Artificial Intelligence (IT3160E)

Than Quang Khoat

khoattq@soict.hust.edu.vn

School of Information and Communication Technology Hanoi University of Science and Technology

Content:

- Introduction of Artificial Intelligence
 - Definition
 - Foundation fields
 - Brief history
 - Successful practical applications
 - Software frameworks and libraries
- Intelligent agent
- Problem solving: Search, Constraint satisfaction
- Logic and reasoning
- Knowledge representation
- Machine learning

AI & DS & Industry 4.0



Some successes: Watson (2011)



IBM's Watson Supercomputer Destroys Humans in Jeopardy

Some successes: GAN (2014)

Enable imagination for a machine



Artificial faces





Goodfellow, Ian, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. "Generative adversarial nets." In *NIPS*, pp. 2672-2680. 2014.

Some successes: AlphaGo (2016)

- AlphaGo of Google beat a top player at Go (cờ vây), 3/2016
 - □ Go is a 2500 year-old game.
 - □ Go is one of the most complex games.
- AlphaGo learns from 30 millions human moves, and plays itself to find new moves.
- It beat Lee Sedol (World champion)
 - http://www.wired.com/2016/03/two-movesalphago-lee-sedol-redefined-future/
 - http://www.nature.com/news/google-ai-algori masters-ancient-game-of-go-1.19234







Some successes: AlphaFold (2021)

Accurate prediction of protein folding





Jumper, John, et al. "Highly accurate protein structure prediction with AlphaFold." Nature 596.7873 (2021): 583-589.

Some successes: ChatGPT (2022)

Human-level Chatting, Writing, QA,...

trời đẹp quá, làm cho tớ bài thơ đi

SO

\$

S OpenAI

Tạm biệt mây mù, đón tới ánh sáng Trời đẹp mới rực, tấm lòng tớ tăng Sức mạnh của mùa xuân, giữa cánh đồng hoa Vẻ đẹp tự nhiên, làm cho tớ mãi mãi mê hoặc.

Lấp lánh tràn đầy niềm vui, nụ cười tràn đầy niềm hạnh phúc Không cần bất kỳ điều gì, chỉ cần trời đẹp để tớ hạnh phúc. Why ChatGPT is about to change how you work, like it or not? - Forbes, Feb. 2, 2023

BUSINESS

Audio Live TV

Log In

ChatGPT passes exams from law and business schools

By Samantha Murphy Kelly, CNN Business

Updated 1:35 PM EST, Thu January 26, 2023

Some successes: ...







Artificial Intelligence

Global impact of AI



Until 2030, AI can create **\$15000 billions**

Impact on Growth of other areas

Insurance	(12.8% - 24.4%) 16.8%	Banking	10.3%
High Tech	15.5%	Consumer packaged goods	9.4%
Pharmaceutical & Medical Products	16.1%	Aerospace & Defense	10.4%
Travel	15.1%	Transport & Logistics	10.1%
Advanced Electronics / Semiconductors	14.9%	Retail	8.7%
Telecom	13.2%	Healthcare Systems and Services	10.0%
Media & Entertainment	12.4%	Agriculture	8.3%

https://www.mckinsey.com/featured-insights/artificial-intelligence/visualizing-the-uses-and-potential-impact-of-ai-and-other-analytics

Definition of AI (1)

- The definitions (views) of Artificial Intelligence (AI) can be categorized in 4 groups:
 - □ (1) Systems that **think like humans** (suy nghĩ như con người)
 - "The exciting new effort to make computers think ... machines with minds, in the full and literal sense." (Haugeland, 1985)
 - "[The automation of] activities that we associate with human thinking, activities such as decision-making, problem solving, learning ..." (Bellman, 1978)
 - □ (2) Systems that **think rationally** (suy nghĩ hợp lý)
 - "The study of mental faculties through the use of computational models." (Charniak and McDermott, 1985)
 - "The study of the computations that make it possible to perceive, reason, and act." (Winston, 1992)

Definition of AI (2)

□ (3) Systems that **act like humans** (hành động như con người)

- "The art of creating machines that perform functions that require intelligence when performed by people." (Kurzweil, 1990)
- "The study of how to make computers do things at which, at the moment, people are better." (Rich and Knight, 1991)
- (4) Systems that act rationally (hanh động hợp lý)
 - "Computational Intelligence is the study of the design of intelligent agents." (Poole et al., 1998)
 - "Al ... is concerned with intelligent behavior in artifacts." (Nilsson, 1998)

Definition of AI (3)

- The definitions (1) and (2) relate to *thinking and inference processes*
- The definitions (3) and (4) relate to *actions*
- The definitions (1) and (3) assess the success (i.e., intelligence) at the level of human intelligence
- The definitions (2) and (4) assess the success (i.e., intelligence) at the level of rationality
 - A system is considered acting rationally if it does its jobs according to what it (the system) knows

 Artificial Intelligence (AI) is the science and engineering of making intelligent machines, especially intelligent computer programs

[John McCarthy, Stanford University, http://www-formal.stanford.edu/jmc/whatisai/node1.html]

Acting humanly: Turing test

Turing (1950), "Computing machinery and intelligence":

- "Can machines think?" \rightarrow "Can machines behave intelligently?"
- Test for intelligent behavior: the Imitation Game



- Predicted that by 2000, a machine might have a 30% chance of surpassing a non-expert person for a Turing test in 5 minutes
- Suggested major components of AI: knowledge, reasoning, language understanding, learning

Acting rationally

- Rational behavior: Doing the right thing
- The right thing: That which is expected to maximize goal achievement, given the available information
- Doesn't necessarily involve thinking
 E.g., blinking reflex
- But *thinking* should be in the service of rational action
- The rationality should take the computation cost into account
 - If the computation resource and time costs are too high, then it is impractical (i.e., not applicable in practice)

Rational agents (1)



- An **agent** is an entity that perceives and acts
- Generally, an agent is a function from percept histories to actions:

$$f: \mathsf{P}^* \to \mathsf{A}$$

Rational agents (2)



- For an environment and a task, we need to find out an agent that has the best performance
- An intelligent agent is the one that can act rationally (i.e., intelligently)
 - Action that helps maximize the achievement of the goal(s), given the perceived information
- Important note: Limits of computation (of the computer) do not allow perfect (optimal) rationality to be achieved
 - \rightarrow Intelligence vs. computation cost (practicality)

Related fields of AI (1)

Philosophy

- Logic
- Methods of reasoning
- Foundations of learning
- Language
- Rationality

Mathematics

- Formal representation and Proof algorithms
- Computation
- Decidable vs. undecidable problems
- Tractable vs. intractable problems (i.e., computational complexity, especially time cost)
- Probability

Related fields of AI (2)

Economics

- Utility function
- Decision making theory

Neuroscience

Natural basis of mental activities

Psychology

- Adaptivity
- Phenomena of perception and motor control
- Experimental techniques (psychophysics, etc.)

Related fields of AI (3)

Computer technology

- Build high-speed computers
- High performance computing
- Control theory
 - Design systems to maximize a certain objective function

Linguistics

- Knowledge representation
- Grammar (of a language)

Brief history of AI (1)

- 1943: McCulloch & Pitts presented the first research on AI, which proposed modeling of two-state (i.e., on/off) artificial neurons
- 1950: The concept of AI was first mentioned by Turing in his article "Computing Machinery and Intelligence"
- 1956: The first workshop (taking place in 2 months) in Dartmouth (USA) discussing the field of AI. AI was born.
- 1952-1969: The initial achievements in AI
- 1950s: First AI programs
 - Samuel's chess program
 - Newell & Simon's logic reasoning program
 - Gelernter's geometric theorem proving program

- John McCarthy
- 1965: Robinson proposed the complete algorithm for logic reasoning



Brief history of AI (2)

- 1966-1973: the first AI winter
 - Al researchers realized the difficulty of computational complexity
 - Artificial neural networks were heavily influenced, and developed very slowly
- 1969-1979: Introduction and early development of knowledge-based systems
- 1980: AI systems and programs were used commercially
- 1980-1988: The emergence of expert systems
- 1986: Artificial neural networks became popularly
- 1987: AI became a scientific field
- 1987-1993: the second AI winter

The main subfields of AI

- Search & Planning: giải quyết vấn đề có mục tiêu bằng tìm kiếm
- Knowledge representation & Reasoning: biểu diễn tri thức và sử dụng chúng hiệu quả
 - Expert systems, Knowledge-based systems, ...
- Machine Learning (1959): giúp máy tính có khả năng học từ dữ liệu
 - □ Reinforcement Learning, Deep Learning, ...
- Machine perception: giúp máy có khả năng nhận thức hình ảnh, ngôn ngữ, ...
 - D Computer Vision, Natural Language Processing, Speech Recognition
- Robotics: máy có khả năng tự hoạt động và ứng xử hiệu quả
- Multi-agent systems: giúp các máy tương tác với nhau

Important achievements in AI (1)

Information retrieval

Virtual assistant: Siri, Google Now, Cortana, Bixby, etc.



- Human-machine communication
 - Voice, Gesture, Natural language understanding, etc.





Artificial Intelligence

Important achievements in AI (2)

Entertainment

Music, Movies, Games, News, Social networks, etc.



Transportation

 Shelf-driving car, Traffic law enforcement, Prediction of demand for car/motorbike ride, etc.







Artificial Intelligence

Coogle DeepMind

Important achievements in AI (3)

- Education and learning
 - Learning materials, Learning path, Knowledge dissemination, etc.





E-commerce

Product/service recommendations, Demand prediction, Promotion campaign, etc.



Artificial Intelligence

Important achievements in AI (4)

System security

 Computer virus detection, Network intrusion detection, Email spam filtering, etc.





Marketing and advertisement



Artificial Intelligence

Successful application fields of AI (1)

E-commerce

- Personalized/target advertisement, Product and service recommendation, etc.
- Entertainment
 - Games, Music, Movies, News, etc.
- Finance
 - Market analysis, Stocks investment, Loan risk estimation, Card fraud detection, etc.
- Manufacturing
 - Defect product detection, Maintenance status prediction, Robots work in production lines, etc.
- Medicine and health
 - Disease diagnostics, Interpretation of x-ray images, Heart rate/brain wave/blood vessel analysis, Micro-surgery robot, etc.

Successful application fields of AI (2)

- Telecommunications
 - Automatic customer support, Data routing and transmission, etc.
- Aeronautics and space
 - Planning the operations of spacecraft, Universe station maintenance prediction, Satellite control, etc.
- Nuclear plant management
 - Problem/risk prediction and warning, etc.
- Military
 - Object recognition and classification, etc.
- And there are many other application fields ...

Software frameworks and libraries (1)

- TensorFlow (www.tensorflow.org)
 - OS: Linux, Mac OS, Windows, Android
 - Languages: Python, C++, Java
- **Caffe** (caffe.berkeleyvision.org)
 - OS: Linux, Mac OS, Windows
 - Languages: Python, Matlab
- **Caffe2** (caffe2.ai), **PyTorch** (pytorch.org)
 - □ In march 2018, Caffe2 and PyTorch were merged in the unified architecture
 - OS: Linux, Mac OS, Windows, iOS, Android, Raspbian
 - Languages: C++, Python
- Keras (keras.io)
 - OS: Linux, Mac OS, Windows
 - Language: Python
- Theano (deeplearning.net/software/Theano)
 - OS: Linux, Mac OS, Windows
 - Language: Python

Software frameworks and libraries (2)

- CNTK (www.microsoft.com/en-us/research/product/cognitivetoolkit/)
 - OS: Windows, Linux
 - Languages: Python, C++, C#
- Deeplearning4j (deeplearning4j.org)
 - OS: Linux, Mac OS, Windows, Android
 - Languages: Java, Scala, Clojure, Python
- Apache Mahout (mahout.apache.org)
 - OS: Any OS with JVM installed
 - Languages: Java, Scala
- Weka (http://www.cs.waikato.ac.nz/ml/weka/)
 - OS: Any OS with JVM installed
 - Language: Java

Open debates about AI (1)

- The ability of AI?
 - Play correctly a table-tennis game?
 - Discover and prove a new mathematical theory?
 - Can converse with one person in 1 hour?
 - Automatically perform a complicated surgery?
 - Instantly translate between bilinguals in a conversation?
 - etc.
- Can computers think like humans?
- Can computers do reasoning?

Open debates about AI (2)

- If computers can replace what is being done by humans, the fewer jobs (unemployed)
- Humans will have too much spare time? (compared to too little, as it is today)
- People feel a loss of their dominant (highest) intelligence
- The use of multiple AI systems can reduce (loose) accountability at work
- Since computers do (and interfere with) many human everyday things, they will feel their privacy is compromised
- The (perfect) success of AI is the end of humanity?

References

- R. E. Bellman. An Introduction to Artificial Intelligence: Can Computers Think? Boyd & Fraser Publishing Company, San Francisco, 1978.
- E. Charniak and D. McDermott. *Introduction to Artificial Intelligence*. Addison-Wesley, Reading, Massachusetts, 1985.
- J. Haugeland. *Artificial Intelligence: The Very Idea*. MIT Press, Cambridge, Massachusetts, 1985.
- R. Kurzweil. The Age of Intelligent Machines. MIT Press, Cambridge, Massachusetts, 1990.
- N. J. Nilsson. *Artificial Intelligence: A New Synthesis*. Morgan Kaufmann, San Mateo, California, 1998.
- D. Poole, A. K. Mackworth, and R. Goebel. *Computational Intelligence: A Logical Approach*. Oxford University Press, Oxford, UK, 1998.
- E. Rich and K. Knight. *Artificial Intelligence* (Second Edition). McGraw-Hill, New York, 1991.
- P. H. Winston. *Artificial Intelligence* (Third Edition). Addison-Wesley, Reading, Massachusetts, 1992.