

Introduction to Artificial Intelligence

Walid Gomaa

**Professor at Egypt Japan University of Science and Technology,
Alexandria, Egypt**





Outline

- What is Intelligence?
- The Turing Test.
- Short History of AI.
- Quantum Computation.
- Deep Learning.
- Applications of AI.
- Challenges.
- Demos.



What is Intelligence?

- “The true sign of intelligence is not knowledge but imagination.”, [Einstein](#).
- “I know that I am intelligent, because I know that I know nothing.”, [Socrates](#).
- Three major theories in the 20th centuries.





Charles Spearman [1904]



- Different types of intelligence.
- **Correlated**: doing well in some sections of IQ test → tend to do well on all sections.
- General intelligence factor g (remains controversial to this day).



Howard Gardner [1983]



- Theory of Multiple Intelligences.
- Eight distinct types of intelligence.
- Needn't be correlated.
- One can emotionally or socially be intelligent, but not necessarily possess analytical intelligence.



Robert Sternberg [1985]




- Triarchic Theory of Intelligence:
 - analytic, creative, and practical.
- More of a cognitive approach rather than **psychometric** (scores from intelligence tests) approach.



The Turing Test




- Aka the **imitation game**.
- Operational definition introduced by Alan Turing in 1950.





Multidisciplinary Field





3 Types of Artificial Intelligence

Artificial Narrow Intelligence (ANI)



Stage-1

Machine Learning


- Specialises in one area and solves one problem



Siri



Alexa



Cortana

Artificial General Intelligence (AGI)



Stage-2

Machine Intelligence

- Refers to a computer that is as smart as a human across the board

Artificial Super Intelligence (ASI)



Stage-3


Machine Consciousness

- An intellect that is much smarter than the best human brains in practically every field




[1943-1955] Gestation of AI

- Alan Turing: the Turing test (article 'Computing Machinery and Intelligence', 1950).
- McCulloch & Pitts [1943]: model of artificial neurons (Boolean circuit model of the brain).
- Hebb [1949]: simple updating rule for modifying the connection strengths between neurons (Hebbian learning).
- Marvin Minsky & Dean Edmonds [early 1950's]: first NN called SNARC.



Hebbian Learning






1956-The Birth of AI

- Dartmouth meeting to adopt AI (John McCarthy, Marvin Minsky, Claude Shannon, Nathaniel Rochester).



- Conjecture: every aspect of intelligence can be precisely described in a way understandable by the machine.
- A. Newell & H. Simon's **Logic Theorist**: reasoning system able to prove logical assertions and mathematical theorems.





[1952-1969] Early Enthusiasm & Great Expectations

- General Problem Solver (GPS): imitate human-like planning.
- Gelernter's Geometry Engine: proving tricky theorems of geometry.
- Checkers: strong amateur levels.
- **Lisp: by McCarthy - dominant AI language for 30 years.**
- M. Minsky: limited domain problems (calculus integration, blocks world, etc).
- Herbert Simon (1965): “Machines will be capable, within twenty years, of doing any work that a man can do”.





[1966-1973] More Realistic

- Early machine translation efforts.
- Neural network research almost disappears.
- H. Simon predicted AI be chess champion and proves a significant mathematical theorem within 10 years (it becomes a reality ~40 years).






[1974-1980] The First AI Winter

- US government canceled all funding for machine translation.
- Funding for neural networks almost disappeared as M. Minsky and Pappert showed that perceptrons (linear classifiers) are not powerful enough.
- 1973: Lighthill report → failure of AI, intractability, combinatorial explosion, only toy problems → UK government cut funding for almost all AI research programs.



[1969-1979] Knowledge-Based Systems

- Domain specific knowledge in narrow areas of expertise.
- Expert systems.
- DENDRAL: inferring molecular structure from the information provided by a mass spectrometer.
- MYCIN: diagnose blood infections. Better than junior doctors.






[1980-present] AI Becomes an Industry

- 1982: first successful expert system RI at the Digital Equipment Corporation
 - configure orders for new computer systems.
 - By 1986, saving about \$40M/year.
 - By 1988, DEC AI group had 40 expert systems deployed.
- 1981: Japan announced the “Fifth Generation” project, 10 years to build intelligent machine running Prolog.
- US responded by forming the Microelectronics and Computer Technology Corporation (MCC) as a research consortium.
- UK: Alvey report countered the Lighthill report and AI funding was reinstated.



[1986-Present] The Return of NN

- Mid 1950s: the back-propagation algorithm.




- Connectionist models as competitors to symbolic and logical approaches.



[1987-1990] Second AI Winter

- 1987: the collapse of the Lisp machines.
- Early 1990's: expert systems too expensive to maintain, difficult to update, not robust, too specialized.
- Early 1990s: “Fifth generation” project goals have not been met → funding cuts.
- Funding cuts also in the US.





[1987-Present] Adopting Scientific Method

- Opening to other areas: information theory, statistical and stochastic analysis, optimization and control, etc.
- To accept a hypothesis:
 - It must be subjected to rigorous experiments.
 - It must be analyzed statistically for its importance.
 - Possibility to replicate experiments and regenerate the results using shared repositories of shared data and code.



[1995-Present] Intelligent Agents


- Internet: widespread applications such as “bots”.
- Integration with other fields such as economics and control theory in order to understand, model, and build intelligent agents (BDI models).
- Looking for a universal algorithms that would learn and act **autonomously** in any environment.





[2001-Present] Very Large Data Sets


- Data-oriented rather than algorithmic/rule oriented.
- Same algorithm: more data gives better performance.
- Learn knowledge from data over hard-coded knowledge.





[2010-Present] Deep Learning Revolution

- Significant advances in machine learning, in particular, deep learning.
- Computer vision, speech recognition, and machine translation are dominated by deep learning.





[Recent] Quantum Computation

- Traditional computing paradigm uses classical physics.
- Quantum computation uses quantum physics for its realizations.



Quantum Computation

- Qubits replaces classical bits.
- Classical bit: either 0 or 1.
- Qubit: superposition of 0 and 1: $a|0\rangle + b|1\rangle$
- Qubit is represented by a wave function.
 - Example: the spin of an electron or the polarity of a photon.

Bit
(Classical Computing)


0



1

Qubit
(Quantum Computing)

0



1



Essential Properties

- **Interference:**
 - constructive/destructive.
 - Exploited in quantum algorithms (David Deutsch).
- **Entanglement:**
 - Multiple particles wave functions can become entangled.
 - Manipulating/measuring one particle necessarily affects the other.





Why Quantum Computation?

- More efficient algorithms.
- Quantum parallelism (machine learning, simulating complex phenomena in physics, chemistry, drug design, developing new materials, etc).
- Breaking classical crypto systems (security/crypt-analysis).
- Building better crypto systems (security/cryptology).
- Simulating quantum systems.



Tycoons

- Google Quantum AI research group.
- Google Bristlecone: 72-qubit processor → approaching “quantum supremacy” - quantum computers can do calculations beyond the reach of today’s fastest supercomputers.
- IBM quantum research: the QISKit (Open Source Quantum Information Software Kit).
- IBM 50-qubit processor.
- D-Wave Systems Inc.: quantum computing Canadian company.




Deep Learning

- Computer vision.
- Speech recognition (moving from hidden Markov models).
- Natural language processing.



Deep Learning Paradigm

Traditional programming





Machine learning





Why Now?

- Abundance of data in digitized form.
- Advancements in high-performance computing (GPUs, TPUs, etc).
- Algorithmic advancements, specially training architectures and optimization methods.






The Philosophy of Neural Networks


- **Symbolic AI**: intelligence is symbol manipulation by rigid logical rules.
- **Connectionism**: concepts in our brains are not represented by symbols, but by patterns of activations.



Nonlinear Transformation




$$\sigma(x) = \frac{1}{1 + e^{-x}}$$



introducing nonlinearity




MultiLayer Perceptron (Fully Connected NN)





Keyword: Representation




- **Representation learning:** automatically learning good features or representations.
- **Deep learning:** automatically learn multiple levels of representations of increasing complexity and abstraction.




Shallow Learning


Classification



Denoising




OCR






Deep Learning






Fully Connected Layer

$32 \times 32 \times 3$ image $\rightarrow 3,072 \times 1$ feature vector





Convolutional Layer






Convolutional Neural Network


CNN (Convent)

Sequence of convolution layers interspersed with activation functions



ReLU: $f(x) = \max\{0, x\}$
introducing nonlinearity
(piecewise linear)

(gives better performance than logistic functions)








Applications of AI

- Science.
- Industry.
- Business.
- Military.



Science

- Prediction of protein folding and Development of new drugs (CADD)
— AlphaFold 2.
- Telesurgery (remote surgery).
- Weather modeling and prediction.
- Dermatologist-level classification of skin cancer (using deep learning, Nature article).





Industry

- Inspection and sorting systems using machine vision.
- Smart factories and the dawn of Industry 4.0.





Business

- Advanced financial systems (predictions of stock market).
- Fraud prevention.
- Chatbots for customer service.
- Risk management.



Military

- Smart bombs.
- Unmanned drone aircrafts.
- Robotic soldiers.
- Decoding of secret codes.





Challenges

- Data bias.
- Vast amounts of data.
- Wartime robots.
- Too specialized.
- Transparency problem.



Data Bias

- Predictive policing in the US.





Vast Amounts of Data

- $\sim 10^5$ more data required than a human being to understand a concept.
- Only big companies (amazon, google, IBM, etc) can afford that.
- Sparsity and scarcity of data in many fields such as healthcare (recognizing tumors in X-ray scans).

Need to develop more efficient DNN that can work with less data.



Wartime Robots

- Establishing responsibility (also applies for autonomous vehicles).
- Lowering the threshold for warfare engagement → violating conditions of just war.
- War zone too complex for robots to discriminate targets.
- Integration with human fighters.
- Difficulties in winning the hearts and minds of the defeated side.
- Proliferation of the technology to terrorist groups.





Too Specialized

- Very efficient in single task, complete ignorance of any other.
- Training for multiple tasks caused interference or forgetting.

True AI requires learning multiple tasks (transfer learning, progressive neural networks, Artificial General Intelligence AGI).




Transparency Problem

- NN are are inscrutable to observers.
- Why do they work? and make these decisions?
- Need to build more coherence into the network from the get-go. Possibly using symbolic logic-based AI.
- Prior knowledge/starting point for understanding the world.



Aggressive Flight (Kumar Lab, UPenn)



Thank You