

AI Science and Engineering: A new scientific discipline?

Prof. Ioannis Pitas

Aristotle University of Thessaloniki

pitas@csd.auth.gr

www.aiia.csd.auth.gr

Version 3.0

AI Science and Engineering

- **What is AI?**
- Statistical Machine Learning
- AI and Human Mind
- Artificial General Intelligence
- AI Science and Engineering?
- University Education on AI
- AI in University Education

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:
 - ***Machine Learning*** (ML),
 - Classical (Symbolic) ***Artificial Intelligence*** (AI)

What is AI?

- Closely related AISE disciplines:
 - **Robotics,**
 - Autonomous Systems,
 - Digital Signal/Image Processing and Analysis,
 - Data Science and Data Analytics
 - **Network Theory.**
- Very useful in defining:
 - Data analysis modes, applications.

What is AI?


- Complementary AISE-related disciplines:
 - Cognitive Science,
 - Neuroscience,
 - Psychology,
 - ***Philosophy, Ethics***
 - Linguistics
 - Sociology.

What is AI?

Data/information/knowledge definitions

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

- **Data are primarily numbers** representing object characteristics (features).
- Passive/active data acquisition.
- Data sampling.
- **Measured in bits.**





What is AI?

Data can have **spatiotemporal structure**:

- 1D temporal signals, e.g., music
- 2D spatial signals: images
- Signals and object features can be represented by **vectors**:

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




What is AI?

Exponential data increase:

- Proliferation of sensors
- Detailed recording of nature and humans
- Sensing automation.

Data volume increase in past decade.



What is AI?

Why we need ever more data?

- To navigate in an ever more complex world.
 - *Why do we need a more complex world?*


Data sustainability:

- HW enabled
- ***Moore's law***
- Data storage constraints
- Data communication constraints.

What is AI?

Unsupervised Machine Learning

- Data clustering:




- Data geometry
- ***Abstraction***
- ***Data compression.***

What is AI?

Supervised Machine Learning

- Learning functions $y = f(\mathbf{x}; \theta)$ from labeled training data $\{(\mathbf{x}_i, y_i), i = 1, \dots, N\}$.
- ***Classification***
- ***Regression.***

- Learning data probability distributions $p(\mathbf{x})$.
 - ***Generative neural networks.***
 - ***Fake data creation.***



What is AI?

Information

- **Notoriously vague definitions**
- My definition: ***Information is the result of the manual or automatic Data Analysis.***

Taxonomy: Data → Information → Knowledge.

Machine Learning/inference produces ***information*** (including metadata).

- ***Information theory/entropy: bits (once more)!***



What is AI?

Concepts and ideas (ιδέες).

- Concepts are specific mental constructs residing in our mind (brain?) that refine and abstract ideas.
- **Concept instances**




Instances of a triangle.

- **Abstraction and generalization:**
 - Simplification and data compression.

What is AI?

Ideas in Philosophy.

- Idealism, materialism, dualism.
- Plato's cave.



What is AI?


Symbolic AI

- A **symbol** (‘Σύμβολο’) is a **comprehensible representation** of an object, idea, concept, action, status, or relationship.
- Symbolic AI mimics and simulates high-level human intelligence and **reasoning**.
- It represents and operates on concepts and their relations through **logic** and **search**.
- **Reasoning** is one of the most complex brain activities.

What is AI?

Knowledge

- It is a familiarity, awareness, or ***understanding of someone or something***:
 - Facts (propositional knowledge),
 - Skills (procedural knowledge),
 - Objects relations (relational knowledge).
- Various knowledge descriptions.




What is AI?

Knowledge is primarily a product of reasoning.

- Is knowledge finite?
- ***Can we measure knowledge?***
- Knowledge increase is linear.

- ***Encyclopedias***
- ***Research publications.***



Global research output (publication) growth.

What is AI?

Current AI revolution:

- *AI means ML, which means Deep Neural Networks*
- Stagnation of symbolic AI
- Resurrection of a dead term: AI

Major breakthrough needed:

- Advancement of symbolic AI
- ***Fusion of Machine Learning and symbolic AI.***

What is AI?

Data/Information society:

- Exponential data growth.
- Data acquisition automation.
- ***Information extraction automation through ML.***

Sustainability?

- More sensors, more processors, Moore's law.
- ***Energy-intensive data and information extraction.***

What is AI?

Knowledge society:

- Exponential knowledge growth.
- Not there yet: ***knowledge production and communication is still manual.***
- Past devastating setbacks in knowledge uptaking:
 - Dark ages (beginning of the medieval times).

What is AI?

Knowledge Sustainability:

- Limitations in brain capacity.
- Solution: **social swarm intelligence**
- Example: collective memory.
- Knowledge communication through **education** is way suboptimal:
 - New education mode needed, stressing **critical thinking** and **abstraction**.
 - **Morphosis**: formation of knowledgeable citizens.
 - **Global education**: diminishing social and regional barriers to education.

- **Unified machine and human learning theories?**

AI Science and Engineering

- What is AI?
- **Statistical Machine Learning**
- AI and Human Mind
- Artificial General Intelligence
- AI Science and Engineering?
- University Education on AI
- AI in University Education

Statistical Machine Learning


Supervised Machine Learning

Learning functions $y = f(\mathbf{x}; \theta)$.

- ***Bayesian Learning/Decision Making:***
 - Learning from Probability distributions $p(\mathbf{x})$.
- ***Statistical Machine Learning:***
 - Learning from labeled training data $\{(\mathbf{x}_i, \mathbf{y}_i), i = 1, \dots, N\}$.

Applications:

- ***Classification***
- ***Regression.***



Statistical Machine Learning




Bayesian Classification

Two-class Maximum Likelihood classification problem:

- Adopt \mathcal{C}_1 , if $r_1(\mathbf{x}) < r_2(\mathbf{x})$ or:

$$\Lambda(\mathbf{x}) = \frac{p(\mathbf{x}|\mathcal{C}_1)}{p(\mathbf{x}|\mathcal{C}_2)} > T_{12}.$$

- Probabilities $p(\mathbf{x}|\mathcal{C}_i)$ are unknown and have to be estimated.



ML decision boundary.

Statistical Machine Learning

Multivariate Gaussian Probability Distribution


- Jointly normal variables X_1, \dots, X_n :

$$f_{\mathbf{X}}(\mathbf{x}) = \frac{1}{(\sqrt{2\pi})^n \det(\mathbf{C})^{\frac{1}{2}}} e^A,$$

$$A = -\frac{1}{2} (\mathbf{x} - \mathbf{m})^T \mathbf{C}^{-1} (\mathbf{x} - \mathbf{m})$$

Parameters $\hat{\mathbf{m}}, \hat{\mathbf{C}}$ to be estimated from data $\{(\mathbf{x}_i, \mathbf{y}_i), i = 1, \dots, N\}$:

- Expected vector: $\mathbf{m} = E\{\mathbf{x}\}$.
- Covariance matrix: $\mathbf{C} = E\{(\mathbf{x} - \mathbf{m})^T (\mathbf{x} - \mathbf{m})\}$.



Statistical Machine Learning



Neural Networks.

Perceptron:

$$\hat{y} = f(z) = f(\mathbf{w}^T \mathbf{x} + b) = f\left(\sum_{i=1}^N w_i x_i + b\right)$$

- f : activation function.
- Simplest form (firing threshold):

$$\mathbf{w}^T \mathbf{x} \geq -b \Rightarrow \mathbf{w}^T \mathbf{x} + b \geq 0.$$



Statistical Machine Learning



Classification is a binary function **prediction** (estimation):

$$\mathbf{y} = f(\mathbf{x}, \mathbf{w}).$$

- **Input.** $\mathbf{x} = [x_1, x_2, \dots, x_n]^T$, e.g., facial 100×80 pixel image.
- **Trainable parameters** (NN weights): $\mathbf{w} = [w_1, w_2, \dots, w_n]^T$.
- **Output.** $\mathbf{y} = [0, 1, 0, \dots, 0]^T$.
- Only the correct facial (person) class label is 1.

Statistical Machine Learning


Multilayer Perceptrons

$$\mathbf{y} = f(\mathbf{x}, \mathbf{w}).$$

- Neural networks training has to minimize an error function $J(\mathbf{w})$.
- Differentiation:

$$\frac{\partial J(\mathbf{w})}{\partial \mathbf{w}} = \mathbf{0}.$$

Input layer ($l = 0$) 1st hidden layer ($l = 1$) 2nd hidden layer ($l = 2$) 3rd hidden layer ($l = 3$) Output layer ($l = L = 4$)



Multilayer perceptron.

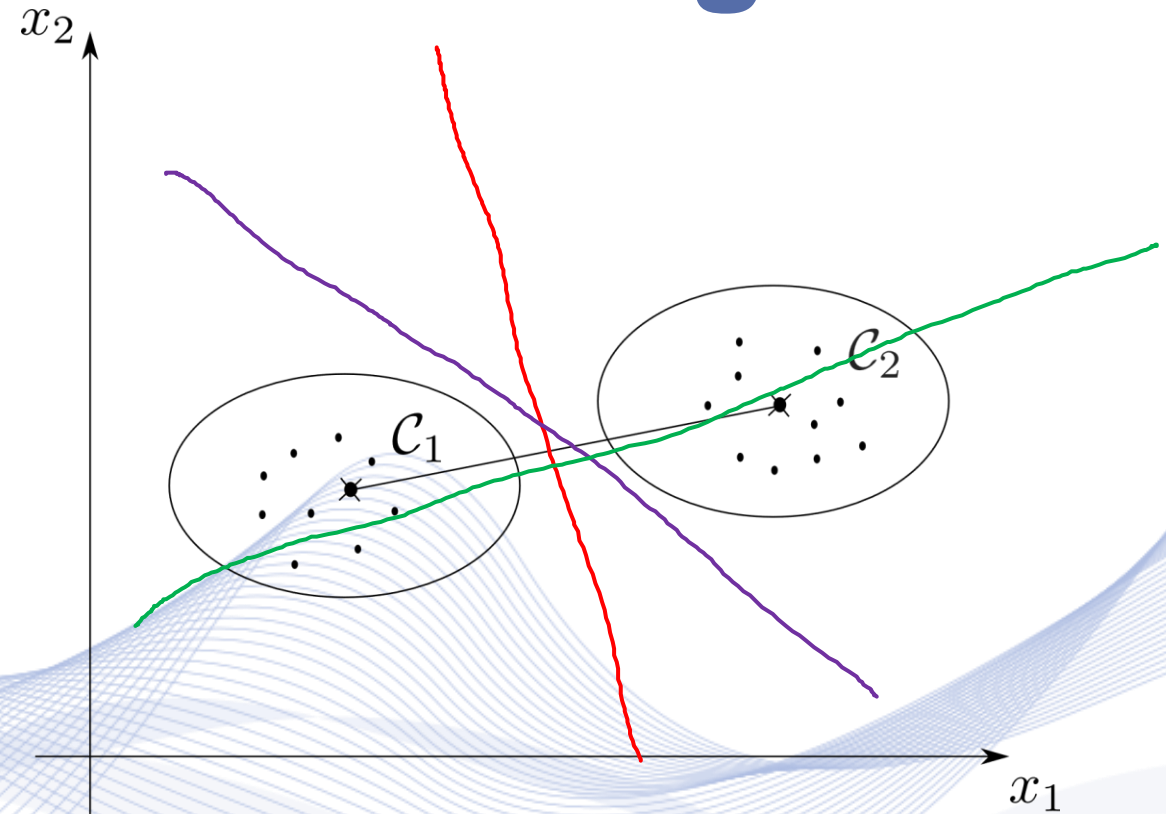

Statistical Machine Learning

Classification.

- ***Training:*** Use training data $\{(\mathbf{y}_i, \mathbf{x}_i)\}$ to find the optimal parameters \mathbf{w} , minimizing the classification error $J(\mathbf{y}_i, \mathbf{x}_i, \mathbf{w})$.
- ***Inference:*** Feed the trained NN with data \mathbf{x} to produce the classification label: $\mathbf{y} = \mathbf{f}(\mathbf{x}, \mathbf{w})$.
- Classification is a special type of ***regression*** (function approximation).

Statistical Machine Learning

2D perceptron.



Separating line: $w_1x_1 + w_2x_2 + b > 0$.

Statistical Machine Learning

2D perceptron.


Decision line: $w_1x_1 + w_2x_2 + b > 0$.

Classification error minimization:

$$J(w_1, w_2, b) = 4.$$

Optimization problem.

- Use gradients to find the minimum!



Statistical Machine Learning

2D perceptron.


Decision line: $w_1x_1 + w_2x_2 + b > 0$.

Classification error minimization:

$$J(w_1, w_2, b) = 2.$$

Optimization problem.

- Use gradients to find the minimum!



Statistical Machine Learning

2D perceptron.

Decision line: $w_1x_1 + w_2x_2 + b > 0$.


Classification error minimization:

$$J(w_1, w_2, b) = 0.$$


$$J(w_1, w_2, b) = 0.$$

Optimization problem.

- **Use gradients (derivatives) to find the minimum!**




Statistical Machine Learning




Perceptron training through classification error
 $J(w_1, w_2)$ minimization.

Statistical Machine Learning




Perceptron training through classification error
 $J(w_1, w_2)$ minimization.

Statistical Machine Learning




Perceptron training through classification error
 $J(w_1, w_2)$ minimization.

Statistical Machine Learning




Perceptron training through classification error
 $J(w_1, w_2)$ minimization.

Statistical Machine Learning




Perceptron training through classification error
 $J(w_1, w_2)$ minimization.

Statistical Machine Learning




Perceptron training through classification error
 $J(w_1, w_2)$ minimization.

Statistical Machine Learning



Perceptron training through classification error
 $J(w_1, w_2)$ minimization.

Statistical Machine Learning




Perceptron training through classification error
 $J(w_1, w_2)$ minimization.

Statistical Machine Learning



Regression is an approximation $\hat{y} = \hat{f}(\mathbf{x}, \mathbf{w})$ of a real-valued function $y = f(\mathbf{x})$.

- **Input:** \mathbf{x} (values in the function domain).
- **Trainable parameters** \mathbf{w} .
- **Output vector** \hat{y} : approximated function values.
- Training and inference.




Statistical Machine Learning



Convolutional Neural Networks (CNN):

- Employ image convolutions in the first layers.
- They may employ fully connected MLPs in the last layers.



Basic CNN structure.

Statistical Machine Learning



CNN image features (vertical image edges).

Transformers and Attention Networks

- CNNs cannot accommodate distant data correlations.
- Attention is essentially a ***statistic correlation*** mechanism to extract meaningful information from data.
- It also provides a data feature diffusion mechanism.



Image patches.

Statistical Machine Learning

Generative learning models

approximate the data generating probability density function (pdf)

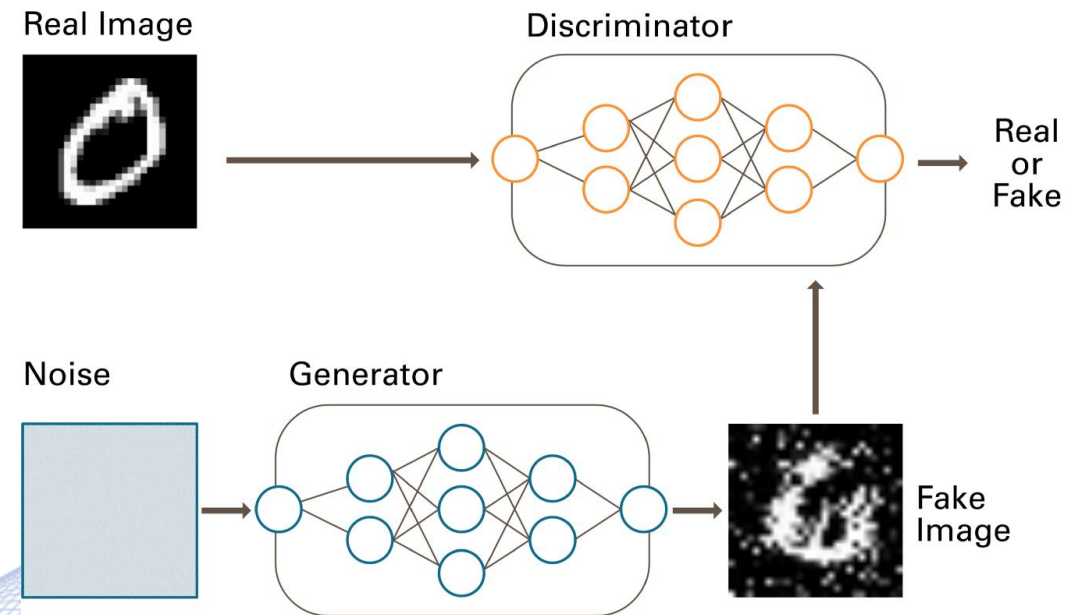
$$\hat{p}(\mathbf{x}) = f(\mathbf{x}; \boldsymbol{\theta}).$$

Generative Adversarial Networks

have:

- A **Generator** function $\hat{y} = G(\mathbf{z}; \boldsymbol{\theta}_G)$.
- A **Discriminator** function $\hat{y} = D(\mathbf{q}; \boldsymbol{\theta}_D)$.

They are one form of Generative AI.



GAN architecture.

Statistical Machine Learning



Generative AI.

- Large Language Models
- Diffusion Models
- Generative Adversarial Networks.

It can be used to create fake data:

- Fake news, images, videos, audio.
- Works of art.




AI Science and Engineering

- What is AI?
- Statistical Machine Learning
- **AI and Human Mind**
- Artificial General Intelligence
- AI Science and Engineering?
- University Education on AI
- AI in University Education


AI and Human Mind

Artificial and Biological neural networks

- Is *network complexity* the basis of both the biological and artificial intelligence?



Biological NN (https://en.wikipedia.org/wiki/Cerebral_cortex)



Multilayer perceptron

AI and Human Mind


Interoception and Physical Intelligence

- **Interoception** is the perception of stimuli from inside our body.
- It supports **homeostasis** (maintenance of functional body equilibrium).
- **It is essential for human (self)consciousness.**
- **Current robots do not have interoception.**
- Closest approximation: **Physical Intelligence** uses distributed sensors to allow robots to live in unstructured environments.
 - Multimodal machine perception: tactile, smell, taste sensors.

AI and Human Mind

Brain-Inspired Computing

- **Computational Neuroscience** creates mathematical models of the brain and nervous systems.
- **Despite advances, no breakthroughs compared to AI revolution.**
- Modeling memory, consciousness, affect etc.
- Major advances expected by Neuroscience and AI/ML fusion.



AI and Human Mind

Intelligent Self-aware systems

- Memory (easy)
- Affect (easy?)
- Consciousness(doable?)
- Real intelligence (difficult?)
- Swarm/social intelligence (doable).

AI and Human Mind

- Intelligent systems can be very useful.
- ***Should we be technophobic?***



AI Science and Engineering

- What is AI?
- Statistical Machine Learning
- AI and Human Mind
- **Artificial General Intelligence**
- AI Science and Engineering?
- University Education on AI
- AI in University Education

Artificial General Intelligence

Is AGI the next step after LLMs?

- Most probably AGI will be VERY different from human intelligence.
 - Airplanes are different than 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 hallucinations 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.



shutterstock.com · 280451036

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.

AI Science and Engineering

- What is AI?
- Statistical Machine Learning
- AI and Human Mind
- Artificial General Intelligence
- AI Science and Engineering?
- University Education on AI
- AI in University Education

AI Science and Engineering: A new scientific discipline?

- **Computer Science** is the study of computation and information.
- **Computer Engineering** is a branch of Engineering that integrates several CSE fields that are required to develop computer hardware and software.
- Traditionally, AI and ML were CSE disciplines.
- Do AI and ML have own scientific methodology?

AI Science and Engineering: A new scientific discipline?

AI/ML: mathematical discipline at the CS/CSE/ECE crossroads.

- **Mathematics lost this topic from the early start (1960s).**
- ECE Departments are much more venturous due to:
 - Early needs in Statistical Communications (1950).
 - Good background in mathematical coursework and Pattern Recognition.
- CS Departments were early starters in the wrong direction:
 - **Symbolic AI**
 - **Weak mathematical coursework.**
- ***Yet CS/CSE is the clear winner in the AI/ML race.***

AI Science and Engineering: A new scientific discipline?

- ***AISE Interdisciplinarity?***
 - AI and Brain/mind studies
 - AI and social studies/engineering.
- **Mature AISE Interdisciplinarity?**
 - **Not there yet!**
- Risks: ***depth vs shallowness.***

AI Science and Engineering: A new scientific discipline?

CSE spawning new disciplines *through specialization*:

- Web science
- Data science
- AI Science and Engineering.
- New scientific methodologies are not **necessarily** essential.
- Poor terminology?
- Past experience: **Physics spawning Engineering disciplines**
 - Electrical Engineering, Mechanical Engineering.

AI Science and Engineering: A new scientific discipline?

AISE background

Lots of mathematics:

- **Analysis/calculus, Optimization**
- Geometry
- Linear Algebra
- Graph Theory
- **Probability theory and statistics**
- Mathematical Logic.

AI Science and Engineering: A new scientific discipline?

AISE background

Classical studies at University and high school level

- Philosophy, ethics, logic
- Linguistic competences.
- Physics?, Biology?
- ***Do we prescribe universal AI scientists?***
- **What about commoners?**

AI Science and Engineering: A new scientific discipline?



Worldwide creation of:

- ***New AI Departments or Schools***
- ***New AI undergraduate studies.***

It seems it is not just a trend.

Many efforts are market driven.

Urgent need for the creation of a good AI Curriculum:

- AIDA AI Curriculum: <https://www.i-aida.org/phd-curriculum/>

AI Science and Engineering: A new scientific discipline?



Changes will be drastic and will come very soon.

Schools of 'Information Science and Engineering' with departments of:

- Computer Science/Informatics,
- Mathematics
- Computer Engineering
- Artificial Intelligence Science and Engineering
- Internet/Web Science.


AI Science and Engineering

- What is AI?
- Statistical Machine Learning
- AI and Human Mind
- Artificial General Intelligence
- AI Science and Engineering?
- **University Education on AI**
- AI in University Education

University Education on AI

- Very many AI MSc and PhD study programs
- AI Schools & Departments (12)
- AI Undergraduate Studies (59)
- ***Developments are mostly demand-driven.***
- Smaller players can be more adventurous in AI studies.


University Education on AI



Countries that offer AI studies.


University Education on AI

Undergraduate AI Programs Distribution




Number of undergraduate AI programs worldwide.

Undergraduate AI Programs Distribution




Global distribution of undergraduate AI studies.

University Education on AI



Distribution of undergraduate AI programs in Europe.

University Education on AI



Geographical distribution of AI undergraduate programs in Europe.

AI Science and Engineering

- What is AI?
- Statistical Machine Learning
- AI and Human Mind
- Artificial General Intelligence
- AI Science and Engineering?
- University Education on AI
- **AI in University Education**

AI in 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 in 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 in University Education

Creation of departments for '***Bio-Science and Engineering***' in Schools of Health Sciences, including:

- Biomedical Engineering, Genetic Engineering and Systems Biology.

Mandatory inclusion of Mathematics and Computer Science courses in all disciplines without exception.

- Simply, one (poor) course in Statistics does not meet the current needs.
- Mandatory courses on AI ***Ethics, Legal and Social Implications*** (ELSI) in all ECE, EE, CS and CSE Curricula.

- It is already partly underway.

I. Pitas, 'AI Science and Society'



4-volume book, 1070+ pages, Amazon 2023.

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

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
pitass@csd.auth.gr**