


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)






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





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

During my PhD I worked in private and public research centers

- Junior Data Scientist at CSH Vienna
- Research Visit at Sony CSL Paris
- Research Visit at Konstanz University





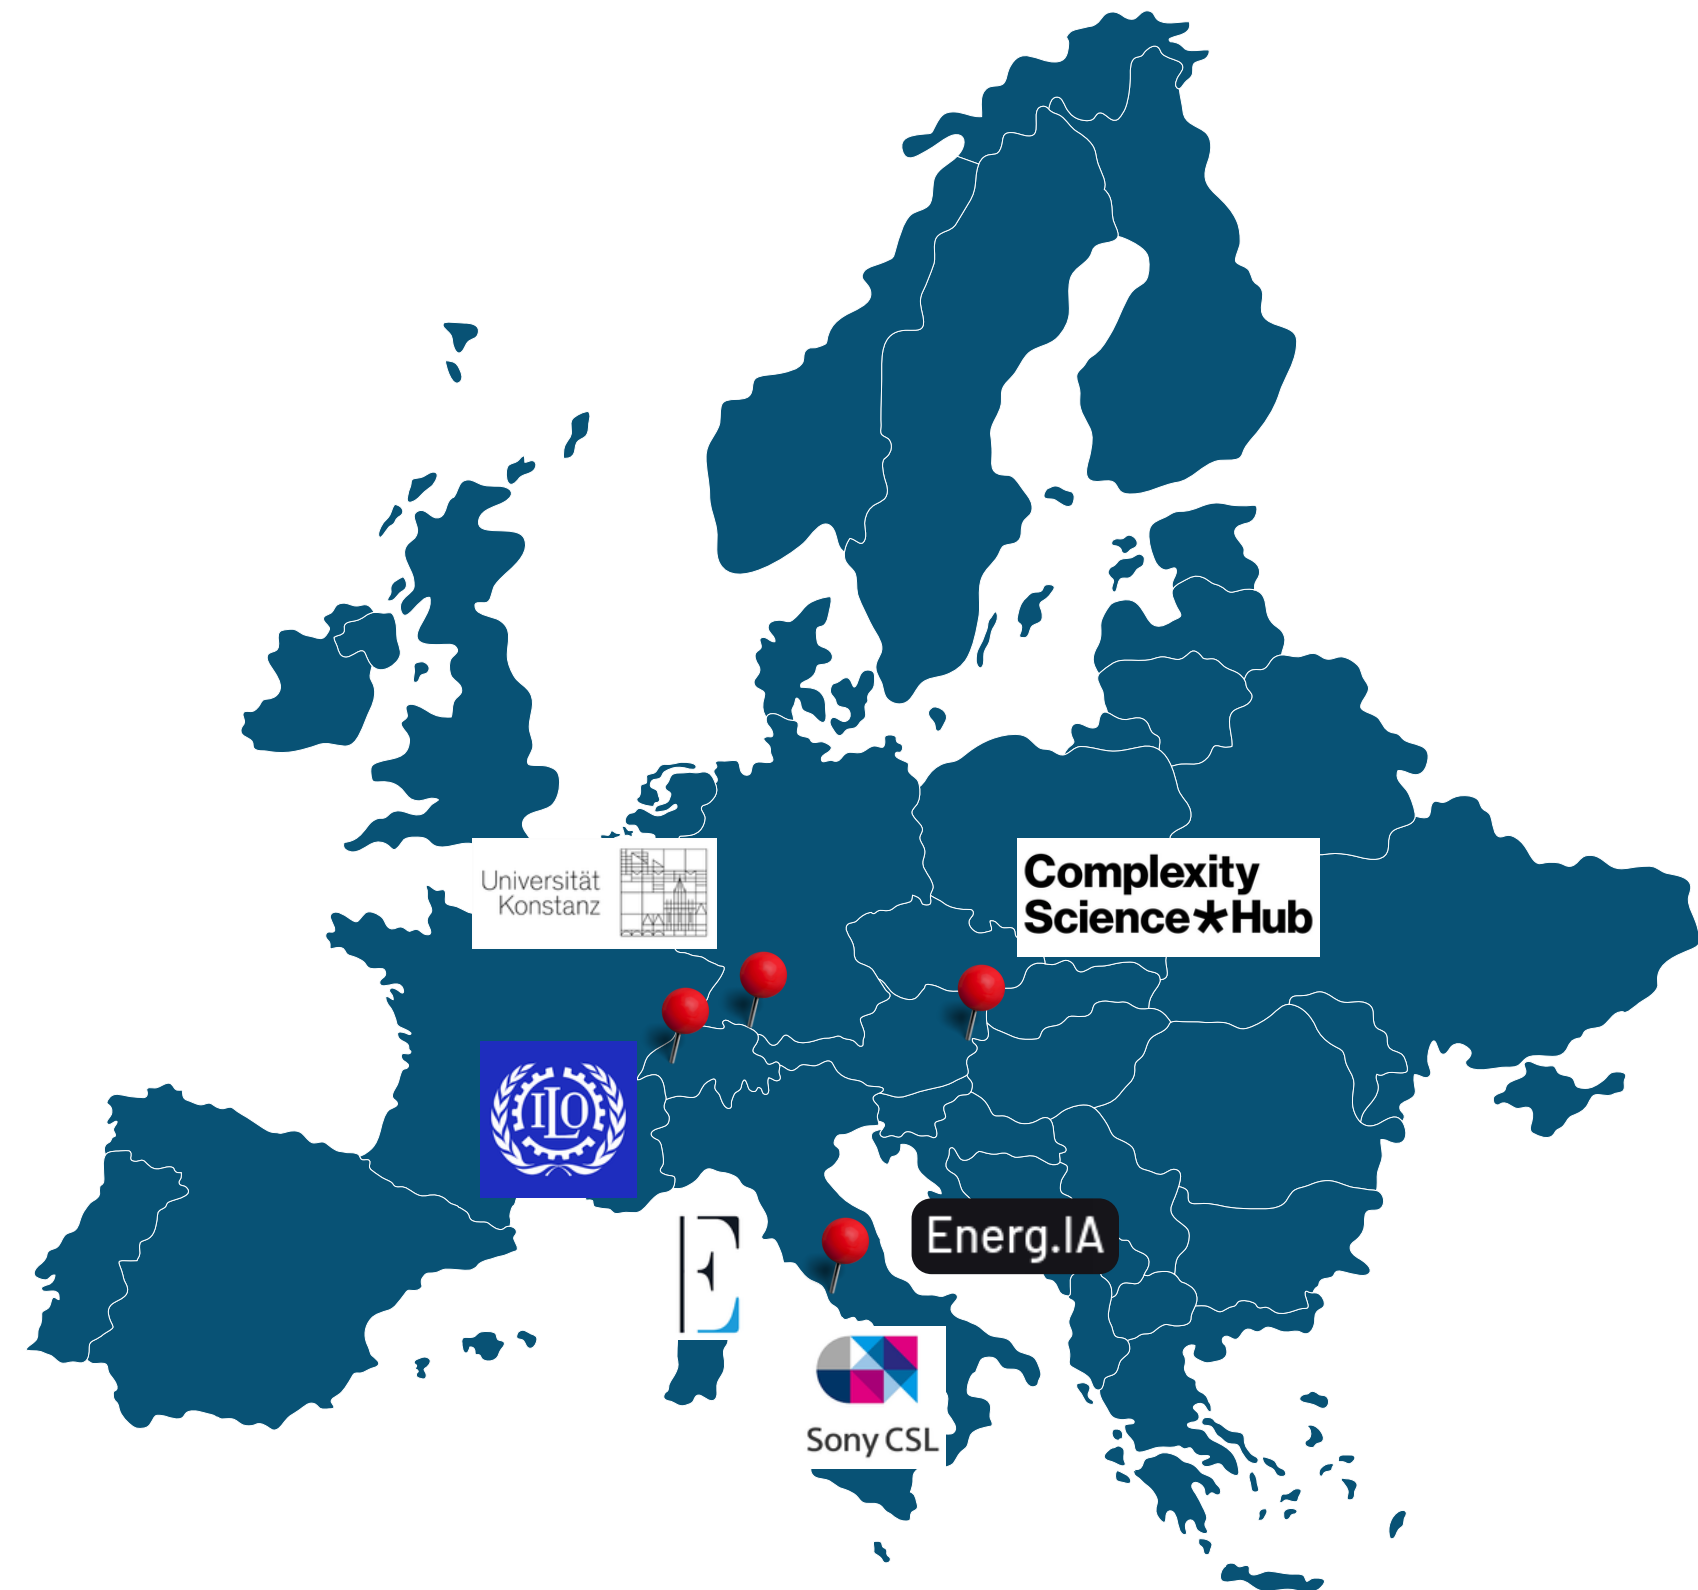
Me in Short


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

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

I also keep collaborating with other institutions

- Junior Research Fellow at CSH Vienna
- Collaborations with CREF and Sony CSL
- Founder and shareholder in AI for green energy company





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
- 3 guest researchers seminars
- 1 students presentation day
- 8 coding sessions

The course will require you to code in python


- you will be provided with some computation resources
 - jupyter notebooks with 10GB GPUs
- you need a github account to access the GPUs
- more details on 17/04 during the coding lab

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




Calendar

Date	Topic	Date	Topic
April 14	Introduction to Neural Networks	April 15	Lab: GPU Server + Machine Learning Basics
April 23	The Multilayer Perceptron	April 24	Lab: Building a MLP with PyTorch
April 28	Regularization of Neural Networks		
May 7	Convolutional Neural Networks	May 8	Lab: Building a CNN to classify clothes
May 14	Graph Neural Networks	May 15	Lab: Introduction to PyTorch Geometric
May 21	Recurrent Neural Networks	May 22	Lab: Time Series + Guest Research Seminar



Calendar

Date	Topic	Date	Topic
May 28	Natural Language Processing		
June 4	Transformer and Attention	June 5	Lab: Working with embeddings
June 25	Large Language Models	June 26	Lab: Fine-tuning LLMs
July 2	Reinforcement Learning and Alignment	July 3	Guest Researcher Seminar
July 9	Image Generation and Multimodality	July 10	Lab: VAE and Diffusion Models
July 16	Students Presentations	July 17	Guest Researcher Seminar



Assignments

The four 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 of github like for ICSS
- the assignments will be assign on the following dates (topics may change)
 - April 24 - Multilayer Perceptron
 - May 15 - Convolutional Neural Network
 - June 5 - Graph Neural Networks/Recurrent Neural Networks
 - June 21 - 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 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 your project on July 17 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-25/>

You can also check last year website

<https://github.com/DLSS-24/DLSS-24/tree/main>



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



Deep Learning





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

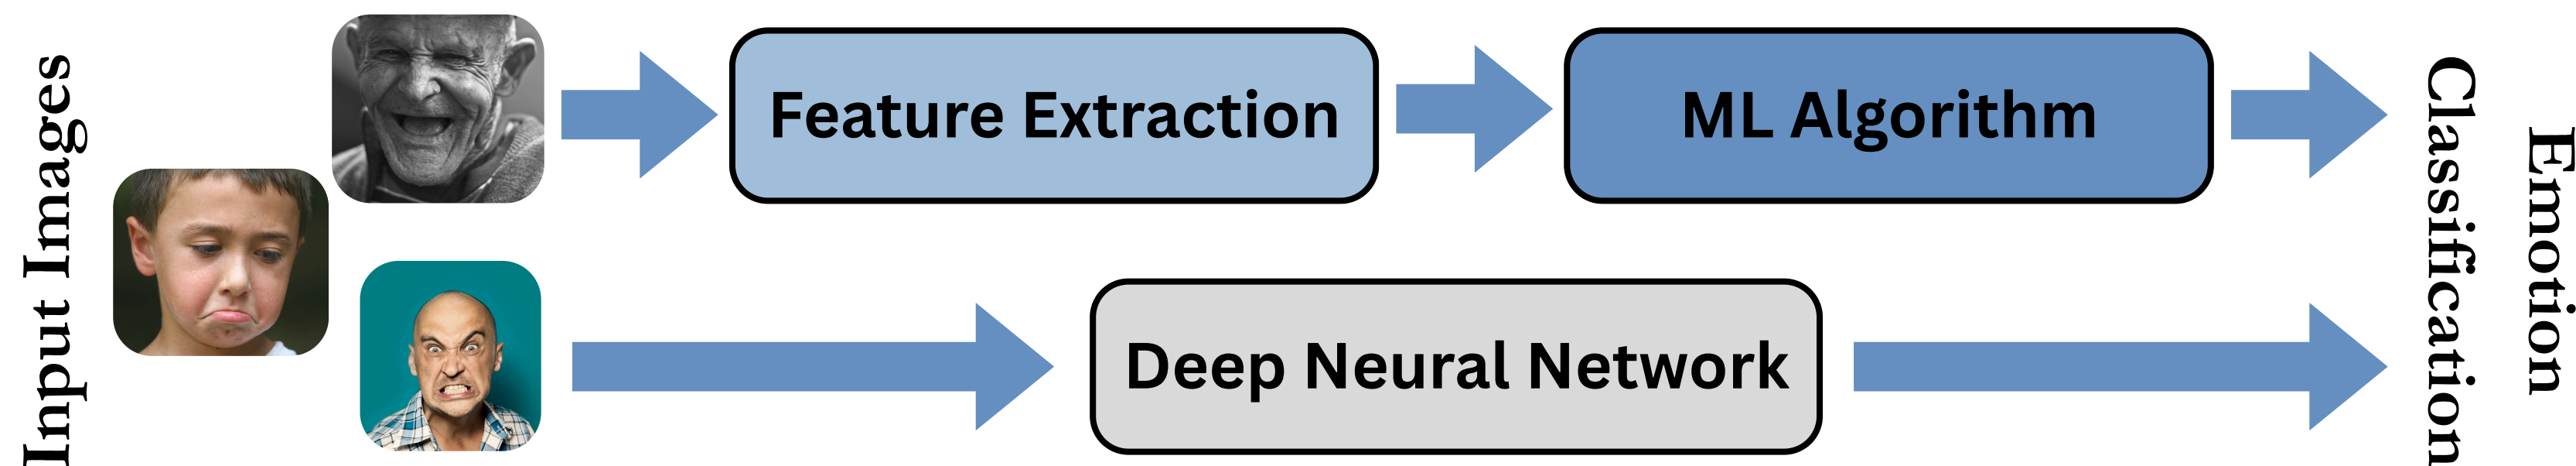
- "Deep" refers to many hidden layers between input and output


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





Deep Learning and Social Sciences







Why Understanding AI is Important

AI is 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



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


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




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

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



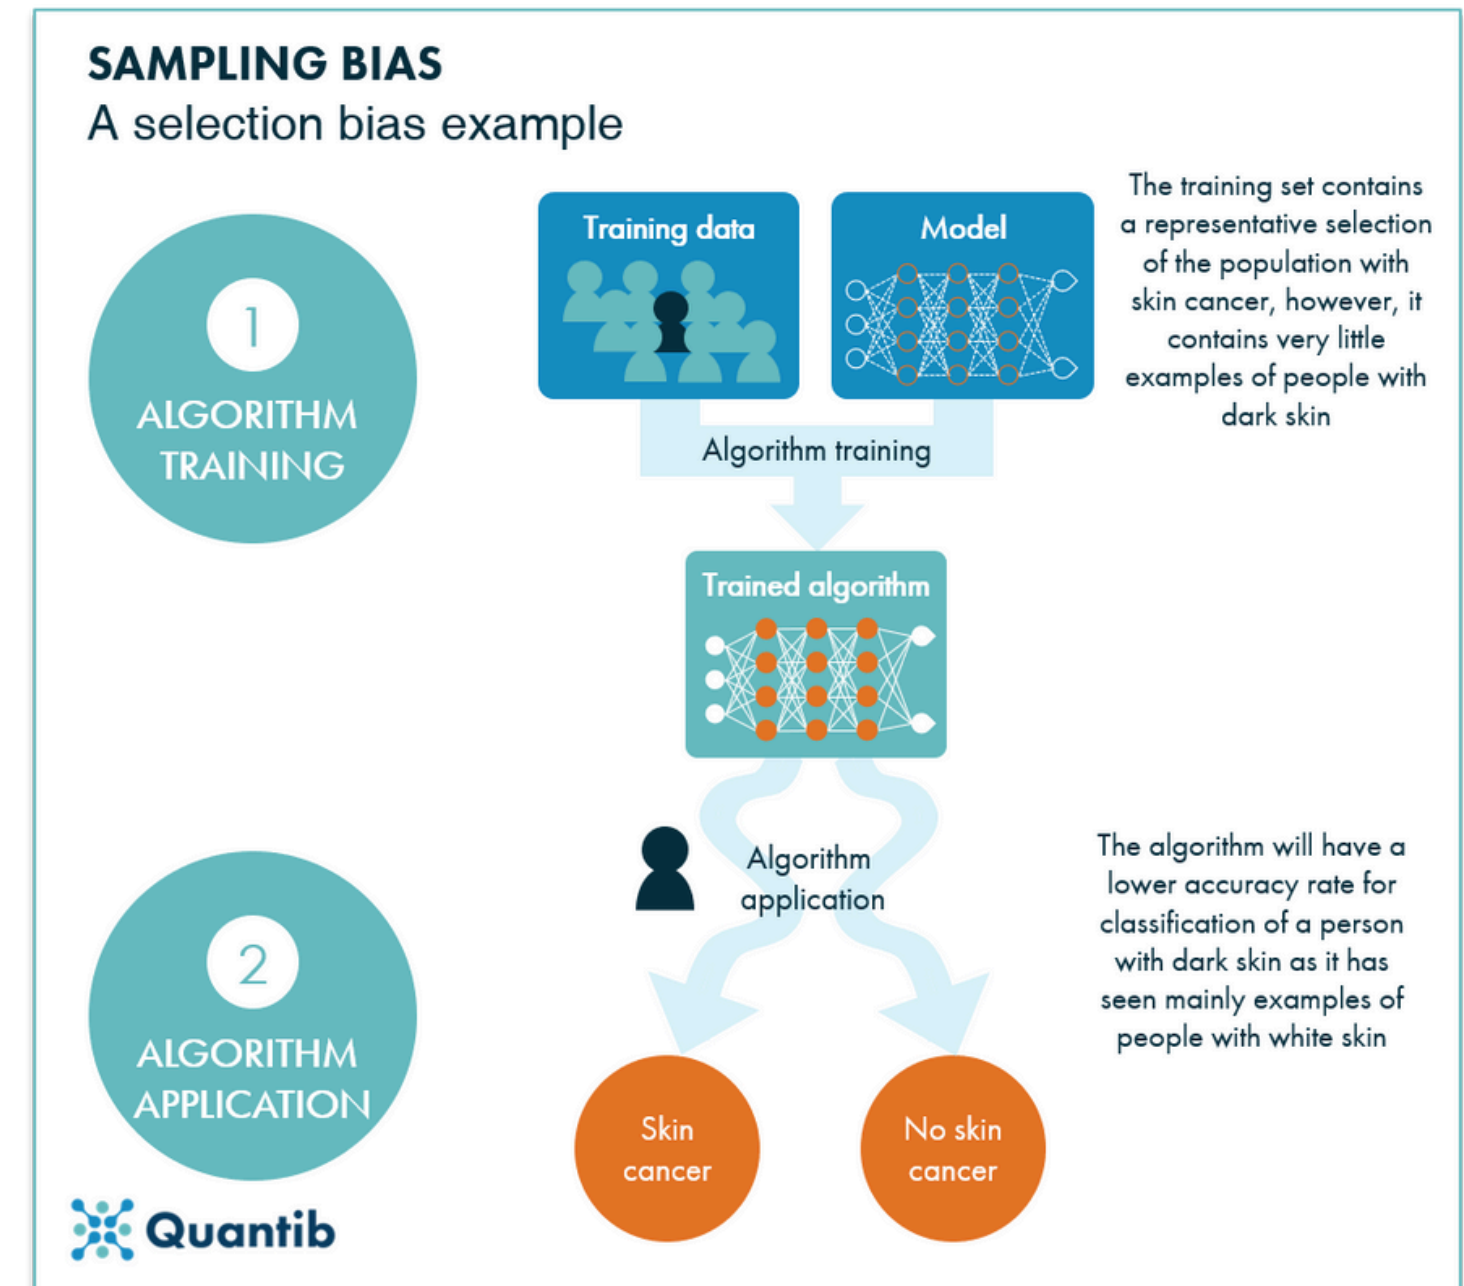
<https://openai.com/index/whisper/>



Ethical Considerations and Challenges

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



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



Summary

Course Info

You can find all the info at <https://giordano-demarzo.github.io/teaching/deep-learning-25/>

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. These techniques enable researchers to extract insights from previously inaccessible data sources including vast text corpora, images, audio recordings, and videos