Universität Konstanz

01 | Introduction to the Course

Giordano De Marzo

https://giordano-demarzo.github.io/

Deep Learning for Social Sciences

Outline

1. Introduction
2. Info about the Course
3. Deep Learning
4. Deep Learning and Social Sciences

Introduction

Me in Short

I'm a Physicist by training

- Master's Degree in Theoretical Physics at Sapienza University (Roma)
- PhD in Physics at Sapienza University, CREF and SSAS

After finishing my PhD I became a PostDoc and Lecturer in David Garcia lab here in Konstanz University. I teach

- Deep Learning (Summer Semester)
- Network Science (Winter Semester)

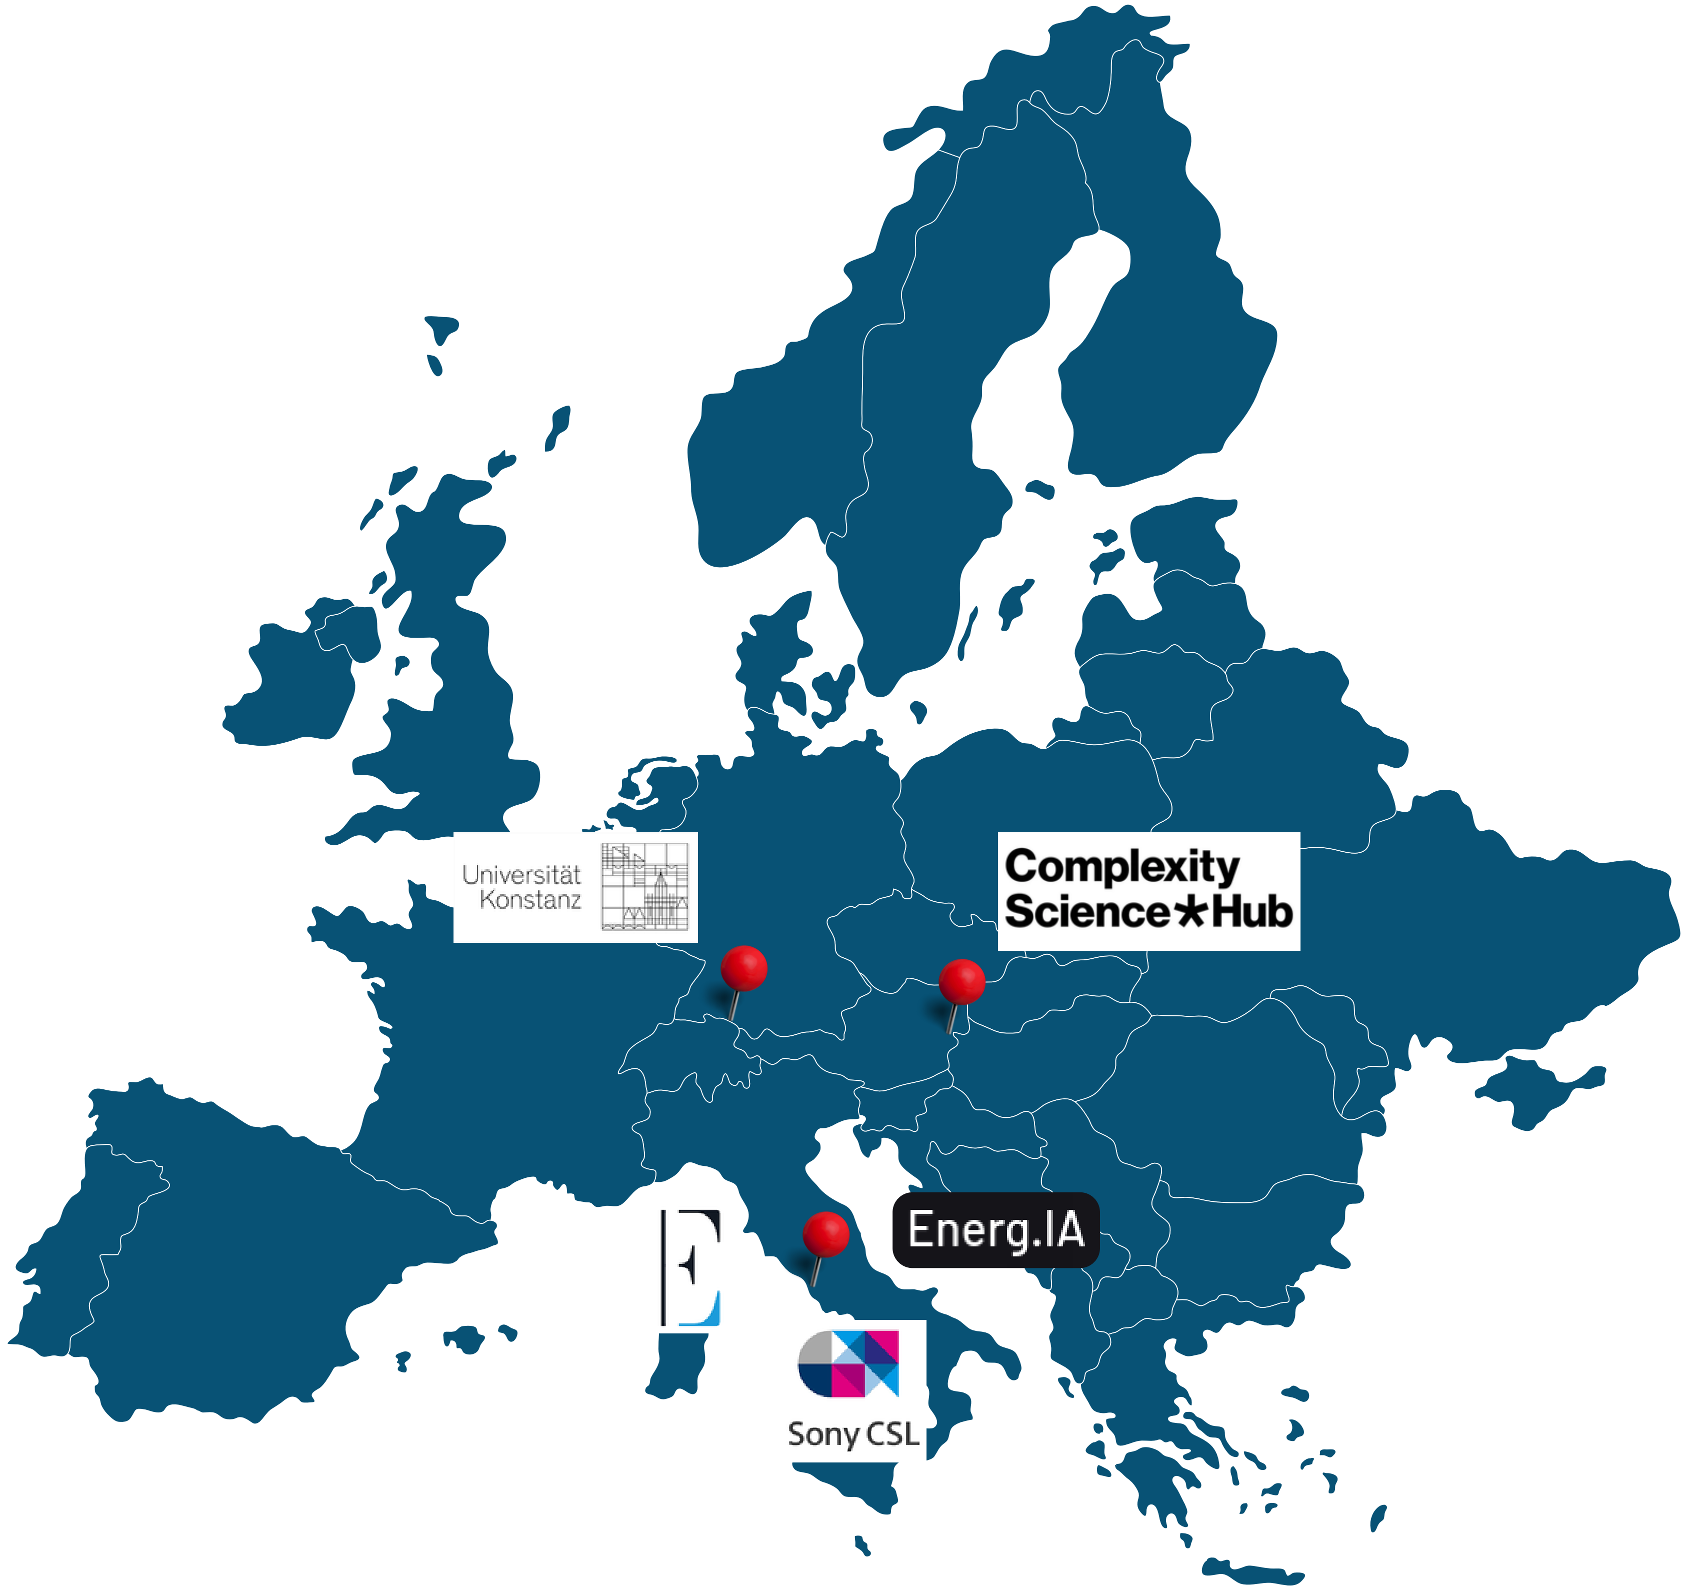

Universität Konstanz

Research Interests

In my research I apply techniques from Statistical Physics, Network Science and Machine Learning to study

- social systems
- economic systems
- artificial intelligence

Here in Konstanz I mostly focus on

- Applications of Large Language Models
- Emergent behavior in AI societies
- Simulating human systems using AI Agents

We use LLMs to forecast technological innovation and understand patents and technologies

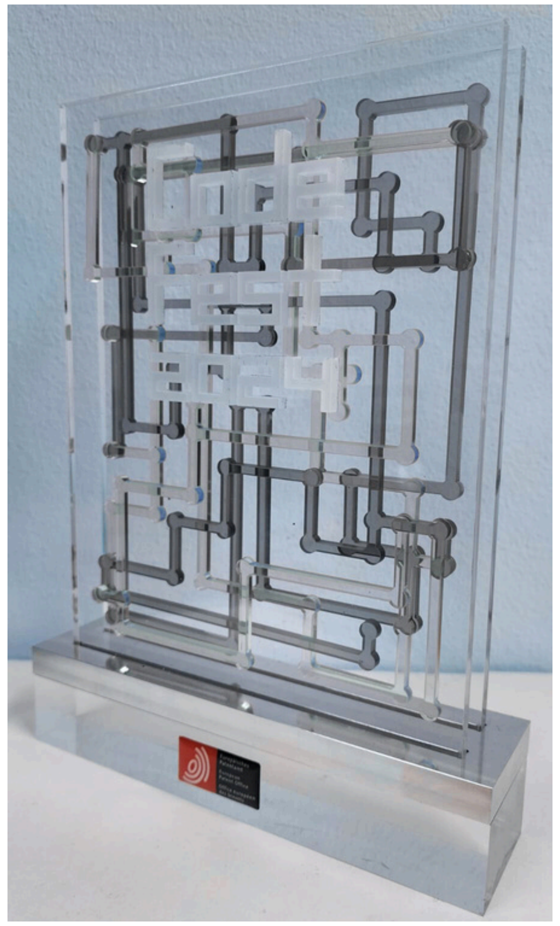

We study how LLMs can coordinate and form socially stable groups

Contacts

You can contact me at

- giordano.de-marzo@uni-konstanz.de

You can also come look for me in my office

- Building **D** Room **339**

If you're interested in my research you can find my latest publications and projects on my website and on google scholar

- https://giordano-demarzo.github.io/
- https://scholar.google.com/citations?user=noLP1ukAAAAJ&hl=it

Who are you?

Some questions for you

- what's your name?
- what's your background?
- do you already know something about machine learning or deep learning?
- why do you want to study deep learning?

Info about the Course

Course Program

This is an introductory course to Deep Learning

- we will only shortly cover the main machine learning concepts
- the core will be on artificial neural networks
- we will cover most modern topics, but we won’t dig into the more technical details
- the course is oriented to practical applications
- there will be not a lot of equations
- focus will be on coding

The course will be divided in

- theory lessons (Wednesday)
- coding sessions (Thursday)

Coding sessions are not mandatory but **strongly recommended**

Logistic

There will be

- 12 theory seminars
- 1 or 2 guest researchers seminars
- 1 students presentation day
- 8 coding sessions

The course will require you to code in python

- we will mostly use google colab
- for the final projects you will have access to more powerful 40GB GPUs
- you need a github account to access these GPUs
- more details during the coding labs

The course grade is composed of the grade of **three assignments** delivered during the semester **(40%)** and a **final project (60%)**

Calendar

Date	Topic	Date	Topic
April 15	Introduction to Machine Learning	April 16	Lab: Colab + Machine Learning Basics
April 22	The Multilayer Perceptron	April 23	
April 29	Training Neural Networks	April 30	Lab: Building a MLP with PyTorch
May 6	Convolutional Neural Networks	May 7	Lab: Building a CNN to classify clothes
May 13	Graph Neural Networks		
May 20	Recurrent Neural Networks	May 21	Lab: Graphs and Time Series

Calendar

Date	Topic	Date	Topic
May 27	Attention and Transformer	May 28	Guest researcher seminar
June 10	Large Language Models	June 11	Lab: Embeddings + Guest researcher seminar
June 17	Fine-tuning LLMs	June 18	Lab: Fine-tuning LLMs
June 24		June 25	
July 01	Image Generation and Multimodality	July 02	Lab: VAE and Diffusion Models
July 08	Reinforcement Learning and Alignment	July 09	Lab: Reinforcement Learning
July 15	Students Presentations	July 16	

Assignments

The three assignment will consists in coding tasks and writing a short report

- you will have 3 weeks for working on the assignment
- you will upload the code and report on github like for ICSS
- the assignments will be released on the following dates
    - April 30 to May 20 - Multilayer Percepetron
    - May 21 to June 17 - Convolutional Neural Network (4 weeks for this assignment)
    - June 18 to July 9 - Large Language Models

Reports should be structured in a scientific format and around 3 pages long

- Title and name
- Short introduction with info about the dataset
- Detailed explanation of the procedure followed
- Conclusions

Final Project

The final project will be similar to the assignments but require more time

- you will work in groups of 3 or 4 people
- I will share a form to ask who you want to work with and on which topic
- Projects will be disclosed at the beginning of June and you can work on them till August 31
- You will present the state of you project on July 15 to receive feedback
- If you need an early grading you can submit the project earlier
- Also in this case you will submit both the code and a report

You can see examples of last year projects and assignments on my website, where I will also regularly upload all the course material

https://giordano-demarzo.github.io/teaching/deep-learning-26/

You can also check last year website

https://giordano-demarzo.github.io/teaching/deep-learning-25/

Course Material

Most of the topics we will cover are contained in the following books

- Understanding Deep Learning by Simon J.D. Prince https://github.com/udlbook/udlbook/releases/download/v1.16/UnderstandingDeepLearning_24_11_23_C.pdf
- The Little Book of Deep Learning by François Fleuret https://fleuret.org/francois/lbdl.html
- Deep Learning Course, a thorough introduction to deep-learning, with examples in the PyTorch framework https://fleuret.org/dlc/
- Alice's Adventures in a differentiable wonderland by Simone Scardapane https://www.sscardapane.it/assets/alice/Alice_book_volume_1.pdf

I will point out during the lesson where you can find useful resources to study and upload relevant material on my website

Policy on ChatGPT and LLMs

LLMs are a powerful tool and it does make no sense to tell you not to use them. However

- if you blindly ask LLMs to code for you, you will probably manage to solve this course assignments, but you won't learn much
- try instead to first write some code and then use LLMs for debugging or improving it
- the better you get, the more you can use LLMs to replace boring coding parts
- in order to do so you need to have a clear plan of what you want to achieve and how

Similar concepts also apply to the reports

- don't generate an entire report from scratch using an LLM
- you can ask it to generate a structure for you, then write a draft of the report
- you can ask to improve your draft

Universität
Konstanz

AI, Machine Learning and Neural Networks

Artificial Intelligence: Systems that can perform tasks requiring human intelligence

- Problem-solving, understanding language, recognizing patterns

Machine Learning: AI systems that learn from data

- Supervised, unsupervised, reinforcement learning

Neural Networks: ML algorithms inspired by the human brain

- Interconnected nodes (neurons) organized in layers
- Learn by adjusting connection strengths (weights)

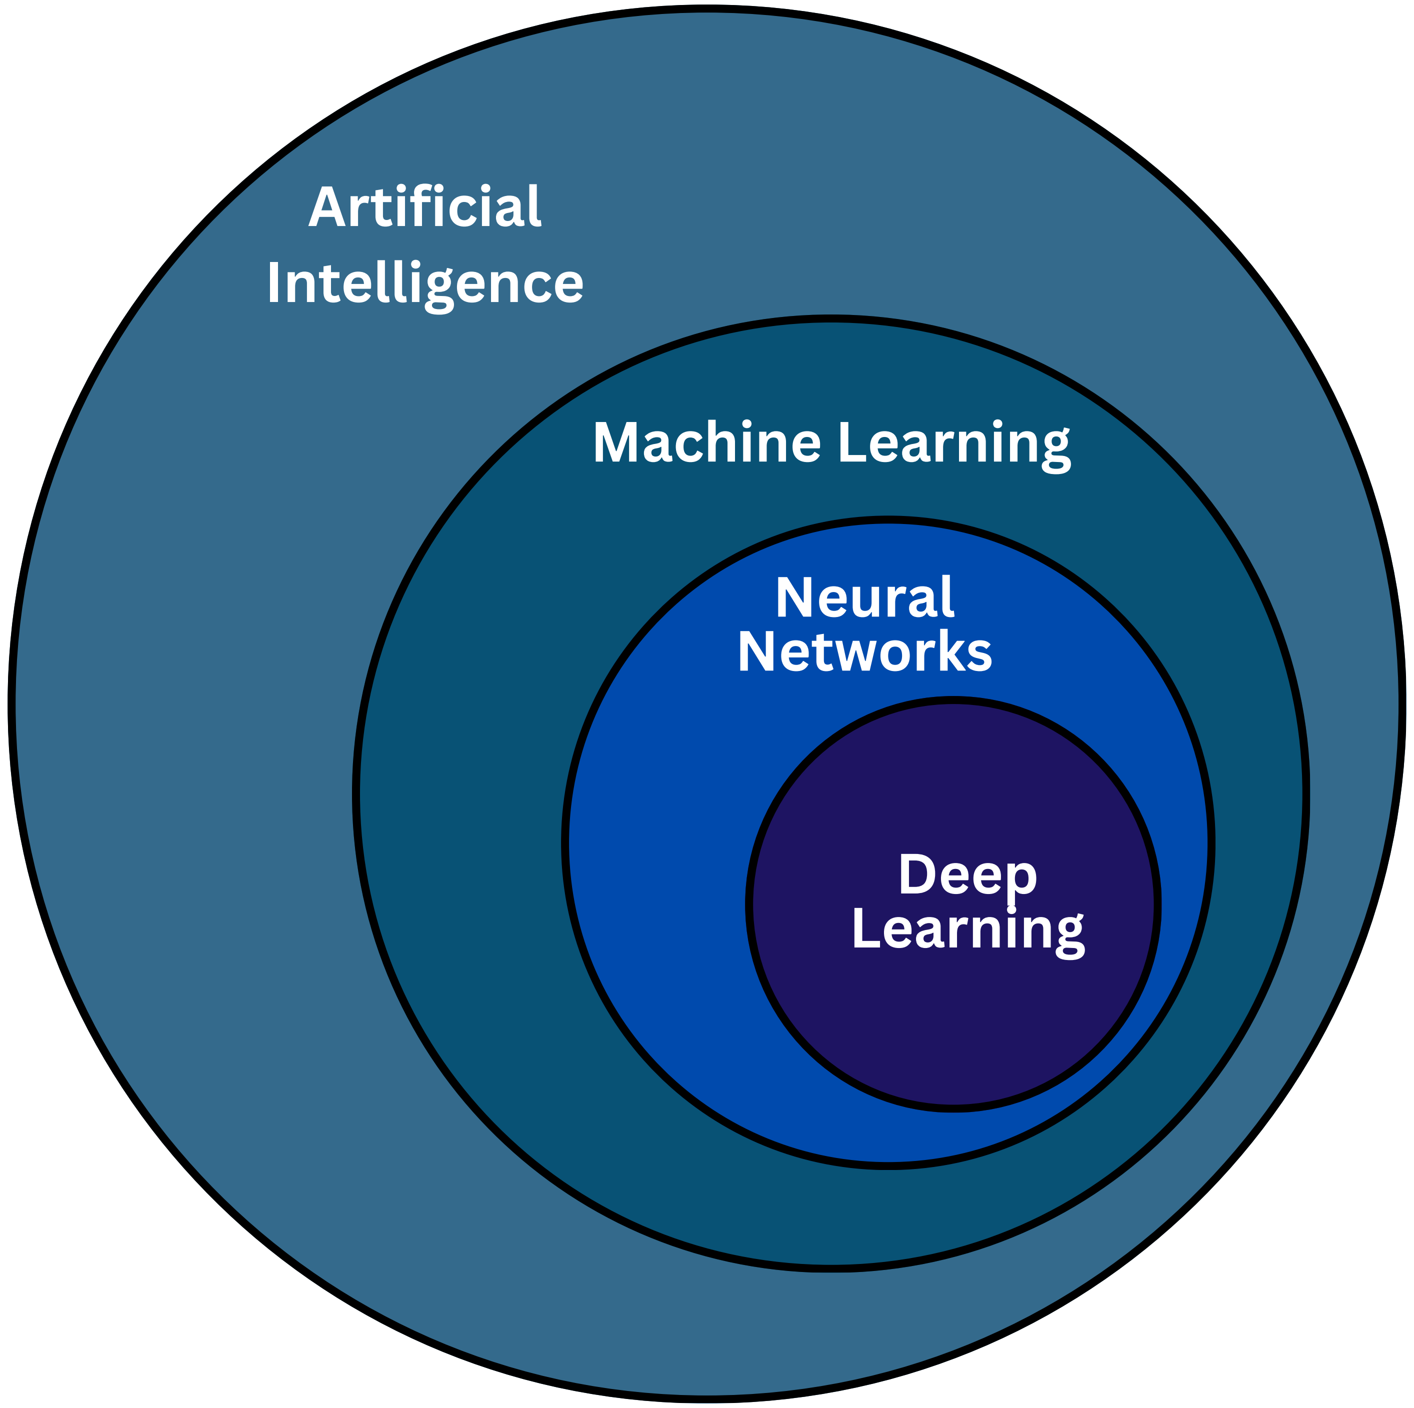

What is Deep Learning?

Definition: Machine learning using neural networks with multiple layers

- "Deep" refers to many hidden layers between input and output

Key characteristics: Automatic feature extraction (no manual feature engineering)

- Hierarchical representation learning
- End-to-end learning from raw data

Traditional ML: Manual feature extraction → Algorithm → Prediction

Deep Learning: Raw Data → Neural Network → Prediction

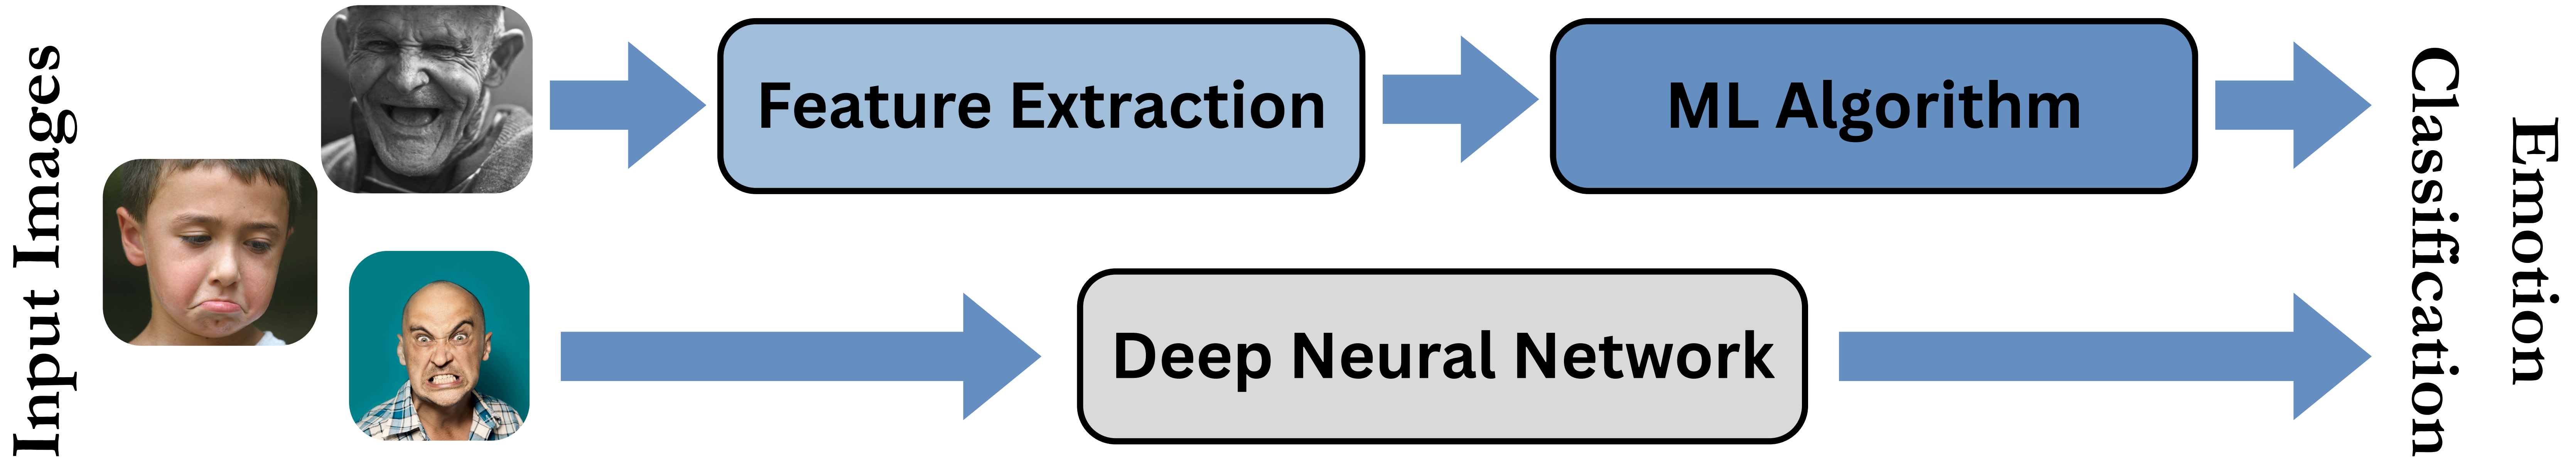

Landmarks in Deep Learning

- 1943: McCulloch & Pitts propose first mathematical model of a neuron
- 1958: Rosenblatt introduces the Perceptron
- 1980s: Backpropagation algorithm for training neural networks
- 1998: LeNet-5 for handwritten digit recognition
- 2012: AlexNet wins ImageNet competition
- 2014: GANs introduced by Goodfellow et al.
- 2017: Transformer architecture revolutionizes NLP
- 2020-present: Large language models
- 2024: Nobel prizes to AI researchers

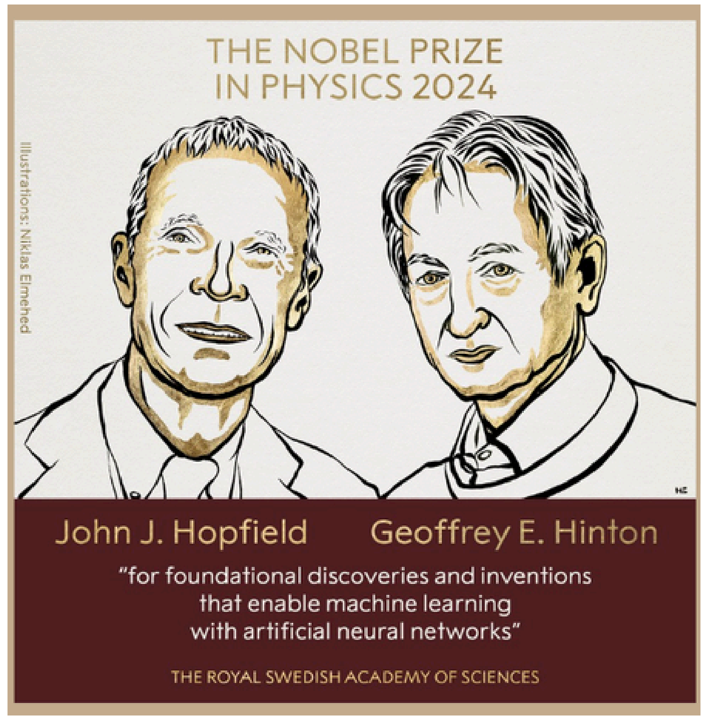

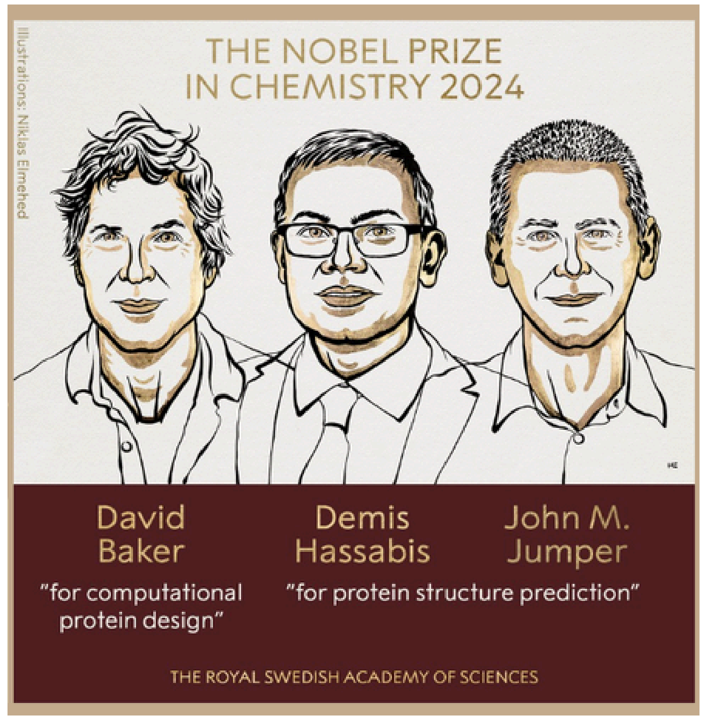

The Rise of Deep Learning

Deep learning has experienced a long journey to become one of the most important technologies of our time

- **From AI winter to renaissance:**
  - Deep Learning was mostly abandoned in between 1990 and 2010
  - Breakthroughs in training techniques (2006-2012)
- **Breakthrough applications in the last 10 years:**
  - Computer vision (object detection, image generation)
  - Natural language processing
  - Game playing (AlphaGo)
  - Protein folding

Google Trends "Deep Learning"

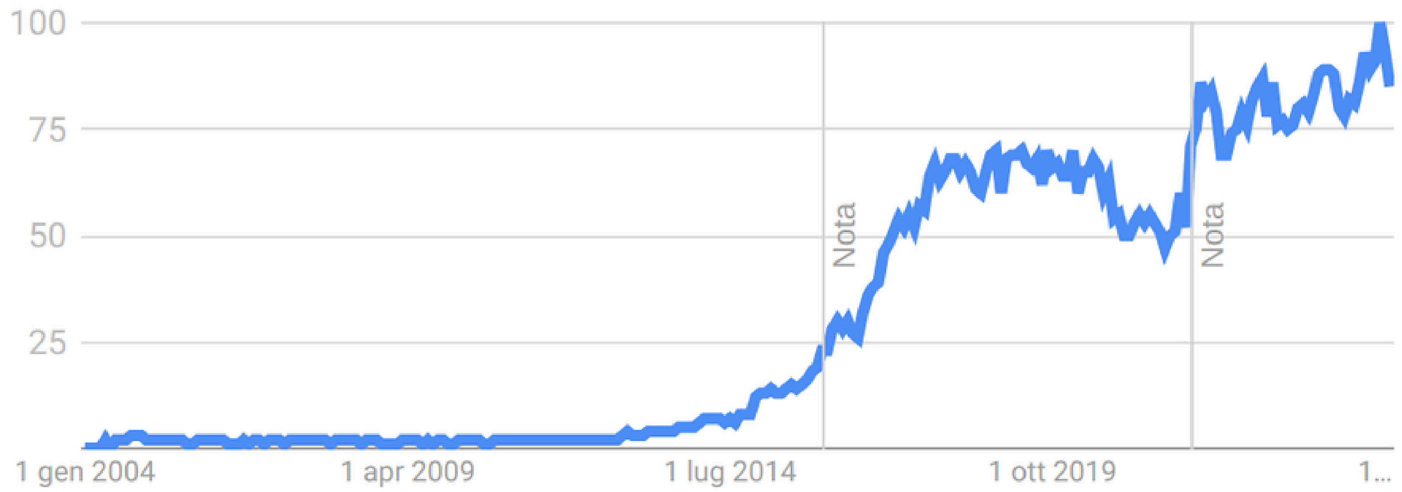

Google Trends "AI"

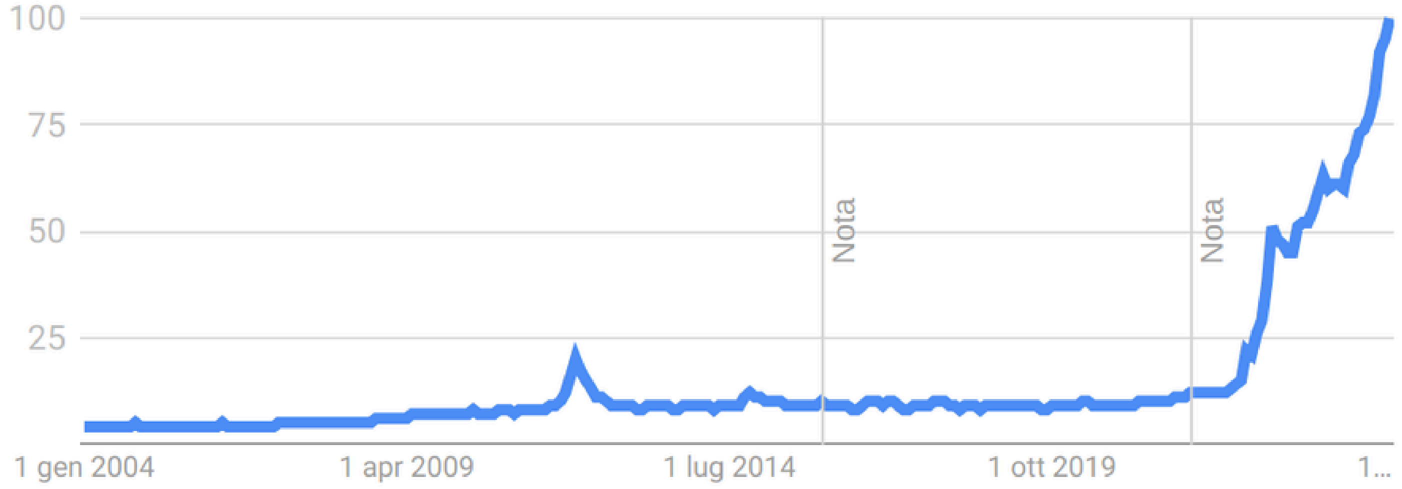

Computing and Big Data

Computational advances enabling deep learning:

- GPUs/TPUs: Massively parallel processing
- Cloud computing: Accessible computing
- Specialized hardware: NVIDIA, Google

Big data as fuel:

- Internet-scale data collection
- ImageNet, Common Crawl, Wikipedia
- User-generated content
- Sensors, IoT, and digital traces

Software frameworks:

- TensorFlow, PyTorch, JAX, Transformers
- Democratizing access to deep learning tools

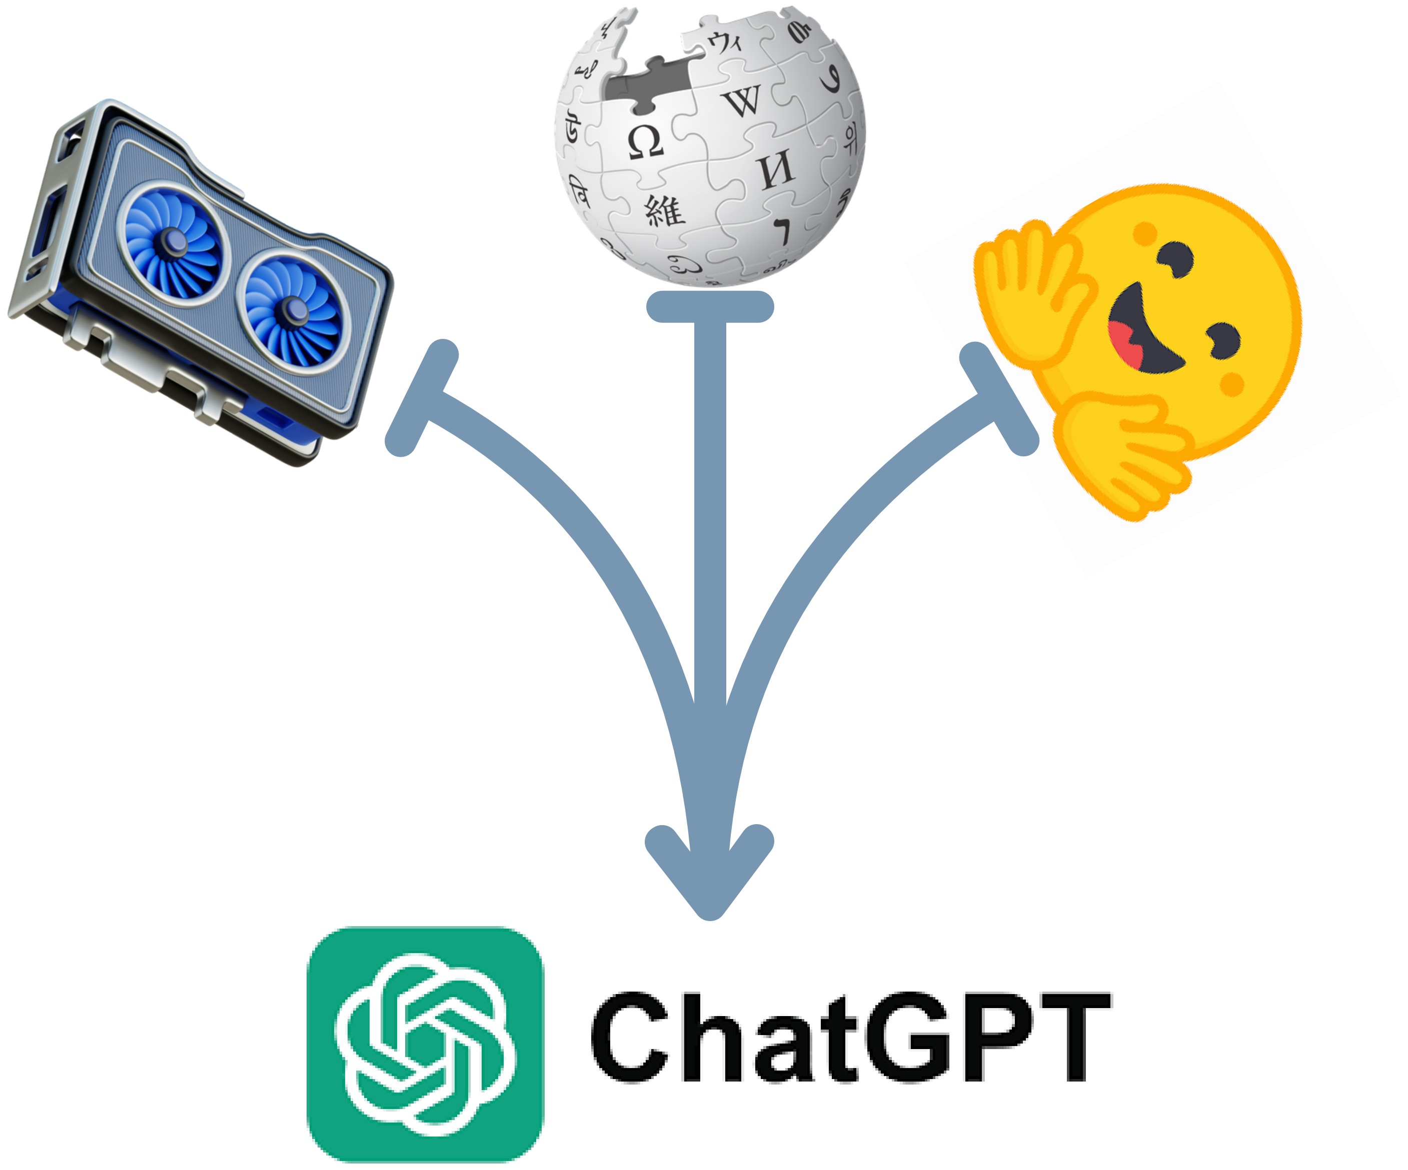

Deep Learning Scaling Laws

More data and bigger models leads to better performance, leading to the scaling laws of deep learning

- **Key scaling dimensions:** Model size (parameters), Dataset size, Computational budget
- **Emergent capabilities:**
  - Large models exhibit unexpected abilities not present in smaller versions
  - Reasoning capabilities

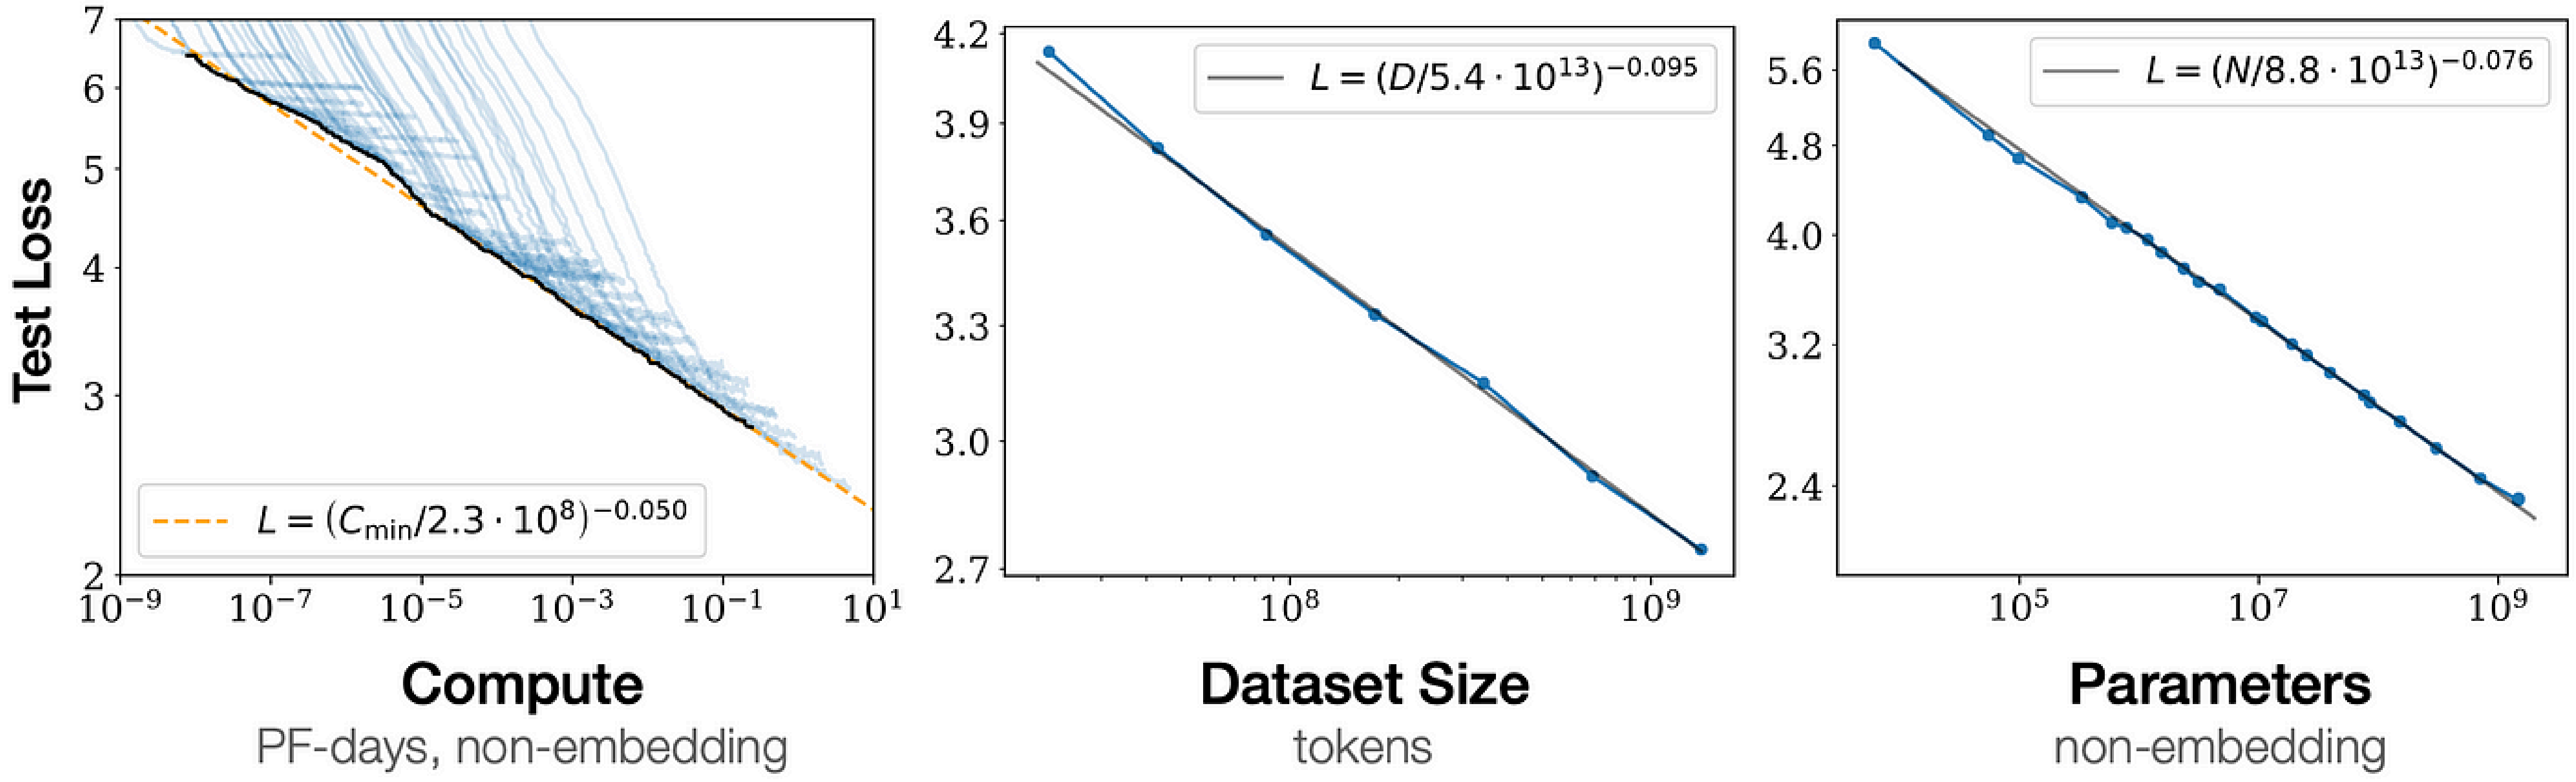

Universität
Konstanz

Deep Learning and Social Sciences

Why Understanding AI is Important

AI is a the transformative technology of our time, reshaping economies, labor markets, and social interactions

- **Geopolitical significance**
  - AI as strategic technology (US-China competition)
  - National AI strategies and investment
  - Regulatory landscape (EU AI Act)
- **Social impact:**
  - Algorithmic bias and fairness concerns
  - Privacy implications
  - Information ecosystem effects (deepfakes, synthetic media)
- **Research opportunity**
  - Unique methodological tools for social science questions

Text Analysis for Social Science

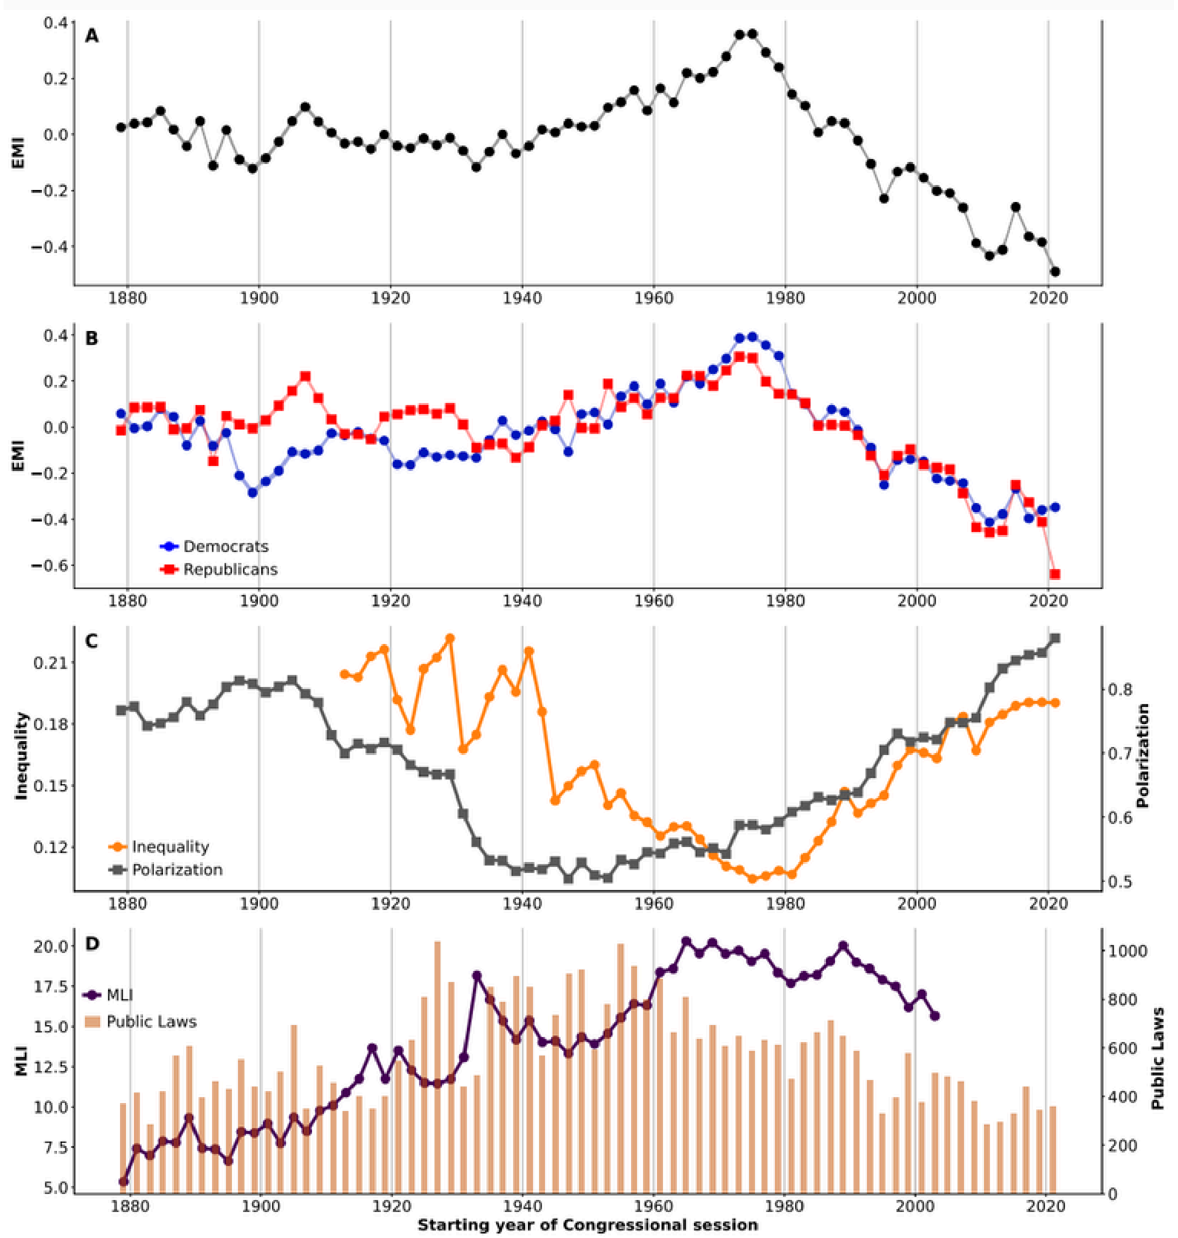

Aroyehun, Segun Taofeek, et al. "Computational analysis of US Congressional speeches reveals a shift from evidence to intuition." arXiv preprint arXiv:2405.07323 (2024).

Natural language processing has revolutionized how social scientists can analyze textual data, enabling research at unprecedented scale and depth

- Natural Language Processing applications:
  - Automated analysis of large text corpora
  - Topic modeling at scale
  - Sentiment analysis of public opinion
- Research examples:
  - Analyzing political polarization in social media
  - Studying media framing across sources
  - Historical text analysis (digitized archives)
  - Policy document analysis

Computer Vision for Social Research

Computer vision techniques enable researchers to extract valuable insights from images and video

- Image analysis applications:
  - Satellite imagery for development indicators
  - Urban planning
  - Visual culture and representation studies
- Research examples:
  - Measuring poverty from satellite imagery
  - Analyzing minorities representation in media
  - Documenting changes in urban landscapes

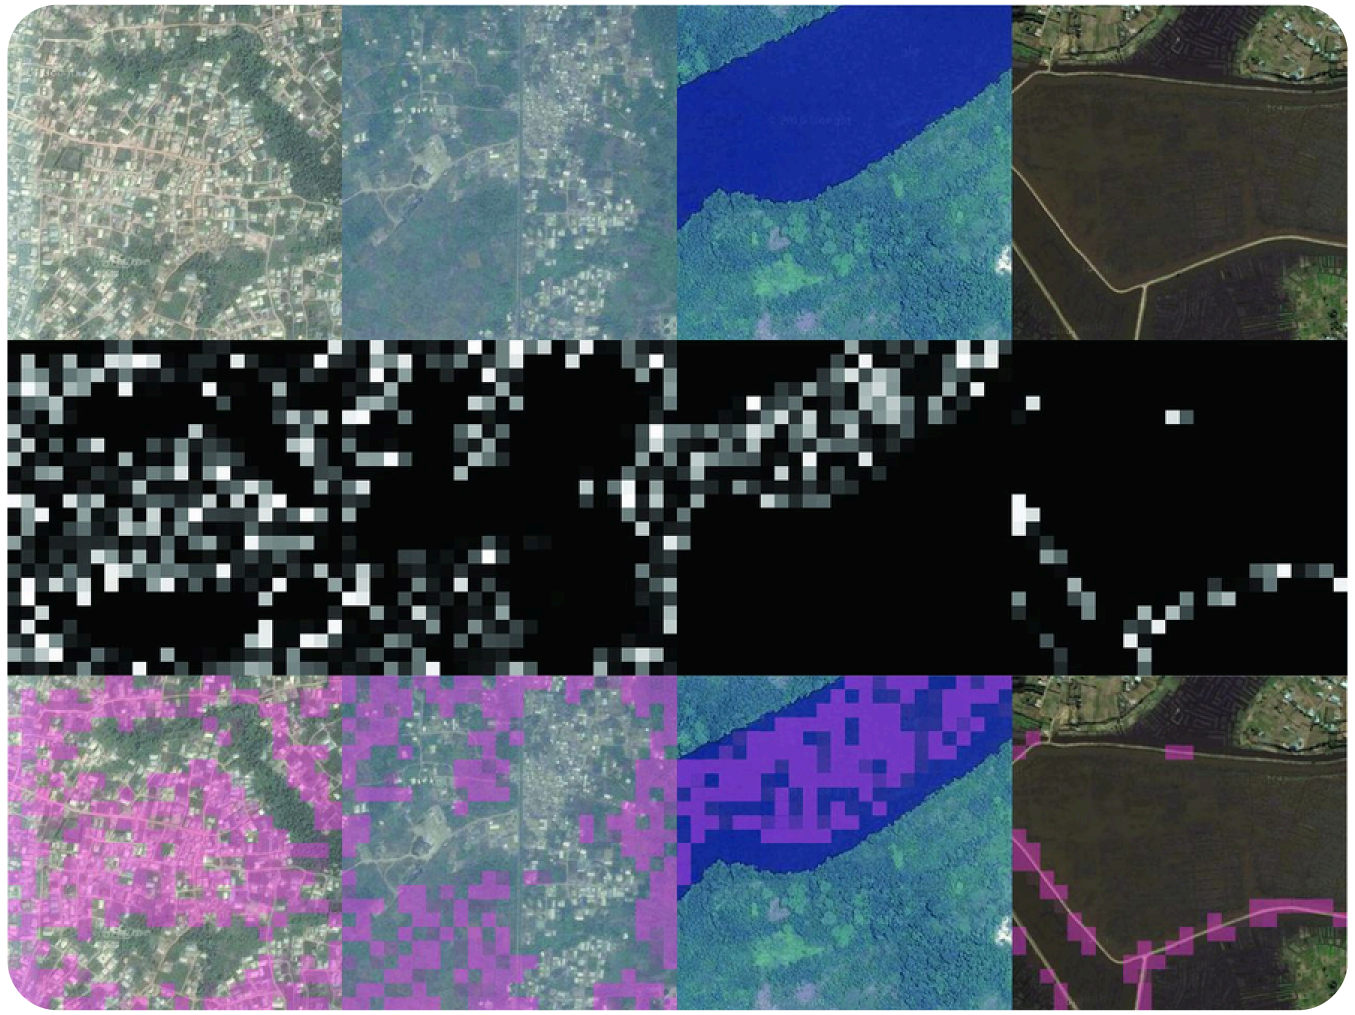

Jean, Neal, et al. "Combining satellite imagery and machine learning to predict poverty." Science 353.6301 (2016): 790-794.

Audio and Video Analysis

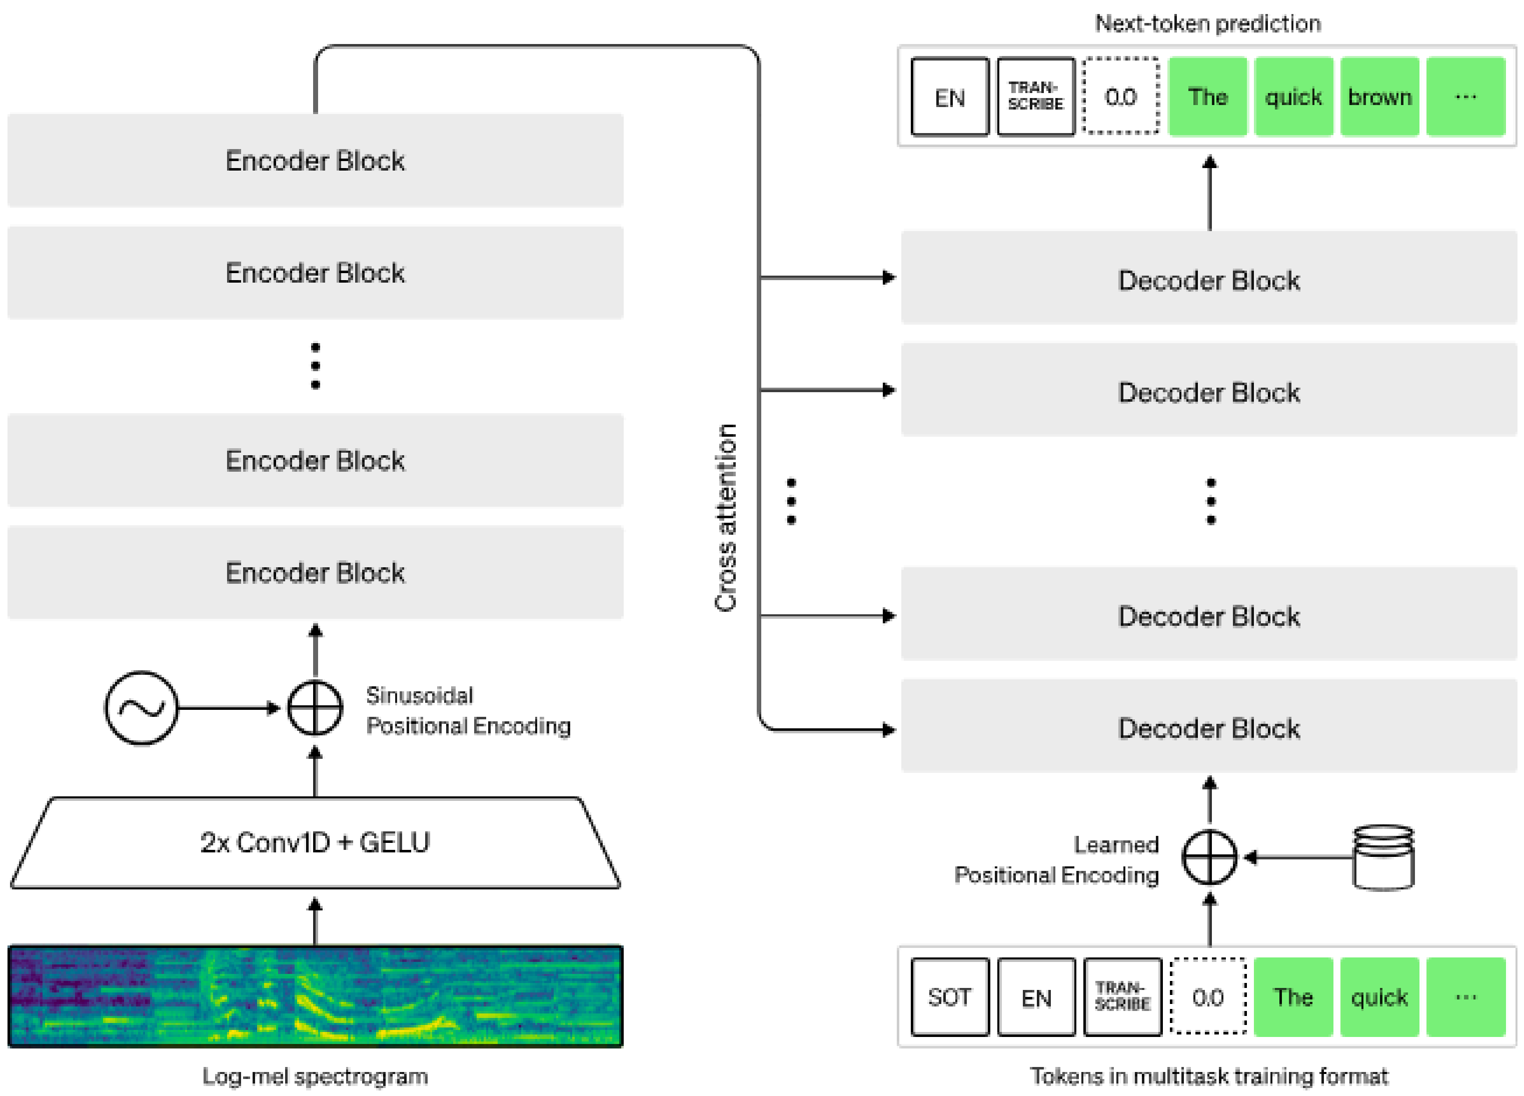

https://openai.com/index/whisper/

Deep learning now enables automated processing of audio and video data sources

- Audio processing applications:
  - Automated interview transcription
  - Speech analysis in political communication
- Video analysis:
  - Analyzing public behavior in urban spaces
  - Documenting social movements
- Research examples:
  - Language use patterns in political debates
  - Non-verbal communication analysis
  - Public space utilization studies

Social Sciences for AI

Most times we talk about how AI can help social sciences, however also humanities and social sciences can help AI

- philosophers help AI societies align their models with human values
- psychology can guide investigation of model behavior
- sociology and social psychology help understand the behavior of group of AIs

We can use what we know from human behavior and experiments we performed to study AI

Ethical Considerations and Challenges

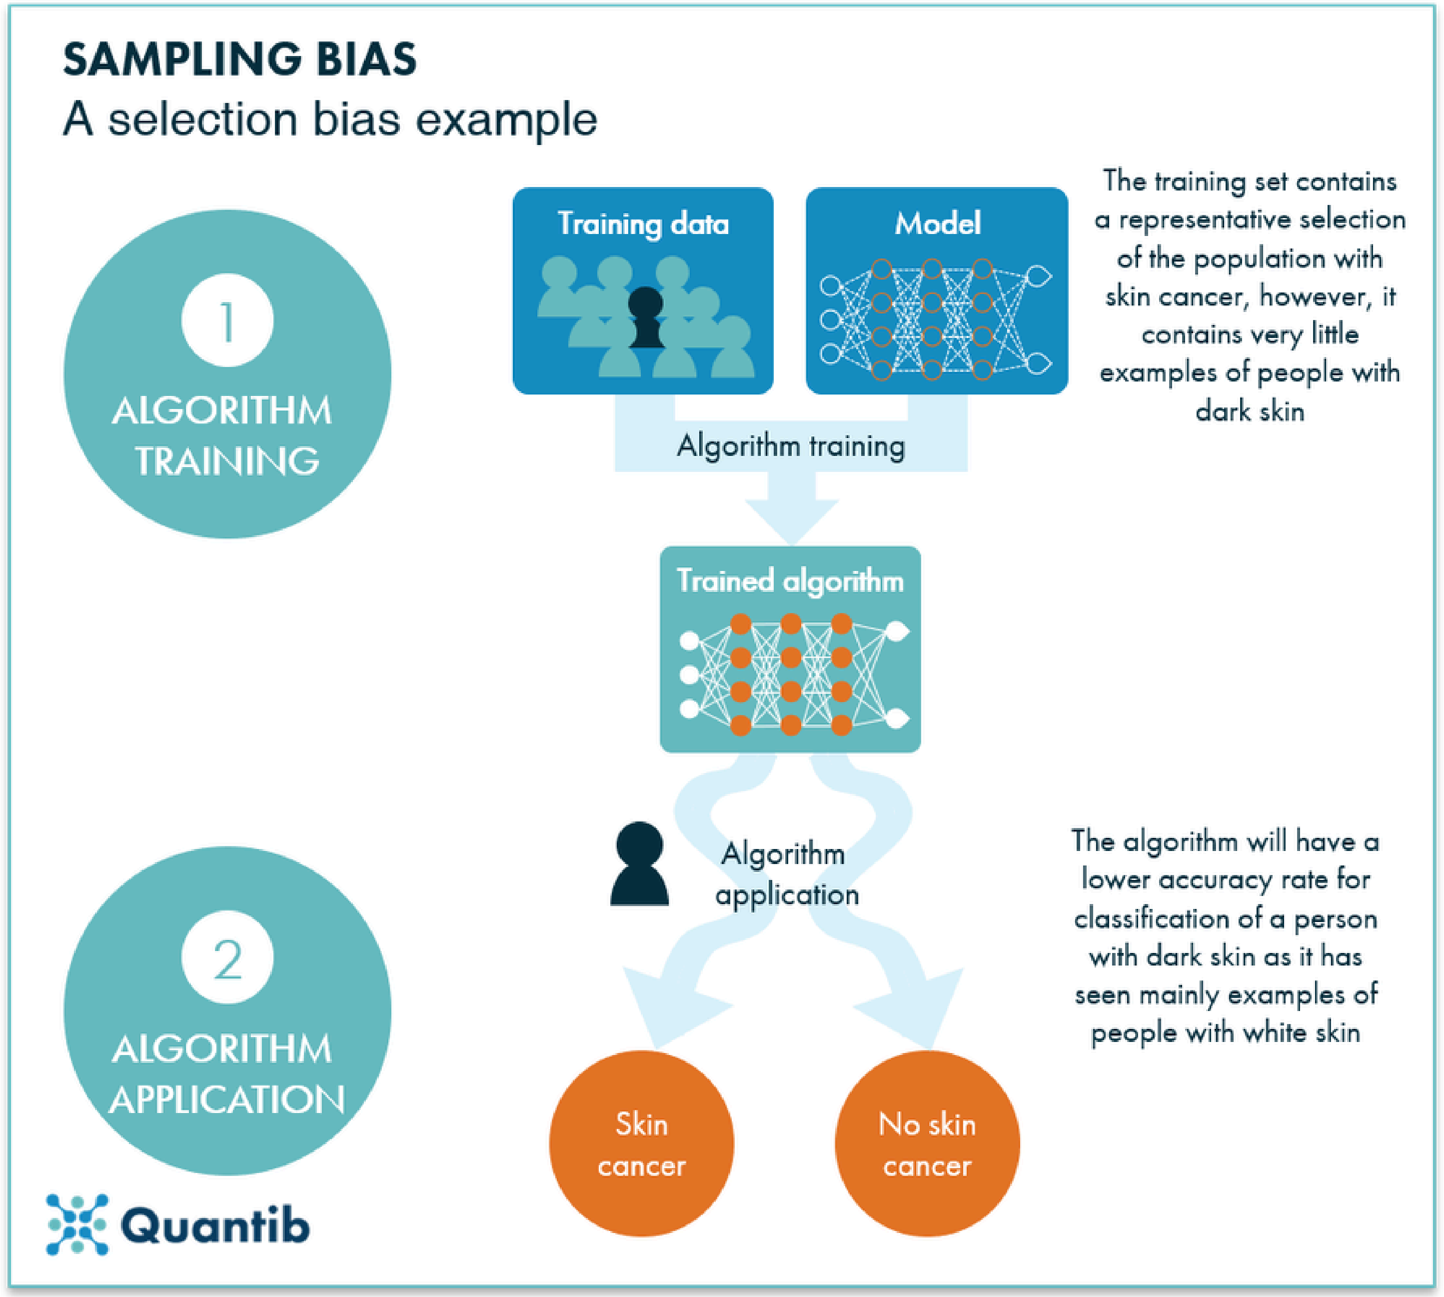

https://www.quantib.com/blog/understanding-the-role-of-ai-bias-in-healthcare

As deep learning tools become widely adopted, we must confront important ethical and methodological questions

- Research ethics:
  - Consent in the age of big data
  - Privacy preservation techniques
- Algorithmic fairness:
  - Bias in training data and models
  - Disparate impacts on different populations
- Methodological challenges:
  - Interpretability vs. performance
  - Reproducibility

Universität
Konstanz

Asch Experiment

Humans tend to conform to the majority

- what happens when the majority is clearly wrong?

Asch investigated this in a famous experiment in the 50s

- participants are asked to determine which comparison line is equal to the reference line
- when many “opponents” provide the wrong answer, participants get fooled

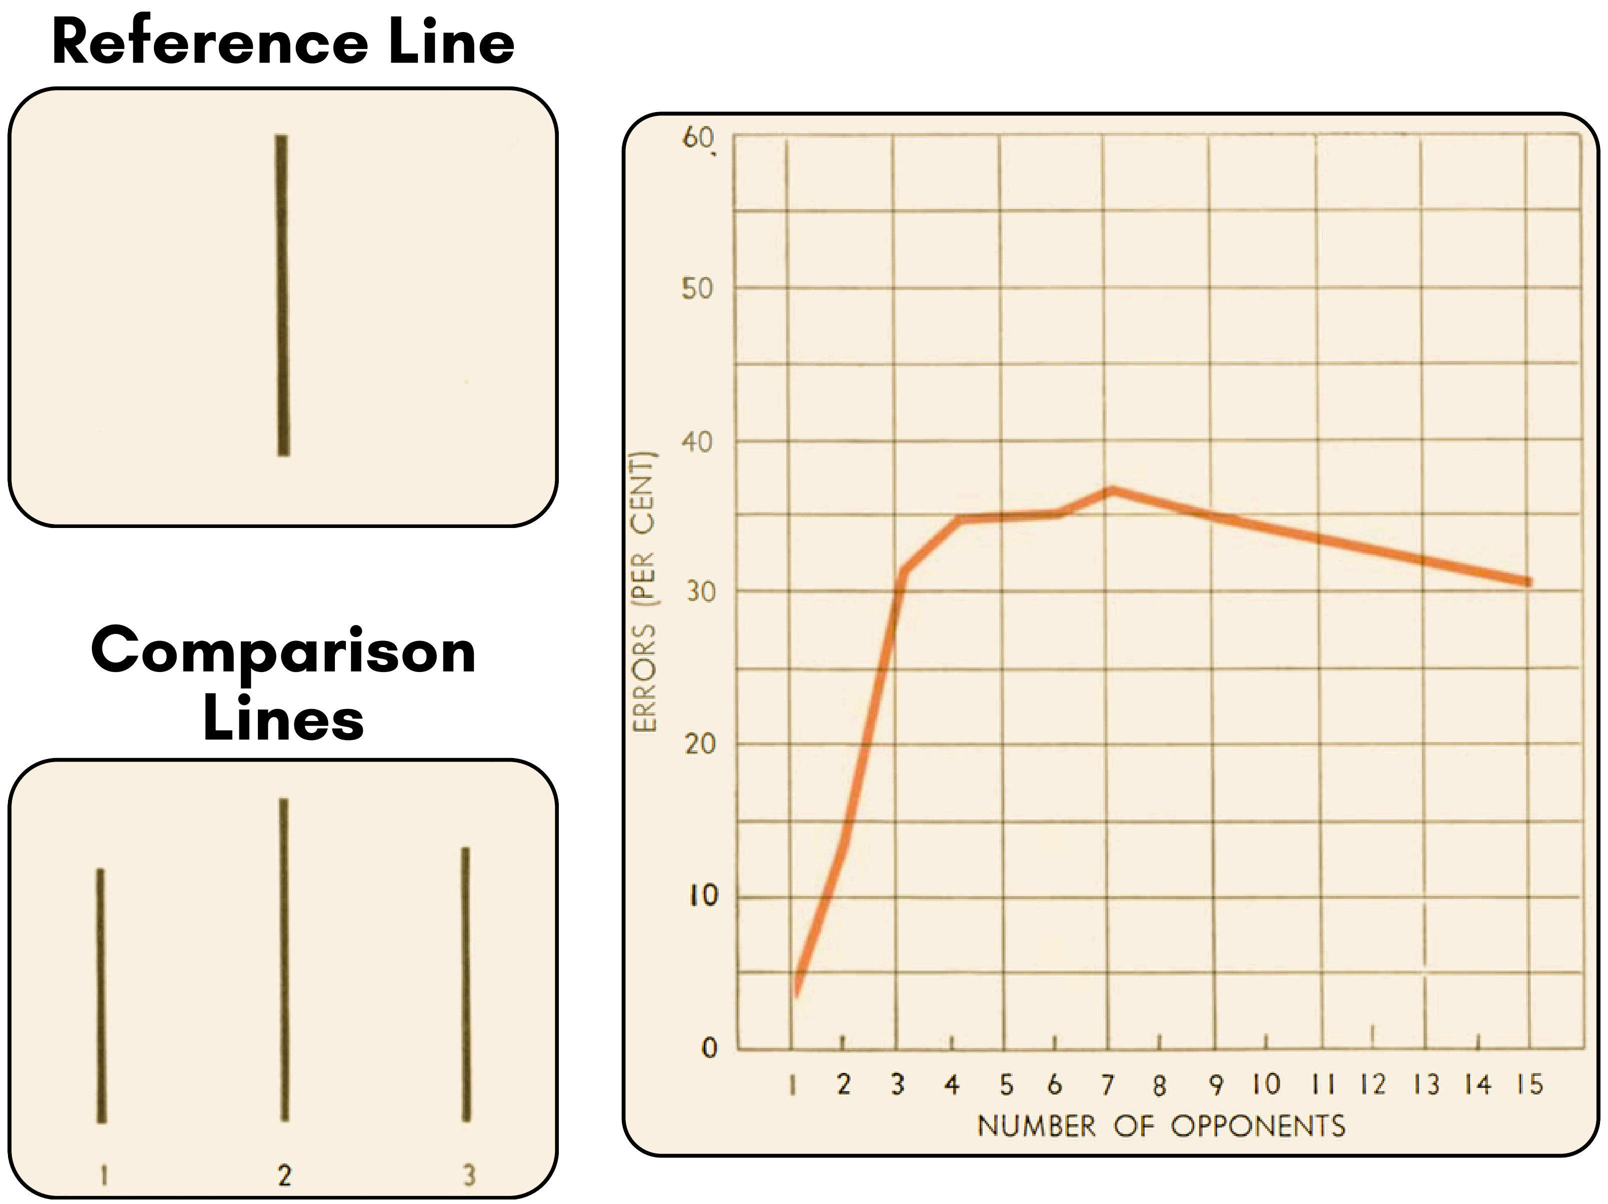

Asch Solomon E. (1955): Opinions and social pressure. Scientific American 193: 31–35.

Universität Konstanz

Conformity in LLMs

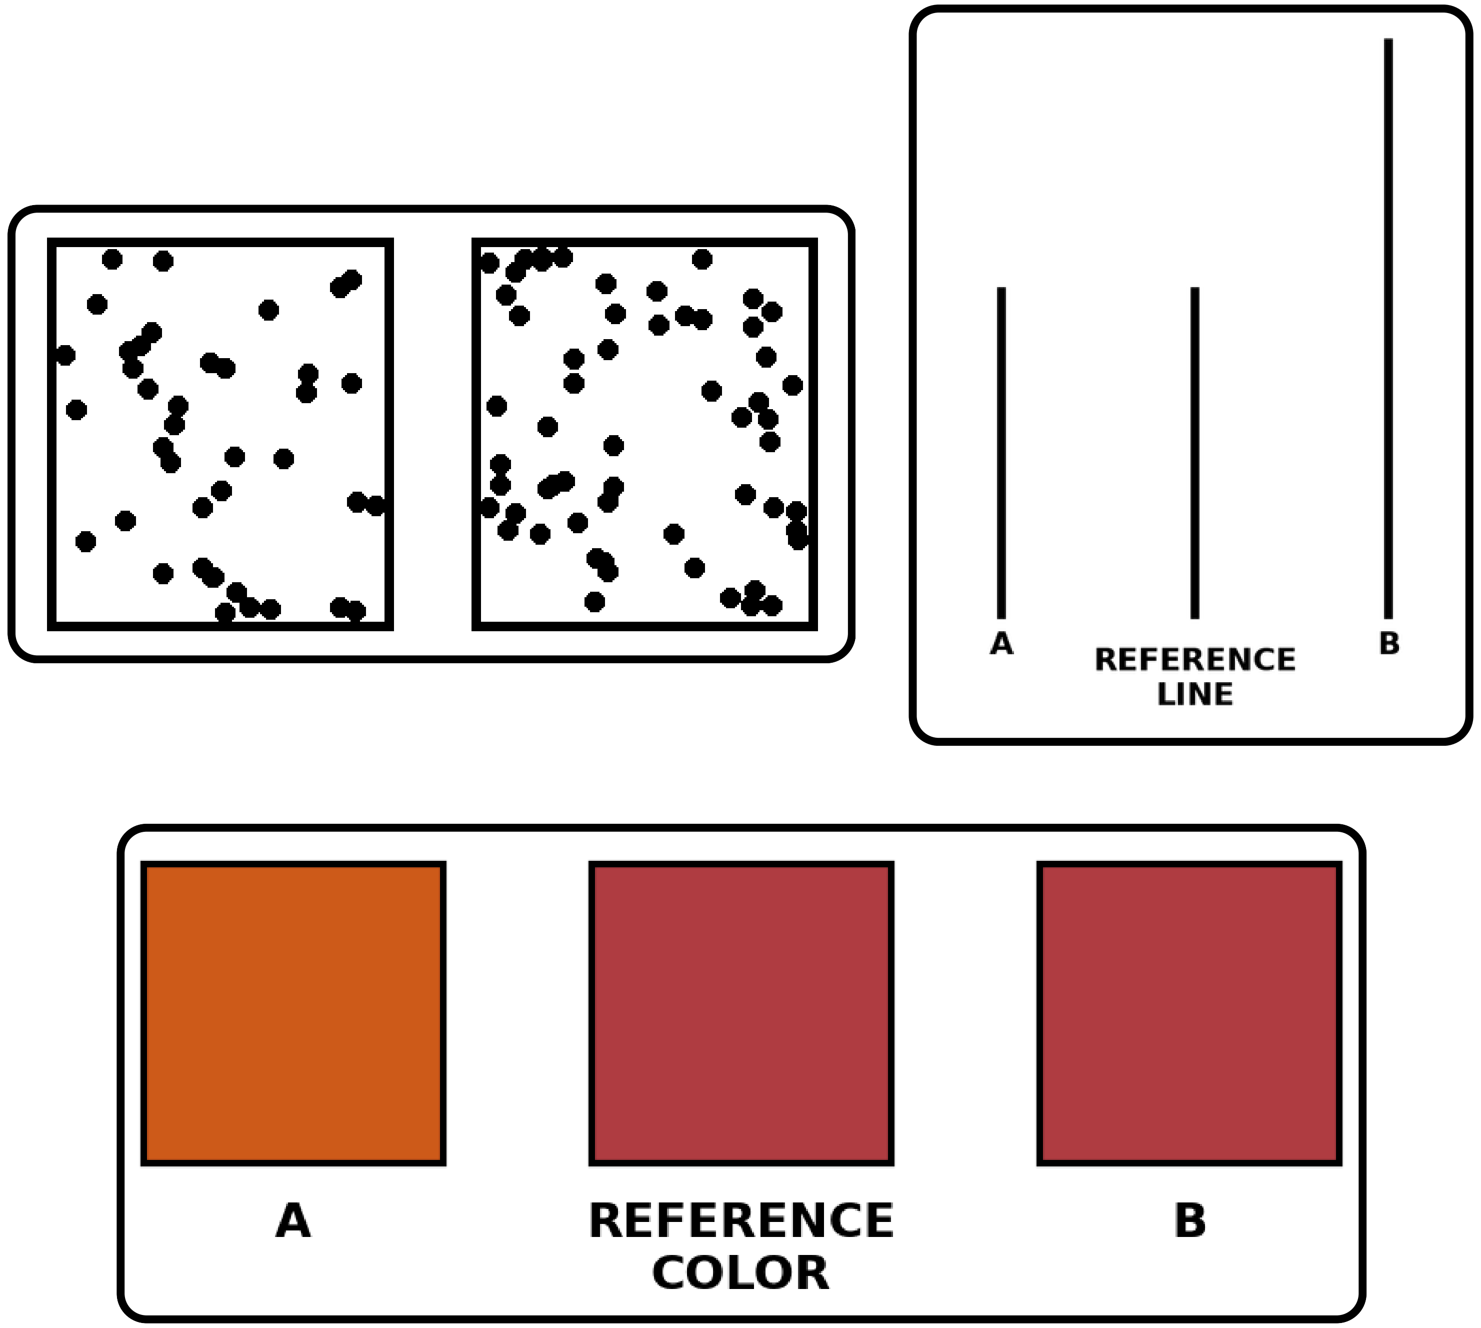

We can use vision LLMs to replicate Asch conformity experiment

- we use three paradigms of conformity experiments in humans
  - dot estimation
  - line measurement
  - color recognition
- there is an option which is clearly correct
- we then simulate confederates who suggest the wrong answer

Social Impact Theory

AI Agents follow the social impact theory by Latané. Conformity is stronger when

- the number of confederates is higher
- the confederates have a form of authority
- confederates are from the same group or nationality

We can use known results for humans to guide the exploration

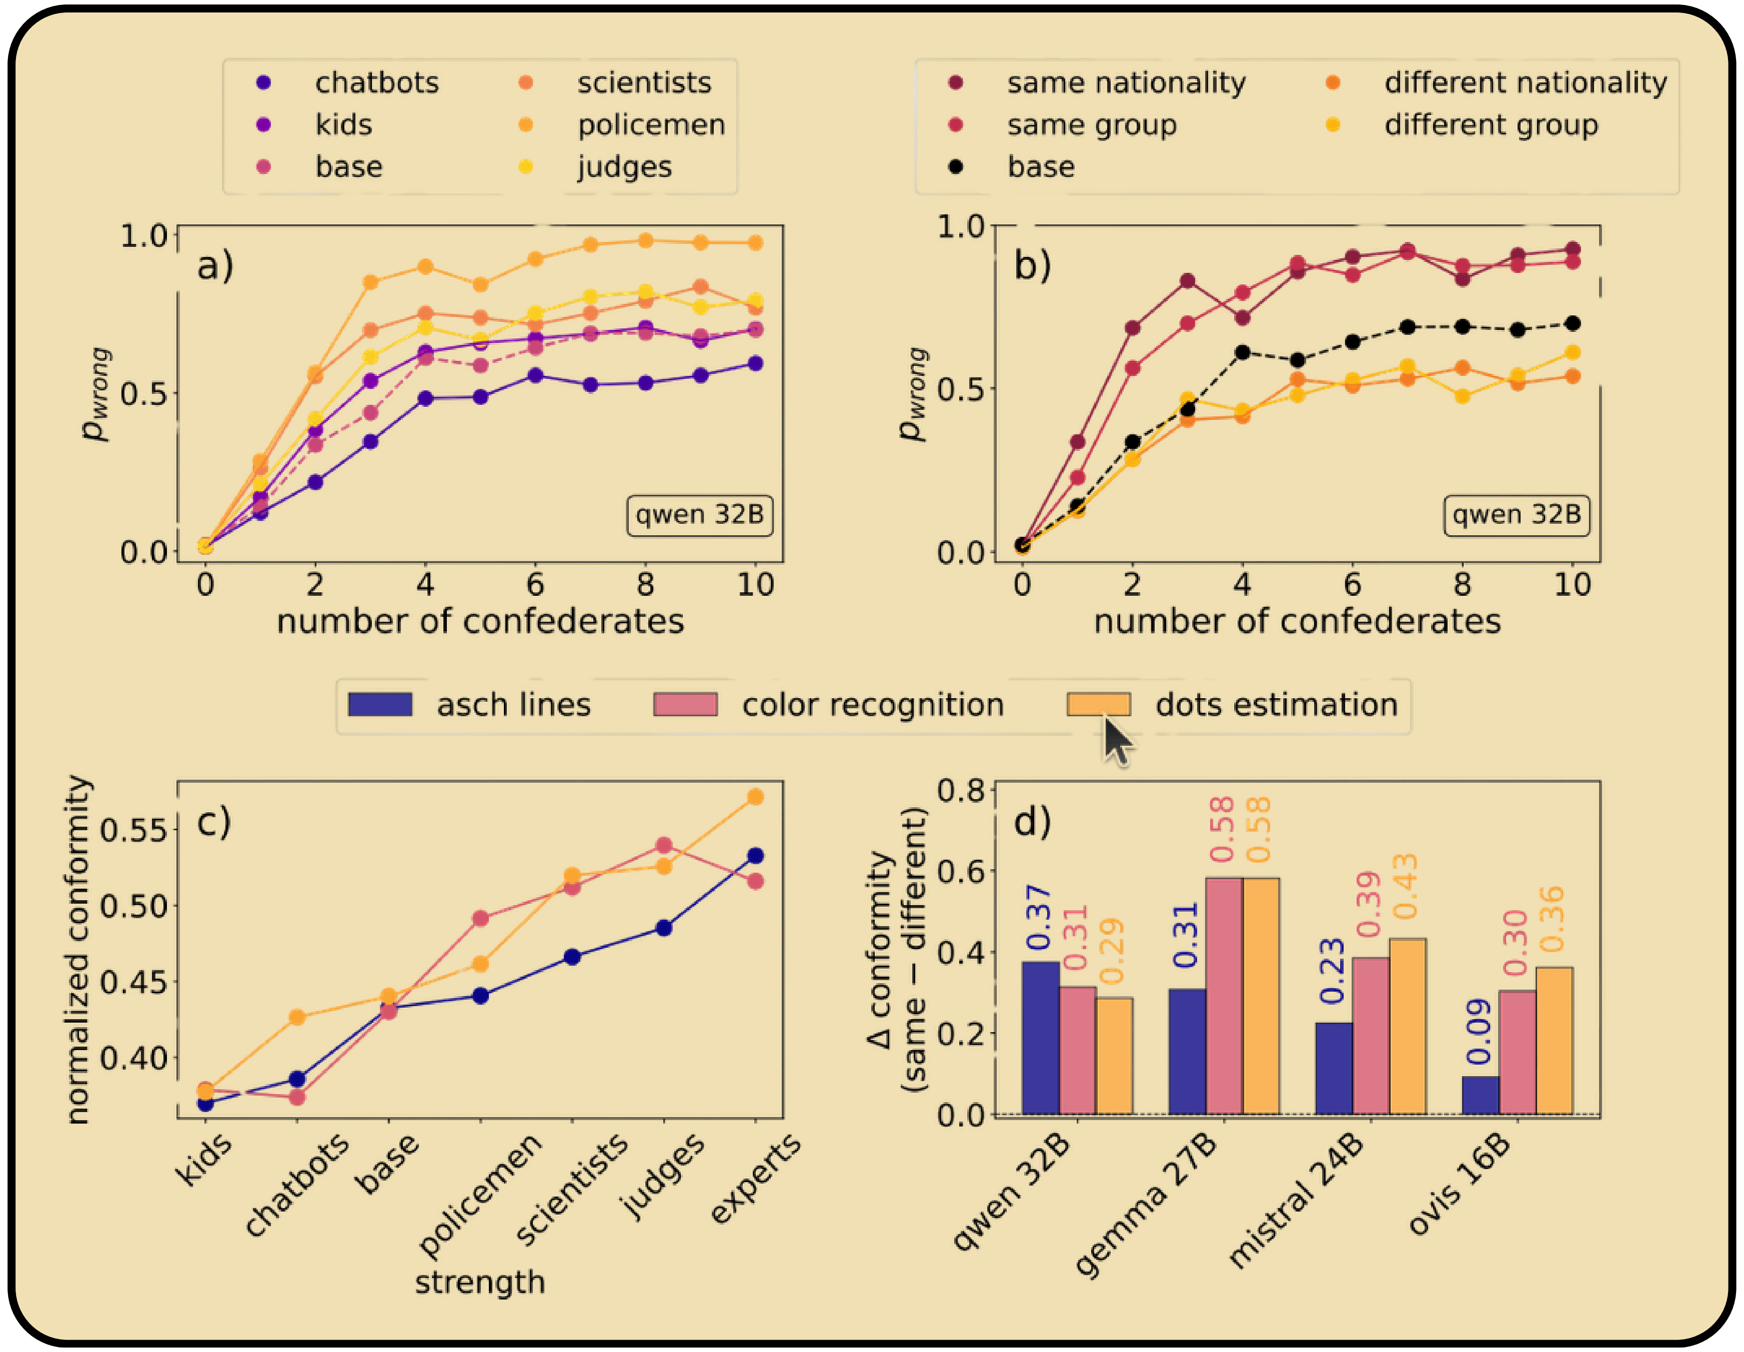

Bellina, Alessandro, Giordano De Marzo, and David Garcia. "Conformity and Social Impact on AI Agents." arXiv preprint arXiv:2601.05384 (2026).

The Social Brain Hypothesis

Humans and primates tend group into societies

- their size is intrinsically limited by the dimension of the neocortex
- the group size grows with the cognitive capabilities of the primate
- for humans this leads to a maximal size of around 150 individuals (Dunbar's numbers)

What about AI Agents?

- are there intrinsic limits to the size of an AI populated society?

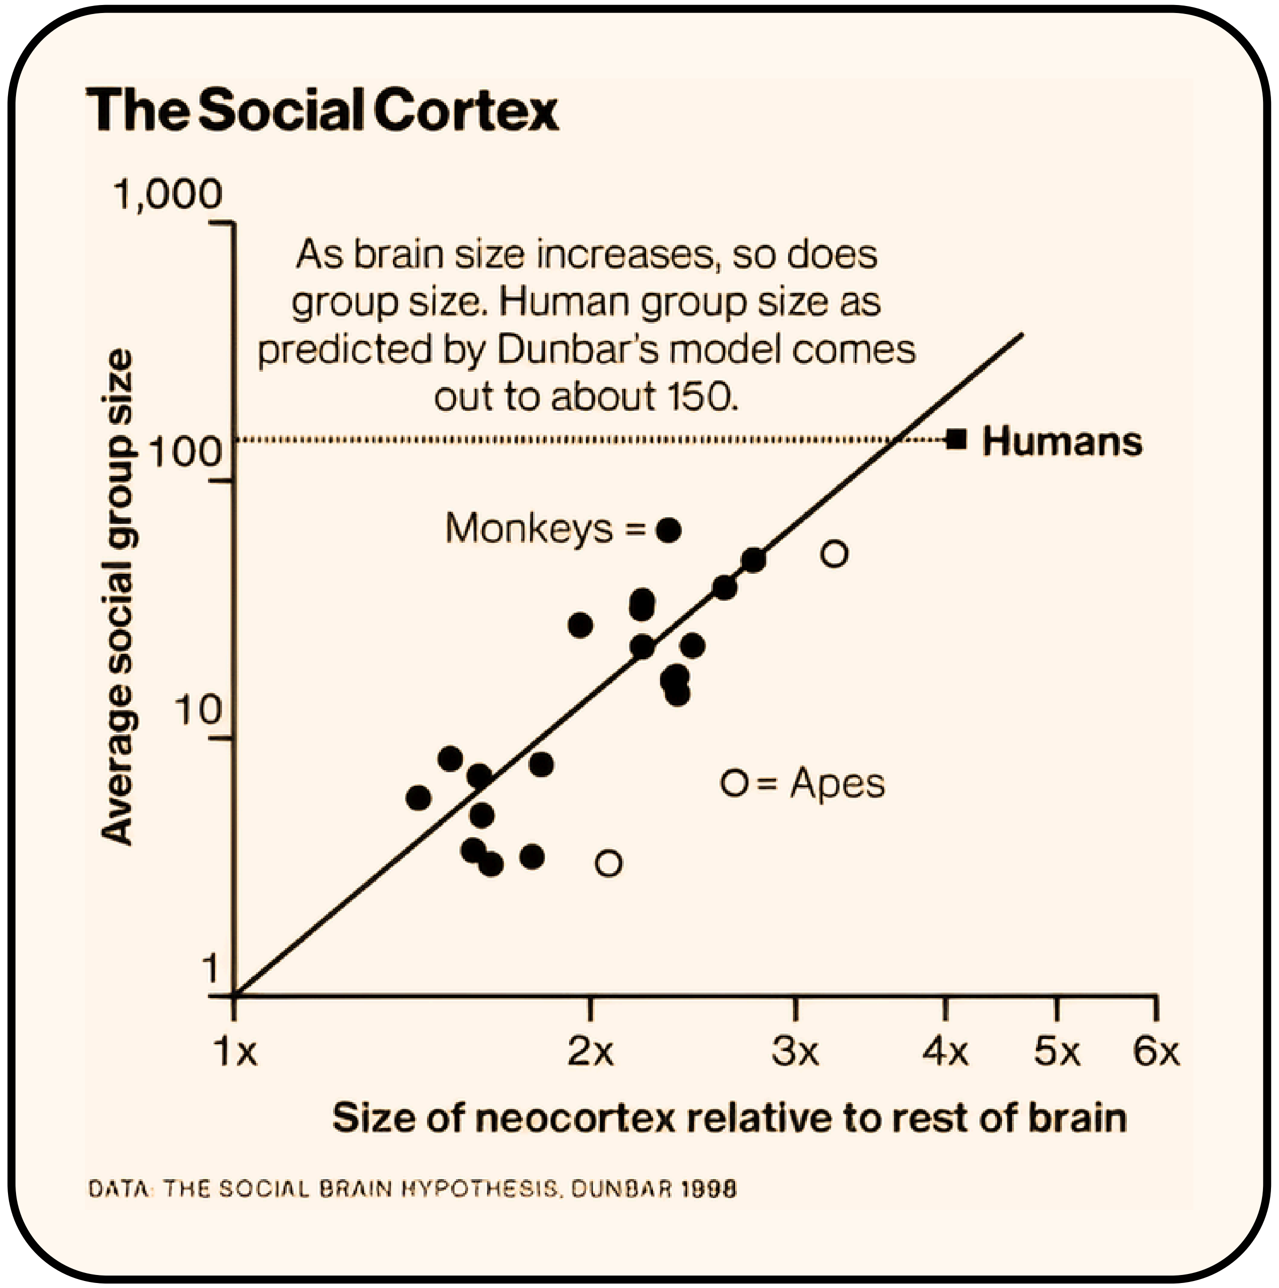

The Social LLM Hypothesis

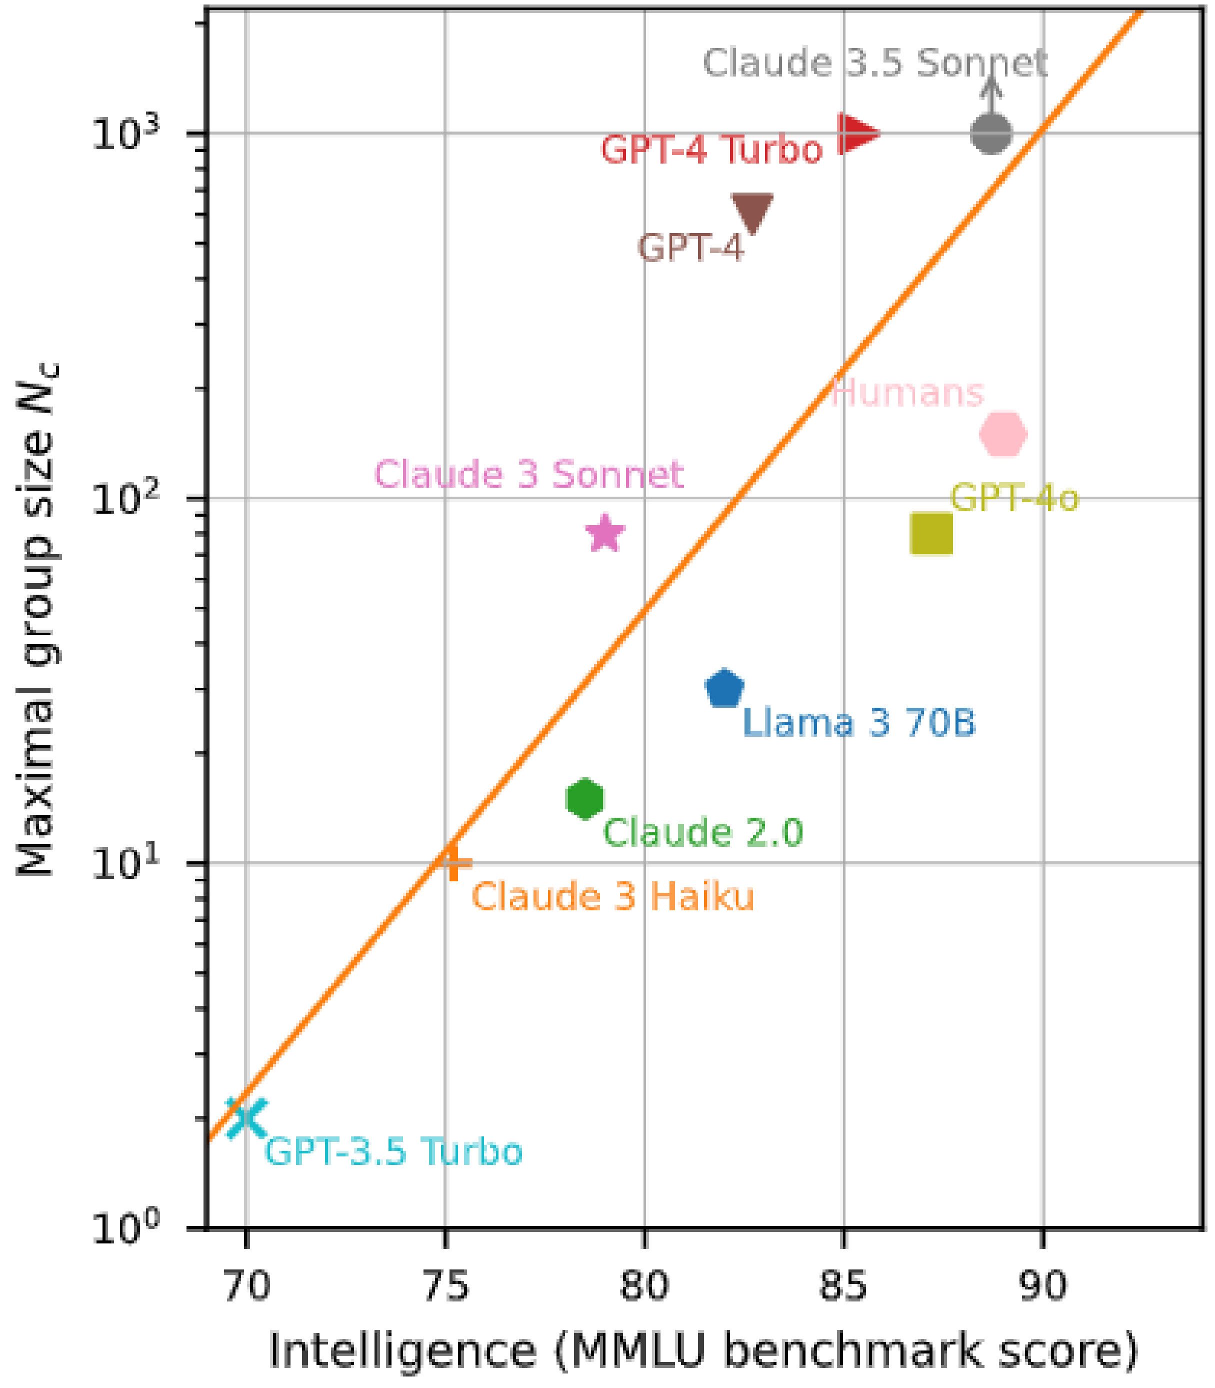

Like primates, also AI Agents have an intrinsic limit on the maximal group size

- using theory from statistical physics we can compute the maximal group size of AI agents
- this maximal group size grows with the capabilities of the model
- the ability to form stable groups is an emergent property

The most advanced models have superhuman coordination capabilities

De Marzo, Giordano, Claudio Castellano, and David Garcia. "AI agents can coordinate beyond human scale." arXiv preprint arXiv:2409.02822 (2024).

Summary

Course Info

You can find all the info at https://giordano-demarzo.github.io/teaching/deep-learning-26/

Deep Learning

Deep learning uses multi-layered neural networks to learn from vast amounts of data. Unlike traditional machine learning, these systems automatically discover patterns and representations without human-engineered features. Algorithmic innovations, massive datasets, and powerful computing has transformed deep learning into a technology central in our society

Deep Learning and Social Sciences

Deep learning offers unprecedented tools to analyze complex human behavior and social phenomena at scale. On the other hand, knowledge from social sciences can help up guide the development of AI models and societies