



# DEEP NETWORK DEVELOPMENT

**Imre Molnár**

PhD student, ELTE, AI Department

✉ [imremolnar@inf.elte.hu](mailto:imremolnar@inf.elte.hu)

🌐 [curiouspercibal.github.io](https://curiouspercibal.github.io)

**Tamás Takács**

PhD student, ELTE, AI Department

✉ [tamastheactual@inf.elte.hu](mailto:tamastheactual@inf.elte.hu)

🌐 [tamastheactual.github.io](https://tamastheactual.github.io)

# Lecture 1.

# Modern AI

---

Budapest, 14<sup>th</sup> February 2025

**1** Course Details

**2** Deep Learning Overview

**3** Modern AI & State of AI

# Lecture 1.

## Course Details

---

Budapest, 14<sup>th</sup> February 2025

**1** Course Details

**2** Deep Learning Overview

**3** Modern AI & State of AI

# Staff

## Lecture and Practice:



**Imre Molnár**

[imremolnar@inf.elte.hu](mailto:imremolnar@inf.elte.hu) ([wkzina@inf.elte.hu](mailto:wkzina@inf.elte.hu))  
<https://curiouspercibal.github.io/>



**Tamás Takács**

[tamastheactual@inf.elte.hu](mailto:tamastheactual@inf.elte.hu) ([cjrnle@inf.elte.hu](mailto:cjrnle@inf.elte.hu))  
<https://tamastheactual.github.io/>

### Demonstrators:



Máté Badó ([blkv0u@inf.elte.hu](mailto:blkv0u@inf.elte.hu))



Samiha Nasser ([wp01th@inf.elte.hu](mailto:wp01th@inf.elte.hu))

### Instructors:



Ádám Fodor ([foauaai@inf.elte.hu](mailto:foauaai@inf.elte.hu))



Zoltán Barta ([dguqkf@inf.elte.hu](mailto:dguqkf@inf.elte.hu))

# Structure

- Classes are every Friday in the South Building:
  - Lectures: 8:00 – 10:00 (O-823 Kitaibel Pál terem)
  - Practices: 10:00 – 12:00 (O-822 Mogyoródi József terem)
- Lecture recordings will be available throughout the semester
- Course syllabus is available on Canvas
- All the materials (lecture slides, practice notebooks, ...) will be uploaded to Canvas

# Prerequisites

Have a basic understanding of:

- Linear Algebra
- Probability Theory
- Programming Skills (for practices)

```
# Python 3: Simple output (with Unicode)
>>> print("Hello, I'm Python!")
Hello, I'm Python!

# Input, assignment
>>> name = input('What is your name?\n')
>>> print('Hi, %s.' % name)
What is your name?
Python
Hi, Python.
```




### Matrix rules

scalar multiplication  $n \begin{bmatrix} a & b & c \\ d & e & f \end{bmatrix} = \begin{bmatrix} na & nb & nc \\ nd & ne & nf \end{bmatrix}$

matrix addition  $\begin{bmatrix} a & b \\ c & d \\ e & f \end{bmatrix} + \begin{bmatrix} g & h \\ i & j \\ k & l \end{bmatrix} = \begin{bmatrix} a+g & b+h \\ c+i & d+j \\ e+k & f+l \end{bmatrix}$

matrix multiplication  $\begin{bmatrix} a & b & c \\ d & e & f \end{bmatrix} \begin{bmatrix} g & h \\ i & j \\ k & l \end{bmatrix} = \begin{bmatrix} ag+bi+ck & ah+bj+cl \\ dg+ei+fk & dh+ej+fl \end{bmatrix}$

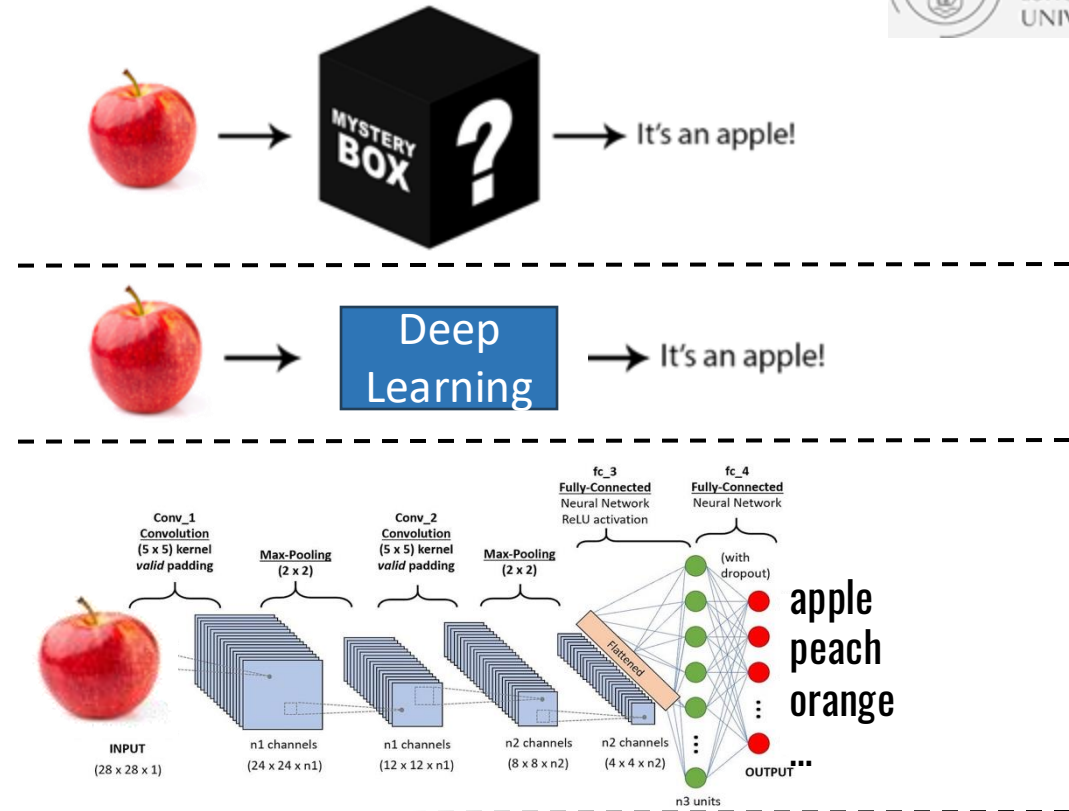
### Probability Line






## Description

- This course provides students with a comprehensive understanding of how **Deep Neural Networks** work:




Level of Understanding





# Description

- The practical part of this course equips students with the necessary skills to implement various Deep Neural Network architectures using a framework called **PyTorch** based on Python programming language;



# Evaluation

- **2 Homeworks** – (announced on **week 3** and **5**)
- **4 Quizzes** – (made available on **week 3, 6, 8, and 10**)
- **1 Assignment** - shared during the **7th week** (tentative) and have submission deadlines (no extension)
  - The assignment **must be defended** (instructors and demonstrators ask questions to verify their knowledge)

### Assignment 1 [Topic TBD]

#### TENTATIVE DATES:

- Deadline: end of week 11 (02.05.2025 11:59 PM)
- Defense: practice of week 12 (09.05.2025)

# Evaluation

- **Lecture Grade:**

$$\text{Lecture} = 0.1 \times (Q1 + Q2 + Q3 + Q4) + 0.6 \times T1$$

- **Q1, Q2, Q3, Q4** – Quiz grades
- **T1** – Theoretical defense grade

- **Practice Grade:**

$$\text{Practice} = 0.2 \times (H1 + H2) + 0.6 \times A1$$

- **H1, H2** - grade from homework 1 and 2
- **A1** - grade from assignment 1 submission and code grade
- To be eligible for the final exam, you must achieve at least a **grade of 2** in both the **Lecture** and **Practice** component
- Further information about the exam will be released later

# Lecture 1.

# Deep Learning Overview

Budapest, 14<sup>th</sup> February 2025

**1** Course Details

**2** Deep Learning Overview

**3** Modern AI & State of AI

# Artificial Intelligence (AI)

**AI:** the science and engineering of making computer systems and algorithms capable of performing tasks that typically require human intelligence.

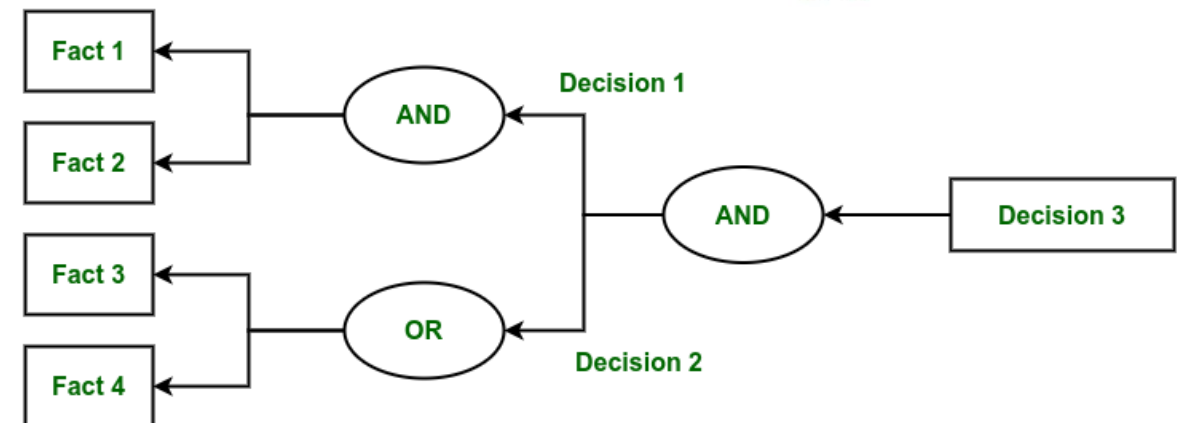
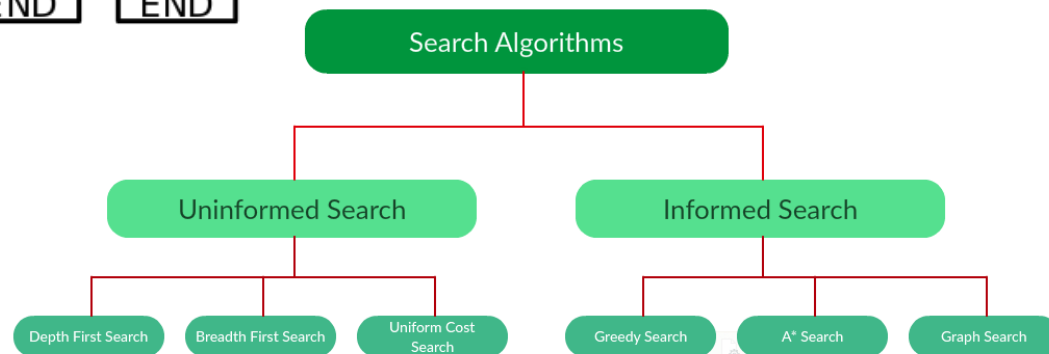
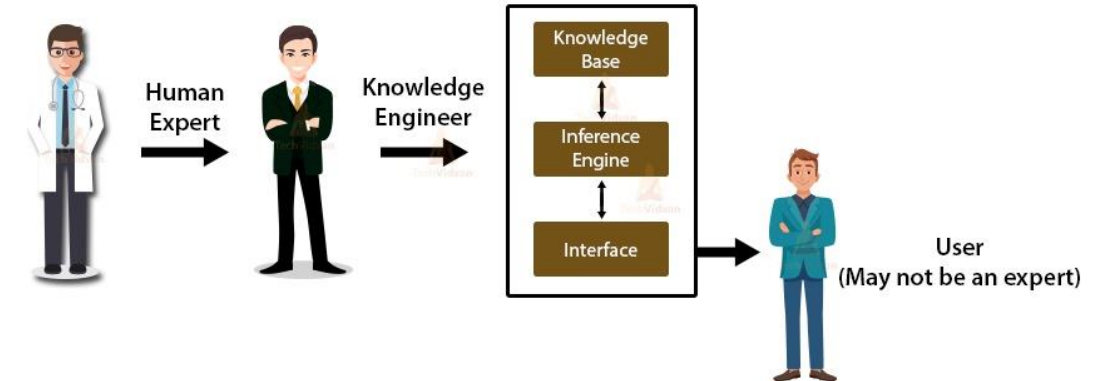
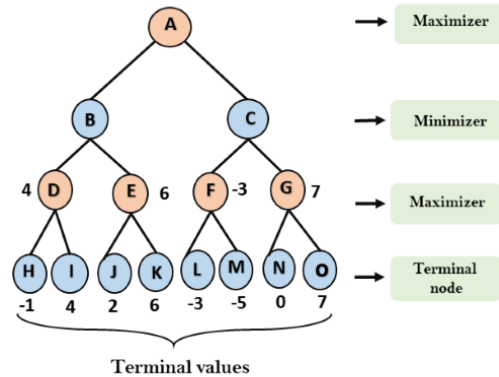



**“Intelligent”  
Machines /  
Systems  
... behave like  
humans  
... talk like  
humans**




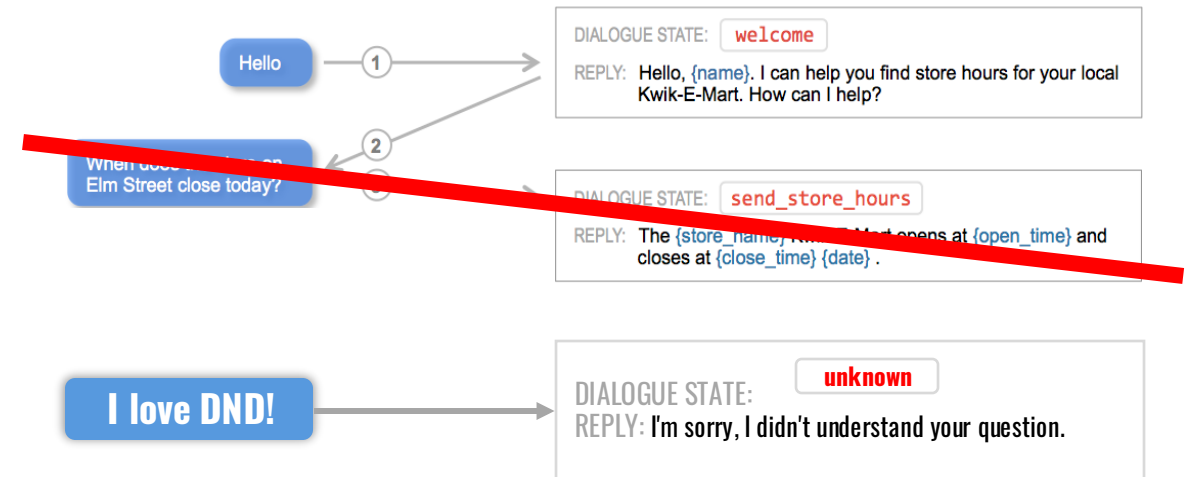
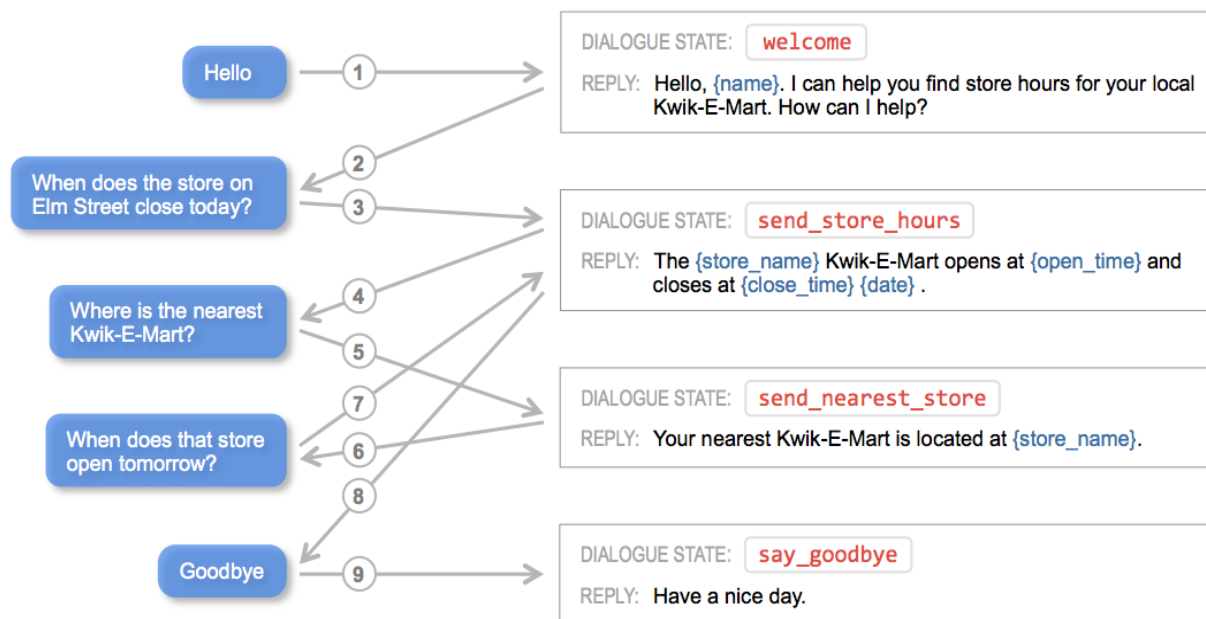
# Artificial Intelligence (AI)

- Traditional AI methods used to be rule-based or goal oriented.
- Expert Systems:** contain a knowledge base and rules for reasoning and decision-making



# Artificial Intelligence (AI)

- Traditional AI methods used to be rule-based or goal oriented.
- Expert Systems:** contain a knowledge base and rules for reasoning and decision-making






# Machine Learning (ML)


**ML:** a subfield of AI that focuses on the development of algorithms and models capable of learning from data and making predictions without being explicitly programmed.



Learn patterns from data

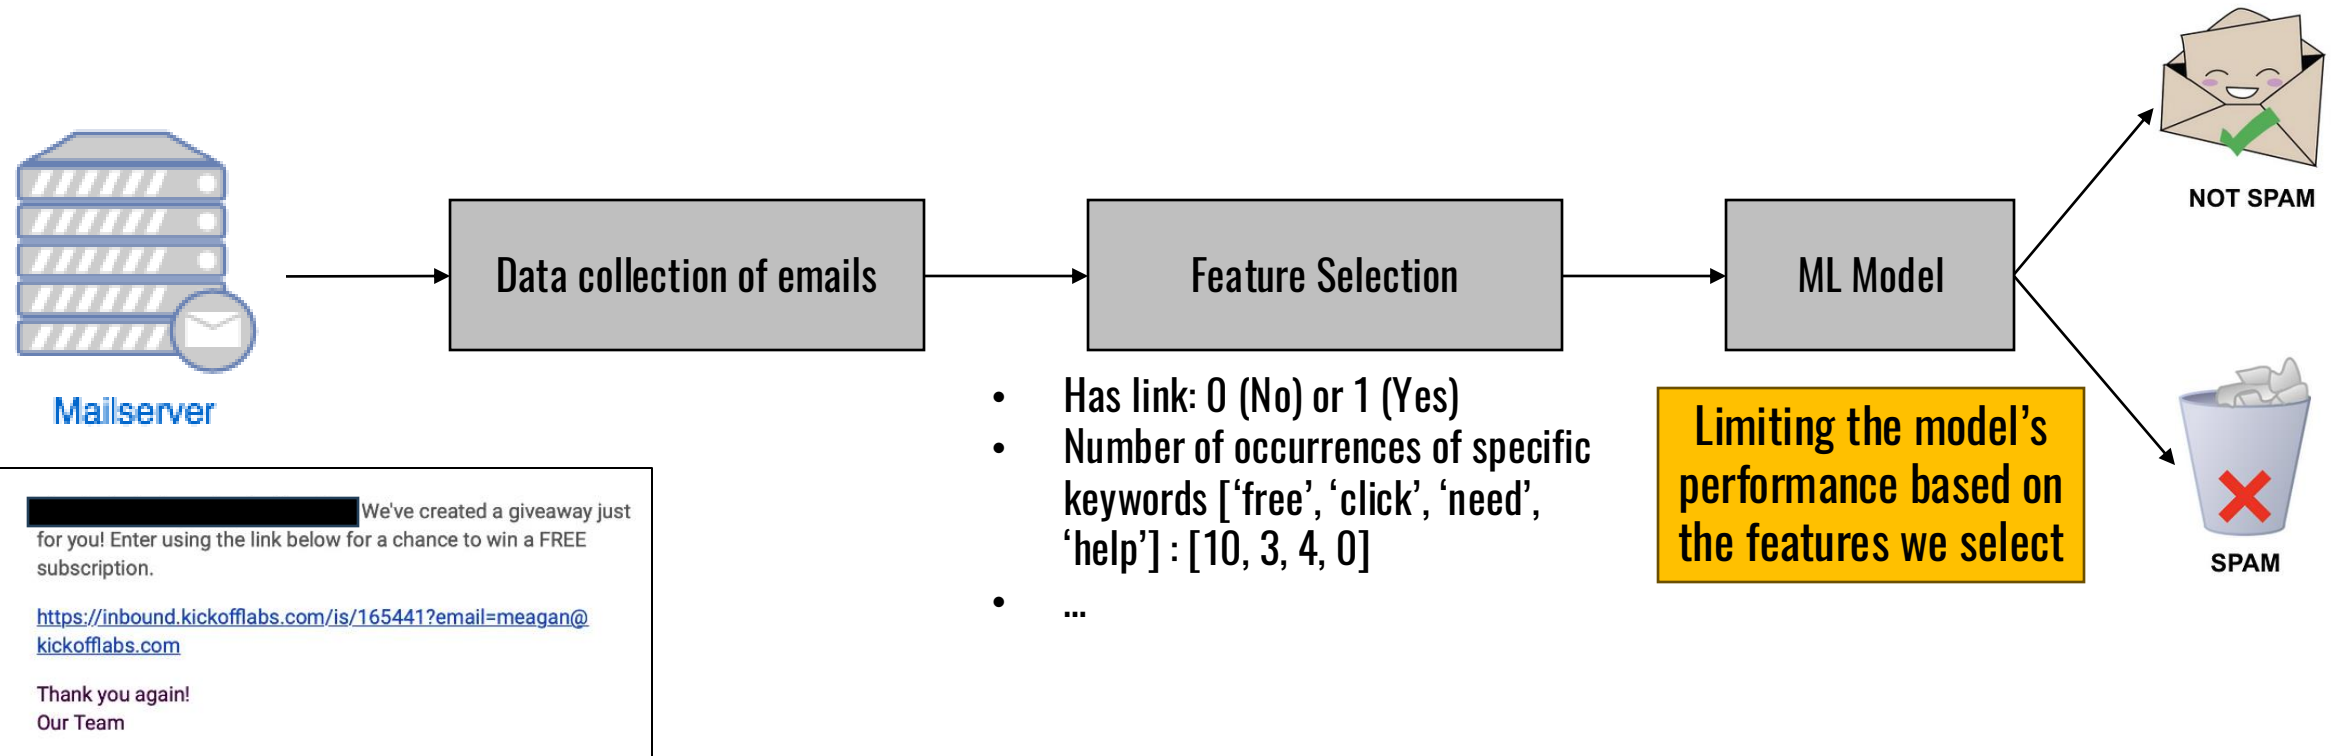


Machine Learning



# Machine Learning (ML)

**ML:** a subfield of AI that focuses on the development of algorithms and models capable of learning from data and making predictions without being explicitly programmed.



# Deep Learning (DL)

**DL:** a subfield of ML that focuses on the development and training of deep neural networks. Deep Learning algorithms leverage these deep architectures to automatically learn hierarchical representations of data, enabling them to capture complex patterns and relationships.

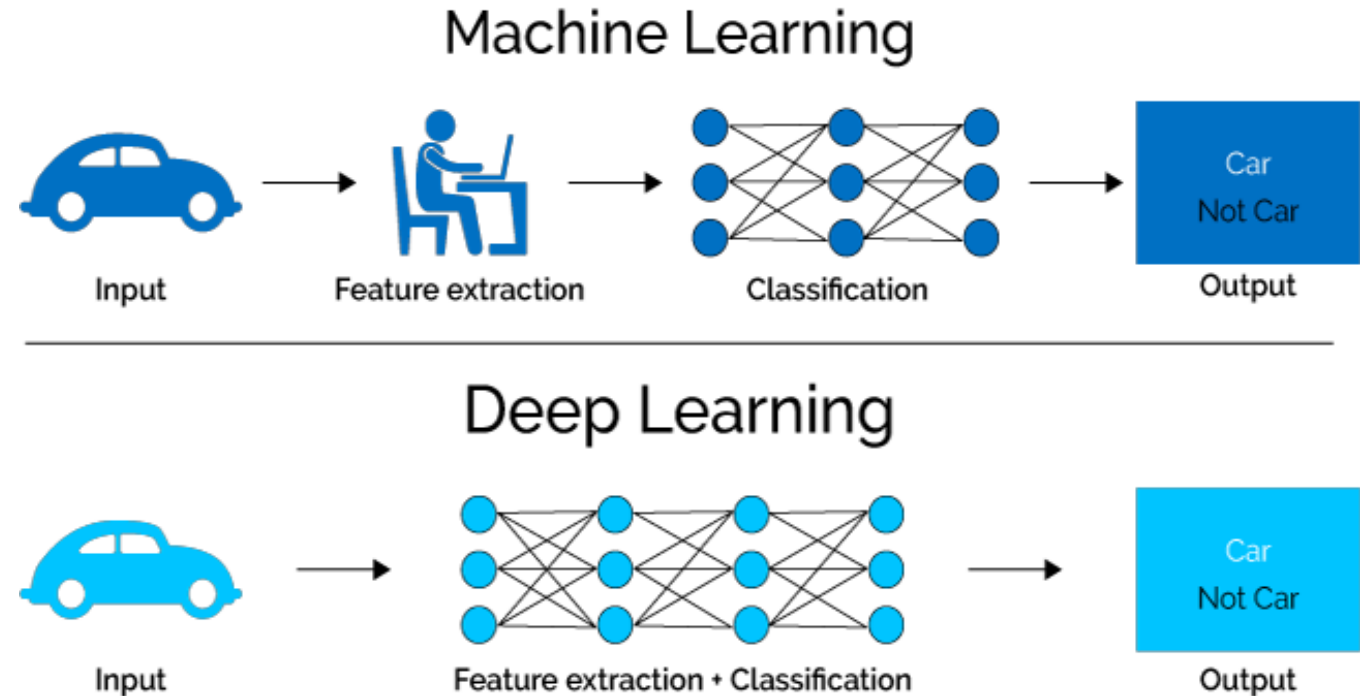



**Increase in computational power (advancements in hardware -> GPU)**





# Deep Learning (DL)

**DL:** a subfield of ML that focuses on the development and training of deep neural networks. Deep Learning algorithms leverage these deep architectures to automatically learn hierarchical representations of data, enabling them to capture complex patterns and relationships.




# Deep Learning (DL)



# Deep Learning (DL)

Deep neural networks learn hierarchical feature representations





45





# Artificial Neural Networks are inspired in Biological Neural Networks

Biological Neuron



Artificial Neuron





# ... but the mechanism is different




# Biological vs Artificial Neuron

Biological Neural Networks (BNNs)



Artificial Neural Networks (ANNs)




- **ANNs** have pre-defined and **fixed architectures** that can only **change the strength of connections** between neurons (weights). **BNNs** can do the former, but also **create new and reorganize existing connections** between neurons (neuroplasticity) and even create new neurons (neurogenesis).

# Biological vs Artificial Neuron

Biological Neural Networks (BNNs)



Artificial Neural Networks (ANNs)

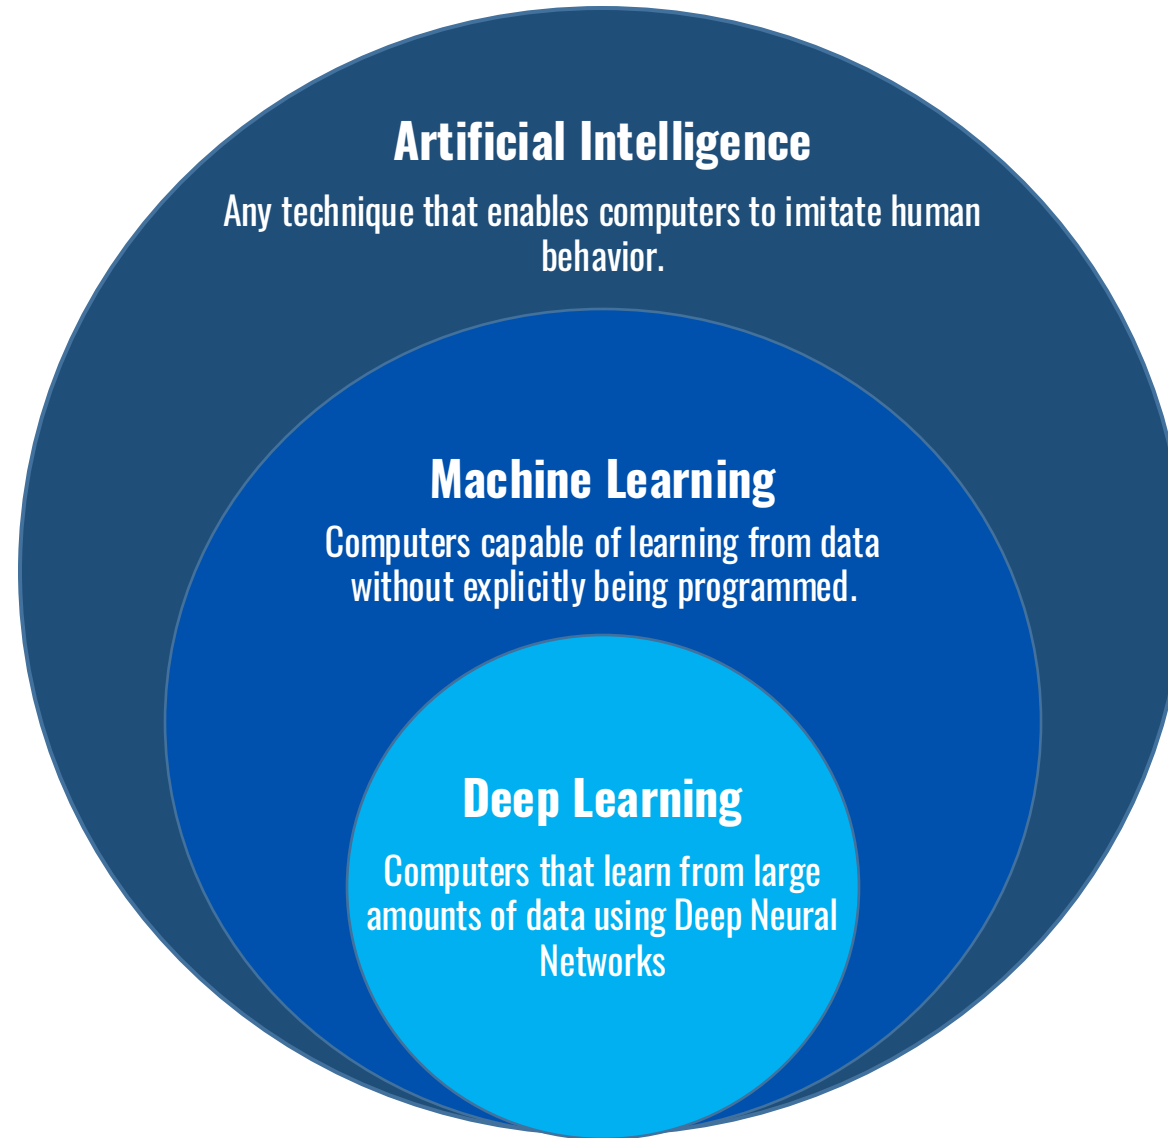


- **ANNs** use **continuous values** with **continuous activation functions** to propagate information throughout the network, whereas **BNNs** function as **firing systems that can be activated or not** (allowing or blocking information to pass).<sup>1</sup>
- In **BNNs**, neurons that **fire together, wire together**. The same **does not hold** for **ANNs**.<sup>2</sup>

1. There are computational models and algorithms inspired by the discrete firing behavior of biological neurons, such as spiking neural networks.

2. Hebbian learning

# Overview



# Lecture 1.

# Modern AI & State of AI

---

Budapest, 14<sup>th</sup> February 2025


**1** Course Details

**2** Deep Learning Overview

**3** Modern AI & State of AI



# Advancements in AI (Deep Learning methods)

## Classification



# Advancements in AI (Deep Learning methods)



## Classification






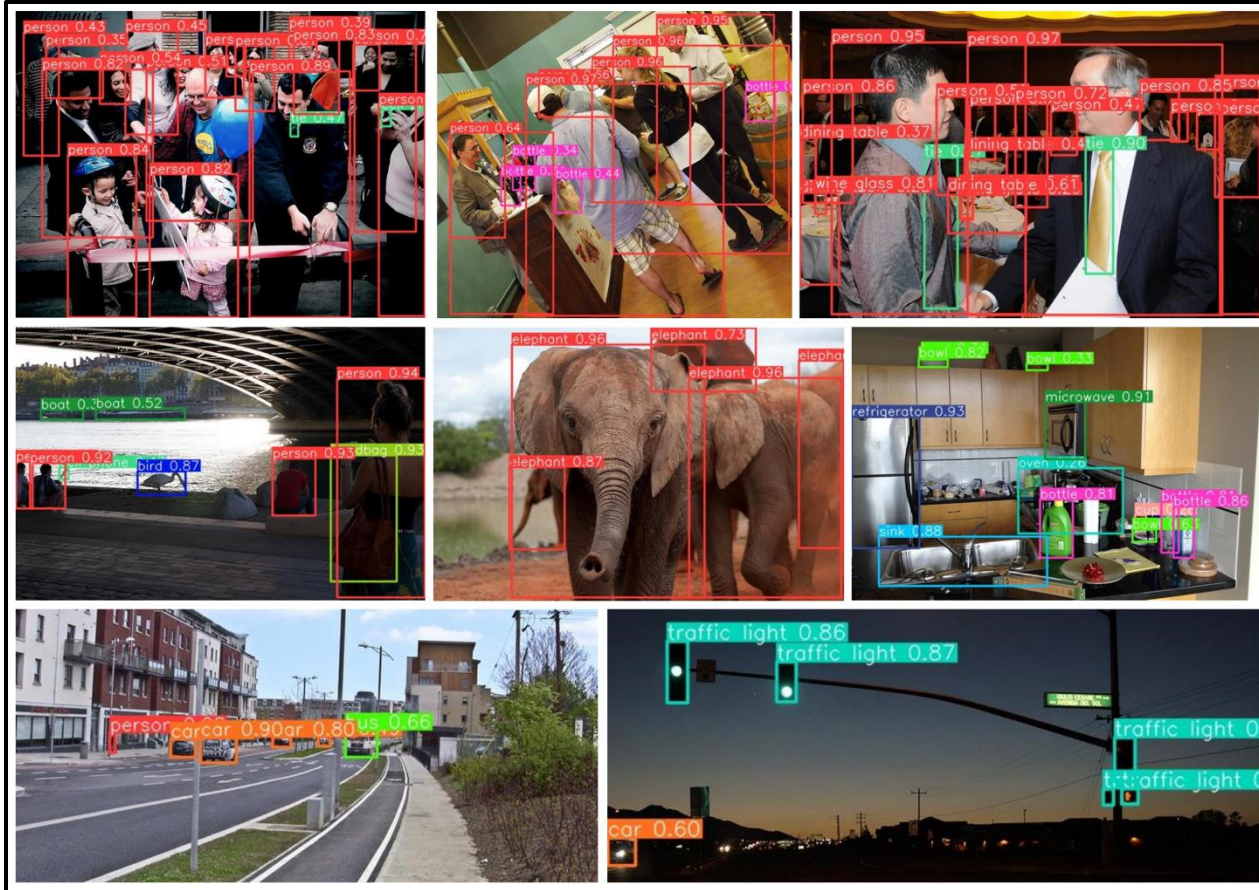
# Advancements in AI (Deep Learning methods)

## Object Detection



# Advancements in AI (Deep Learning methods)

## Object Detection [1]




[1] Wang, A., Chen, H., Liu, L., Chen, K., Lin, Z., Han, J., & Ding, G. (2024). YOLOv10: Real-Time End-to-End Object Detection. arXiv [Cs.CV]. Retrieved from <http://arxiv.org/abs/2405.14458>




# Advancements in AI (Deep Learning methods)

## Image Segmentation




Semantic Segmentation



Instance Segmentation




Semantic Segmentation



Instance Segmentation

# Advancements in AI (Deep Learning methods)

## Image Segmentation: Segment Anything 1 & 2 (SAM) [2]



- SAM segments any type of object (even if it was not explicitly trained on that object)
- SAM 2 is built on top of SAM with the addition of memory for tracking objects overtime.

[2] Ravi, N., Gabeur, V., Hu, Y.-T., Hu, R., Ryali, C., Ma, T., ... Feichtenhofer, C. (2024). SAM 2: Segment Anything in Images and Videos. arXiv [Cs.CV]. Retrieved from <http://arxiv.org/abs/2408.00714>



# Advancements in AI (Deep Learning methods)

## Text-to-Image



**Stable Diffusion 1.2 (2022)**



**FLUX.1-schnell (2024)**




**FLUX.1-dev-LoRa (2024)**

## Advancements in AI (Deep Learning methods)

### Text-to-Image: DALL-E 2


<https://openai.com/dall-e-2/>



[3] Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., & Chen, M. (2022). Hierarchical Text-Conditional Image Generation with CLIP Latents. arXiv [Cs.CV]. Retrieved from <http://arxiv.org/abs/2204.06125>

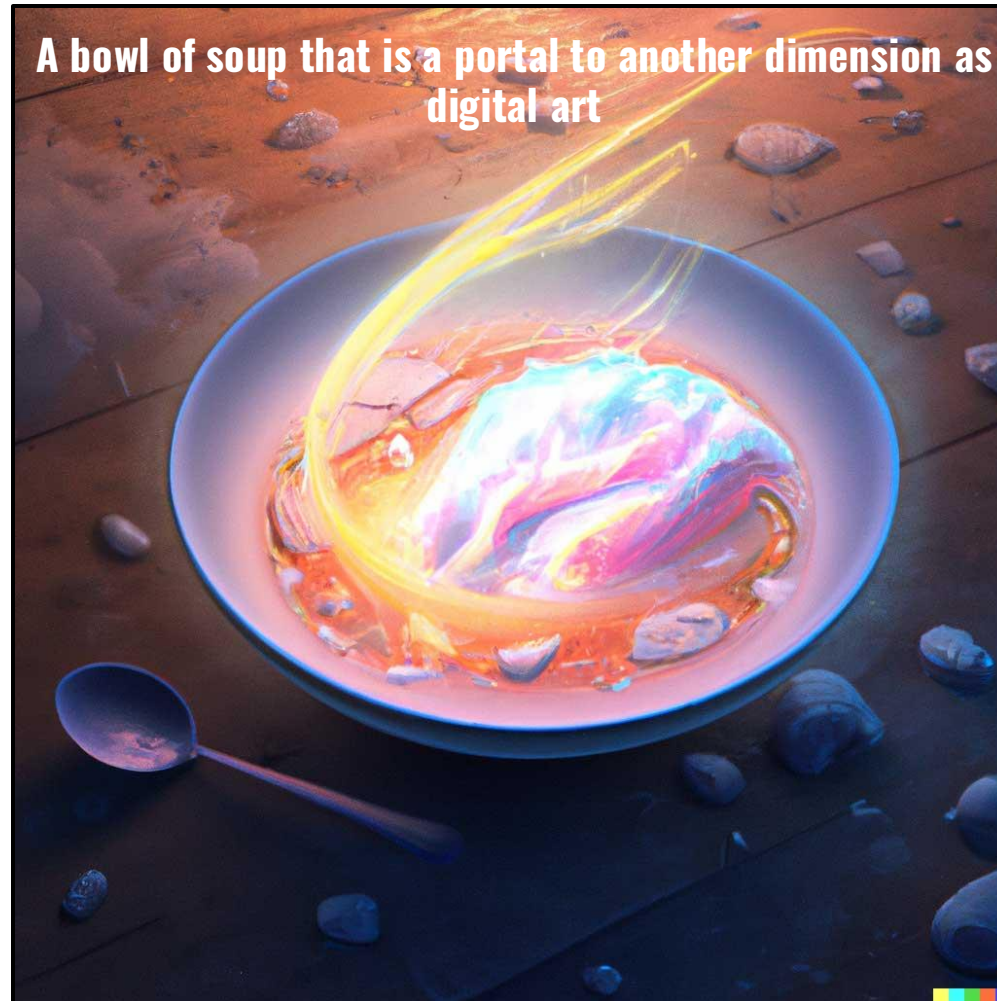
# Advancements in AI (Deep Learning methods)

**An astronaut riding a horse in photo realistic style**





# Advancements in AI (Deep Learning methods)



# Advancements in AI (Deep Learning methods)

## Text-to-Video: SORA

<https://openai.com/index/sora/>



# Modern AI categories


## What is Artificial Intelligence?

Traditional version – the science and engineering of making machines that simulate human intelligence and improve themselves.

# Modern AI categories

## What is Artificial Intelligence?

The goal of traditional AI is to pass the Turing test.



- What happens if a computer (algorithm) passes the **Turing test**?
- Is it intelligent?
- Is it considered **Artificial General Intelligence (AGI)**?

## Modern AI categories

“In this paper, we argue that it is critical for the AI research community to explicitly reflect on what we mean by “AGI,” and aspire to quantify attributes like the performance, generality, and autonomy of AI systems” [5]

About Turing Test: “in practice, the test often highlights the ease of fooling people (Weizenbaum, 1966; Wikipedia, 2023a) rather than the “intelligence” of the machine. Given that modern LLMs pass some framings of the Turing Test, it seems clear that this criteria is insufficient for operationalizing or benchmarking AGI”



originally published Nov. 2023; updated Jan. 2024

### Levels of AGI: Operationalizing Progress on the Path to AGI

Meredith Ringel Morris<sup>1</sup>, Jascha Sohl-dickstein<sup>1</sup>, Noah Fiedel<sup>1</sup>, Tris Warkentin<sup>1</sup>, Allan Dafoe<sup>1</sup>, Aleksandra Faust<sup>1</sup>, Clement Farabet<sup>1</sup> and Shane Legg<sup>1</sup>

<sup>1</sup>Google DeepMind

We propose a framework for classifying the capabilities and behavior of Artificial General Intelligence (AGI) models and their precursors. This framework introduces levels of AGI performance, generality, and autonomy. It is our hope that this framework will be useful in an analogous way to the levels of autonomous driving, by providing a common language to compare models, assess risks, and measure progress along the path to AGI. To develop our framework, we analyze existing definitions of AGI, and distill six principles that a useful ontology for AGI should satisfy. These principles include focusing on capabilities rather than mechanisms; separately evaluating generality and performance; and defining stages along the path toward AGI, rather than focusing on the endpoint. With these principles in mind, we propose “Levels of AGI” based on depth (performance) and breadth (generality) of capabilities, and reflect on how current systems fit into this ontology. We discuss the challenging requirements for future benchmarks that quantify the behavior and capabilities of AGI models against these levels. Finally, we discuss how these levels of AGI interact with deployment considerations such as autonomy and risk, and emphasize the importance of carefully selecting Human-AI Interaction paradigms for responsible and safe deployment of highly capable AI systems.

**Keywords:** AI, AGI, Artificial General Intelligence, General AI, Human-Level AI, HLAI, ASI, frontier models, benchmarking, metrics, AI safety, AI risk, autonomous systems, Human-AI Interaction

2v2 [cs.AI] 5 Jan 2024

[5] Morris, M. R., Sohl-dickstein, J., Fiedel, N., Warkentin, T., Dafoe, A., Faust, A., ... Legg, S. (2024). Levels of AGI for Operationalizing Progress on the Path to AGI. arXiv [Cs.AI]. Retrieved from <http://arxiv.org/abs/2311.02462>



# Modern AI categories

## What is Artificial Intelligence?

The goal of traditional AI is to pass the Turing test.



There are many claims that ChatGPT / GPT-4 has passed the Turing test (not necessarily true)

- **Is ChatGPT intelligent?**
- **Has ChatGPT achieved AGI?**



## Modern AI categories

~~AGI remains a distant goal~~, so let us now focus on the current state of AI

AI has changed and is being applied in new ways.

**Futuristic**

**Skeptical**


**Political**

**Enterprise**



## Futuristic AI


A technology that extends human capabilities and aids people in solving real world problems.




**Brain Tumor detection**



**Autonomous driving**



**Prosthetics**




**Brain Computer Interface**

**... makes our  
lives better**


**... enhances  
humans**

# Futuristic AI

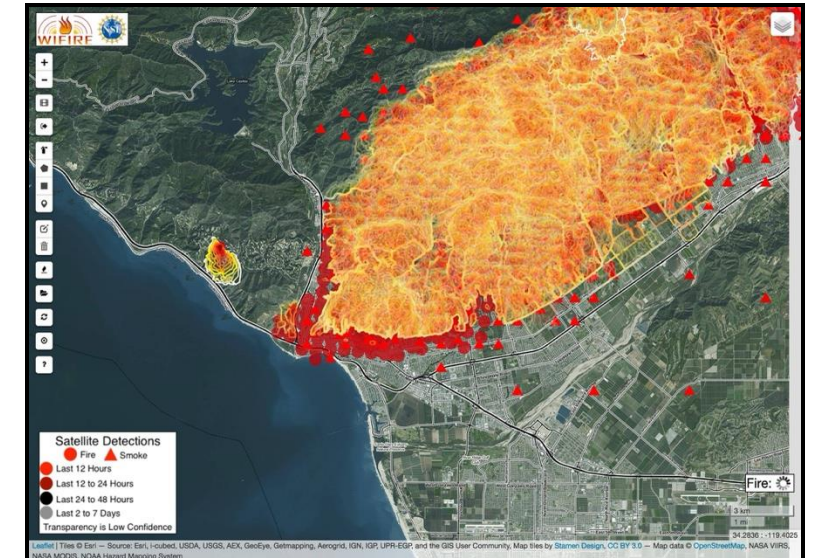
A technology that extends human capabilities and aids people in solving real world problems.



Monitoring wildlife



Monitoring ocean





Monitoring forests

# ... saves the world

## Futuristic AI (Examples)

Brain tumour detection

<https://huggingface.co/spaces/sindhoorar/brain-tumor-classifier>






# Futuristic AI (Examples)

The most advanced Medical AI system: **Med-PaLM 2** (Google)

<https://sites.research.google/med-palm/>

Can you write me a report analyzing this chest X-ray?




**Findings:**

- Devices: None.
- Lungs: No pneumothorax. No substantial pleural effusion. Lungs appear clear.
- Cardiomeastinal: Normal heart size. Mediastinal contours within normal limits.
- Other: No acute skeletal abnormality.

**Impression:**  
No active disease seen in chest.

Enter a question here

For multiple modalities:  
**Med-PaLM M** (Google) [6]



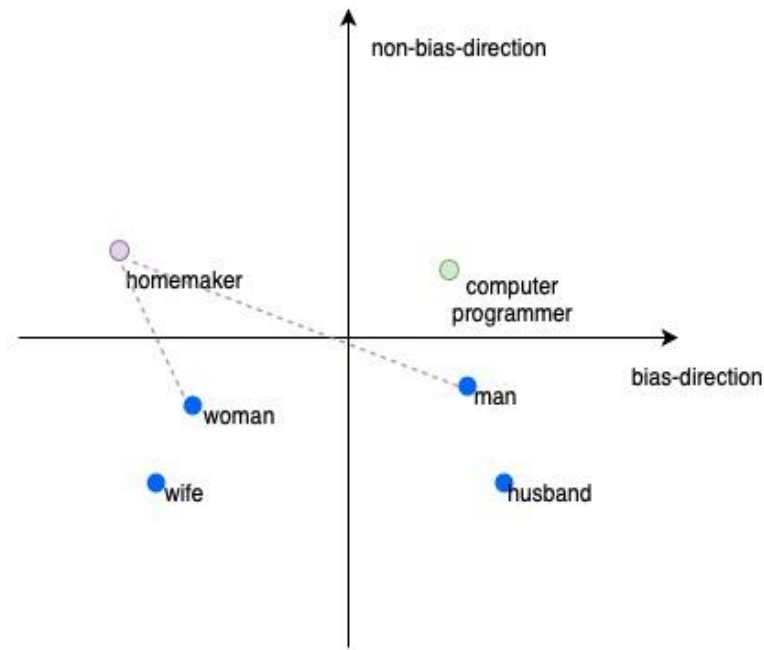
[6] Tu, T., Azizi, S., Driess, D., Schaekermann, M., Amin, M., Chang, P.-C., ... Natarajan, V. (2023). Towards Generalist Biomedical AI. arXiv [Cs.CL]. Retrieved from <http://arxiv.org/abs/2307.14334>

## Skeptical AI

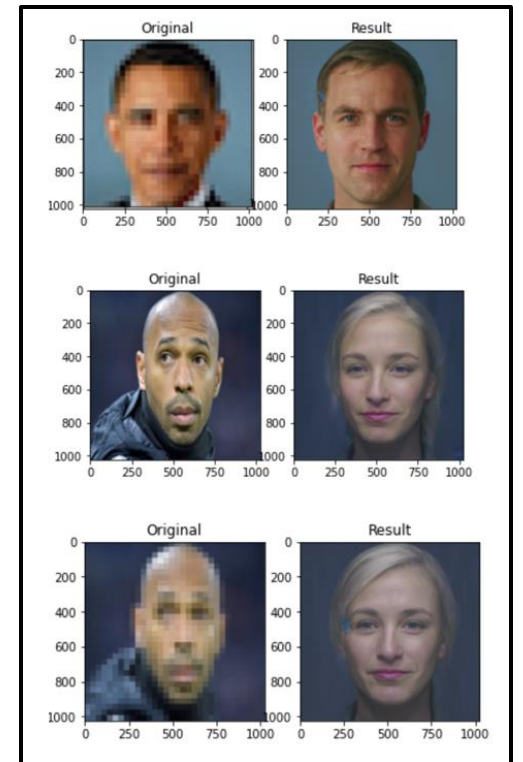
Any AI system that is biased, unsafe, insecure and is used as a resource for inequality and war.



...  
**manipulates  
information**



... **biased** ...





## Skeptical AI (Examples)

**Diffusion Bias Explorer** (<https://huggingface.co/spaces/society-ethics/DiffusionBiasExplorer>)

A method that shows the bias in text-to-image methods


Ambitious CEO  
Ambitious Receptionist  
Ambitious Lawyer  
Sensitive/Emotional Lawyer  
Scientist  
Teacher  
...  
Why?



## Skeptical AI (Examples)

The previous example was simple, but it can get messy:


<https://futurism.com/the-byte/lawyer-chatgpt-court>



## Skeptical AI (Examples)

### DeepFakes:


- Misinformation - <https://www.youtube.com/watch?v=cQ54GDm1eL0&t=13s>
- Not always skeptical: Entertainment/Tribute - <https://www.youtube.com/watch?v=uAPUkgeiFVY>



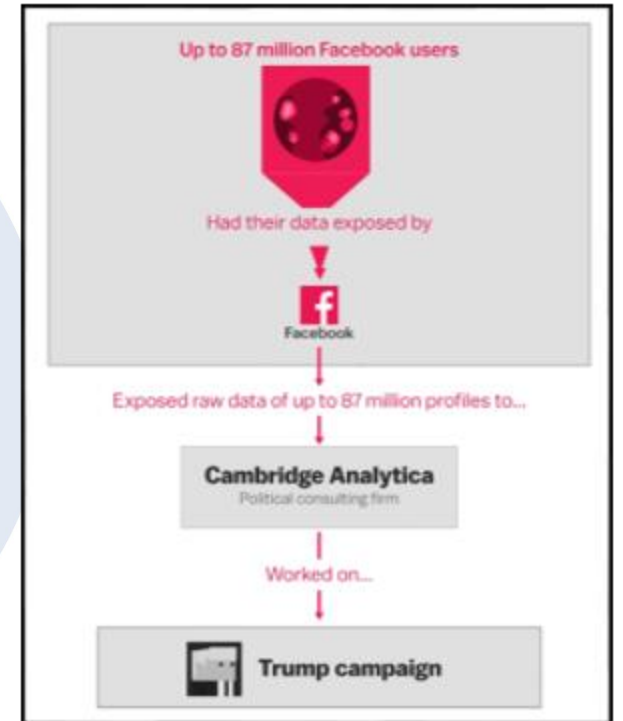


## Political AI

A tool for controlling people's decision making towards the government's benefit.



**... take advantage of our disadvantages**

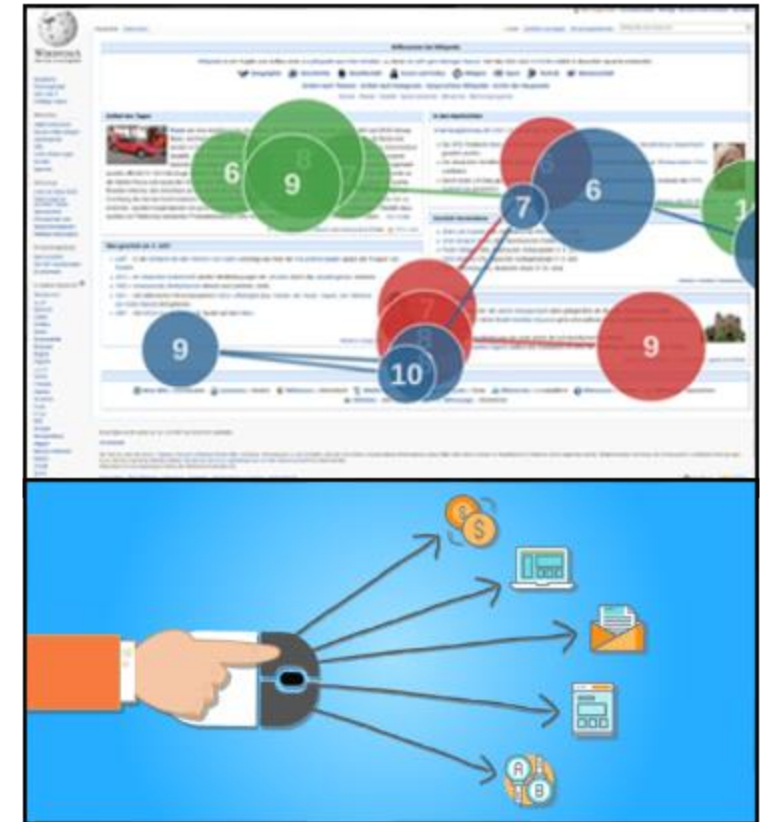


# Enterprise AI

A data-driven tool to maximize profit.



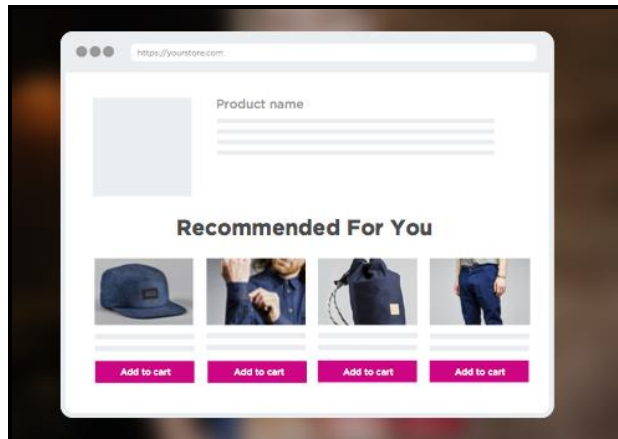
**... data is  
everything  
(no such thing  
as free service)**



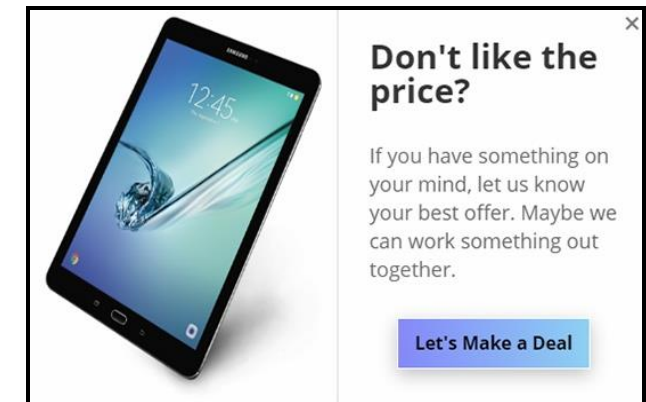
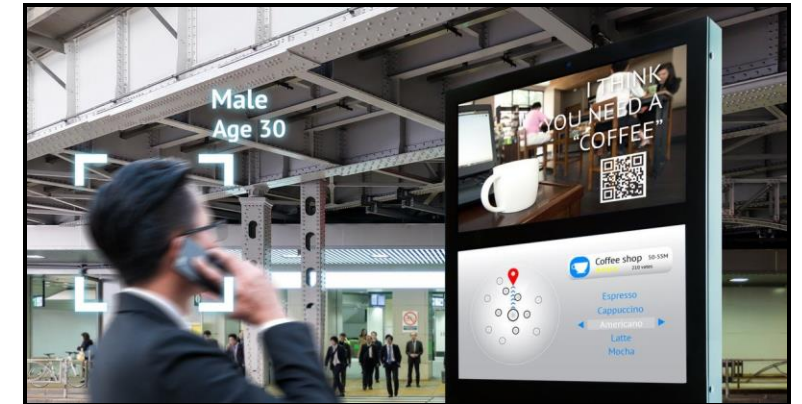


# Enterprise AI

A data-driven tool to maximize profit.



**The more they  
know about  
you, the more  
they profit...**



## Enterprise AI (Examples)

### Data needed for training an AI model

- Can sell trained model as a service


### Data produced by trained model

- Can use produced data to improve model
- Can sell produced data



## Enterprise AI (Examples)

A data-driven tool to maximize profit.



## Job automation

## Enterprise AI (Examples)

Cognition labs **Devin**



## Job automation



# Summary

**FUTURISTIC AI:** The goal is to extend human life.



**POLITICAL AI:** The goal is to control humans.



**SKEPTICAL AI:** The goal is to have objectives that are misaligned with human objectives.



**ENTERPRISE AI:** The goal is to maximize profit.



# Ethical Considerations

- This course also emphasizes ethical considerations in AI development, ensuring that students not only learn the technical aspects of **Deep Learning** but also understand its broader impact on society.



# More Speech and Fewer Mistakes (07.01.2025)



# Highlights

- Replace fact checkers with **Community Notes**
  - They are too politically biased
  - Destroyed more trust than they have created
- Filters that scan policy violations
  - Made a lot of mistakes and took down content that is shouldn't have
- Europe has an ever-increasing number of laws institutionalizing censorship and making it difficult to build anything innovative




# EU AI Act

“The AI Act ... is the first-ever comprehensive legal framework on AI worldwide. The aim of the rules is to foster trustworthy AI in Europe”

## Why do we need rules on AI?

- “While most AI systems pose limited to no risk [...] certain AI systems create risks that we must address”
- “For example, it is often not possible to find out why an AI system has made a decision or prediction and taken a particular action.”  
...




[7] <https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai>

# EU AI Act

## Unacceptable risk

1. **harmful AI-based manipulation and deception**
2. harmful AI-based exploitation of vulnerabilities
3. social scoring
4. **Individual criminal offence risk assessment or prediction**
5. untargeted scraping of the internet or CCTV material to create or expand facial recognition databases
6. **emotion recognition in workplaces and education institutions**
7. biometric categorisation to deduce certain protected characteristics
8. real-time remote biometric identification for law enforcement purposes in publicly accessible spaces




[7] <https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai>

# EU AI Act

## High risk

- AI safety components in critical infrastructures (e.g. **transport**), the failure of which could put the life and health of citizens at risk
- AI solutions used in education institutions, that may determine the access to education and course of someone's professional life (e.g. **scoring of exams**)
- AI tools for employment, management of workers and access to self-employment (e.g. **CV-sorting software for recruitment**)
- Certain AI use-cases utilised to give access to essential private and public services (e.g. **credit scoring denying citizens opportunity to obtain a loan**)
- And many more ...



[7] <https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai>

## Summary

# What version will you create?



## The future is in your hands ...



---

# Resources

## Books:

- Courville, Goodfellow, Bengio: Deep Learning  
Freely available: <https://www.deeplearningbook.org/>
- Zhang, Aston and Lipton, Zachary C. and Li, Mu and Smola, Alexander J.: Dive into Deep Learning  
Freely available: <https://d2l.ai/>

## Courses:

- Deep Learning specialization by Andrew NG
- <https://www.coursera.org/specializations/deep-learning>

# That's all for today!

