

E016350 - Artificial Intelligence

Lecture 1

Introduction

Part 1: Survey of AI and Course Overview

Aleksandra Pizurica

Ghent University
Fall 2024



Aleksandra Pizurica
Dr. Ir., Full professor
Statistical image modelling and AI
Ghent University (UGent)

- Research Group Artificial Intelligence and Sparse Modelling (GAIM), Department Telecommunications and Information processing (TELIN); <https://gaim.ugent.be/>
- Research: Statistical image modeling; restoration and analysis of high-dimensional data; machine learning; sparse coding; compressed sensing; deep learning; remote sensing; medical imaging; art investigation
- Teaching courses at UGent: **Artificial intelligence**; **Computer graphics**, **Probabilistic graphical models**
- E-mail: aleksandra.pizurica@ugent.be
W: <https://telin.ugent.be/~sanja>

Course staff - Teaching Assistants



Dr. Nicolas Nadisic
Asst. prof., *Head TA*



Nicolas Vercheval
Head TA



Yoann Arhant



Roel Huysentruyt



Xianlu Li



Ide Van den Borre



Niels Vyncke

Contact: ai@lists.ugent.be



TELIN Welcome to the Department of Telecommunications and Information Processing

TE

DDCM

DIGCOM

GAIM

IPI

SMACS

DDCM

Database Document and Content Management

DIGCOM

Digital Communications

GAIM

Group for Artificial Intelligence and Sparse Modelling

IPI

Image Processing and Interpretation

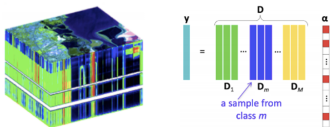
SMACS

Stochastic Modeling and Analysis of Communication Systems

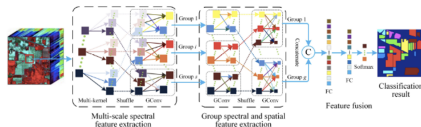
GAIM: Group for Artificial Intelligence and Sparse Modelling

Welcome to GAIM, the Group for Artificial Intelligence and Sparse Modelling in the department TELIN of the Faculty of Engineering and Architecture at Ghent University

Sparse modelling in high-dimensional data analysis



Deep learning in sensor fusion and HSI analysis

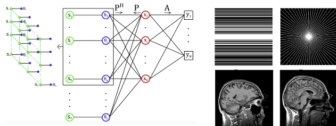


Signal & image processing – machine learning – information theory

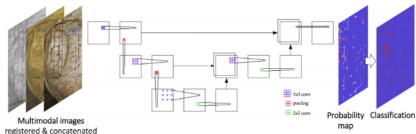
GAIM: Group for Artificial Intelligence and Sparse Modelling

Welcome to GAIM, the Group for Artificial Intelligence and Sparse Modelling in the department TELIN of the Faculty of Engineering and Architecture at Ghent University

Graphical models in sparse image recovery



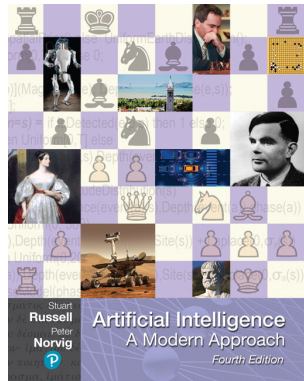
Deep learning in digital painting analysis



Signal & image processing – machine learning – information theory

Study material

- Slides & Lecture notes (available on Ufora)
- Theory exercises (available on Ufora)
- Recommended book:
S. Russel and P. Norvig:
Artificial Intelligence - A Modern Approach
(4th edition; denoted as [R&N] in the slides)



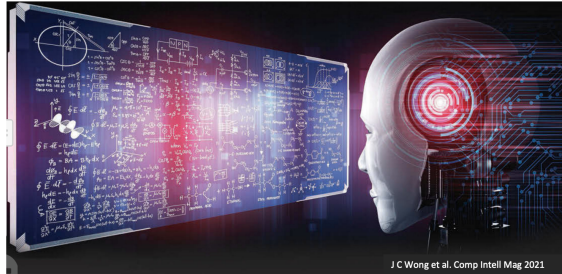
Lectures, lab sessions and exam

- Theory lectures: Tuesdays 13:00 - 15:45
Aud Magnel, Campus Ardoyen
- Exercises: Fridays 8:30 - 11:15
Leslokaal Rudolf E. Richter, Volta,
Campus Ardoyen, (Building 131)
 - ▶ Theory exercises
 - ▶ Computer exercises
 - ▶ Projects
 - ★ Facilitate understanding of the theory
 - ★ Involve Python programming
 - ★ Machine learning tasks
- Guest lecture (advanced topics)

Examination:

- Written exam:
60% of the total score
 - ▶ Theory
 - ▶ Problem solving
- Projects:
30% of the total score
 - ▶ Code, reports
- Homework:
10% of the total score
 - ▶ Computer exercises

Rise of Artificial Intelligence



J C Wong et al. Comp Intell Mag 2021



For the first time, a drug created by artificial intelligence will be tested in humans

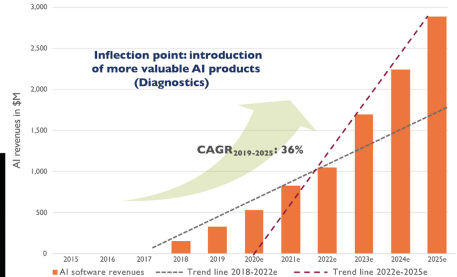
April 30, 2020

In "Artificial Intelligence"



The global evolution forecasts of artificial intelligence revenues for medical imaging 2015-2025

(Source: Artificial Intelligence for Medical Imaging 2020 report, Yole Développement, 2020)



Rise of Artificial Intelligence



ChatGPT reaches 100 million users two months after launch

Unprecedented take-up may make AI chatbot the fastest-growing consumer internet app ever, analysts say



📷 ChatGPT is owned by Microsoft-backed company OpenAI. Photograph: Pavlo Gonchar/Sopa Images/Rex/Shutterstock

Rise of Artificial Intelligence



I'd like to see the rock formations in Cappadocia, Turkey. Can you plan a 3-day trip around that?



Of course! Exploring the unique rock formations in Cappadocia, Turkey, is a fantastic experience. Here's a suggested 3-day itinerary to help you make the most of your trip:



Day 1: Arrival in Cappadocia

- Arrive at either Nevşehir Kapadokya Airport or Kayseri Airport, both of which serve Cappadocia.
- Check into your chosen accommodation in the towns of Göreme, Uçhisar, or Ürgüp. These towns are centrally located and offer easy access to the main attractions.
- After settling in, start your exploration by visiting the Göreme Open-Air Museum. This UNESCO World Heritage site features numerous rock-cut churches with stunning frescoes.
- Enjoy dinner at a local restaurant in Göreme and sample some traditional Turkish cuisine.

Day 2: Hot Air Balloon Ride and Cappadocia Exploration

- Start your day early with a hot air balloon ride over Cappadocia. This is a must-do activity to witness the incredible landscape from above. Be sure to book in advance.

Rise of Artificial Intelligence



Show me a code snippet of a website's sticky header in CSS and JavaScript.



Certainly! A sticky header is a common feature on websites where the header remains visible at the top of the page as you scroll down. Here's a code snippet for creating a sticky header using HTML, CSS, and JavaScript:



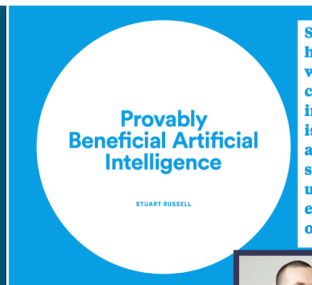
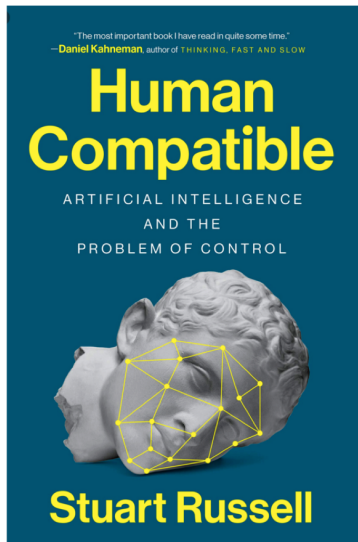
HTML (index.html):

html

Copy code

```
<!DOCTYPE html>
<html lang="en">
<head>
  <meta charset="UTF-8">
  <meta name="viewport" content="width=device-width, initial-scale=1.0">
  <link rel="stylesheet" href="styles.css">
  <title>Sticky Header Example</title>
</head>
<body>
```

Provably beneficial AI



Should we be concerned about long-term risks to humanity from superintelligent AI? If so, what can we do about it? While some in the mainstream AI community dismiss these concerns, I will argue instead that a fundamental reorientation of the field is required. Instead of building systems that optimize arbitrary objectives, we need to learn how to build systems that will, in fact, be provably beneficial for us. I will show that it is useful to imbue systems with explicit uncertainty concerning the true objectives of the humans they are designed to help.



future
of life
An Open Letter
RESEARCH PRIORITIES FOR ROBUST AND
BENEFICIAL ARTIFICIAL INTELLIGENCE



Preferences θ
Acts roughly according to θ

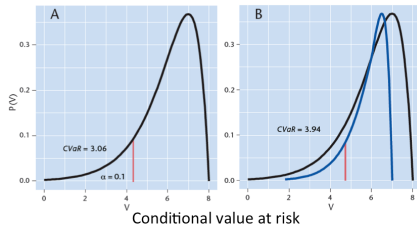
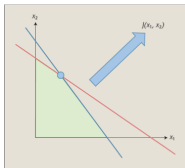


Maximize unknown human θ
Prior $P(\theta)$

OpenAI



Robust AI



Every AI system acts without having a **complete** and **correct** model of the world.

- known unknowns
- unknown unknowns

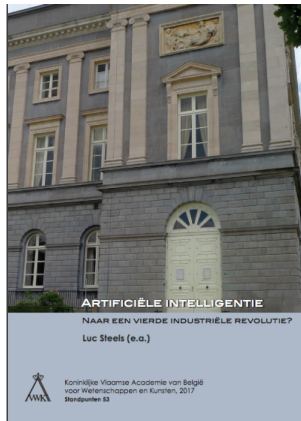
Think of high-stakes applications, such as robotic surgery, self-driving cars, stock trading or autonomous weapons.

Standard decision making policy: maximize expected total reward.

More conservative: **Conditional Value at Risk** (CVaR) – the expected value of a given percentage of the worst possible outcomes.

T.G. Dietterich. Steps Toward Robust Artificial Intelligence. AI Magazine, Fall 2017.

Rising awareness of AI: Benefits and risks



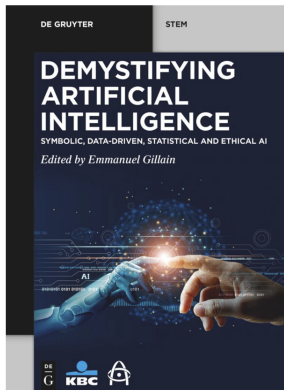
Position paper of the Royal Flemish Academy of Belgium (KVAB) about AI (2017): **Artificial Intelligence – Towards a fourth industrial revolution**

Luc Steels, Bettina Berendt, Aleksandra Pizurica, Dirk Van Dyck, and Joos Vandewalle

<https://www.kvab.be/en/standpunten/artificial-intelligence>

“We believe that AI can be a strong and positive force for our society, if the necessary measures are taken to spread the research, development and knowledge of AI widely. But we are at the same time worried about premature applications or misuse. We must become aware of the limits of AI and take action so that AI is used in a safe and reliable way, to everyone’s advantage.”

Many efforts to 'demystify' AI



Open Access

Published by De Gruyter 2024

Demystifying Artificial Intelligence

Symbolic, Data-Driven, Statistical and Ethical AI

Edited by: Emmanuel Gillain

In the series De Gruyter STEM

<https://doi.org/10.1515/9783111426143>

The theory of the core chapters in this book was contributed by professors from Belgian Universities specializing in different fields of AI: Yves Deville, Bart Bogaerts, Isabelle Linden, Aleksandra Pizurica, Hendrik Blockeel, Walter Daelemans, Erik Mannens

Ghent University has a strong tradition in artificial intelligence and data science related research. Many research teams are working on foundational aspects of AI and are applying the latest AI techniques in a wide range of application areas. In total this accounts for about 50 professors, 50 post docs and 200 PhD students and it results in collaboration with over 300 local and international companies.

[Home](#) > [Research](#)

Research

UGent.AI is advancing the AI research field in a number of foundational areas. In addition the AI research is applied to a wide range of important application fields.

Foundational areas

Automating Data Science

Big Data Analytics and HPC

Deep Learning and Representation Learning

Distributed and Embedded Machine Learning

Explainable AI

(Exploratory) Data Mining

Human-Centered AI

Linked Data and Semantics

Machine Learning

Multimodality and Data Fusion

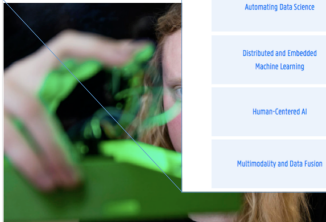
Non-conventional computing paradigms

Probabilistic modelling, statistics and uncertainty

PEOPLE



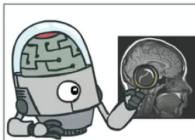
RESEARCH



What is AI?

The science of making machines that:

Think like people



Think rationally



Act like people



Act rationally



Illustration credit: D. Klein and P. Abbeel: Intro to AI, <http://ai.berkeley.edu>

Thinking like people: Cognitive science

Determine how humans think

- Introspection (catch our own thoughts as they go by)
- Psychological experiments (observing a person in action)
- Brain imaging (observing the brain in action)

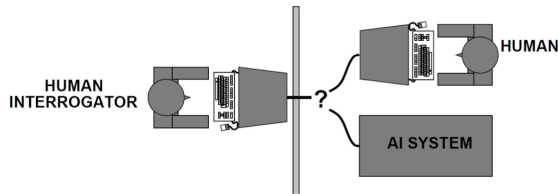
Cognitive science – theories of human mind

based on experimental psychology & computer models

- Scientific theories of brain activities at different levels of abstraction
- Validation:
 - ▶ Predicting and testing behavior of human subjects (top-down)
 - ▶ Identification from neurological data (bottom-up) – cognitive neuroscience

Acting humanly: The Turing test

The Turing test (Alan Turing, 1950) was designed to provide a satisfactory operational definition of intelligence



Suggested major components of AI:

- Knowledge representation (store what it hears or knows)
- Automated reasoning (use the stored info to draw conclusions)
- Machine learning (adapt to new scenarios; detect and extrapolate patterns)
- Language processing (e.g., communicate in English or another language)

Extension - total Turing test includes computer vision and robotics

Thinking rationally: Laws of Thought

Aristotle (384 – 322 BC): codifying the “right thinking” (irrefutable reasoning)

sylogisms - argument structures that always yield correct conclusions when given the right premises. E.g.,

“Socrates is a man; all men are mortal; therefore, Socrates is mortal”

Studying these laws of thought initiated the field of **logic**.

Direct line through mathematics and philosophy to modern AI

- From the 19th century: a precise notation for statements about all kinds of objects and relations among them
- By 1965: programs for solving “in principle” any problem in logistic notation

Logicist tradition in AI is still present. Problems:

- Stating **informal knowledge** in the formal terms of logistic notation
- Computational resources (solvable “in principle” \neq solvable in practice)

Acting rationally

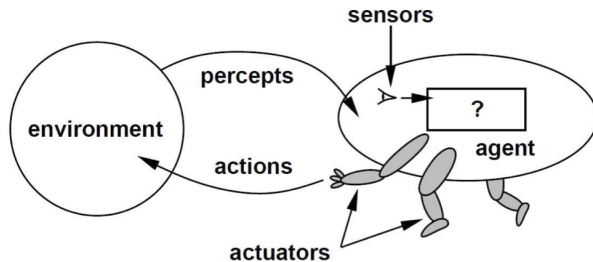
Rational behaviour: doing the right thing

- i.e, doing what 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

Important advantages over the other three categories of AI definitions:

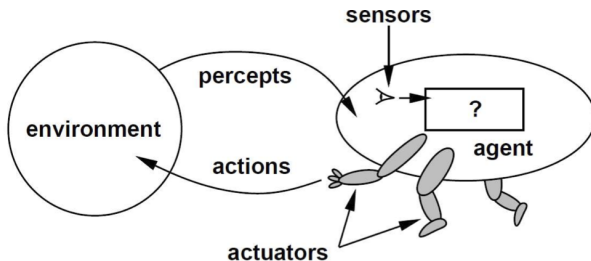
- More general
e.g. all the skills needed for Turing test are needed to act rationally; also, "laws of thought" are only one of possible mechanisms for acting rationally
- Much better suited for formal mathematical description

Intelligent agents



Agent – an entity that perceives the environment through sensors and acts upon it through actuators.

Intelligent agents



Agent – an entity that perceives the environment through sensors and acts upon it through actuators.

Artificial Intelligence deals with the design of **rational agents**, i.e., the design of agents that act to achieve best expected outcome given all information available.

The foundations of AI

Philosophy	logic, methods of reasoning mind as physical system foundations of learning, language, rationality
Mathematics	formal representation and proof algorithms, computation, (un)decidability, (in)tractability probability
Psychology	adaptation phenomena of perception and motor control experimental techniques (psychophysics, etc.)
Economics	formal theory of rational decisions
Linguistics	knowledge representation grammar
Neuroscience	plastic physical substrate for mental activity
Control theory	homeostatic systems, stability simple optimal agent designs

History of AI - Some milestones

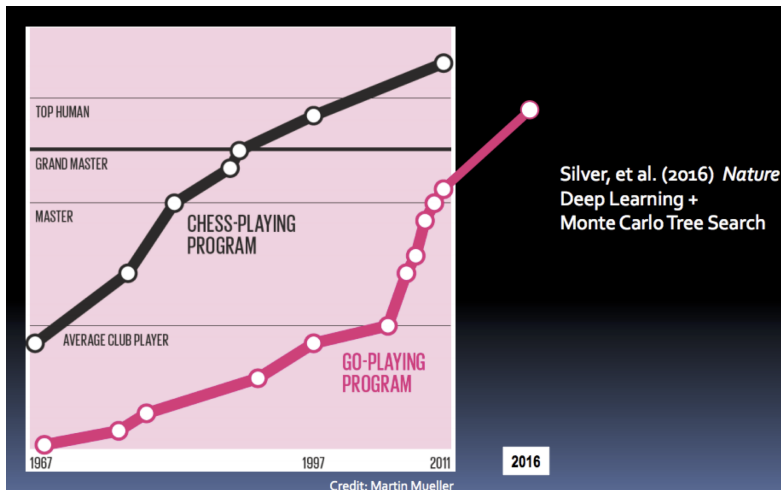
1943	McCulloch & Pitts Neuron: Boolean circuit model of brain
1950	Turing's "Computing Machinery and Intelligence"
1950s	Early AI programs, (incl. Checkers program, Logic Theorist, Geometry Engine)
1956	Dartmouth meeting: "Artificial Intelligence" adopted
1965	Robinson's complete algorithm for logical reasoning
1966–74	AI discovers computational complexity Neural network research almost disappears
1980–88	Expert systems industry booms
1988–93	Expert systems industry busts: "AI Winter"
1985–95	Neural nets return to popularity; Convolutional Neural Nets (CNNs) appear
1995–	Agents, agents, everywhere ...
2003–	Human-level AI back on the agenda
2012–	Deep neural networks (AlexNet, VGGNet, ResNet, Inception ...)
2014–19	Beginnings of generative AI models, attention mechanisms, transformers
2020–	Transformers everywhere, ChatGPT booming, diffusion models...
2023–	The generative AI race (Microsoft's Bing, Google's Bard, OpenAI's GPT-4 ...)

Some important milestones in game playing



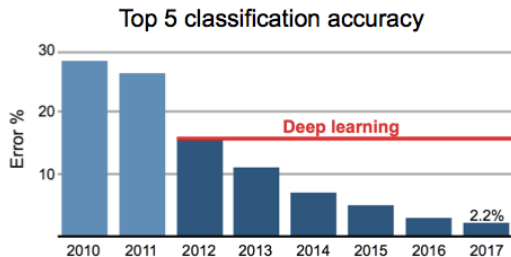
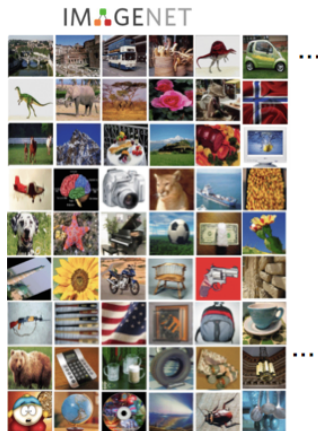
March 2016: AlphaGo defeated a 9-dan master Lee Sedol in the game of Go

AI evolution in game playing



T.G. Dietterich: Steps toward robust artificial intelligence. *AI Magazine*, Fall 2017.

State of the art in AI: Image classification



ImageNet database contains millions of annotated images.

Top-5 error: the percentage of cases where the classifier did not include the correct class among its 5 top guesses.

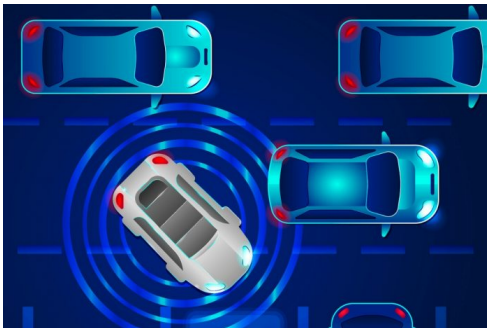
State of the art in AI: Face recognition



State of the art in AI: Autonomous vehicles

“Highway pilots” for hands-free driving

Predictions for up to 75% autonomous vehicles by 2040



Baidu launches Beijing robotaxi's. ITS International, September 2020.

State of the art in AI: Space exploration

Autonomous planning and scheduling in space exploration

- Beginnings: NASA's remote agent program (1999) *"It's one small step in the history of space flight. But it was one giant leap for computer-kind, with a state of the art artificial intelligence system being given primary command of a spacecraft"* <https://ti.arc.nasa.gov/tech/asr/groups/planning-and-scheduling/remote-agent/>
- Charting unmarked galaxies, supernovas, stars, blackholes, and studying cosmic events that would otherwise go unnoticed, tracking asteroids
- Facilitating life of astronauts
- Mars exploration - hunt for life in the solar system



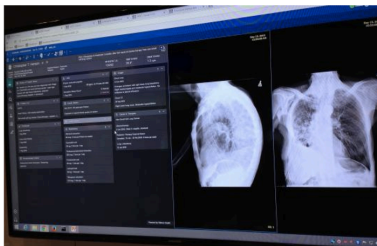
M. Anderson: AI Seeks ET: Machine Learning Powers Hunt for Life in the Solar System. IEEE Spectrum, July 2020.

State of the art in AI: Healthcare

FEATURE | ARTIFICIAL INTELLIGENCE | FEBRUARY 24, 2017 | DAVE FORNELL

How Artificial Intelligence Will Change Medical Imaging

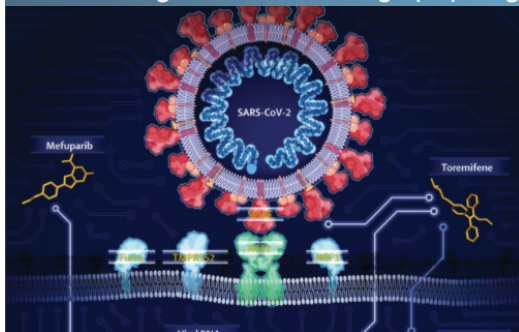
Machine learning software will serve as a very experienced clinical assistant, augmenting the doctor and making workflow more efficient



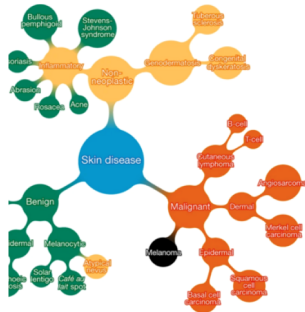
THE LANCET
Digital Health

Published: September 18, 2020 • DOI: [https://doi.org/10.1016/S2589-7500\(20\)30192-8](https://doi.org/10.1016/S2589-7500(20)30192-8)

Artificial intelligence in COVID-19 drug repurposing



Some recently featured topics



A. Esteva et al. Dermatologist-level classification of skin cancer with deep neural networks. Nature, 542, June 2017.

Some recently featured topics



Feature | Biomedical | Biomedical Devices

24 Sep 2018 | 19:00 GMT

AI Could Provide Moment-by-Moment Nursing for a Hospital's Sickest Patients

In the intensive care unit, artificial intelligence can keep watch at a patient's bedside

By Behnood Gholami, Wassim M. Haddad and James M. Bailey

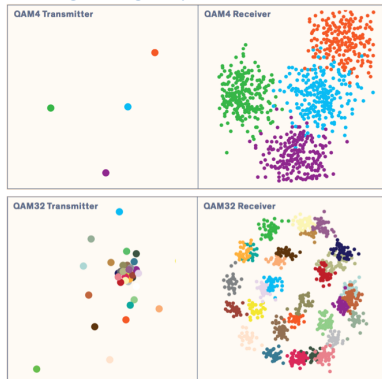


State of the Art in AI: Wireless communications



AI-enabled Spectrum Access System

Signals designed by channel autoencoders



M. Yao et al. Artificial Intelligence Defined 5G Radio Access networks.

IEEE Communications Magazine, March 2019.

J. Downey et al. Machine learning remakes radio - In the future, AIs – not humans – will design our wireless signals. IEEE Spectrum, May 2020.

AI entering all spheres of our life

Home / Enterprise Architecture / Artificial intelligence to make most business decisions by 2035

Artificial intelligence to make most business decisions by 2035



By **SIBAHLE MALINGA**, ITWeb senior news journalist.
Johannesburg, 22 Sep 2020

Read time: 3min 50sec

IN THIS ARTICLE ▾



More than half (57%) of surveyed company executives believe [artificial intelligence](#) (AI) will make most business decisions by 2035, and potentially eliminate the need for senior management teams.

Some recently featured topics: Composing text and poems

AI writing newspaper articles, essays, poems...



had Hagen

Here's a stanza from a sonnet written by William Shakespeare:

*Shall I compare thee to a summer's day?
thou art more lovely and more temperate:
rough winds do shake the darling buds of May,
and summer's lease hath all too short a date*

And here's one written by Deep-speare, an artificial intelligence program that we trained to write sonnets:

*Yet in a circle pallid as it flow,
by this bright sun, that with his light display,
roll'd from the sands, and half the buds of snow,
and calmly on him shall infold away*

“Deep-speare” crafted Shakespearean verse that most readers couldn't distinguish from human-written poems. (J. Lau et al. IEEE Spectrum, May 2020)

Some recently featured topics: Image synthesis

Synthesizing realistic images of non-existing people, animals, plants and arbitrary objects and scenes.



Random images of imagined people: <https://thispersondoesnotexist.com/>

T. Karras, S. Laine and T. Aila (NVIDIA): A Style-Based Generator Architecture for Generative Adversarial Networks (2019).

Examples of AI revolutionizing image generation and manipulation



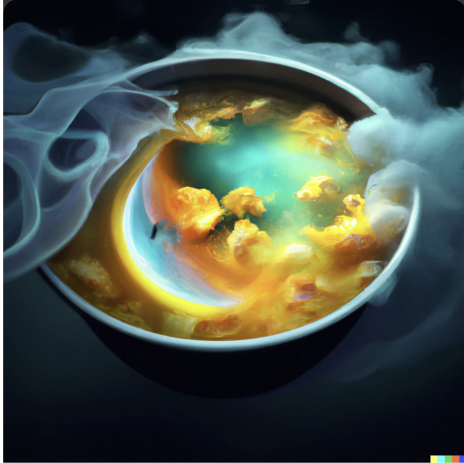
An abstract oil painting of a river



An oil painting by Matisse of a humanoid robot playing chess

Generated by DALL-E, OpenAI

Examples of AI revolutionizing image generation and manipulation



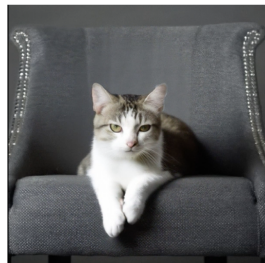
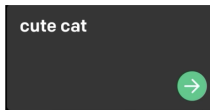
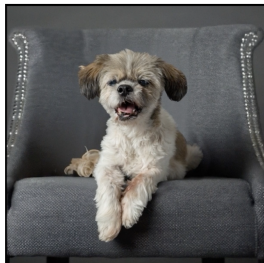
A bowl of soup that is also a portal to another dimension, digital art



A Shiba Inu dog wearing a beret and black turtleneck

Generated by DALL-E, OpenAI

Examples of AI revolutionizing image generation and manipulation



Generated by DALL-E, OpenAI

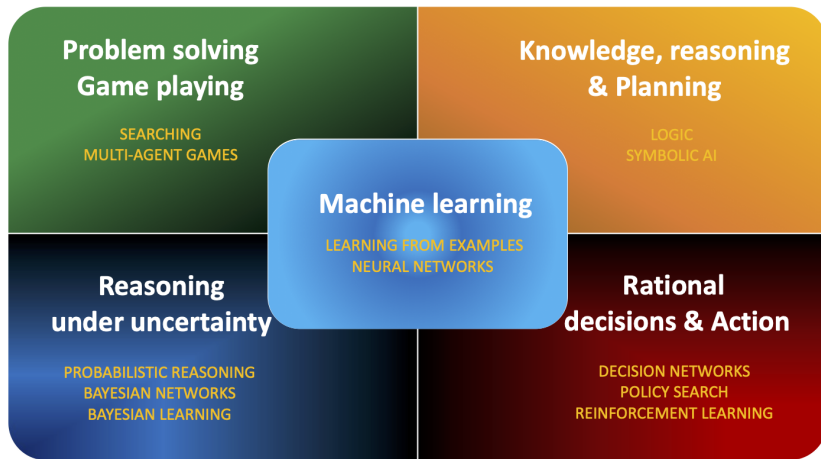
Examples of AI revolutionizing image generation and manipulation



Generated by DALL-E, OpenAI

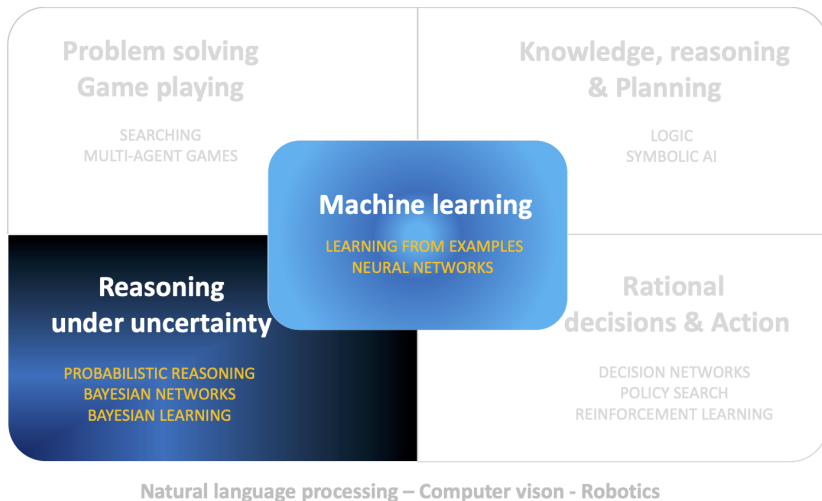
DALL-E is based on the GPT-3 architecture, but trained to generate images from textual descriptions instead of just text.

Main parts of modern AI

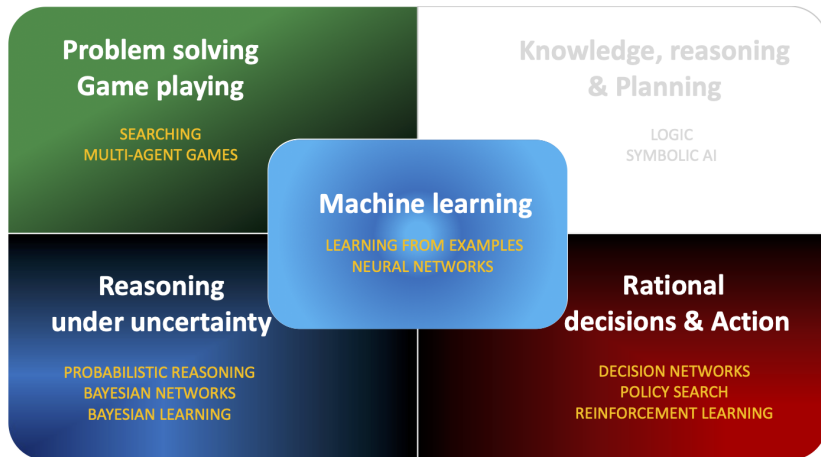


Natural language processing – Computer vision - Robotics

This course – 3 credit version



This course – 6 credit version



Natural language processing – Computer vision - Robotics

Contents of this course (1/3)

- Module 1: Intro to AI
 - ▶ Foundations and history of AI
 - ▶ Fundamental machine learning concepts
 - ★ Feedback in machine learning
 - ★ The concept of supervised and semi-supervised learning
 - ★ Model training and testing, overfitting
- Module 2: Machine Learning (ML) basics – Regression and Classification
 - ▶ Logistic regression
 - ▶ Decision trees
 - ▶ Classification and Clustering
 - ▶ Best practices in data-driven models
 - ▶ White-box models and parameter estimation
 - ▶ Black-box models (Perceptron and neural networks)

Contents of this course (2/3)

- Module 3: Reasoning under uncertainty and Bayesian ML
 - ▶ Bayesian reasoning and learning
 - ▶ Naive Bayes classifier
 - ▶ Bayesian networks and inference
 - ▶ Design of experiment
- Module 4: Societal context
 - ▶ Ethics in AI
 - ▶ Example applications
 - ▶ Benefits and limitations

Contents of this course (3/3)

- Module 5: Search problems
 - ▶ Informed search; local search
 - ▶ Games (minimax, expectimax)
- Module 6: Decisions and actions
 - ▶ Rationality, Decision networks
 - ▶ Markov Decision Processes (MDP)
 - ▶ Reinforcement Learning (RL)
- Module 7: More advanced topics
 - ▶ Reasoning over time; Prediction; Viterbi algorithm
 - ▶ Fundamentals on Hidden Markov Models and Dynamic Decision networks
 - ▶ Examples from Robotics and/or Computer vision and/or NLP

Related Courses

A **non-exhaustive list** of some of the courses that **build on this course**:

- **More advanced AI/ML** courses such as E061330 (Machine Learning), E061350 (Deep Generative Models) and E016340 (Probabilistic Graphical Models)
- Focused on **specific application domains**: E016712 (Computer graphics), E061370 (Data Visualization for and with AI), E019380 (Intelligent Robot Manipulation), E061341 (Natural Language Processing), E018230 (Recommender systems), E019370 (Robotics)

Cross-fertilization (through **computational neuroscience**):

- E010382 (Neuro-Engineering Science), E010620 (Computational Neurophysiology)

Some future prospects of AI research

- Understanding the mechanisms behind deep learning
 - ▶ The number of layers, input and output neurons and filter responses are determined through experiments that require expert knowledge
 - ▶ Rigorous mathematical models to enable performance guarantees are lacking
- Unsupervised learning
 - ▶ There are much more unlabelled data; humans learn by observing too
- Combination of machine learning and complex reasoning
 - ▶ Includes efficient inference methods in graphical models
- Priorities for beneficial AI - building AI for the people
 - ▶ Dealing with model uncertainties and with imprecisely stated goals and preferences
 - ▶ Dealing with erroneous labels, degraded data and adversarial attacks
 - ▶ Optimizing AI's economic impact
 - ▶ Social, moral, ethics & legal considerations