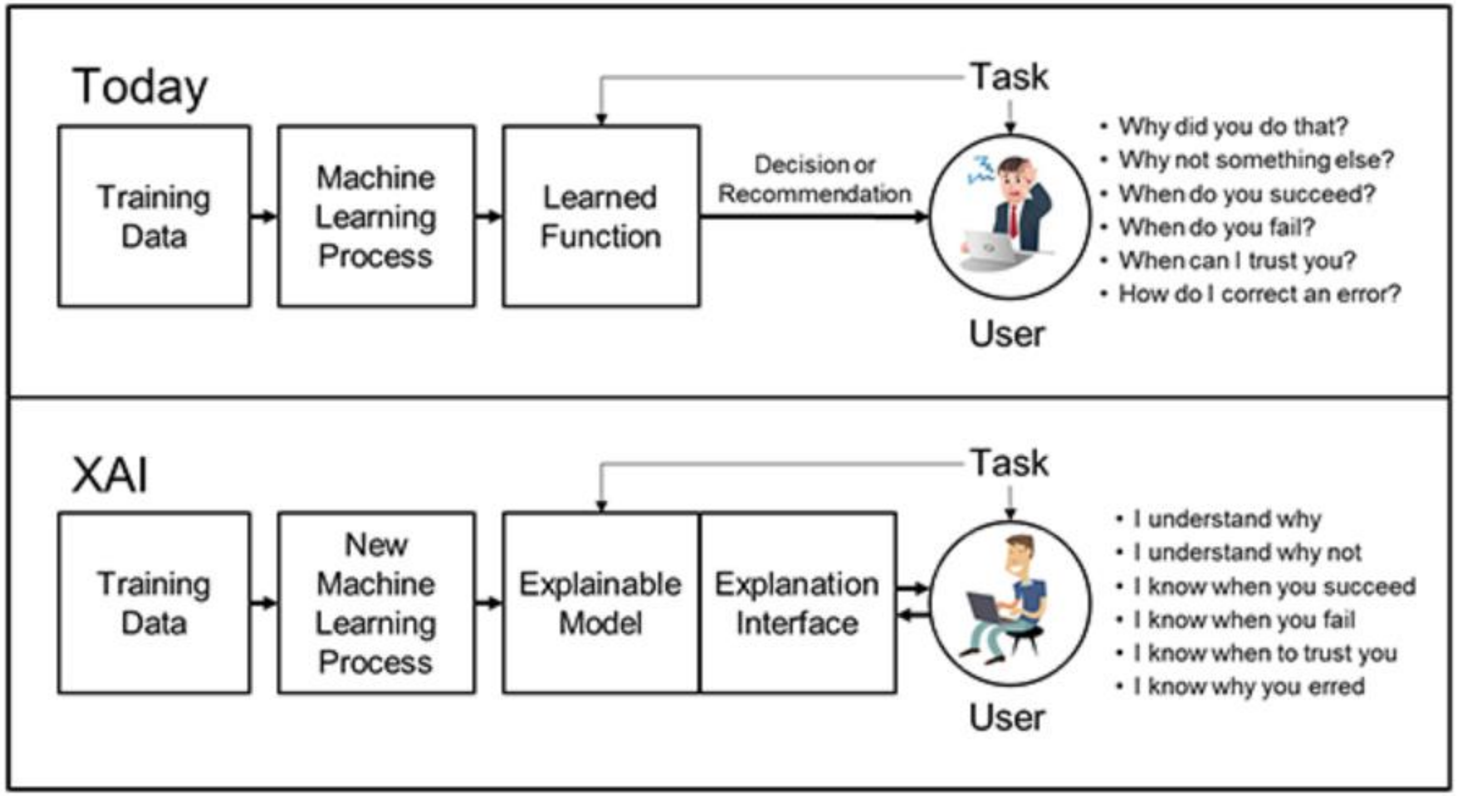

DEFENSE ADVANCED RESEARCH PROJECTS AGENCY

https://www.darpa.mil/program/explainable-artificial-intelligence

Third wave of AI?

- Boosting the implementation of 4 features of AI

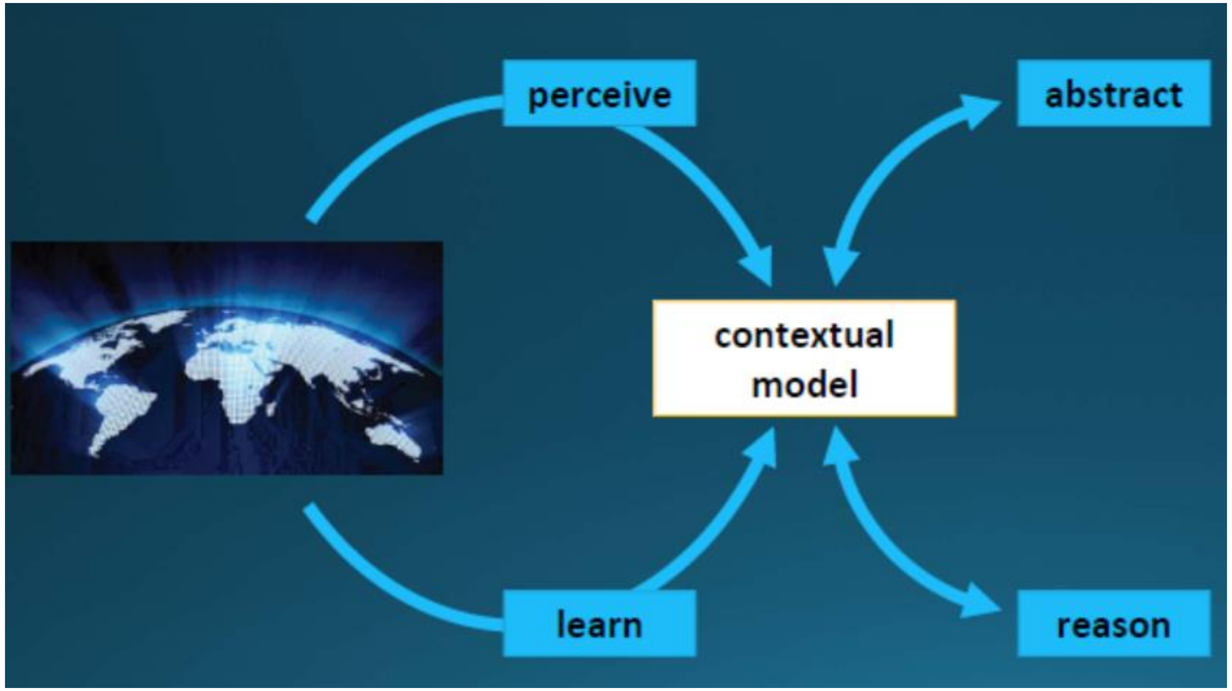

1st wave

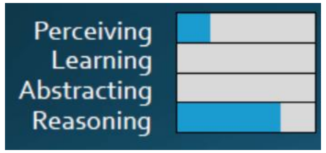

2nd wave

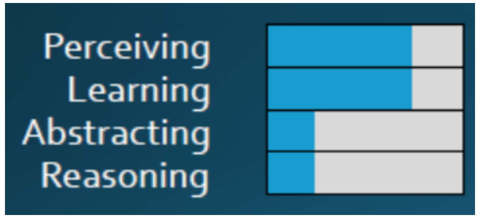

3rd wave

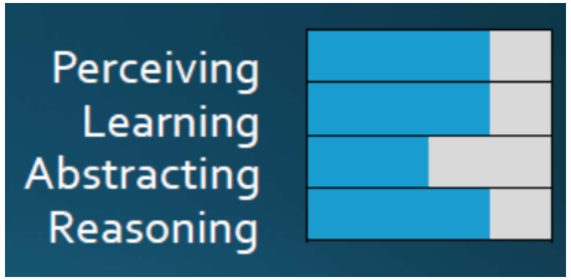

https://www.darpa.mil/program/explainable-artificial-intelligence

Other paradigms: distributed AI

- Distributed Problem Solving (DPS)
  - **sub-algorithms that collaborate**
  - **Ej.: blackboard systems**
- Parallel AI (PAI)
  - **Split algorithms to improve performance**
  - **Typically multi-processors**
- Multi-agent systems (MAS)
  - **Each component «thinks» autonomously**

Other paradigms

- **Distributed Artificial Intelligence**

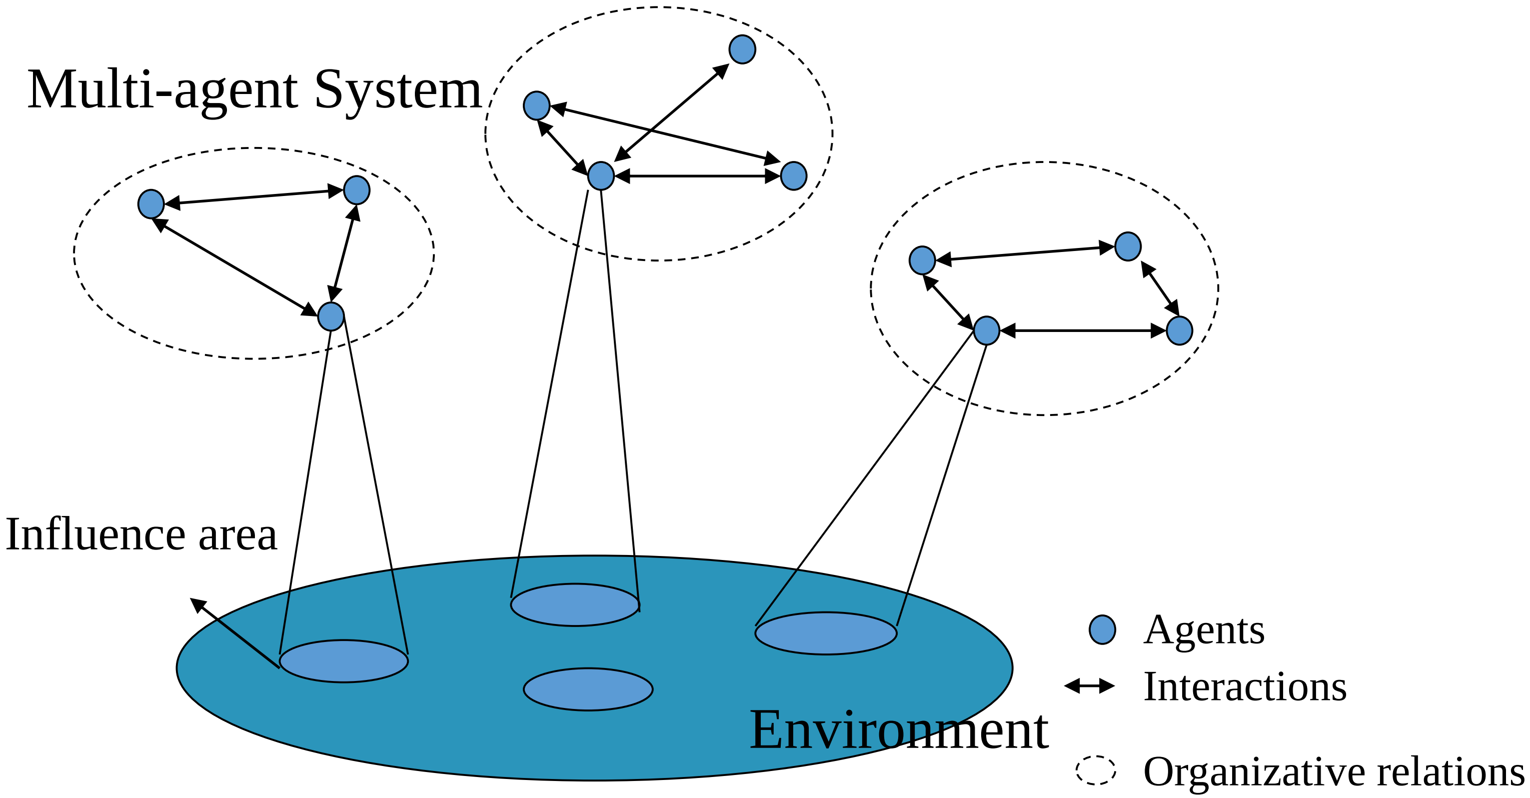

Cognitive agents

- Cognitive agents. BDI

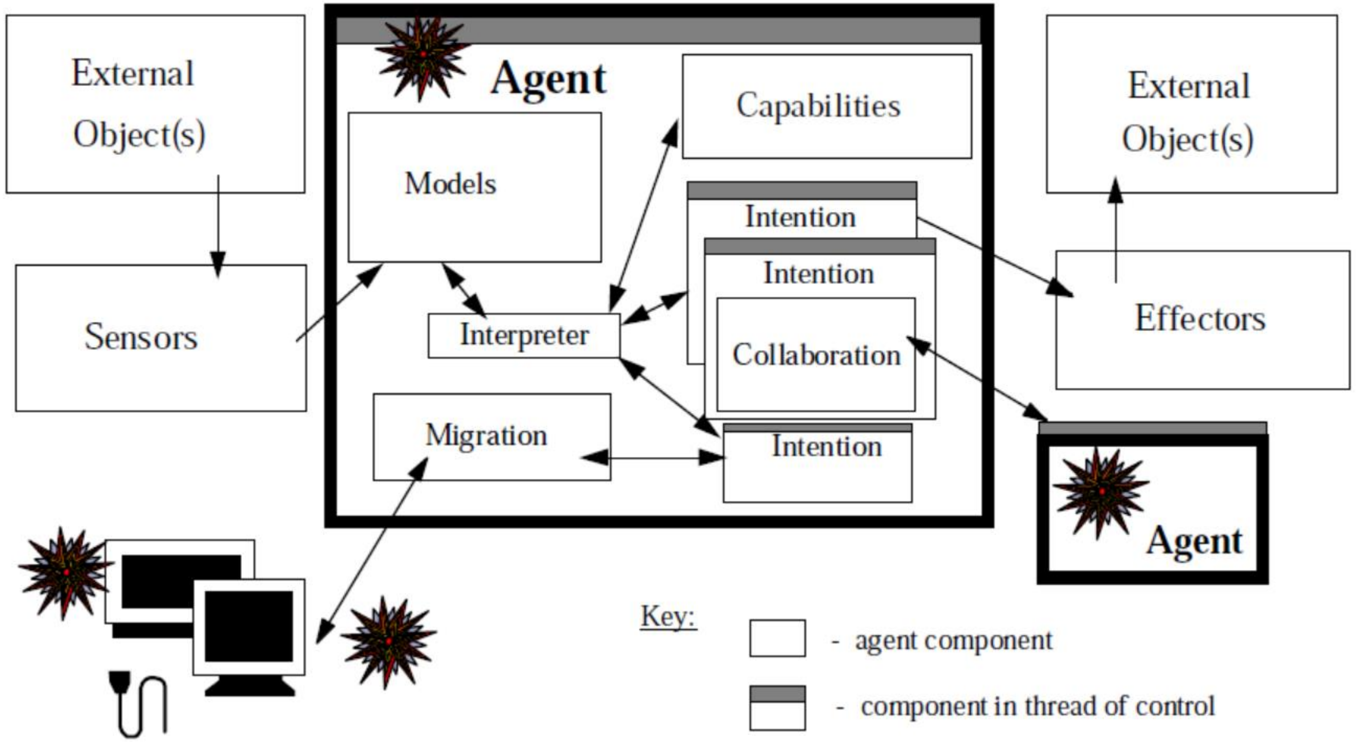

[kendall] Kendall, Pathak, Murali Krishna and Suresh, “The Layered Agent Pattern Language”

Cognitive agents

[Wooldridge]

- **Weak agent**
  - **i) autonomous**
  - **ii) reactive**
  - **iii) pro- active**
  - **iv) collective behaviour**

- **Strong agent**, adds one or more of the following:
  - **v) mentalistic notions (beliefs, goals, plans, and intentions)**
  - **vi) rationality**
  - **vii) veracity**

[DARPA]

Cognitive systems: Can reason and use knowledge

- **Can learn from experience and improve performance**
- **Explain itself and be told what to do**
- **Can be aware of its own capabilities and reflect on itself**
- **Can interact with their environment**
- **Can respond robustly to unpredicted changes**
- **Can achieve human-like performance in activities requiring context-specifi knowledge**

About robustness

- Natural question with any distributed AI solution
- Robust design is the primary source of complexity motivated by biological and technological systems [Reynolds]
  - Self-Organized Criticality (SOC) **"complexity emerges in systems that are otherwise internally homogenous and simple"**
  - Highly Optimized Tolerance (HOT) **"complexity associated with intricately designed or highly evolved systems"**

AI paradigms in ATM

- Connections with ATM automation roadmap

1st Scenario: Local optimization: AI solutions embedded within the machine, improving the capability of the machine to support the human to resolve problems.

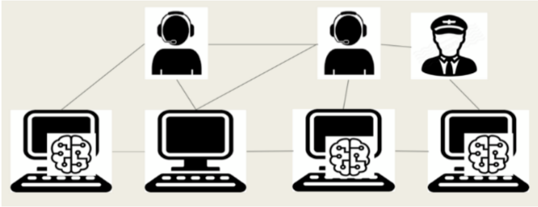

2nd Scenario: Holistic cognitive support to ATM: Human remains in control, but the system gains a certain level of autonomy, informing the human user about its actions.

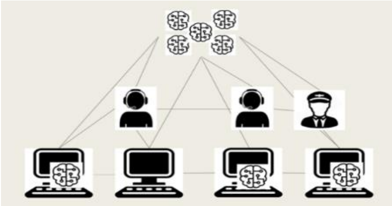

3rd Scenario: Autonomous ATM: Complete change of human and machine roles: Technical systems autonomously decide and execute actions, and all ATM functionality is based on M2M interactions.

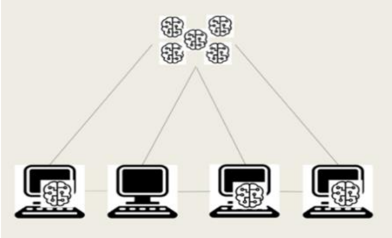

Conclusions

- Remarkabel success of AI/DL paradigms in certain challenges can be now transferred to diverse problems. Ex.: driverless cars
- Most of current successes can be described as “narrow” domains with clearly defined tasks (mainly classification, reinforcement learning, etc)
- A current challenge of AI is explainability/transparency and intrinsic dependence on training sets: critical in many applications
- Distributed systems open new challenges, analysis of emergent behaviour and design for robust performance under uncertainty