AI and Linguistic Studies

Prof. Ioannis Pitas
Aristotle University of Thessaloniki
pitas@csd.auth.gr
www.aiia.csd.auth.gr
Version 1.0

AI and Linguistic Studies

- **What is AI?**
- Machine Learning
- Natural Language Processing
- Large Language Models
- GPT and ChatGPT
- LLMs and AI in Education
- AI and University Education
- Citizen Morphosis
- Artificial General Intelligence

Artificial Intelligence & Information Analysis Lab

What is AI?

- ***AI Science and Engineering*** (AISE) is the interdisciplinary, scientific study and engineering of ***Artificial Systems*** that mimic and/or surpass ***human intelligence*** in information analysis and ***human interaction*** with the world.

- Core AISE disciplines are:
  - Classical (Symbolic) ***Artificial Intelligence*** (AI),
  - ***Machine Learning*** (ML).

Symbolic AI

Concepts and ideas (ιδέες).

- Concepts are specific mental constructs residing in our mind (brain?) that refine and abstract ideas.
- Examples: ‘Triangle’, ‘Freedom’, ‘Love’.
- ***Concept definition***: Triangle consists of three points connected by 3 straight line segments.

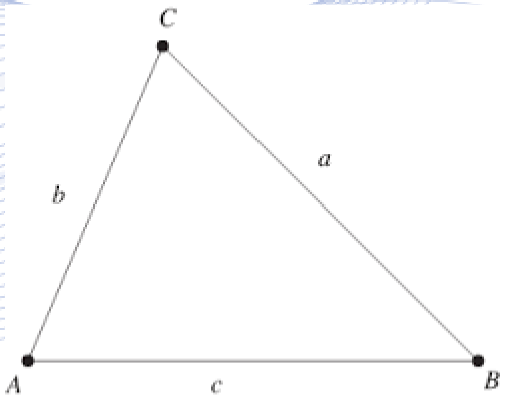

Triangle.

Symbolic AI

Ideas in Philosophy.

- Plato’s cave.
- ***Idealism***: reality is a reflection of ideas.
- ***Materialism***: ideas are shadows of matter on itself (brain).

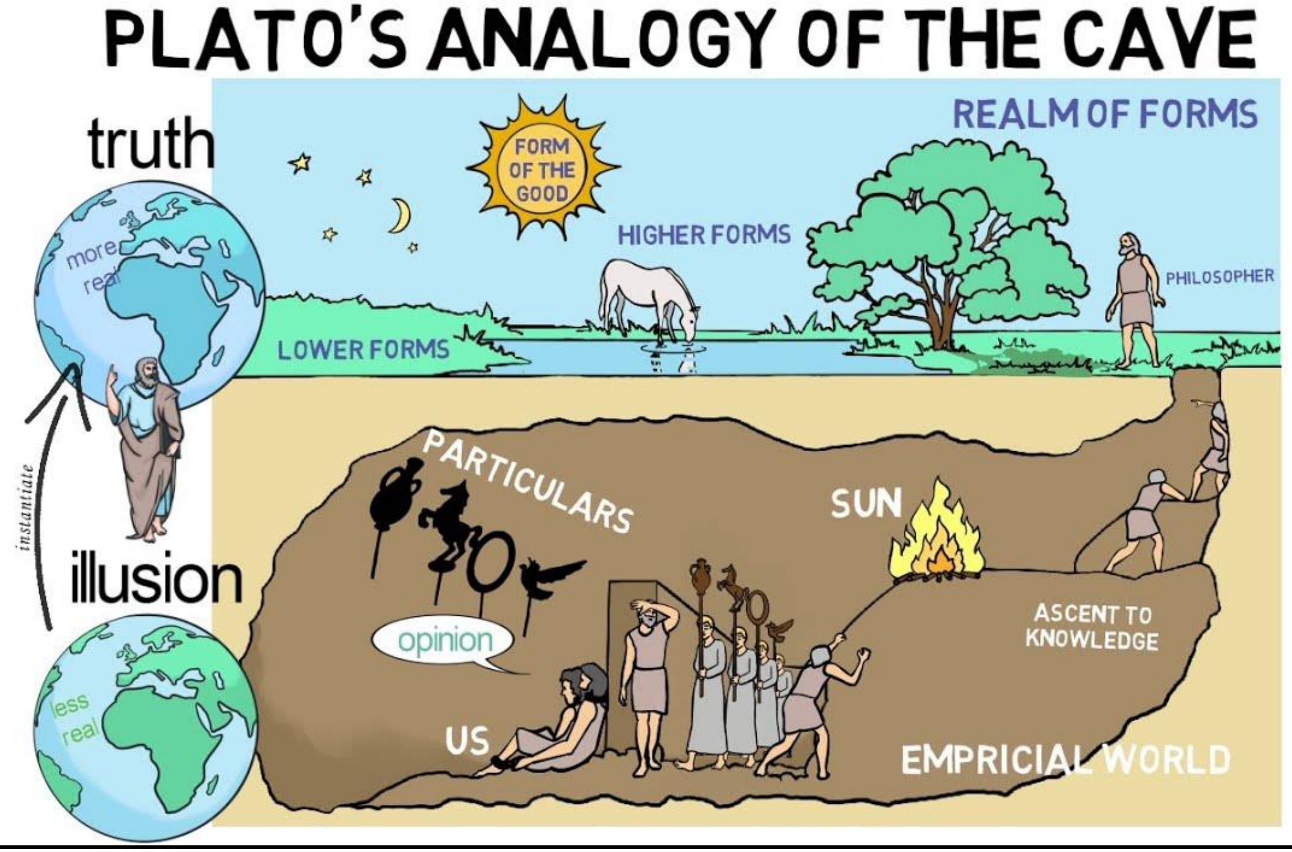

Symbolic AI

- ***Symbolic AI*** operates on concepts and their relations though ***logic*** and ***search***.
- It mimics and simulates high-level human intelligence and ***reasoning***.

- ***Reasoning*** is one of the most complex brain activities.

- ***Symbolic AI*** employs Mathematical Logic.

Symbolic AI

- Examples:

‘If somebody has high fever and coughs, she/he has flu.’

‘If I turn left, I may enter the opposite lane.’

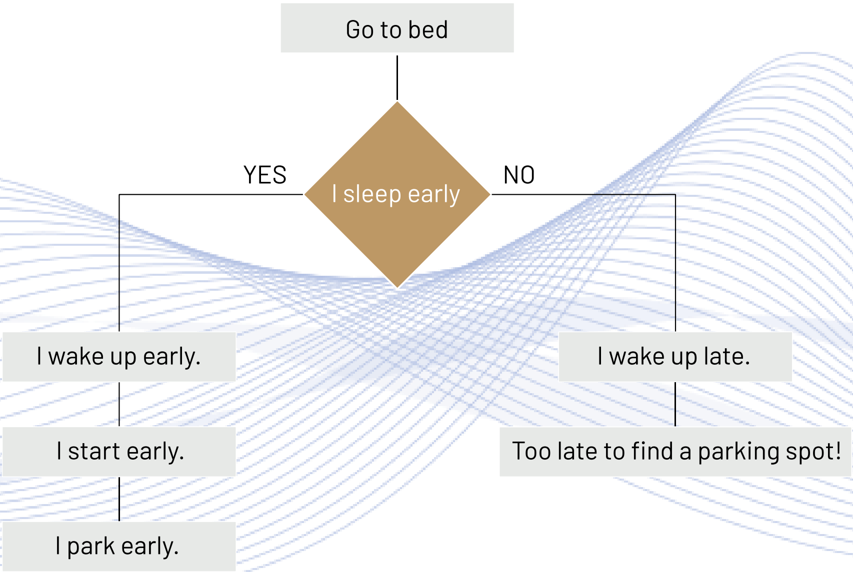

- ***Symbolic AI failed to deliver!***

AI and Linguistic Studies

- What is AI?
- **Machine Learning**
- Natural Language Processing
- Large Language Models
- GPT and ChatGPT
- LLMs and AI in Education
- AI and University Education
- Citizen Morphosis
- Artificial General Intelligence

Machine Learning

- ***Data***: measured quantities related to nature and/or human activities.

- ***Data are primarily numbers*** representing object characteristics (***features***).
- ***Measured in bits.***

- ***Data can be organized in vectors.***

Machine Learning

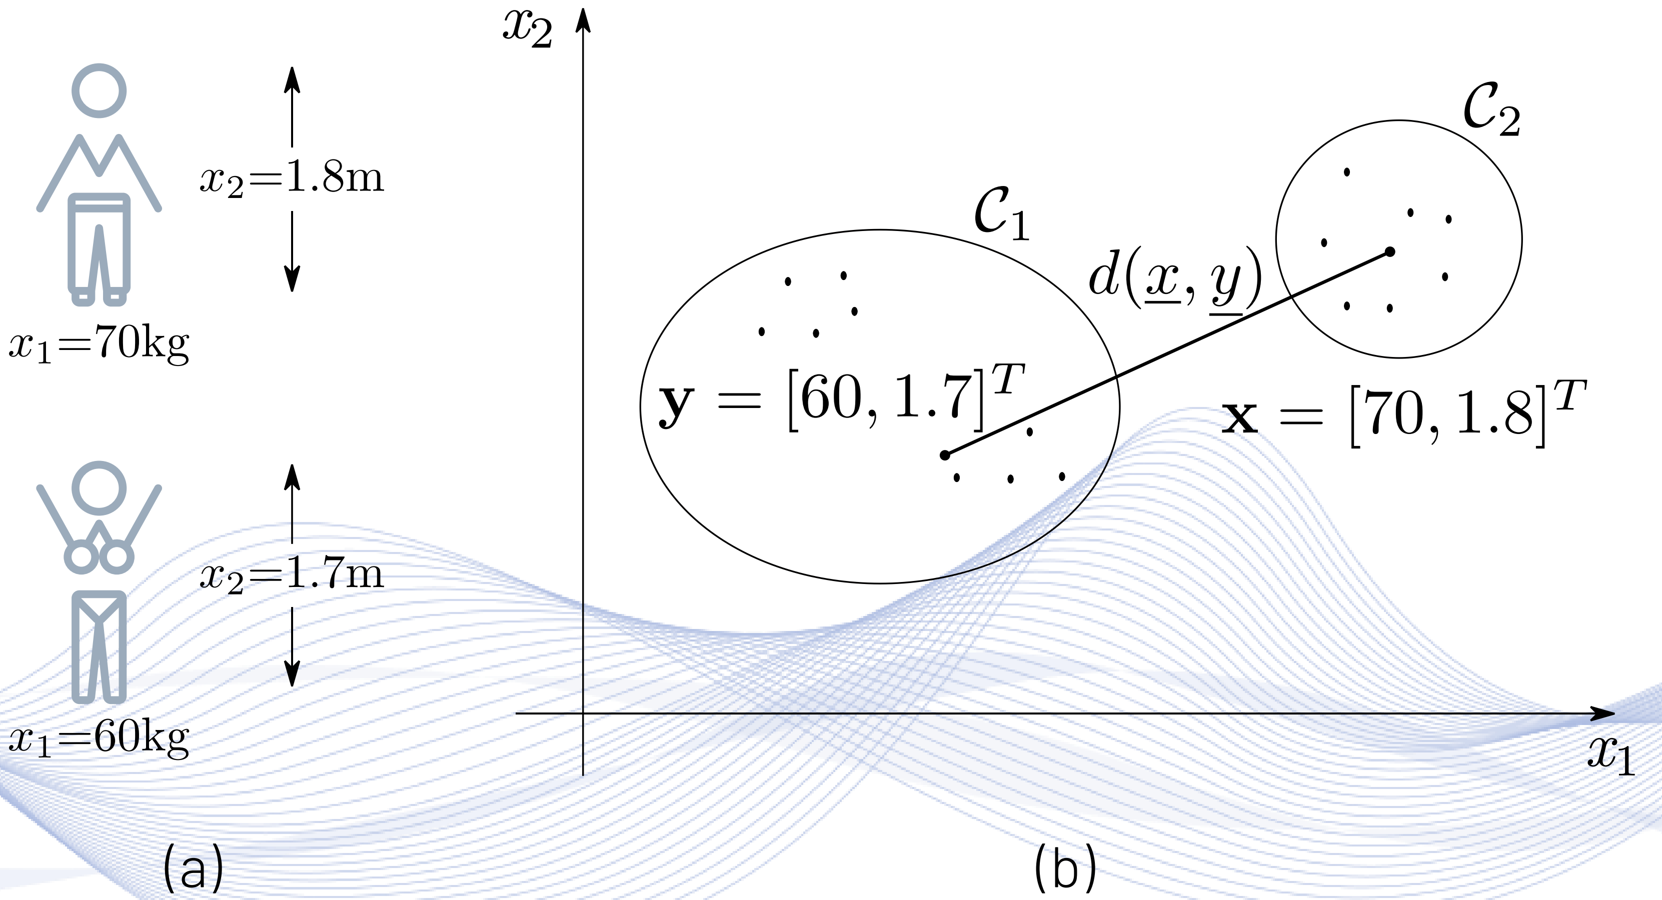

Measuring humans and producing their weight and height vectors.

Machine Learning

Data can have ***spatiotemporal structure***:

- 1D temporal signals, e.g., music
- 2D spatial signals: images

- Data features can be represented by ***vectors***:

$$\mathbf{x}^T = [x_1, x_2, \ldots, x_n].$$

Machine Learning algorithms learn from data vectors.

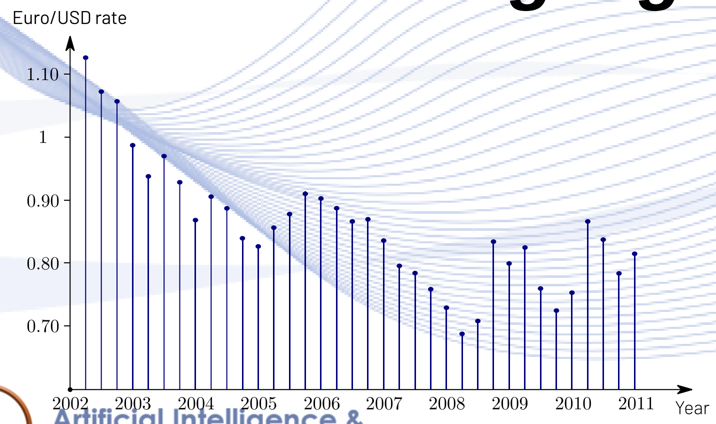

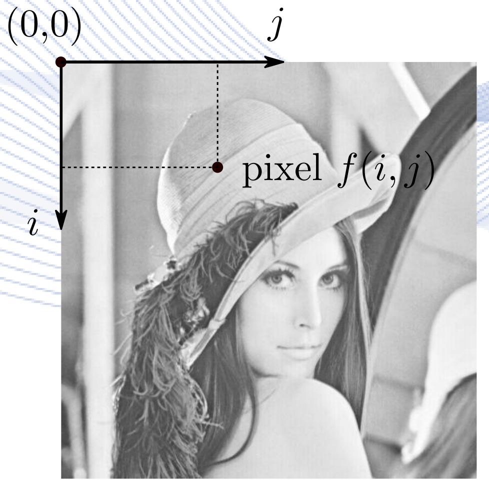

Machine Learning

- Graphs can represent relations of historical actors.

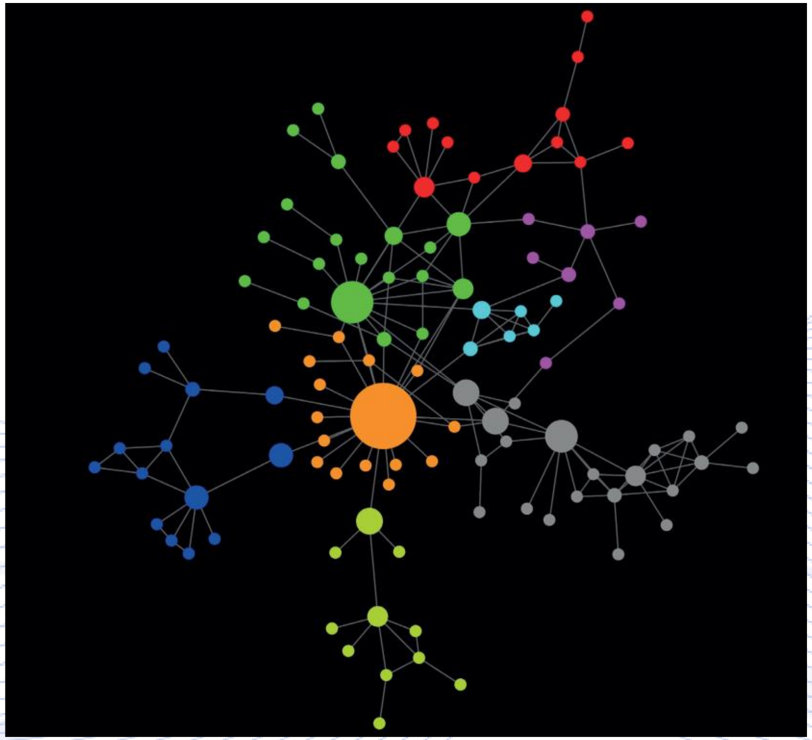

Clusters of the Byzantine nobility in the civil war period 1321-1328 AD.

Machine Learning

Citizen communities (graphs)

- Citizens are graph nodes connected by relations (graph edges):
  - friendship, political affiliation, etc.

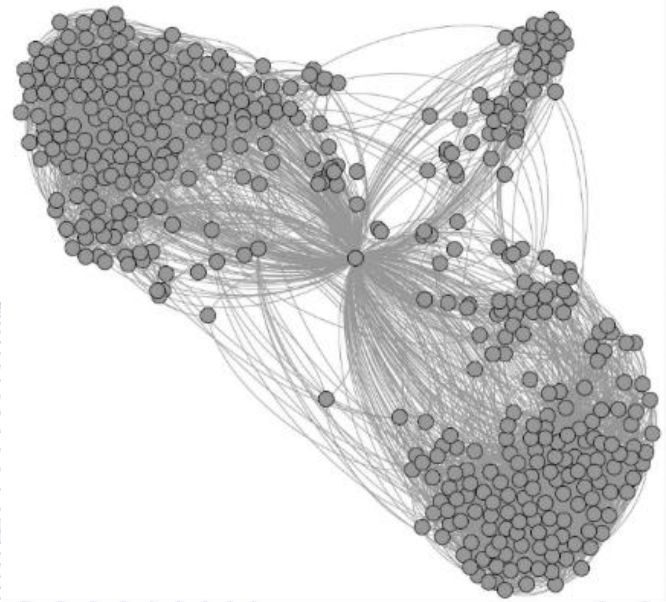

Facebook friendship graph.

Generative AI

Machine Learning Algorithms that learn data and produce new data.

- ***Large Language Models***
  - Text production
- ***Generative Adversarial Networks, Diffusion Models***
  - Multimedia content creation (images, video, audio, computer graphics)

GAN-generated video.

AI and Linguistic Studies

- What is AI?
- Machine Learning
- **Natural Language Processing**
- Large Language Models
- GPT and ChatGPT
- LLMs and AI in Education
- AI and University Education
- Citizen Morphosis
- Artificial General Intelligence

Artificial Intelligence & Information Analysis Lab

Natural Language Processing

Word embeddings: ***Word2Vec*** (example)

Two-layer NN trained to reconstruct linguistic context of words.

- Training is performed with pairs of context-target words.
- 2 training variations.

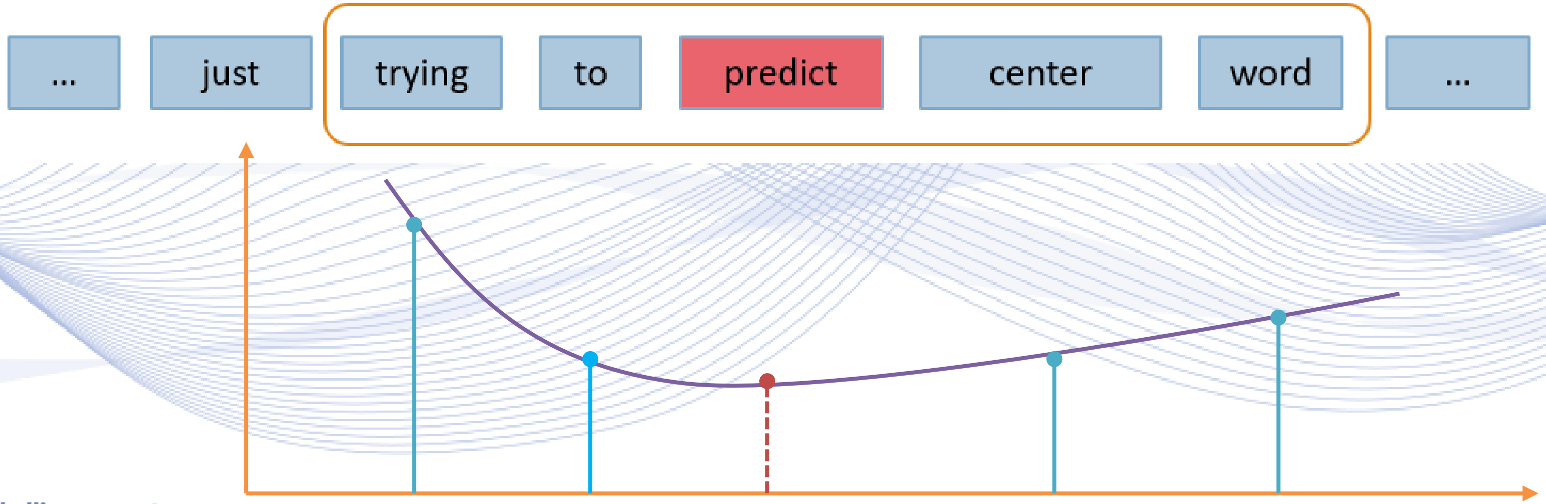

Natural Language Processing

Visualization of word prediction in 2D space

A sentence can be visualized as a curve in the vectorial space over time, connecting all its word embeddings.

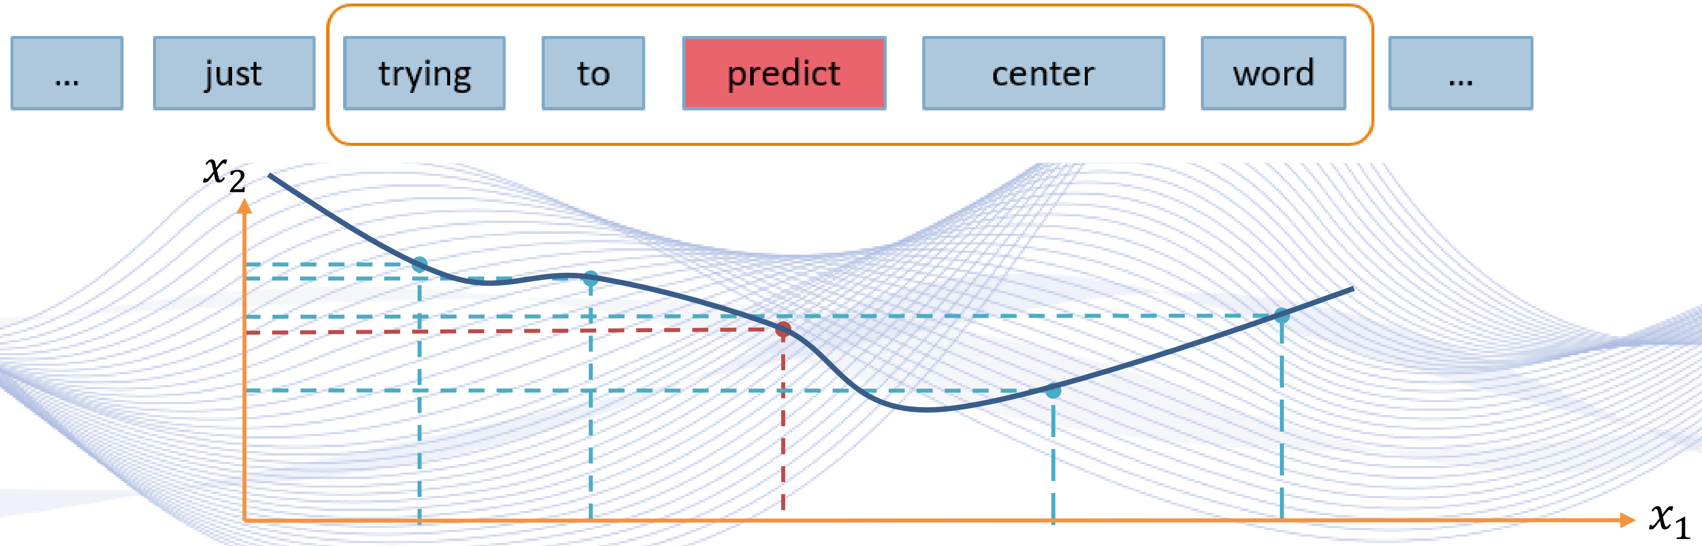

Word trajectory in a 2D vectorial space $[x_1, x_2]$.

Natural Language Processing

Vectors representing words

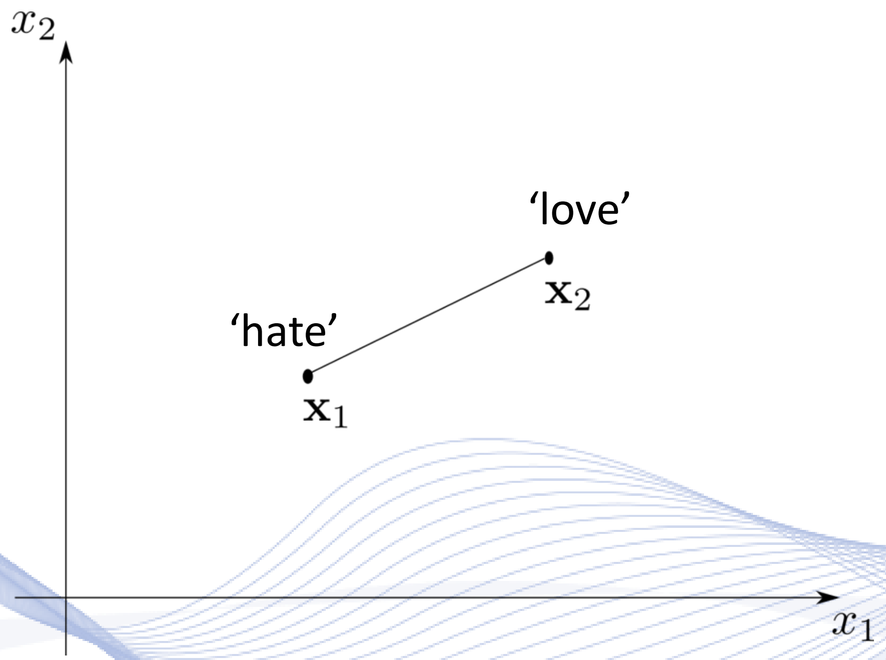

Distance between two words.

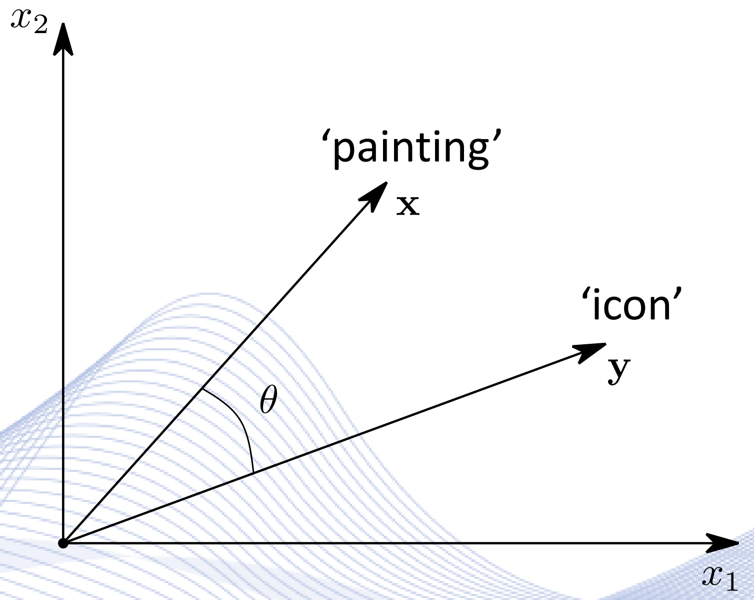

Semantic similarity between two words.

Natural Language Processing

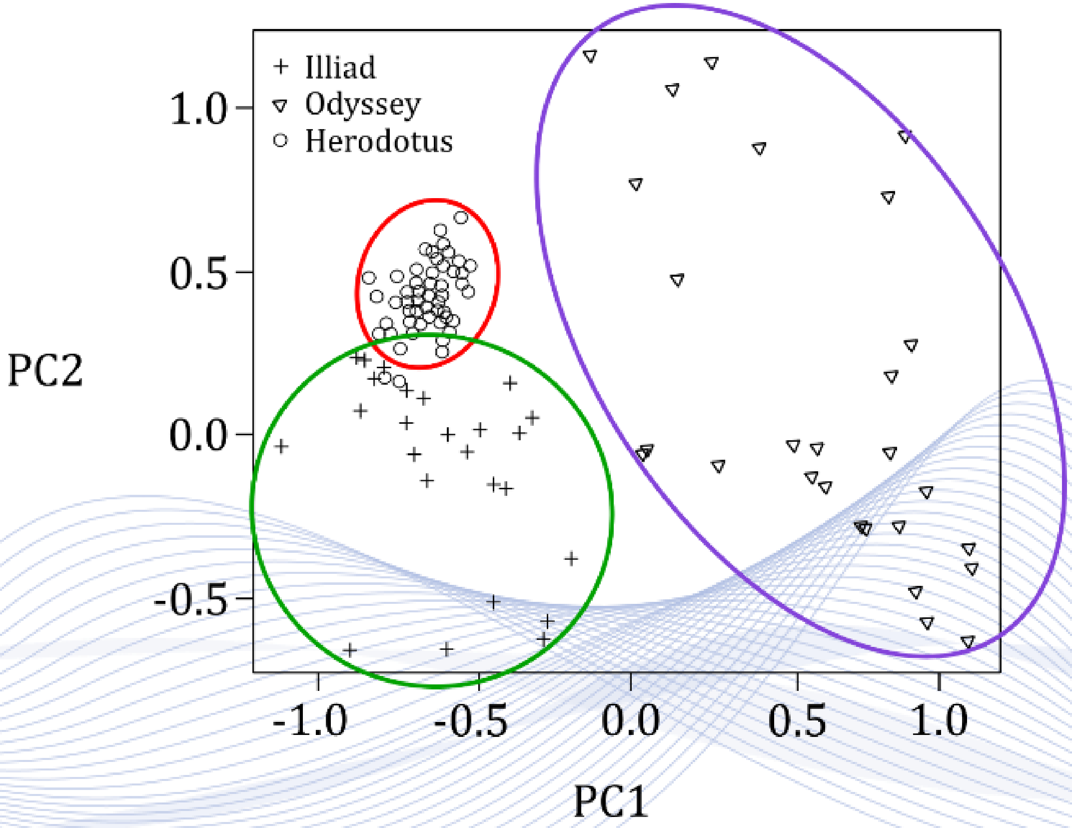

Representing texts by vectors:

Principal component analysis of Homer's Iliad and Odyssey.

Artificial Intelligence & Information Analysis Lab

Natural Language Processing

Text Classification

- Does text $\mathbf{x}$ belong to class $\mathcal{C}_1$ or class $\mathcal{C}_2$?

Examples:

- Text sentiment analysis
  - Is the text 'sad' or 'joyful'?
- Author recognition
  - Is this Epistle authored by St. Paul or not.

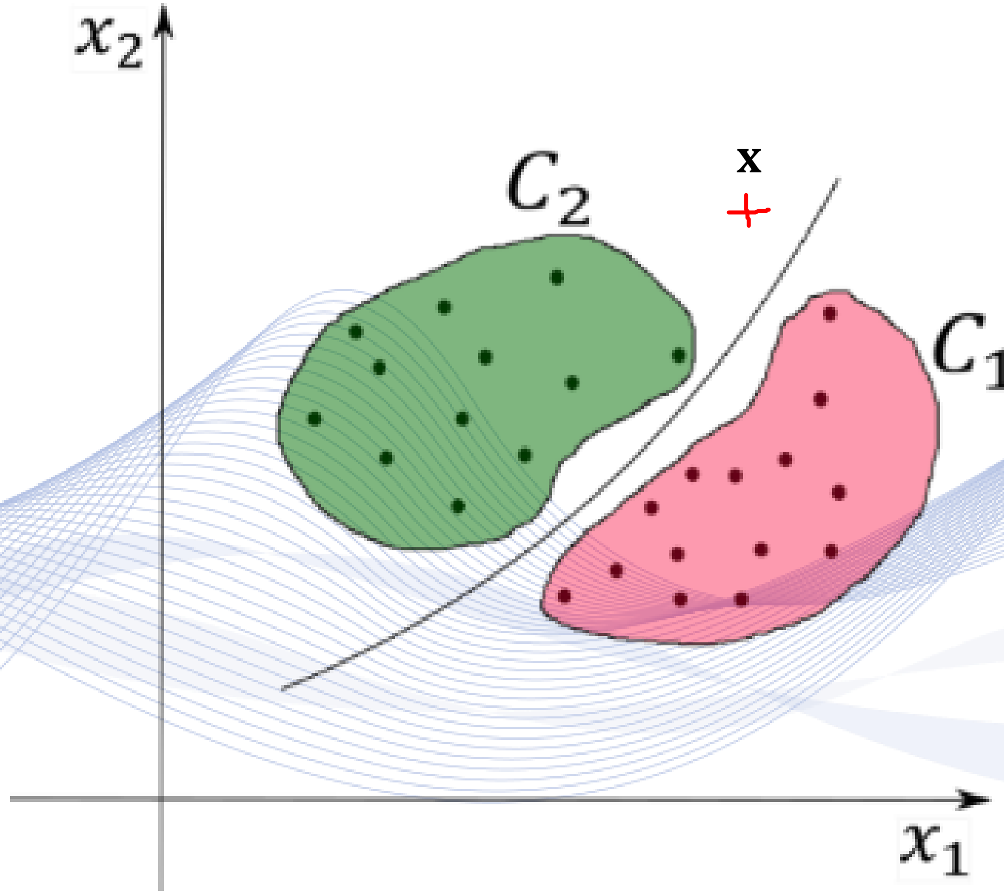

Natural Language Processing

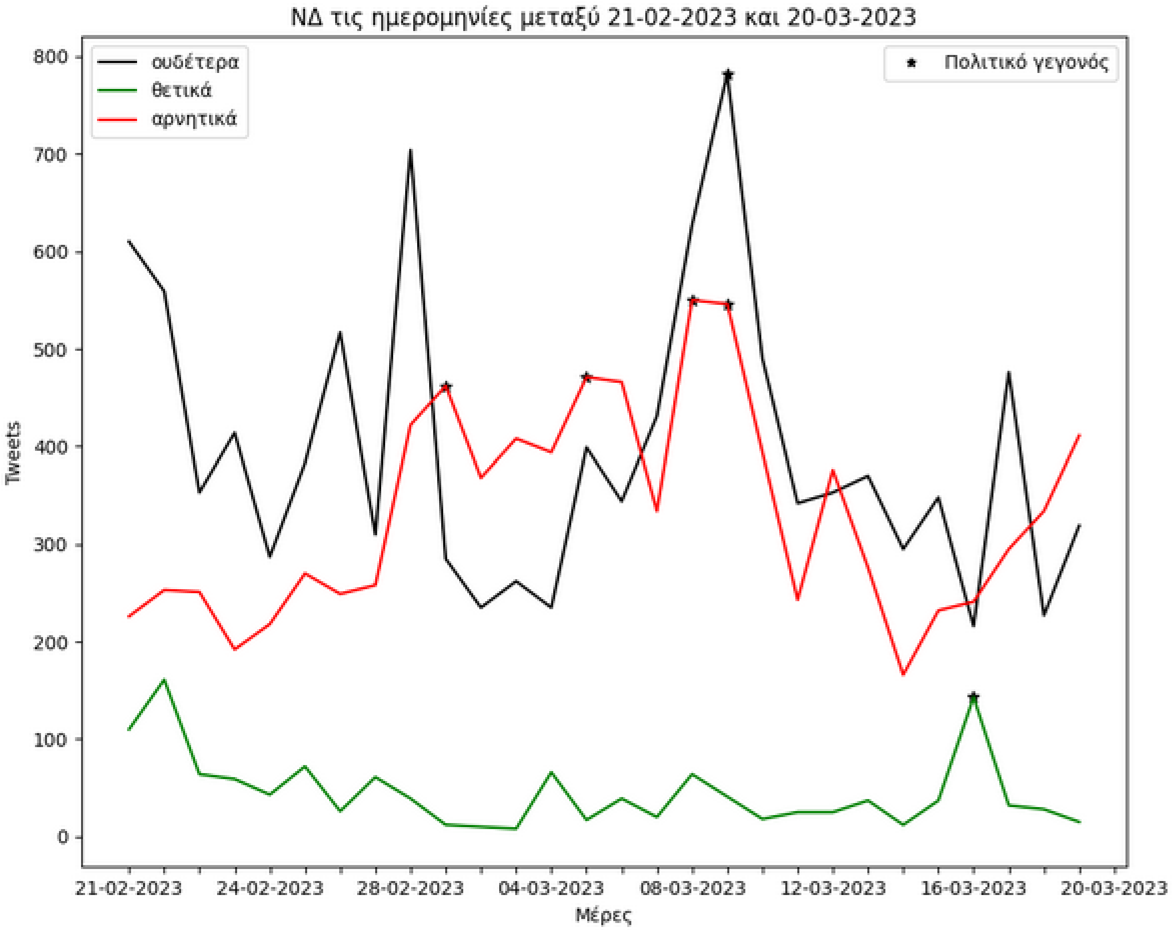

Sentiment Analysis of Greek Political Tweets.

AI and Linguistic Studies

- What is AI?
- Machine Learning
- Natural Language Processing
- Large Language Models
- GPT and ChatGPT
- LLMs and AI in Education
- AI and University Education
- Citizen Morphosis
- Artificial General Intelligence

Large Language Models

- ChatGPT is a ***Large Language Model*** (***LLM***) that is fine-tuned from a ***Generative Pre-Trained Transformer-3.5*** (***GPT-3.5***) LLM series, produced by OpenAI.

- An LLM is a ***Deep Neural Network*** (***DNN***) trained to generate smooth text similar to the human-generated one.
- The fine-tuning of the GPT-3.5 is performed using supervised and reinforcement learning with human feedback [OPE2023] .

Large Language Models

The building blocks of LLMs are [AJI2023] :

- ***Tokenization***: transforming a text in a series of tokens, e.g.,:
  - ***sub-words, words***.
- Text compression, in order to minimize the size of the encoded token, while retaining the ability to represent well text sequences.

- ***Vector embedding***: Token representation by vectors capturing their semantic meaning in a high-dimensional space.
- Vector embeddings are processed by the NN and are learned during the training.

Large Language Models

Transformers provide data representations based on statistical correlations of input elements (NLP tokens).

- They comprise of the ***encoder*** and ***decoder***.
- ***Self-attention*** weighs the importance of input or output tokens.
- ***Cross-attention*** cross-correlates input and output tokens.

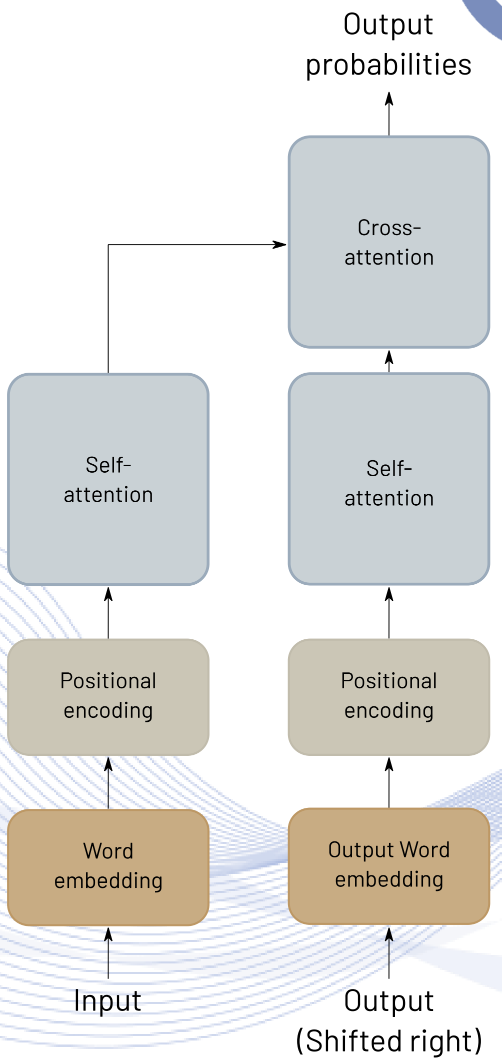

Transformer architecture.

Large Language Models

Transformers

- ***Transformers*** comprise of the encoder and decoder and use the self-attention mechanism to weigh the importance of input elements [VAS2017] .
- GPT-3.5 is a fine-tuned model of the GPT-3, which is a Transformer DNN.

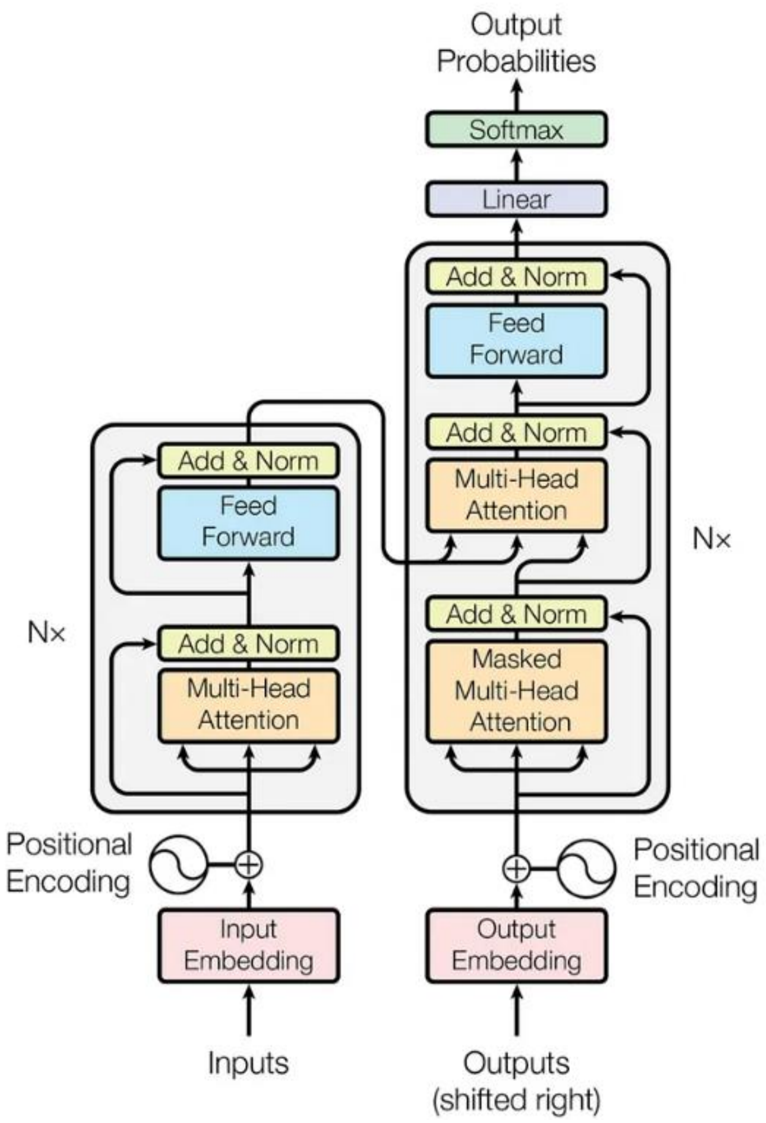

Transformer architecture [VAS2017].

Large Language Models

LLM training and text production example:

- LLMs' reply to the query '***What is the capital of Spain?***' would be '***Madrid***' rather than '***death penalty***', since:
- a) they encountered this semantic association (Spain, Madrid, capital) too many times in their training corpora.
- b) the learned association (Spain, country) helps them disambiguate the meaning of the query word 'capital'.
- ***Such statistical associations may occasionally be out of context, or semantically wrong or completely fabricated.***

Large Language Models

LLM training and text production:

- LLMs search for text patterns and correlations in huge amounts of training data and produce statistically probable output (text).
- They become increasingly better in learning word predictions and relations.
- This is an essential feature in outputting smooth 'human-like' text.
- ***Is Language all we need?***

Large Language Models

LLMs have high expressive and abstraction power.

They are mathematical functions

$$\mathbf{y} = \boldsymbol{f}(\mathbf{x}; \boldsymbol{\theta})$$

that learn parameters $\boldsymbol{\theta}$ from labeled training text data $(\mathbf{x}, \mathbf{y})$.

- Their power is in the ***huge number*** of parameters in $\boldsymbol{\theta}$.
- ***Special case of Generative AI.***
- Huge expressive and abstraction power compared to classical linguistic approaches.
- Non-explainable operation (so far).

AI and Linguistic Studies

- What is AI?
- Machine Learning
- Natural Language Processing
- Large Language Models
- **GPT and ChatGPT**
- LLMs and AI in Education
- AI and University Education
- Citizen Morphosis
- Artificial General Intelligence

GPT

- The ***Generative Pre-Trained Transformer*** (***GPT***) is a ***decoder-only Transformer*** model that generates one token at a time [RAD2018].

- Semi-supervised training:

a) Unsupervised pre-training.

b) Supervised fine-tuning.

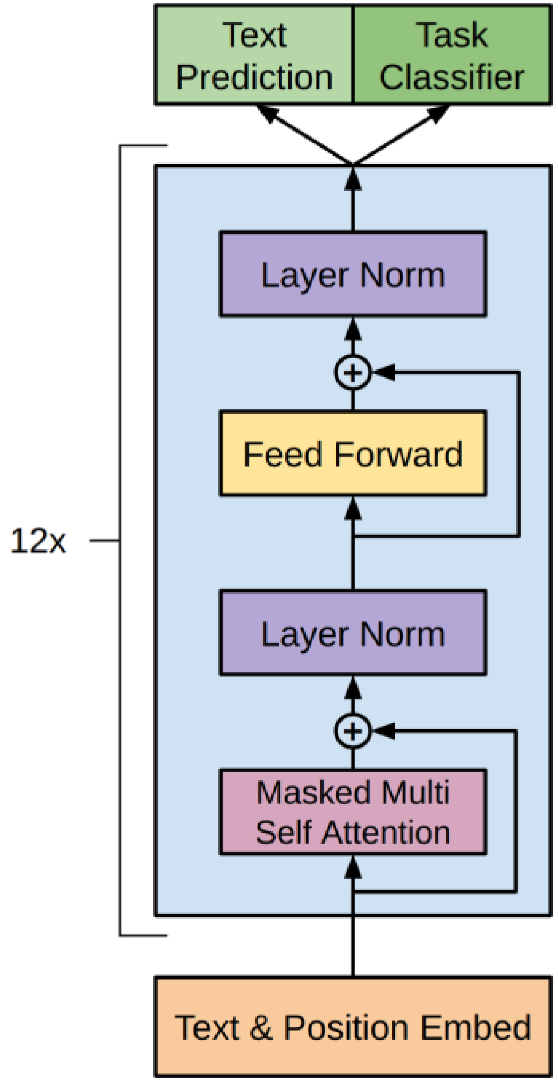

GPT architecture [RAD2018].

GPT Training stages

Unsupervised Pre-training stage:

- *Training dataset*: BooksCorpus [ZHU2015].
- *Objective*: Standard language modelling [RAD2018].

Fine-tuning stage:

- *Training dataset*: a labelled dataset corresponding to the fine-tuning task
- *Objective*: GPT model parameters adaptation to the supervised target task and language modelling [RAD2018].

ChatGPT Reward Model

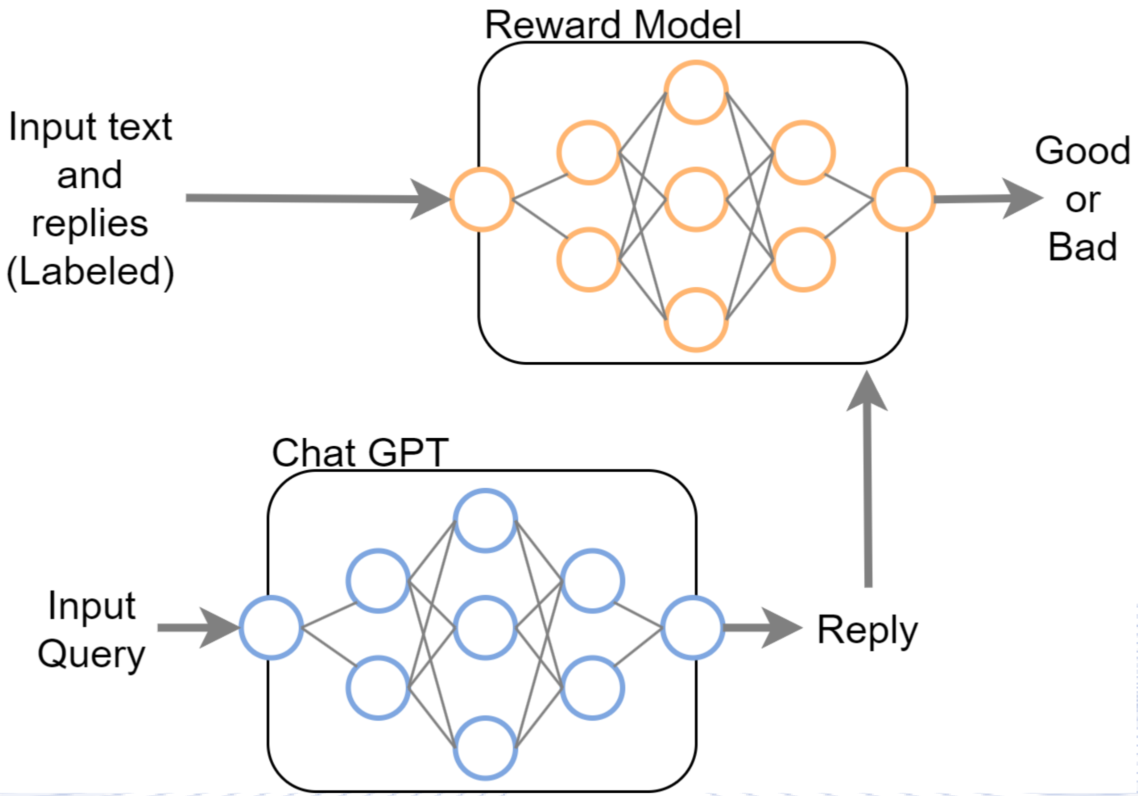

ChatGPT fine tuning methodology.

GPT-4

- GPT-4 is a large multimodal model

Input: Both images and text

Output: Text

- Trained on next word prediction using public and licensed data.
- Fine-tuned through ***Reinforcement Learning with Human Feedback*** (RLHF) in order to align the models output with the user's intent [OP2023].
- Models capabilities originate from the pre-training process and not the RLHF [OP2023].

GPT-4

- GPT-4 exhibits human-level performance on various professional and academic benchmarks [OP2023].

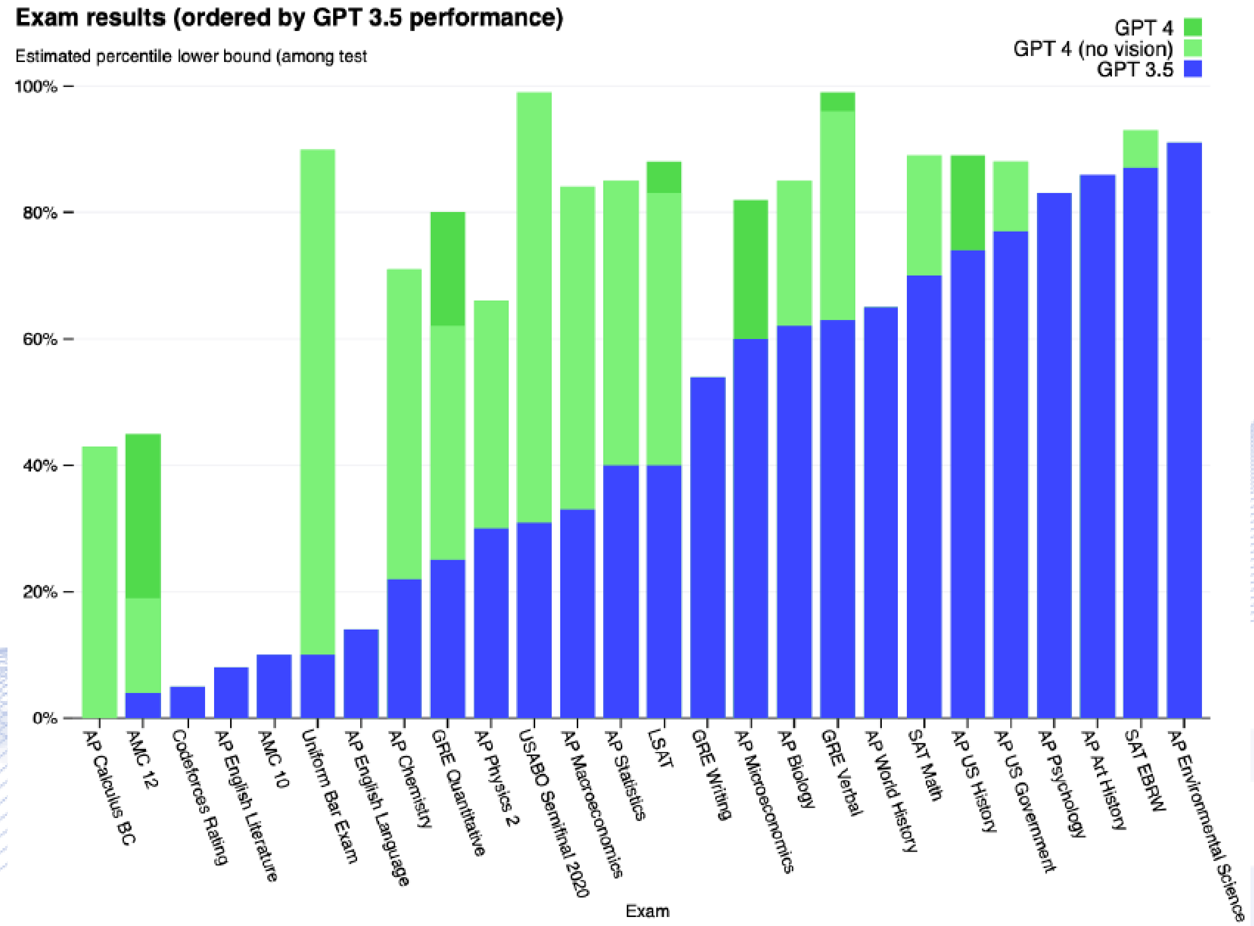

GPT performance on academic and professional exam [OP2023].

GPT-4 Capabilities

	GPT-4 Evaluated few-shot	GPT-3.5 Evaluated few-shot	LM SOTA Best external LM evaluated few-shot	SOTA Best external model (incl. benchmark-specific tuning)
MMLU [49] Multiple-choice questions in 57 subjects (professional & academic)	**86.4%** 5-shot	70.0% 5-shot	70.7% 5-shot U-PaLM [50]	75.2% 5-shot Flan-PaLM [51]
HellaSwag [52] Commonsense reasoning around everyday events	**95.3%** 10-shot	85.5% 10-shot	84.2% LLaMA (validation set) [28]	85.6 ALUM [53]
AI2 Reasoning Challenge (ARC) [54] Grade-school multiple choice science questions. Challenge-set.	**96.3%** 25-shot	85.2% 25-shot	85.2% 8-shot PaLM [55]	86.5% ST-MOE [18]
WinoGrande [56] Commonsense reasoning around pronoun resolution	**87.5%** 5-shot	81.6% 5-shot	85.1% 5-shot PaLM [3]	85.1% 5-shot PaLM [3]
HumanEval [43] Python coding tasks	**67.0%** 0-shot	48.1% 0-shot	26.2% 0-shot PaLM [3]	65.8% CodeT + GPT-3.5 [57]
DROP [58] (F1 score) Reading comprehension & arithmetic.	80.9 3-shot	64.1 3-shot	70.8 1-shot PaLM [3]	**88.4** QDGAT [59]
GSM-8K [60] Grade-school mathematics questions	**92.0%*** 5-shot chain-of-thought	57.1% 5-shot	58.8% 8-shot Minerva [61]	87.3% Chinchilla + SFT+ORM-RL, ORM reranking [62]

Performance of GPT-4 on academic benchmarks [OP2023].

GPT-4 Limitations

GPT-4 suffers from the same limitations as the previous GPT models [OP2023]:

- Hallucinations.
- Bias in its output text.
- Lack knowledge past 2021 and doesn’t learn from its experience.
- There is still a risk of generating harmful advice, buggy code and inaccurate information. This risk has been reduced compared to older models through additional signal in the RLHF.

Artificial Intelligence & Information Analysis Lab

ChatGPT Capabilities

ChatGPT ***text processing*** capabilities:

- ***Translation***: chatGPT performs well translating in English [BAN2023].
- ***Summarization***: Adequate results (similar to GPT3). However, it is outperformed by SOTA works [BAN2023].
- ***Question Answering***: Near perfect scores [BAN2023].
- ***Sentiment Analysis***: It outperforms supervised SOTA works [SCA2022] and zero-shot multilingual LLM [CAH2022] (evaluation metric: F1 score) [BAN2023].

ChatGPT Capabilities

- ***Dialogue tasks***: ChatGPT generates high quality fluent human-like responses [BAN2023].
- ***Misinformation detection***: ChatGPT detected misinformation at 92% and 73.33% accuracy on covid-scientific and covid-social datasets, containing scientific and social claims related to Covid-19 accordingly [BAN2023].
- ***Code understanding and generation***: ChatGPT achieved higher score on the LinkedIn Python skills assessment than 85% of humans [CFTE].

ChatGPT Limitations

- ChatGPTs responses sometimes sound plausible, while they are ***incorrect or nonsensical*** [OPE2023].
- ChatGPT responses are sensitive to tweaks in input phrasing and prompt repetition [OPE2023].
- Training data bias causes ***excessively verbose responses*** and overuse of certain phrases [OPE2023].
- In translation, it still lacks excellent ability to successfully translate English in other languages [BAN2023].

ChatGPT Limitations

- In the case of an ambiguous query, the model ***guesses what the user intended to say***, rather than ask for clarifying questions [OPE2023].
- ChatGPT sometimes responds to ***harmful instructions or outputs biased answers***.
  - The Moderation API is used to flag certain types of unsafe content [OPE2023].
- ChatGPT has a limited understanding of ***low-resource languages***, due to low training data volume [BAN2023].

ChatGPT Limitations

ChatGPT hallucinations

- Reward functions can induce ChatGPT into hallucinating facts, rather than admitting ignorance.
- Hallucinations can become even more serious when ***human-in-the-loop*** LLM retraining or fine-tuning is employed.
- Users can trigger hallucinated replies, e.g., that 'the Pope is a pop singer', as the LLM thinks it maximizes its reward.

ChatGPT Limitations

ChatGPT hallucinations

- Humans make such judgement errors as well:
  - Sensory illusions, wild children's imagination.
- The human mind creates ***mental images*** of the world that map reality, yet are completely artificial, real, but different from reality.
- ***Arts can be considered as a form of creative expressed hallucination***.

ChatGPT Limitations

ChatGPT hallucinations

- In principle, ***Generative AI fabricates imaginary outputs***.
- They may deviate from the training data and ‘common human sense’.
- Depending on their ***social use***, we can call them Art or Fake data or Hallucinations.

ChatGPT: Questionmarks

- ***Does ChatGPT have access to external resources? No.***
  - Yet, if suitably trained ChatGPT can provide lots of factual information.
  - If not, what is its ***knowledge storage capacity***?
- ***Should LLMs have access to external resources? Yes.***
  - Knowledge graphs? Algebraic computations (Symbolic Algebra)?
  - This combination has great potential, e.g., in search.

ChatGPT: Questionmarks

- ***Can LLMs provide hints on how human memory works?***
  - Associative memories, Hopfield networks.
  - CNNs can store some training data information.
  - Transformer-based LLMs are based on ***statistical associations***.

- ***Relation between human imagination and ChatGPT hallucination***?
- Kids are particularly good at fabricating facts or stories.

ChatGPT: Questionmarks

- ***Does ChatGPT have explicit reasoning mechanisms***?
  - No, it has been trained as a pure language model.
  - However, its replies ***show*** some reasoning capabilities.
- ***'Text is all we need' to learn reasoning***?
  - Language/text contain many examples of reasoning.
  - Reasoning as a result of learning-by-examples?
  - ***If proven, it is a Nobel-level breakthrough.***
    - It can reconcile Machine Learning and Symbolic AI.

ChatGPT: Questionmarks

Does ChatGPT have explicit reasoning mechanisms?

- Humans learn from their mothers, relatives, and peers how to think, based on countless everyday discussions.
- An eventual LLM ‘inference by example’ capacity may hint towards ways that ***humans learn to think***.

ChatGPT: Questionmarks

Causal, approximate reasoning?

- LLM output (statistical event cross-association):

  ‘It has repeatedly been observed (or better, has been found in the literature) that plants thrive, when the sun shines’.

- Causal argumentation:

  ‘Plants thrive when the sun shines, because they use sunlight in their photosynthesis’.

ChatGPT Questionmarks

- ***Do LLM/ChatGPT have abstraction mechanisms?***
  - Their internal structure and functionalities are unknown.
  - Clustering and concept creation? Rule creation?

- ***Can ChatGPT provide explicit language modelling***?
  - Derivation of ***grammar and syntax rules.***

- ***ChatGPT explainability?***

ChatGPT Questionmarks

- ***Do LLMs/ChatGPT have affect?***
  - Absolutely not in the human sense.
  - Yet, it is a disgrace that they can create such an impression to unsuspecting public, when texting like ‘I love you’.
  - Machines are very good in understanding certain affect signals, e.g., ***facial expressions***.

LLM criticism

- ‘*Human intelligence can work well with few data*’ (Chomsky, 2023) [CHO2023]: ***completely wrong***.
- The contrary is true: both machine and human learning require massive training, in terms of data, architecture complexity and energy needs.
- ***Is it possible that similar laws govern both machine and human learning?***

LLM criticism

Criticism:

- ‘*Current LLMs have many deficiencies*’,
- ‘*They do just massive plagiarism*’,
- ‘*They know nothing about particular domains*’,
- ‘*They are not multimodal, e.g., supporting visual perception*’ (except GPT-4).

- ***Completely wrong claims***. LLMs are only at the start. Great advances are expected.
- Such nihilistic criticism is similar to the ill-fated criticism of Rosenblatt’s perceptron by Minsky and Papert that led to the AI winter at the end of the 1960s.

AI and Linguistic Studies

- What is AI?
- Machine Learning
- Natural Language Processing
- Large Language Models
- GPT and ChatGPT
- **LLMs and AI in Education**
- AI and University Education
- Citizen Morphosis
- Artificial General Intelligence

LLMs and AI in Education

***ChatGPT opportunities*.**

- LLMs can be used as a new education tool with massive impact in education.
- We have to research how to best use it.

- Its interaction with other teaching modes must be researched.
- Can it be used to trigger creative thinking, while speeding up tedious processes?

LLMs and AI in Education

IT and AI opportunities in education.

- What is the impact of IT and AI in teaching Mathematics?
- What is the impact of LLMs in teaching languages?
- What is the impact of Deep Arts in Arts Schools?

- What is the ***long-term impact of IT and AI*** in human memory?
- Will brain be 'restructured' to be, e.g., devoted more to thinking tasks than to memory?
- Can we observe such findings from historical records?

LLMs and AI in Education

UNESCO guidelines [MIA2023].

- Promote inclusion, equity, linguistic and cultural diversity.
- Protect human agency.
- Monitor and validate GenAI systems for education.
- ***Develop AI competencies including GenAI-related skills for learners.***
- Build capacity for teachers and researchers to make proper use of GenAI.
- Promote plural opinions and plural expressions of ideas.
- Test locally relevant application models and build a cumulative evidence base.
- Review long-term implications in intersectoral and interdisciplinary manner.

- ***Less than 10% of 450 schools/universities had policies on GenAI (2023).***

LLMs and AI in Education

Restrictive/regulated use of LLMs in education.

- Plagiarism tools to detect LLM-produced documents.
- Extreme caution when examining student projects
  - ***Very effort–intensive on Professors and students.***
- Extra caution in distance learning environments.
  - ***Return to old close student-Professor relations.***

- Imposition of minimal age to use LLM tools.

Artificial Intelligence & Information Analysis Lab

ChatGPT in Education

- ChatGPT can change the way we search and retrieve information.
- It has the capacity to help students reply to scientific questions.

- ChatGPT changes:
  - student project execution and examination.
  - educational exams.

ChatGPT in Education

‘Scientific’ capacity of ChatGPT:

- Good at replying factual questions on known topics.
- It has certain capacity to reply mathematical questions.
- It can solve programming exercises very well (e.g., in Python).
- ***Currently, it can neither process nor output diagrams and figures.***

ChatGPT in Education

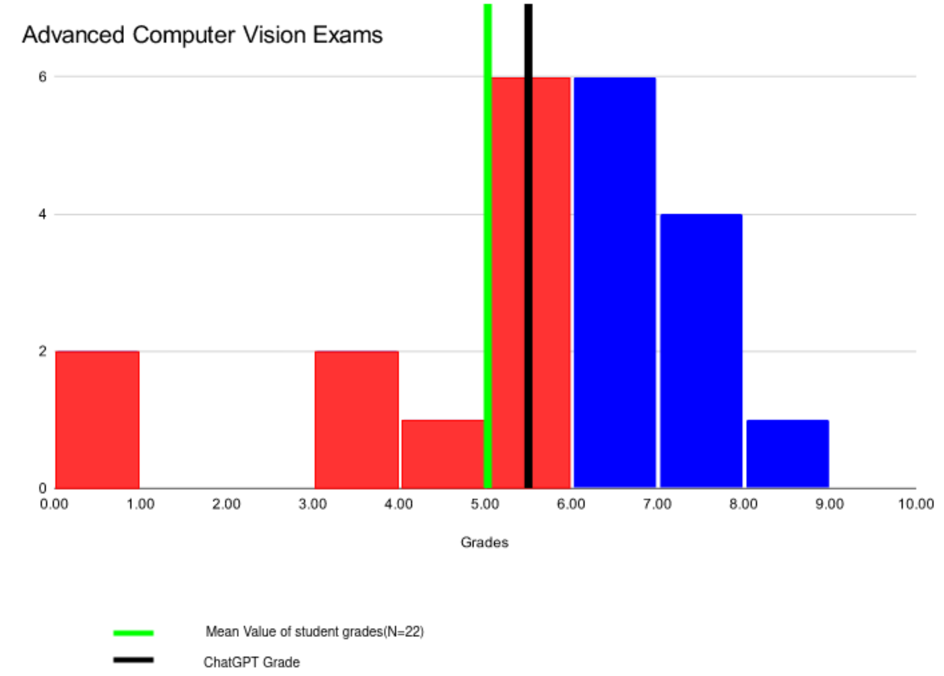

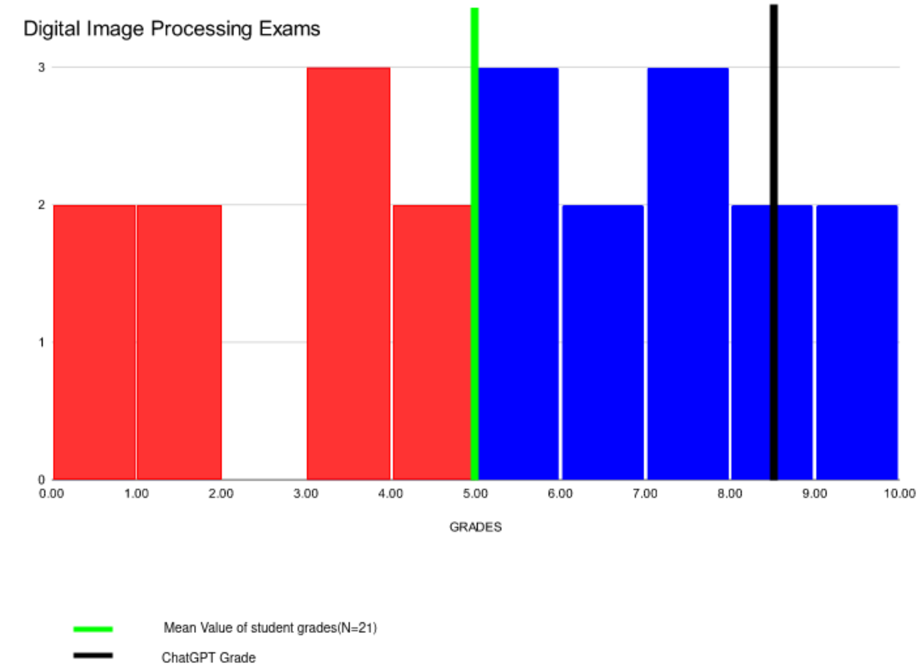

ChatGPT in CS/ECE exams: very good score in mathematical questions.

Artificial Intelligence & Information Analysis Lab

AI and Linguistic Studies

- What is AI?
- Machine Learning
- Natural Language Processing
- Large Language Models
- GPT and ChatGPT
- LLMs and AI in Education
- **AI and University Education**
- Citizen Morphosis
- Artificial General Intelligence

AI and University Education

- Currently, the Humanities face the greatest pressure from LLMs and AI.
- The ***mathematization of classical subjects*** (e.g., Linguistics, Sociology) has advanced significantly.
- Alternative? Creation of departments for '***Philological/Linguistic Engineering***' or '***Social Engineering***' in Science/Engineering Schools.

AI and University Education

Is the combination of Liberal studies and AI doable?

- The distinction between Humanities and Natural Sciences/Engineering persists in most countries.
- Women prefer Humanities than Natural Sciences/Engineering.
  - Only 25-30% of engineers are women.

AI and University Education

The distinction between Humanities and Natural Sciences/Engineering has deep historical roots.

- Humanities were meant to be for the gentry (ruling class) [BER1946].
- They are very old disciplines.
- Emphasis on character rather than knowledge:
  - ‘καλός κ’ αγαθός’ (in Ancient Greece)
  - Liberal education of ‘Piano and French’ style.

AI and University Education

The distinction between Humanities and Natural Sciences/Engineering has deep historical roots.

- Natural Sciences/Engineering are much younger (16-19th centuries).
- They facilitated the industrial revolution.
- They are much closer to profit making and burgeois ideology.

AI and University Education

Does the distinction between Humanities and Natural Sciences/Engineering have biological roots?

Contrasting arguments:

- There is no evidence that women are worse than men in mathematics.
- Women tend to have inclination to humanities, even if they do well in Mathematics (D. Kimura).
- Most people do not perform well in both linguistic and mathematical tests.

Exception: Few people are excellent in all disciplines.

AI and University Education

New Language Theory and Linguistic Methodologies

- Understanding of LLM performance.
- Development of new methodologies in Linguistic studies
  - Complement grammar, syntax, etymology etc.
- Teaching of this new theory and methodology.

- ***We are just at the start!***

AI and University Education

Past experience: from Humanities to Mathematics

- Transition from Aristotelian Logic to Mathematical Logic
  - Boolean Algebra and (19th century)
  - Foundations of Computers (mid 20th century)
- Mathematical Logic is essential tool in Symbolic AI (1960-1980)

- Too bad Symbolic AI failed to deliver so far.

AI and University Education

Creation of Departments for ‘***Mind and Social Science and Engineering***' in Schools of Arts and Humanities.

- Groundbreaking proposal.
- ***Departments of Digital Humanities*** is another good solution.
- The exact name or form is not important, as long as it serves the transfer of mathematical and programming skills to arts and humanities students.

AI and University Education

Alternatives:

- Introduction of 2-3 obligatory Mathematics and Compuer Science courses in each Liberal Discipline.

- Double BSc/MSc degrees ‘X+AI’
  - X: any Liberal Discipline (major).
  - AI minor

AI and University Education

Essential CS courses for AI education (minor in Liberal Studies):

- Mathematical Analysis
- Linear Algebra
- Probabilities and Statistics
- Signals and Systems
- Programming
- Machine Learning/Pattern Recognition
- Neural Networks
- Natural Language Processing.

AI and Linguistic Studies

- What is AI?
- Machine Learning
- Natural Language Processing
- Large Language Models
- GPT and ChatGPT
- LLMs and AI in Education
- AI and University Education
- **Citizen Morphosis**
- Artificial General Intelligence

Citizen Morphosis

Information and Knowledge Society

- Information society: exponential increase of data/information, linear increase of knowledge.
- Knowledge society: exponential increase of knowledge?
- AI, IT and ***citizen morphosis*** are our only hope to have a smooth transition from the current Information Society to a Knowledge Society.
- Else, humanity may face a catastrophic social implosion, if proven unable to advance and pass knowledge to new generations (see ***start of Medieval Times***).

Citizen Morphosis

Citizen morphosis (rather than education) emphasizes the need for conscious citizens:

- with ***critical thinking, communication precision skills, imagination, and emotional intelligence;***
- being able to understand, adapt, and ultimately harness the tremendous new technological and economic possibilities and employment prospects.
- Such a level of education is sought after today in many job positions internationally.

Citizen Morphosis

Major overhaul of education at all levels to master knowledge development and uptaking needs.

- The need for such education permeates all levels of education and all social strata.
- A ***1/3-2/3 society***, where 1/3 of the population understands and benefits from scientific progress, while the remaining 2/3 lags, being impoverished and technophobic, is simply not sustainable.
- Need to educate women, minorities and Global South to improve the global education level.

Citizen Morphosis

The ***basic AI and IT concepts*** are simple and can be taught at all educational levels:

- Data clustering, similarity, classification etc.
- Educational readjustment for their teaching by ***rearranging the curriculum of Mathematics and Informatics***.
- A (partial) mathematization of education is inevitable.
- It is not certain that it is feasible, given the traditional separation of the sciences and the humanities.

Citizen Morphosis

- ***Classical studies*** are also an ideal tool for developing critical thinking and precision.
- They provide a solid basis for ***Ethics, Legal and Social Implications*** (ELSI) knowledge.

AI and Linguistic Studies

- What is AI?
- Machine Learning
- Natural Language Processing
- Large Language Models
- GPT and ChatGPT
- LLMs and AI in Education
- AI and University Education
- Citizen Morphosis
- **Artificial General Intelligence**

Artificial General Intelligence

Is AGI the next step after LLMs?

- A deeper understanding of LLM operation is needed.
- The exact GPT4 architecture and parameters (transformer network weights) are a well-kept corporate secret.
- A deep LLM functionality understanding would be difficult, even if LLMs were open, due to their immense complexity.
- Neuroscience did not advance enough to understand brain and human intelligence.

Artificial General Intelligence

Is AGI the next step after LLMs?

- Most probably AGI will be VERY different from human intelligence.
  - Airplanes are different then birds, yet they obey the same laws of Physics.
- The physical substrate of AI and human intelligence is very different.
  - Robots have very limited but different physical intelligence.
  - Things may change by developing biological robots.
- ***Life evolution by-design*** than through physical selection.
- Massive ***human-machine symbiosis*** at various levels.

Artificial General Intelligence

Is AGI the next step after LLMs?

- Will AGI be any different from human intelligence from a behavioral point of view that is worth talking about?
- Today ***too many*** commoners cannot make the difference.
- The phenomenon is intensified by:
  - Lack of proper education.
  - Access of machines remotely.
  - Unwise claims and behavior of AI agents to the general public, e.g.,:
  - AI halucinations being misunderstood as imagination.
  - False claims of sentiments (internal affect states) by machines.

Artificial General Intelligence

Layman’s technophobia

- ***Fear of the unknown*** as commoners cannot understand AI.
- Machines appear to be intelligent and possibly better at that than the humans themselves.
- They are ***massively better*** in certain tasks, e.g., computations, memory/retrieval.
- Machines appear to be ***sentient***.
- Humans are awed by ChatGPT ‘intelligence’ much more than by other Generative AI methods, e.g., Deep Arts.
- ***Any technophobia can be socially destructive.***

Artificial General Intelligence

Scientific technophobia

- Very recent trend: scientists fearing the unknown.

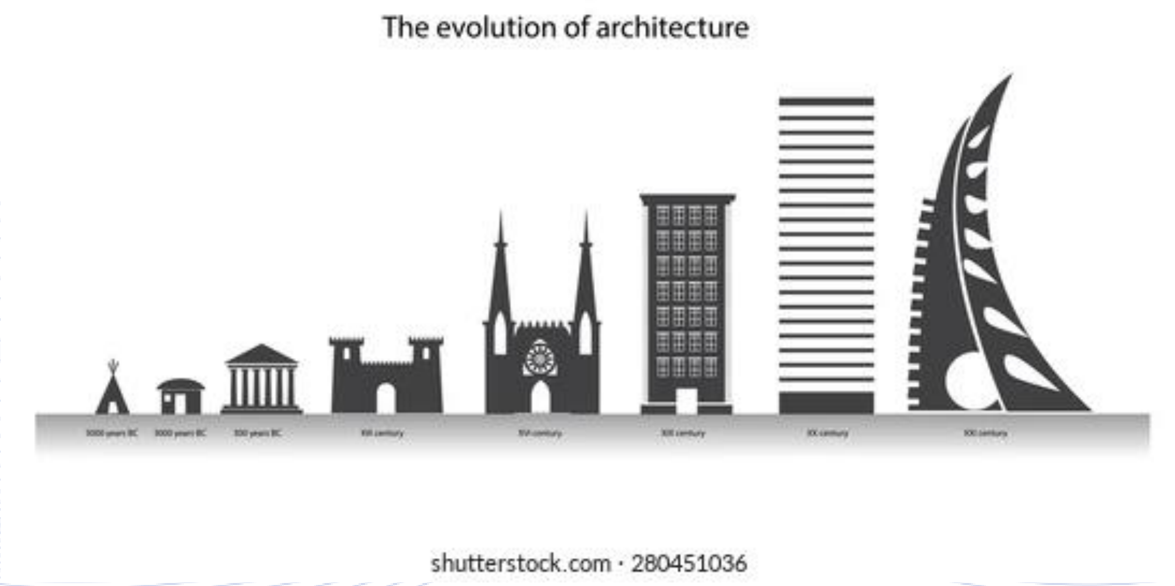

Parable: AI and the tower of Babel.

Artificial General Intelligence

Can AI be stopped or delayed?

- ***AI is the response of humanity to a global society and physical world of ever-increasing complexity.***
- The physical and social complexity increase processes are ***very deep and seeming relentless***.
- ***AI is a blessing, but it can become a curse.***
- Political, ethical, and regulatory concerns cannot and should not stop AI research [FUT2023].
- Scientific technophobia leads nowhere [NYT2023].

Artificial General Intelligence

Can AI be stopped or delayed?

- ***AI research can and should become more open, democratic, scientific and ethical.***
- Simple AI regulatory examples:
  - AI system registry,
  - Clear indication that somebody converses with a machine.
- AI deployment should be regulated and can be temporarily delayed.
  - Geopolitical aspects must be dealt by international cooperation.

Bibliography

[1] I. Pitas, “Artificial Intelligence Science and Society Part A: Introduction to AI Science and Information Technology“, Amazon/Kindle Direct Publishing, 2022,

https://www.amazon.com/dp/9609156460?ref_=pe_3052080_397514860

[2] I. Pitas, “Artificial Intelligence Science and Society Part B: AI Science, Mind and Humans“, Amazon/Kindle Direct Publishing, 2022, https://www.amazon.com/dp/9609156479?ref_=pe_3052080_397514860

[3] I. Pitas, “Artificial Intelligence Science and Society Part C: AI Science and Society”, Amazon/Kindle Direct Publishing, 2022, https://www.amazon.com/dp/9609156487?ref_=pe_3052080_397514860

[4] I. Pitas, “Artificial Intelligence Science and Society Part D: AI Science and the Environment“, Amazon/Kindle Direct Publishing, 2022,

https://www.amazon.com/dp/9609156495?ref_=pe_3052080_397514860

Bibliography

[PIT2023a] Ioannis Pitas, CVML short course, “AI Science and Engineering and its Impact on the Society”, https://icarus.csd.auth.gr/introduction-to-ai-science-and-engineering-and-its-impact-on-the-society-and-the-environment/

[PIT2022] Ioannis Pitas, “AI Science and Engineering: A new scientific discipline?”, https://icarus.csd.auth.gr/chatgtp-in-education/

[PIT2023b] Ioannis Pitas, “ChatGPT in education”, http://icarus.csd.auth.gr/ai-science-and-engineering-a-new-scientific-discipline/

[PIT2023c] I. Pitas, “Artificial intelligence is not the new Tower of Babel. We must beware of technophobia instead”, Euronews, 8/5/2023, https://www.euronews.com/2023/05/08/artificial-intelligence-is-not-the-new-tower-of-babel-we-should-beware-of-technophobia-ins

[CHO2023] N. Chomsky, I. Roberts, J. Watumull, “The false promise of ChatGPT”, The New York Times, 8th March 2023.

[FUT2023] “Pause Giant AI Experiments: An Open Letter”, https://futureoflife.org/open-letter/pause-giant-ai-experiments/, 2023.

[NYT2023] The Godfather of A.I.’ Leaves Google and Warns of Danger Ahead, The NY Times, https://www.nytimes.com/2023/05/01/technology/ai-google-chatbot-engineer-quits-hinton.html

Bibliography

[OPE2023] OpenAI. “CHATGPT: Optimizing Language Models for Dialogue.” OpenAI, 2 Feb. 2023, https://openai.com/blog/chatgpt/.

[OP2023] OpenAI. “GPT-4 Technical Report.” OpenAI, 15 March 2023

[VAS2017] Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. "Attention is all you need." Advances in neural information processing systems 30 (2017).

[RAD2018] Radford, Alec, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. "Improving language understanding by generative pre-training." (2018).

[RAD2019] Radford, Alec, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. "Language models are unsupervised multitask learners." OpenAI blog 1, no. 8 (2019): 9.

[BRO2020] Brown, Tom, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D. Kaplan, Prafulla Dhariwal, Arvind Neelakantan et al. "Language models are few-shot learners." Advances in neural information processing systems 33 (2020): 1877-1901.

Bibliography

[BAN2023] Bang, Yejin, Samuel Cahyawijaya, Nayeon Lee, Wenliang Dai, Dan Su, Bryan Wilie, Holy Lovenia et al. "A Multitask, Multilingual, Multimodal Evaluation of ChatGPT on Reasoning, Hallucination, and Interactivity." arXiv preprint arXiv:2302.04023 (2023).

[ZHU2015] Y. Zhu et al., "Aligning Books and Movies: Towards Story-Like Visual Explanations by Watching Movies and Reading Books," 2015 IEEE International Conference on Computer Vision (ICCV), Santiago, Chile, 2015, pp. 19-27, doi: 10.1109/ICCV.2015.11.

[FOU] The Common Crawl Foundation. Common crawl. URL http://commoncrawl.org.

[ALE2019] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners, 2019.

[JAR2020] Yukun Zhu, Ryan Kiros, Rich Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. Aligning books and movies: Towards story-like visual explanations by watching movies and reading books. 2015 IEEE International Conference on Computer Vision (ICCV), Dec 2015. doi:10.1109/iccv.2015.11.

Bibliography

[CFTE] CHATGPT outperforms 85% of the 4 million programmers assessed on python on linkedin (2022) CFTE. Available at: https://courses.cfte.education/chatgpt-outperforms-4mprogrammers/#:~:text=Some%20of%20its%20abilities%20include,large%20portion%20of%20our%20work

[SCA2022] Scao, T.L., Fan, A., Akiki, C., Pavlick, E., Ilić, S., Hesslow, D., Castagné, R., Luccioni, A.S., Yvon, F., Gallé, M. and Tow, J., 2022. Bloom: A 176b-parameter open-access multilingual language model. arXiv preprint arXiv:2211.05100.

[CAH2022] Cahyawijaya, S., Lovenia, H., Aji, A.F., Winata, G.I., Wilie, B., Mahendra, R., Wibisono, C., Romadhony, A., Vincentio, K., Koto, F. and Santoso, J., 2022. NusaCrowd: Open Source Initiative for Indonesian NLP Resources. arXiv preprint arXiv:2212.09648.

[MIA2023] F. Miao, W. Holmes, “Guidance for generative AI in education and research”, UNESCO 2023.

[BER1946] J.D. Bernal, ‘Science and the Humanities’, lecture at Birkbeck College, 1946.

Q & A

Thank you very much for your attention!

More material in
http://icarus.csd.auth.gr/cvml-web-lecture-series/

Contact: Prof. I. Pitas
pitas@csd.auth.gr