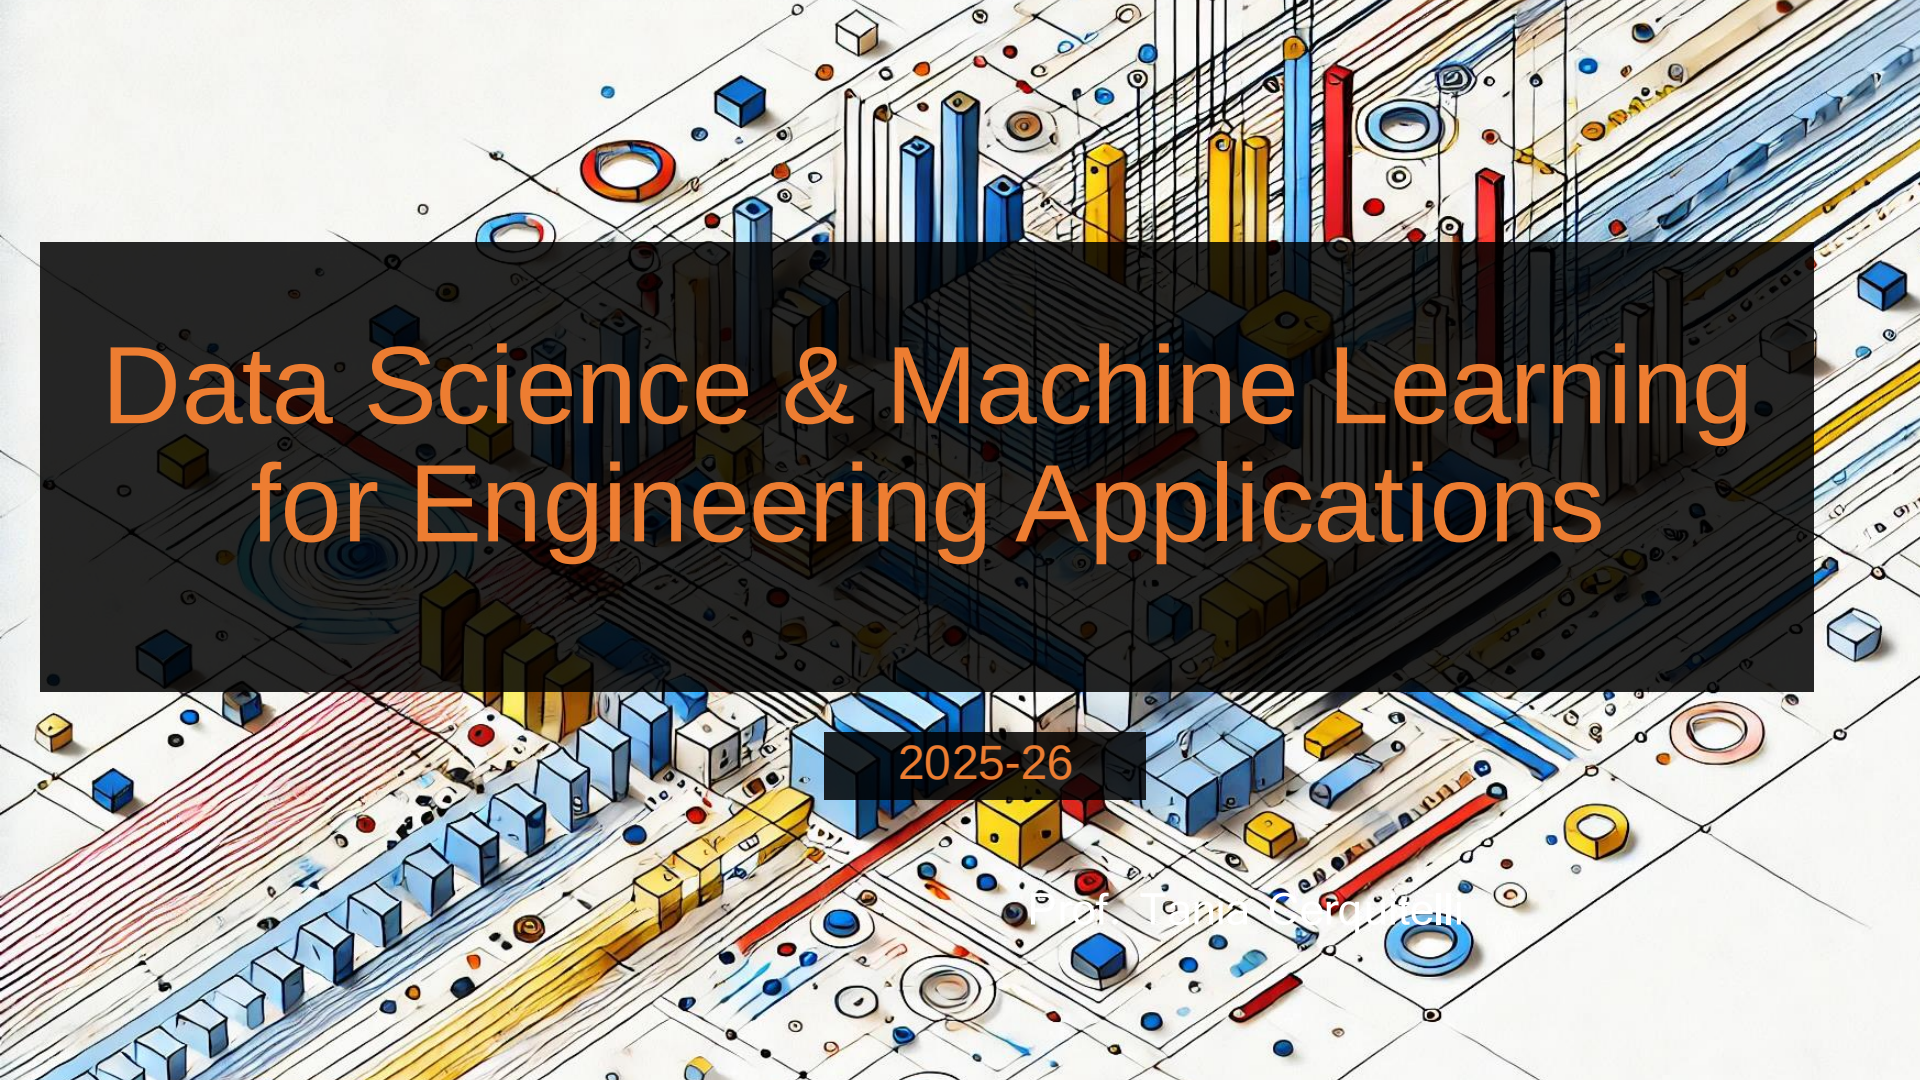

Data Science & Machine Learning for Engineering Applications

2025-26

Prof. Jaan Kalda, Tallinn

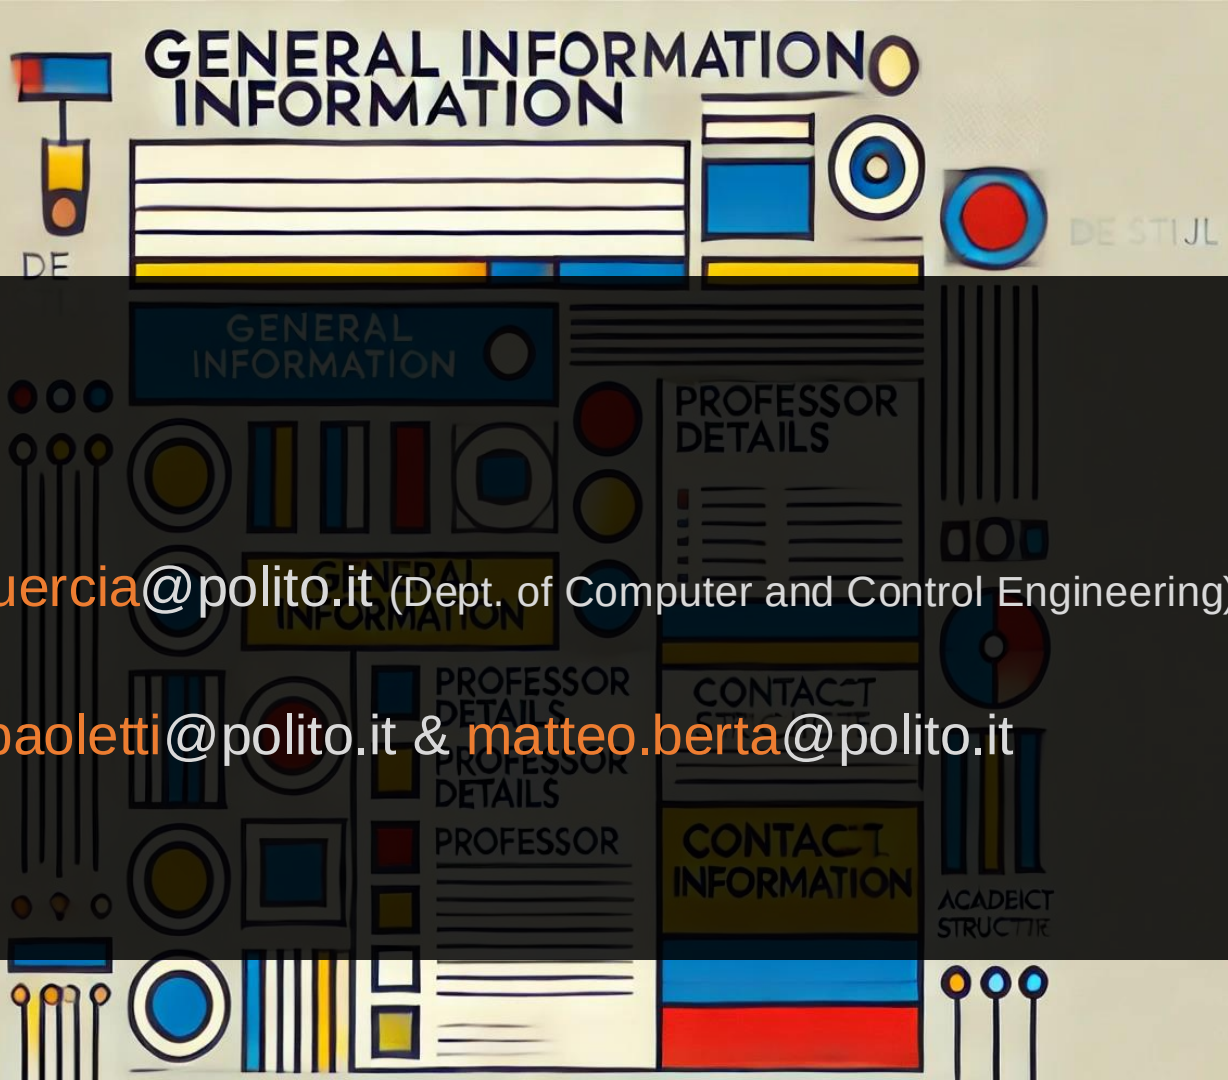

6 Credits

Prof daniele.quercia@polito.it (Dept. of Computer and Control Engineering)

TAs giordano.paoletti@polito.it & matteo.berta@polito.it


POLITO
POLITECNICO DI TORINO

UNIVERSITY
COLLEGE
LONDON

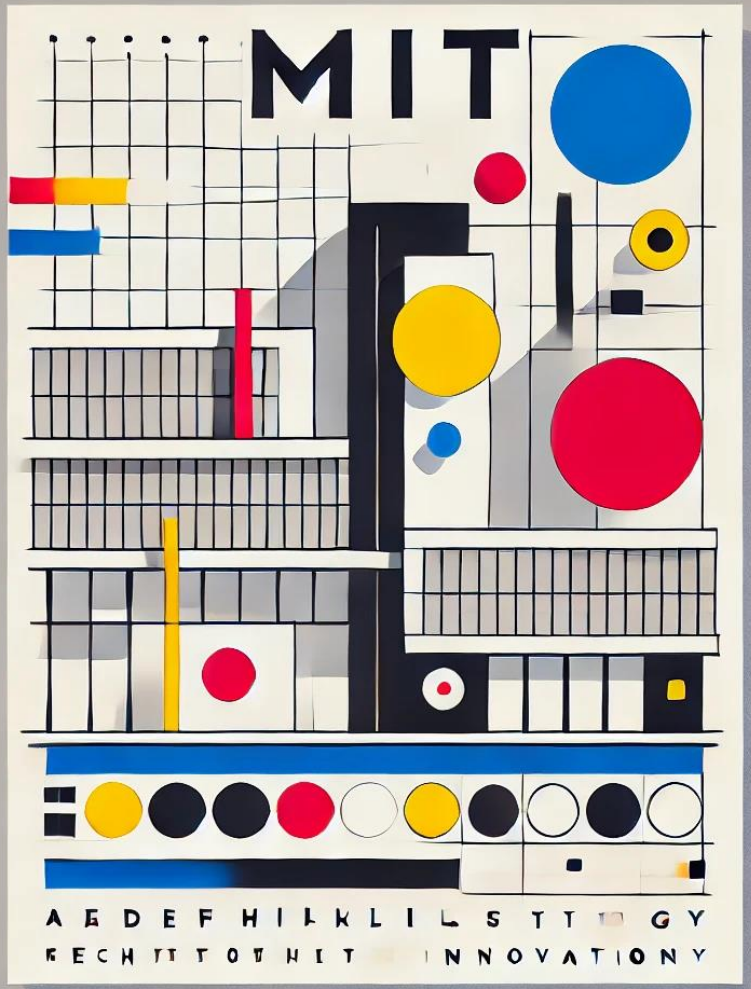

MIT

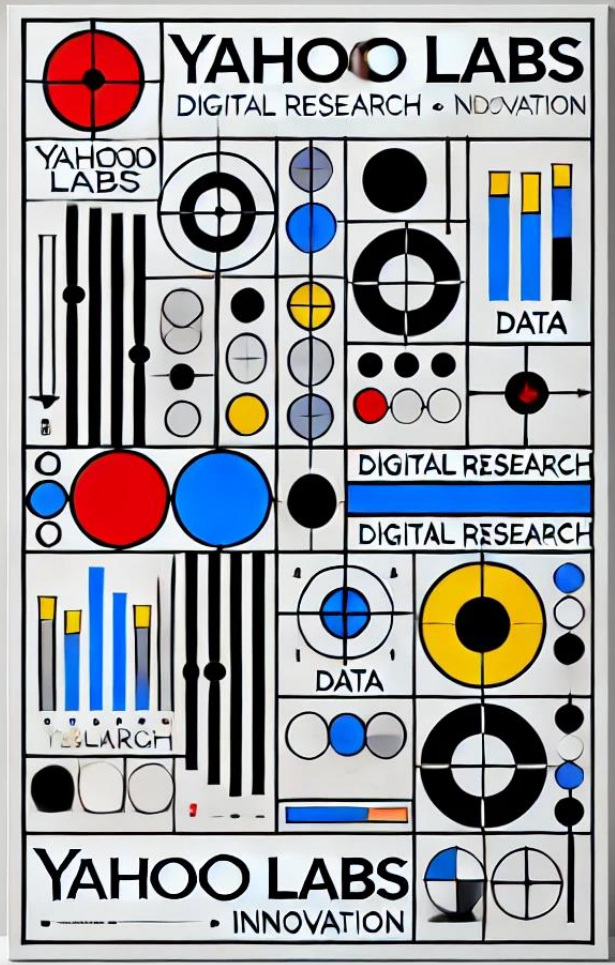
YAHOO LABS
DIGITAL RESEARCH • INNOVATION
DATA
DIGITAL RESEARCH
DIGITAL RESEARCH
DATA
YAHOO LABS
INNOVATION

CAMBRIDGE
UNIVERSITY

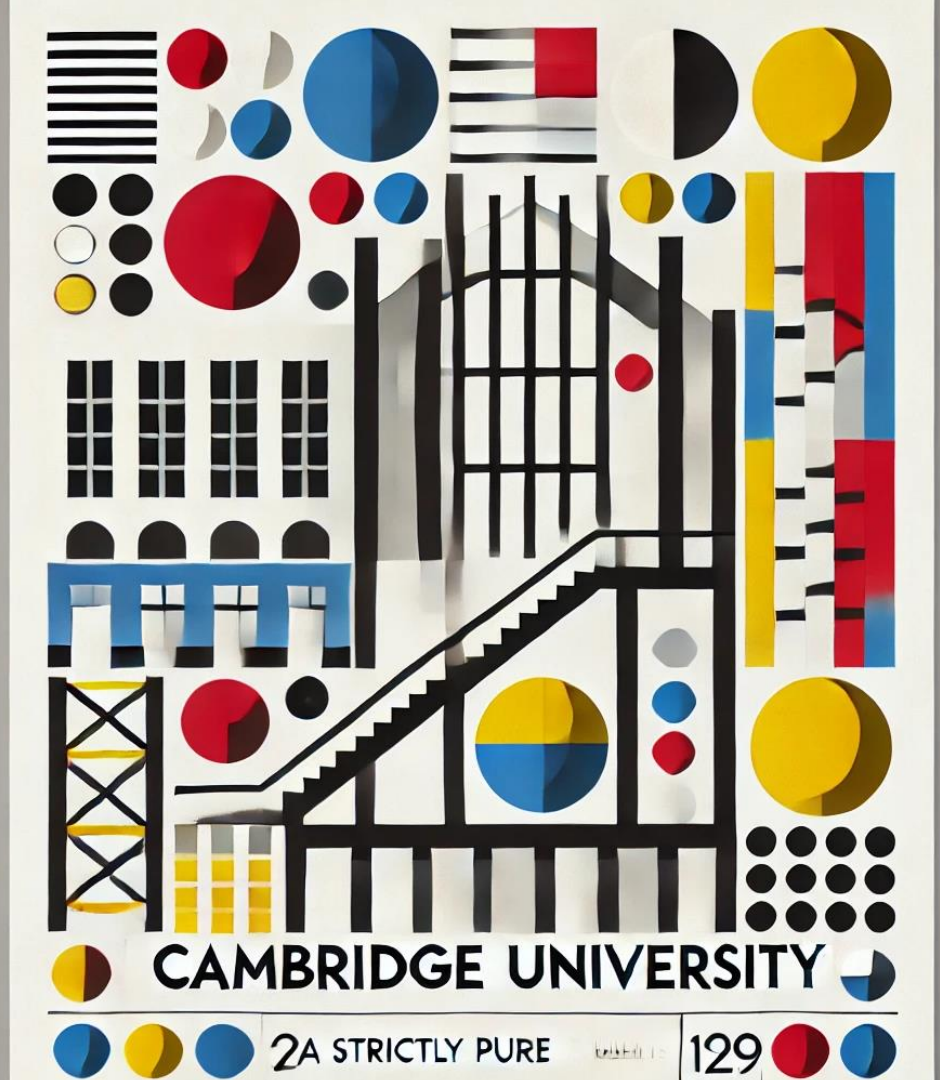

CAMBRIDGE
UNIVERSITY

NOKIA Bell Labs

Inventing the Future X Network

Creation of Bell Labs
The engineering departments of the American Telephone and Telegraph Company (AT&T) and Western Electric were consolidated into Bell Telephone Laboratories. Their mission was to research and design communication technologies for the rapidly expanding telephone network and to explore fundamental areas of science that could shape the future of the industry. Over the years, many cornerstone technologies of modern society have been invented at Bell Labs and 8 Nobel Prizes have been awarded to its researchers

1925

1930's

NOBEL PRIZE

1937
Electron Diffraction
Demonstrating wave nature of matter

1940's

1948
"A Mathematical Theory of Communications"
By showing that all communications channels – of any type – have a fundamental capacity limit, Claude E Shannon founded the field of information theory

1950's

1954
Solar cells

NOBEL PRIZE

1956
Transistor
To replace the vacuum tube, Bardeen, Brattain and Shockley created a working point-contact transistor. This basic building block for all digital products is the foundation for our information society

1958
LASER
In their 1958 paper, Schawlow and his brother-in-law Townes described in detail a proof of concept for the LASER. The laser enables a wide variety of applications: fiber-optic communications, digital storage, barcode scanners, precision surgery and industrial cutting tools

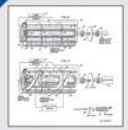

1960's

1962
Telstar
Transatlantic live TV broadcast via satellite

1970's

1973
UNIX and C Language
Thompson and Ritchie's elegant design made it an immediate hit with the programming community when it was released in 1974. UNIX would later on become the Internet's foundation

1976
Fiber Optic Network
First demonstration of 45 Mbit/s transmission

NOBEL PRIZE

1977
Electronic Structure of Magnets and Glasses

NOBEL PRIZE

1978
Cosmic Microwave Background Radiation
Pioneering work on radio communications using the Holmdel Horn Antenna provides support for the Big Bang Theory

1978
Commercial Cellular Network
Invention of the cellular concept and creation of the first commercial network

1980's

1980
Demonstration of DSP
Large-scale integrated circuit for digital signal processing

1995
Commercial DWDM
Pioneering work on wavelength multiplexing in optical fibers

1990's

NOBEL PRIZE

1997
Laser-Based Cooling and Trapping of Atoms
To understand the fundamental limits of materials and matter

NOBEL PRIZE

1998
Fractional Quantum Hall Effect
Discovery of a novel collective quantum fluid state of matter

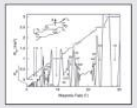

1998
Wireless MIMO Spatial Multiplexing
Invention of wireless transmission based on multiple spatial paths

1995
Integrated ADSL Chip
After co-inventing ADSL technology, follow-up innovations like vectoring continued to generate world records for high speed data transfer over copper telephone lines, fueling the Internet

2000's

2006
Software Defined Routing
Predecessor of Software Defined Networks (SDN)

NOBEL PRIZE

2009
CCD
Boyle and Smith's picture phone research realized the enormous potential of the Charge Coupled Device as an imaging device, leading to the invention of the digital photo, video cameras, scanners, satellite surveillance and ultra-sensitive astronomical telescopes

2009
Coherent 100G Optics
Invention of the future of high speed optical communications with coherent processing

2009
World's first standard compliant LTE call

2010's

2011
lightRadio Cube
First demonstration of building block of future small cell wireless networks

NOBEL PRIZE

2014
Fluorescence Microscopy
Ground-breaking work on sub-wavelength optical microscopy leads to super-resolution microscopy at cellular level

2014
XG-FAST
First demonstration of 10 Gbps over copper telephone wires

2015
GreenTouch
International consortium delivers new technologies to improve energy efficiency in wireless networks by more than 10,000X

2015
Optical MIMO-SDM
Pioneering work on utilizing the spatial dimension in fiber, showing greater than 10X increase in optical network capacity

2015
The Future X Network: A Nokia Bell Labs Perspective
First Nokia Bell Labs book written

2016
5G Massive Connectivity
First demonstration of 1M simultaneous, ultra-low latency connections in a single cell for 5G and IoT

The Future

Nokia Bell Labs continues to solve the great industry challenges, producing disruptive innovations for the next phase of human existence

Daniele Quercia

Eager to join forces on urban computing and responsible AI

Sanja Šćepanović

Let's make AI for Earth Observation and AI for Public Health responsible together

Marios Constantinides

Open to collaborate on projects about Responsible AI and Future of Work

Adam Ke Zhou

Keen on teaming up for ethical and transparent NLP collaboration

new hire

Ali Septiandri

Open to collaborative opportunities in the fields of Healthcare and Responsible AI

new hire

Edyta Bogucka

Seeking partners to make AI responsible through data visualization and design

Do you want to join us?

We welcome applications from outstanding PhD students and scholars to join us in 2024 as summer interns, self-funded visitors or project collaborators

Apply →

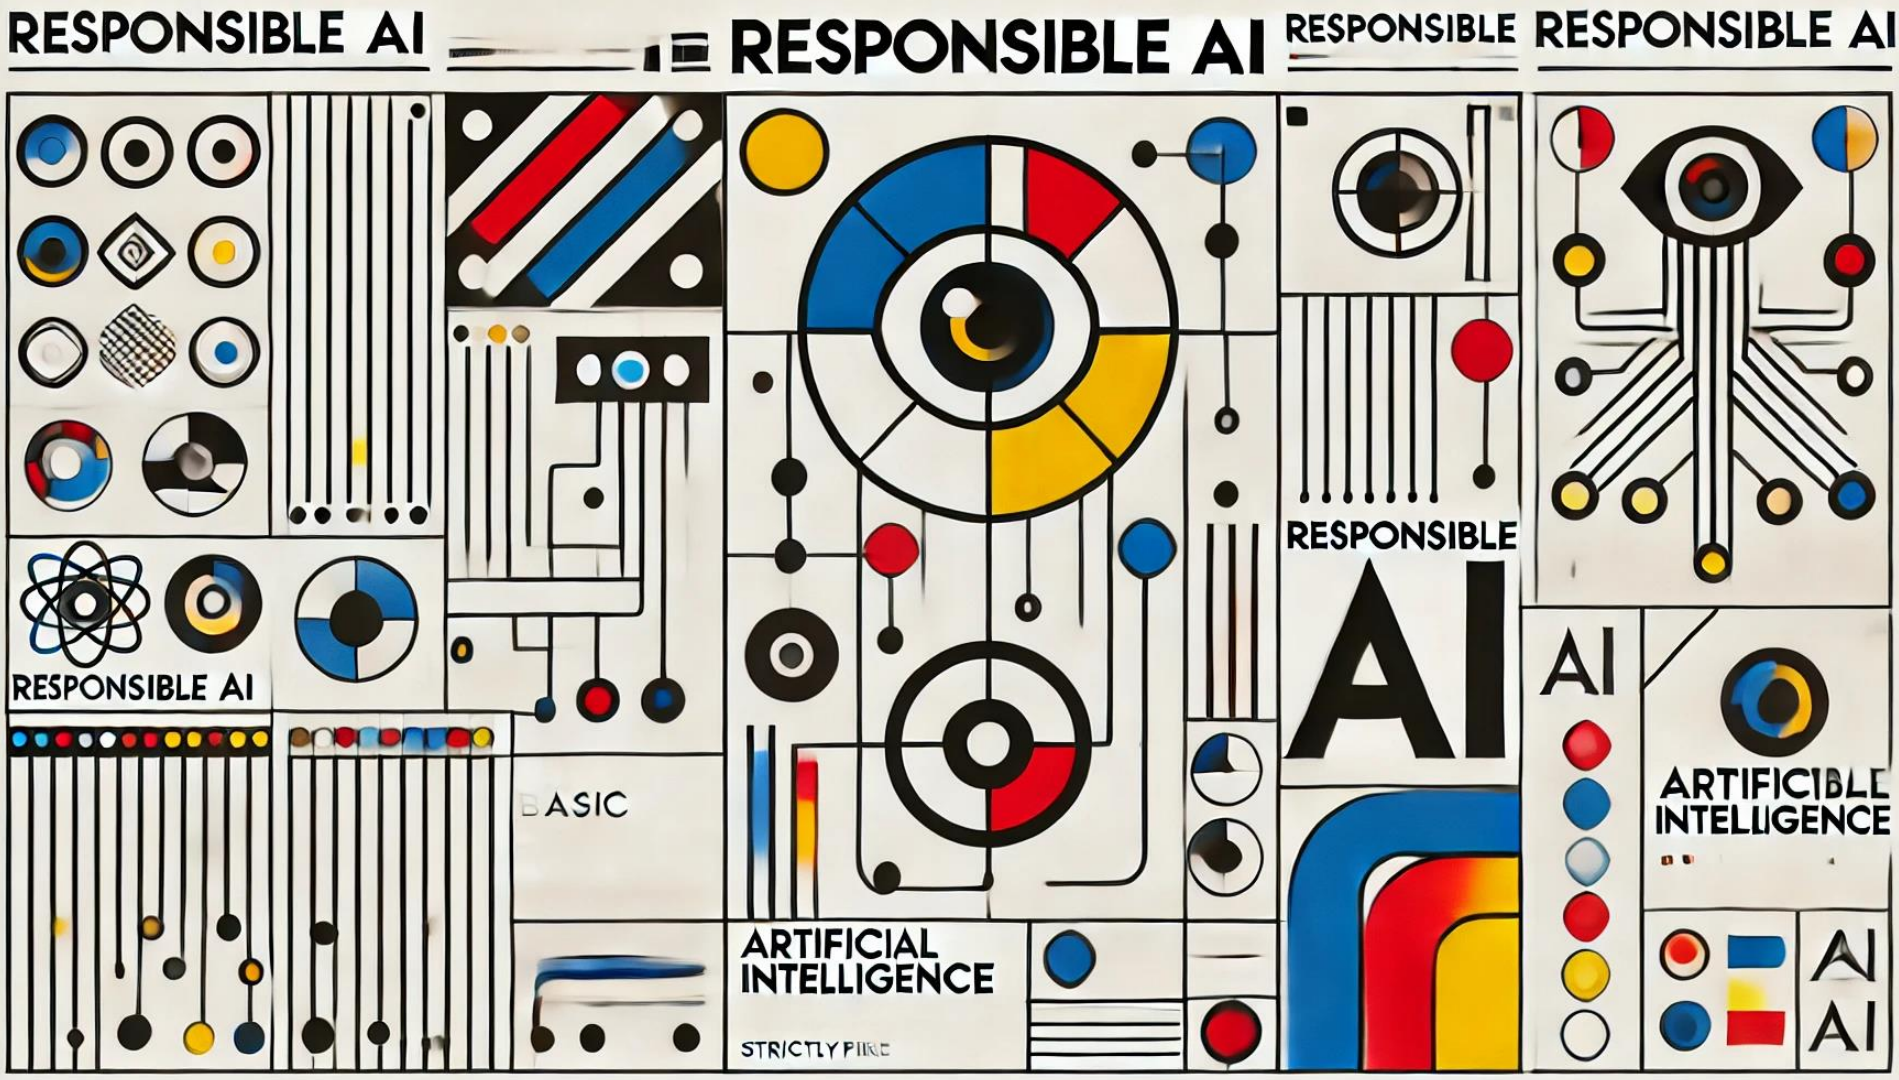
RESPONSIBLE AI
RESPONSIBLE AI
RESPONSIBLE
RESPONSIBLE AI
RESPONSIBLE AI
BASIC
ARTIFICIAL INTELLIGENCE
RESPONSIBLE
AI
AI
ARTIFICIBLE INTELLIGENCE
AI
AI

Data Science
& Machine Learning

Data Science & Machine Learning

What's This About?

Data Science and Machine Learning (ML) are how computers learn from data -like how your phone predicts your next text, but on a much larger scale. This course teaches you to **turn raw data into useful insights** and apply it to engineering problems.

What Will You Learn?

1. The **data science process** – collecting, cleaning, and making sense of data.
2. How to use **Python** and top **data science libraries** (because no one does this by hand anymore).
3. Key **machine learning and deep learning algorithms** – the things that make AI work.

Data Science
& Machine Learning

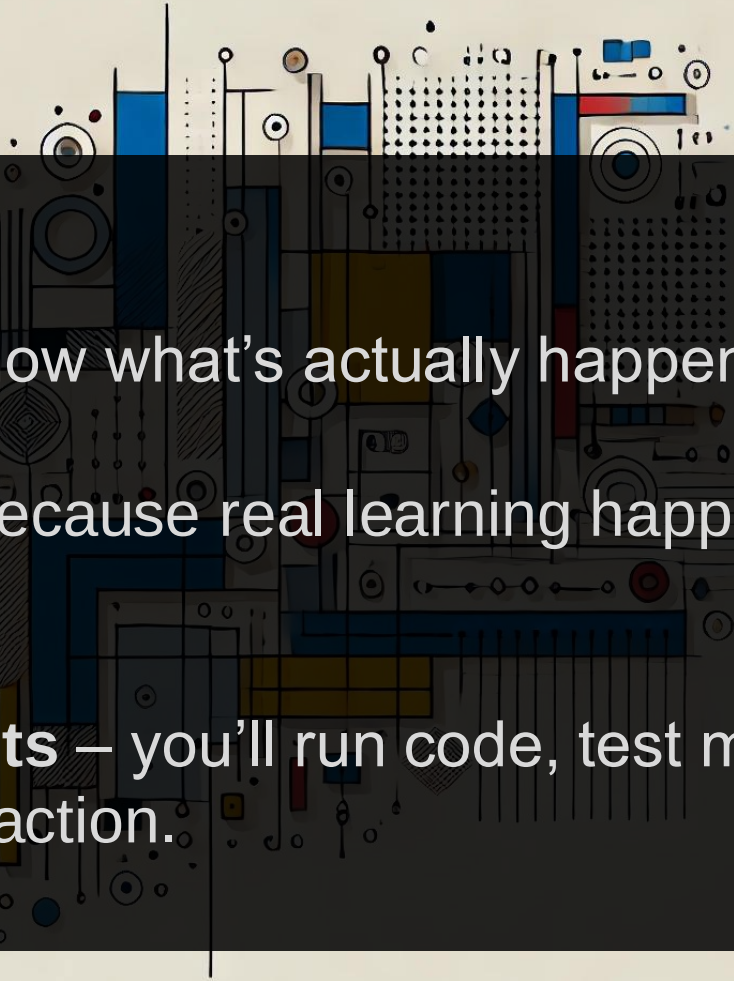

How's It Taught?

1. Theory – so you know what's actually happening.

2. Hands-on labs – because real learning happens when you *do* it.

3. Python experiments – you'll run code, test models, and see how ML works in action.

Data Science
& Machine Learning

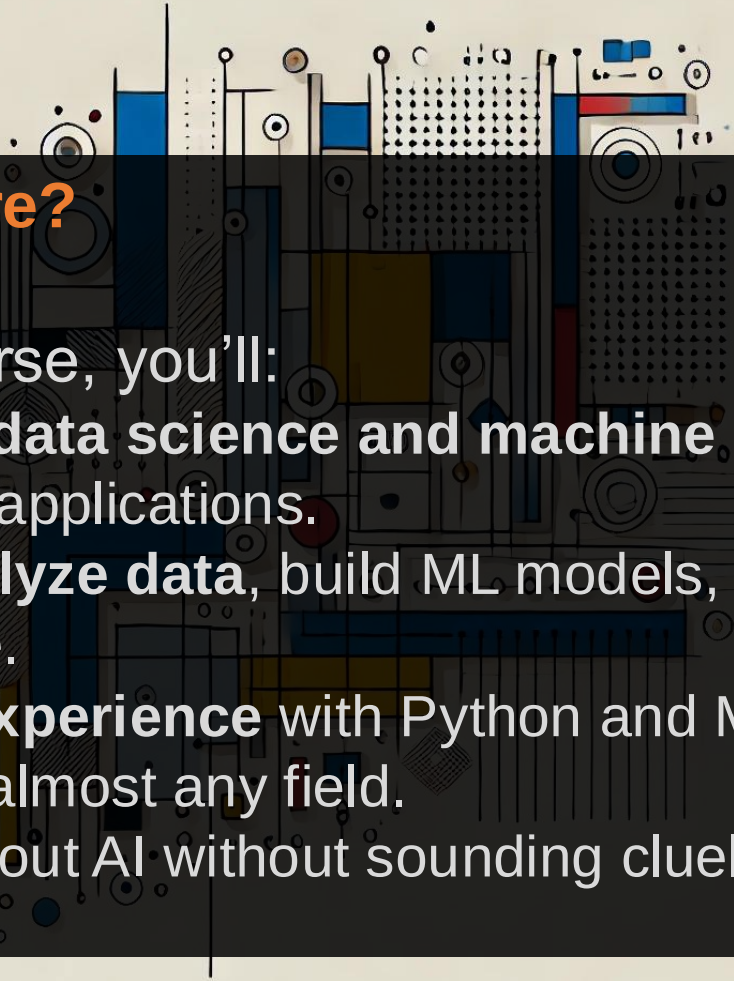

Why Should You Care?

By the end of this course, you'll:

1. Understand how **data science and machine learning** power real-world applications.
2. Know how to **analyze data**, build ML models, and evaluate their performance.
3. Have **practical experience** with Python and ML libraries—valuable skills in almost any field.
4. Be able to talk about AI without sounding clueless.

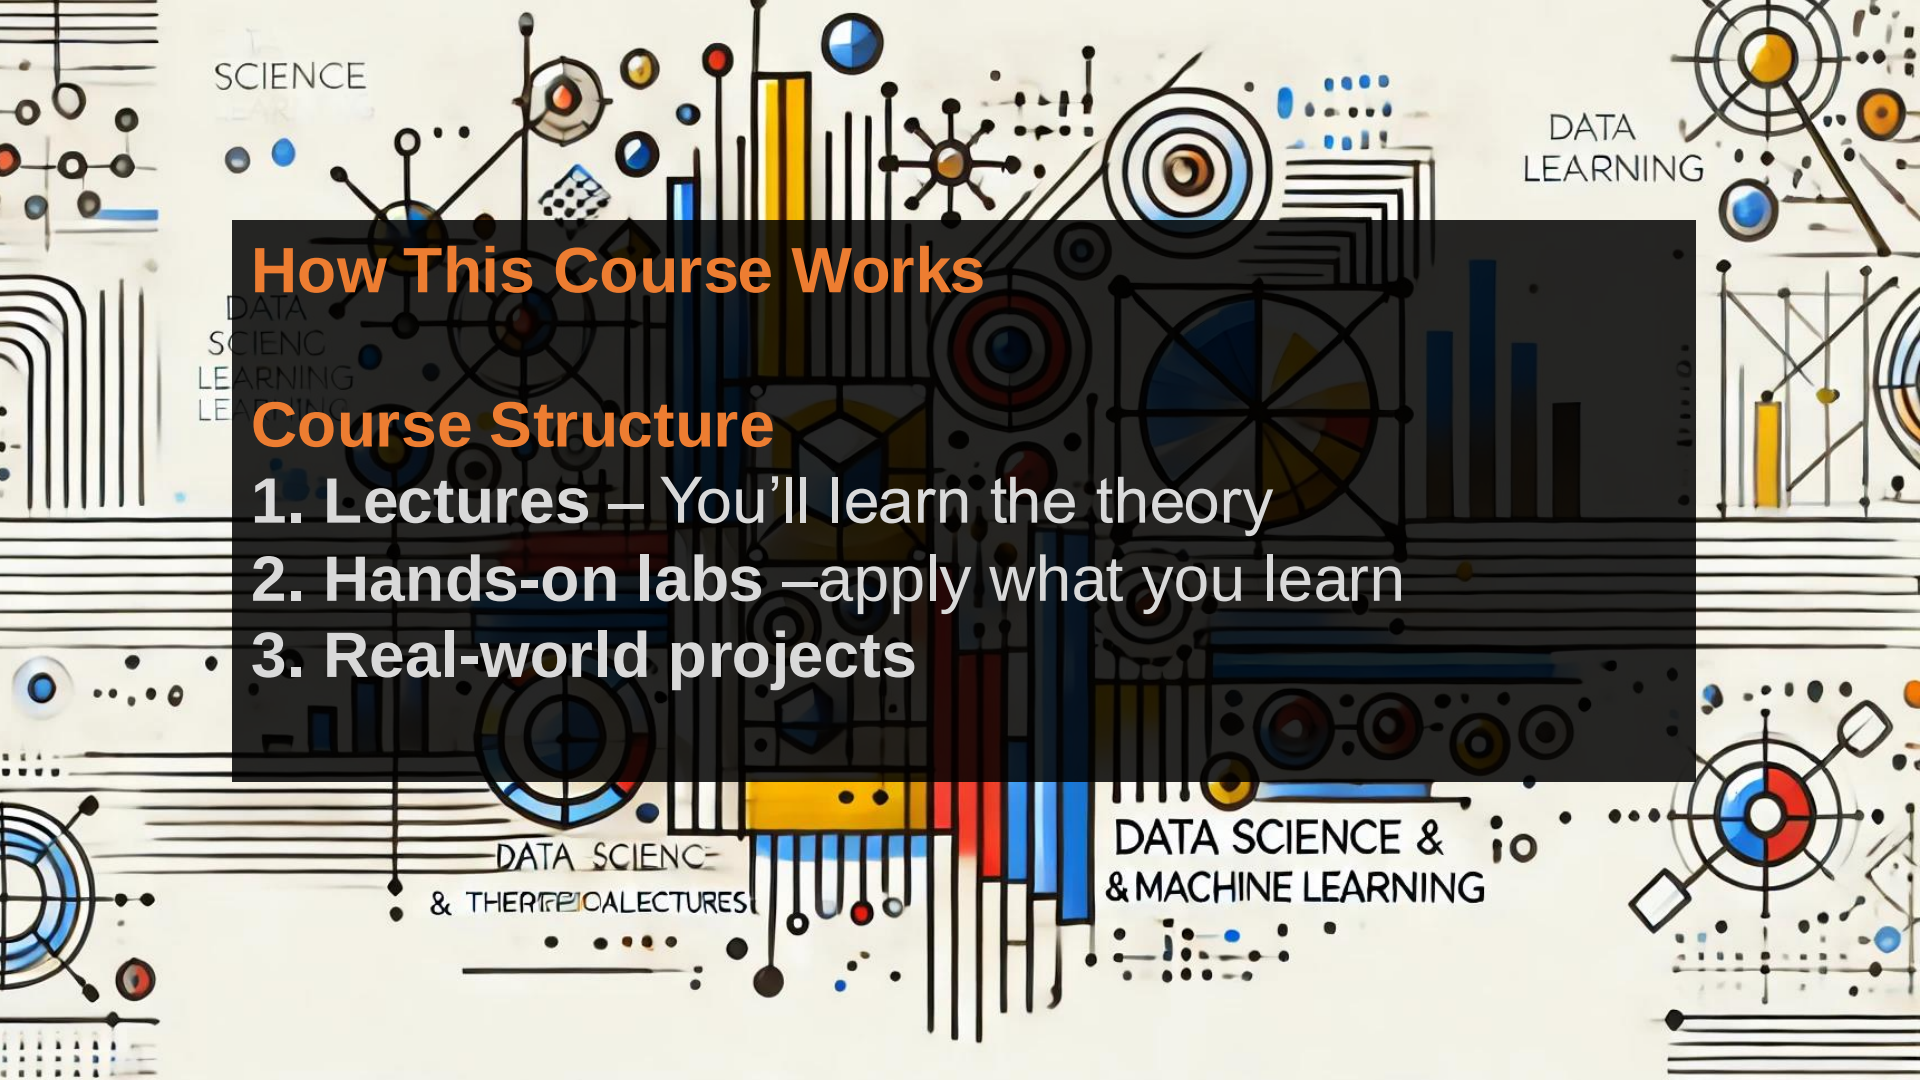

How This Course Works

Course Structure

1. **Lectures** – You'll learn the theory
2. **Hands-on labs** –apply what you learn
3. **Real-world projects**

How This Course Works

How You'll Be Graded

1. **Homework (4 points)** – Small exercises to make sure you can apply what you've learned.

2. **Group Project (20 points)** – Work with a team to build a full **data science process** for a real engineering problem. Your project needs to be accurate, well-documented, and not fall apart under scrutiny.

3. **Written Exam (10 points)** – Multiple-choice questions No notes, no phones—just you and your brain.

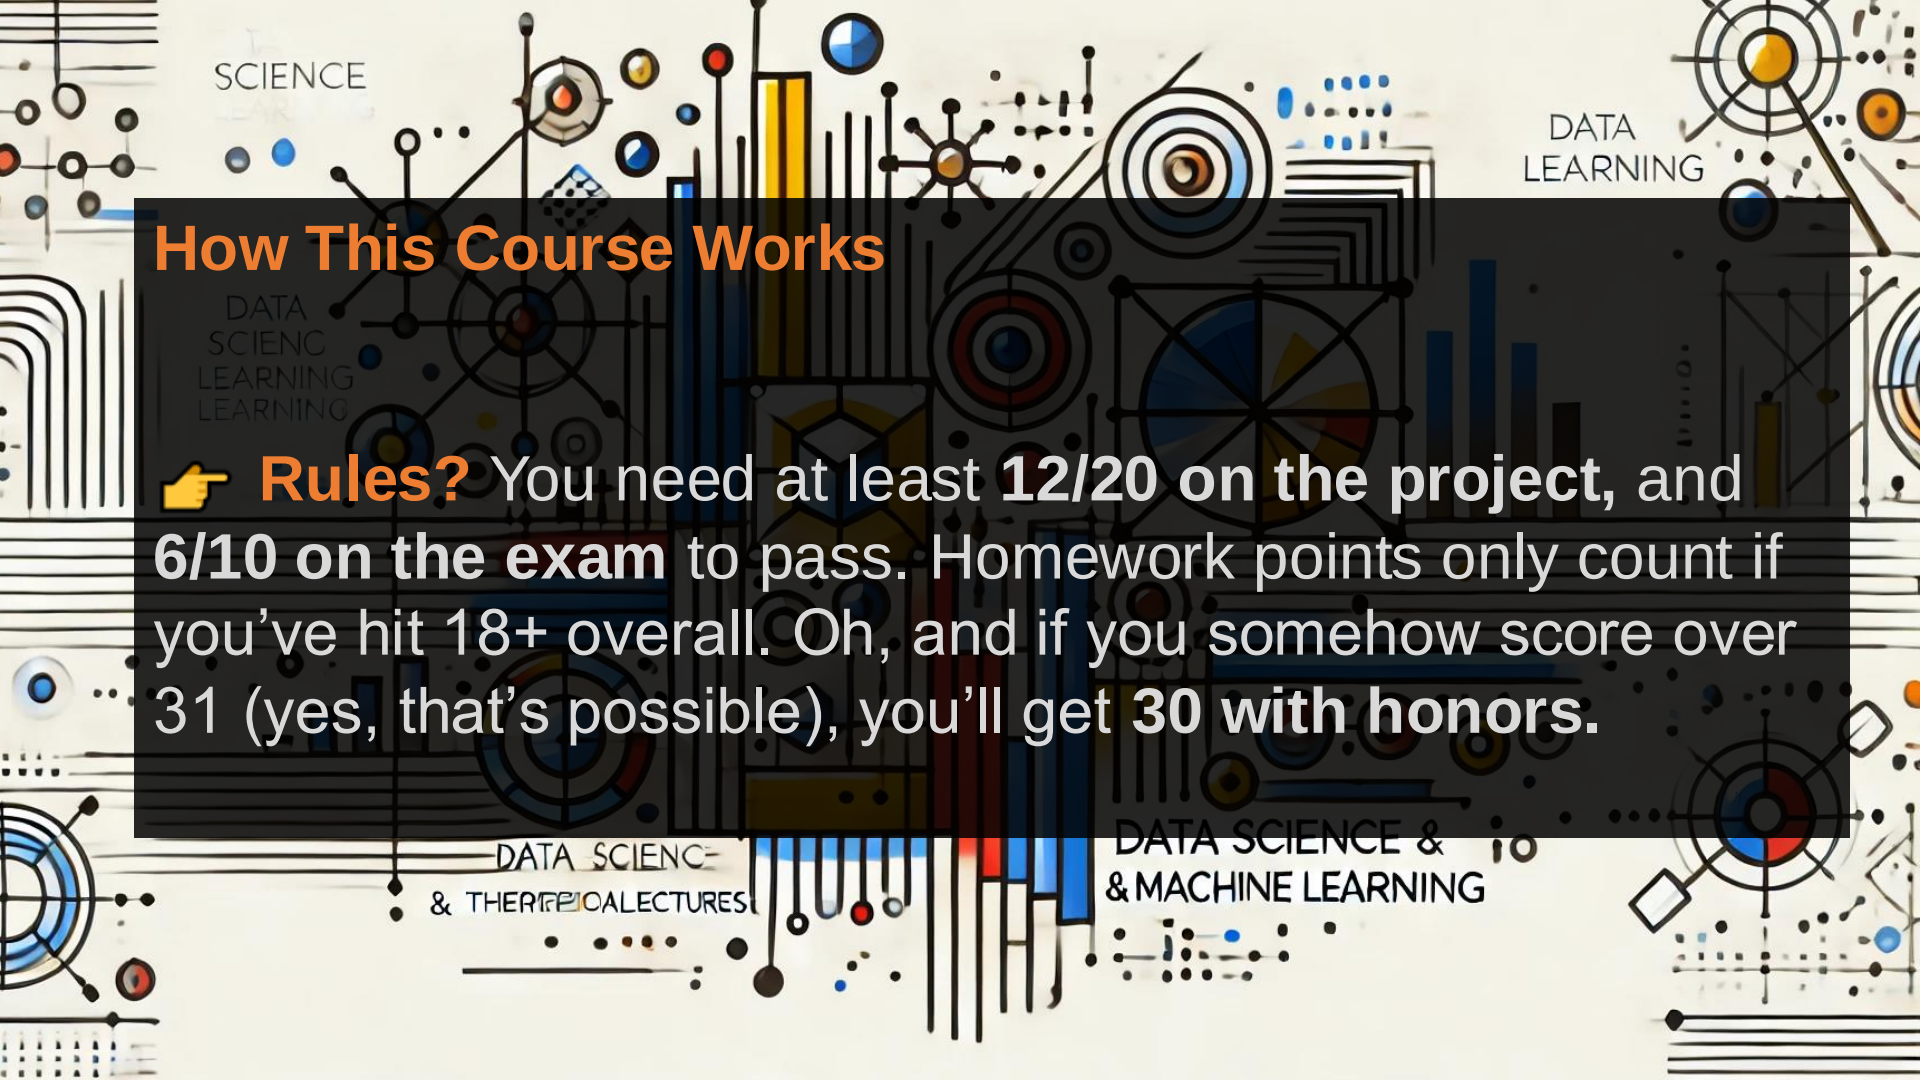

How This Course Works

👉 **Rules?** You need at least **12/20 on the project,** and **6/10 on the exam** to pass. Homework points only count if you've hit 18+ overall. Oh, and if you somehow score over 31 (yes, that's possible), you'll get **30 with honors.**

Course Topics: What You'll Actually Learn

This course is about learning how to **make data work for you**. Here's what we'll cover:

1/ The Data Science Process

A. How to collect, clean, and transform data (because raw data is a mess).
B. Feature engineering – picking the right details that actually matter.

Course Topics: What You'll Actually Learn

2/ Data Science Algorithms

A. **Classification** – Teaching machines to put things into categories.
B. **Clustering** – Grouping things together based on similarities (like Netflix recommending "quirky indie films" when all you watch is Marvel).
C. **Association rules** – Finding hidden patterns in data (think: "People who buy chips also buy salsa").

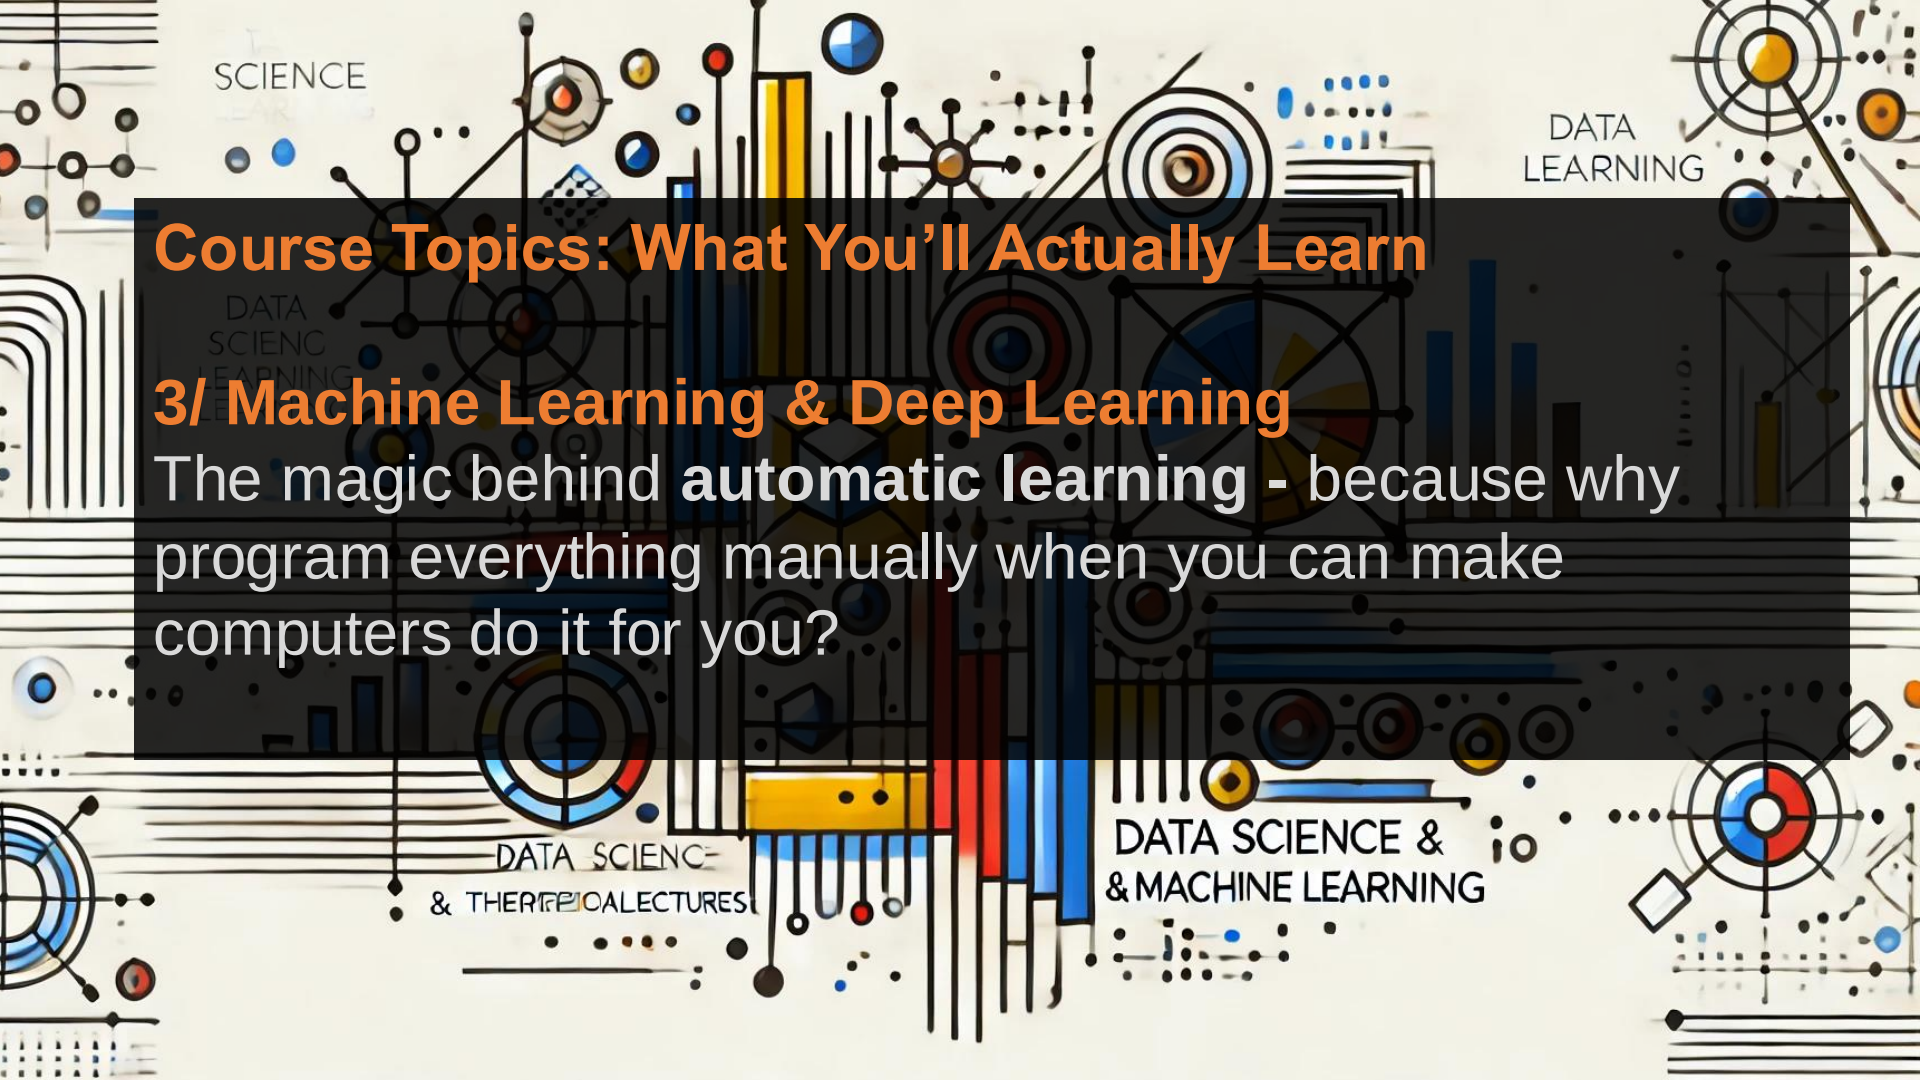

Course Topics: What You'll Actually Learn

3/ Machine Learning & Deep Learning

The magic behind **automatic learning -** because why program everything manually when you can make computers do it for you?

Course Topics: What You'll Actually Learn

4/ Python & Data Science Libraries

A. Learn **Python**, the language everyone in AI and data science swears by.
B. Use top libraries like **scikit-learn** - so you don't have to reinvent the wheel.

Course Topics: What You'll Actually Learn

5/ Real-World Applications & Case Studies

A. You'll design and build a complete **data science process**, applying ML & Deep Learning to real engineering problems.
B. Basically, you'll take what you've learned and make it actually *useful*.

How You'll Be Graded (a.k.a. Proving You Actually Learned Something)

The Breakdown. Your final grade is based on three things:

1. **Homework (4 points)** – Small hands-on tasks to practice **Python, data science, and ML algorithms**.

2. **Group Project (20 points)** – Work in a team to design and implement a full **data science process** for an actual engineering problem. You'll write a report explaining your choices and prove your model works.

3. **Written Exam (10 points)** – Multiple-choice questions covering **data mining, machine learning, and deep learning**.

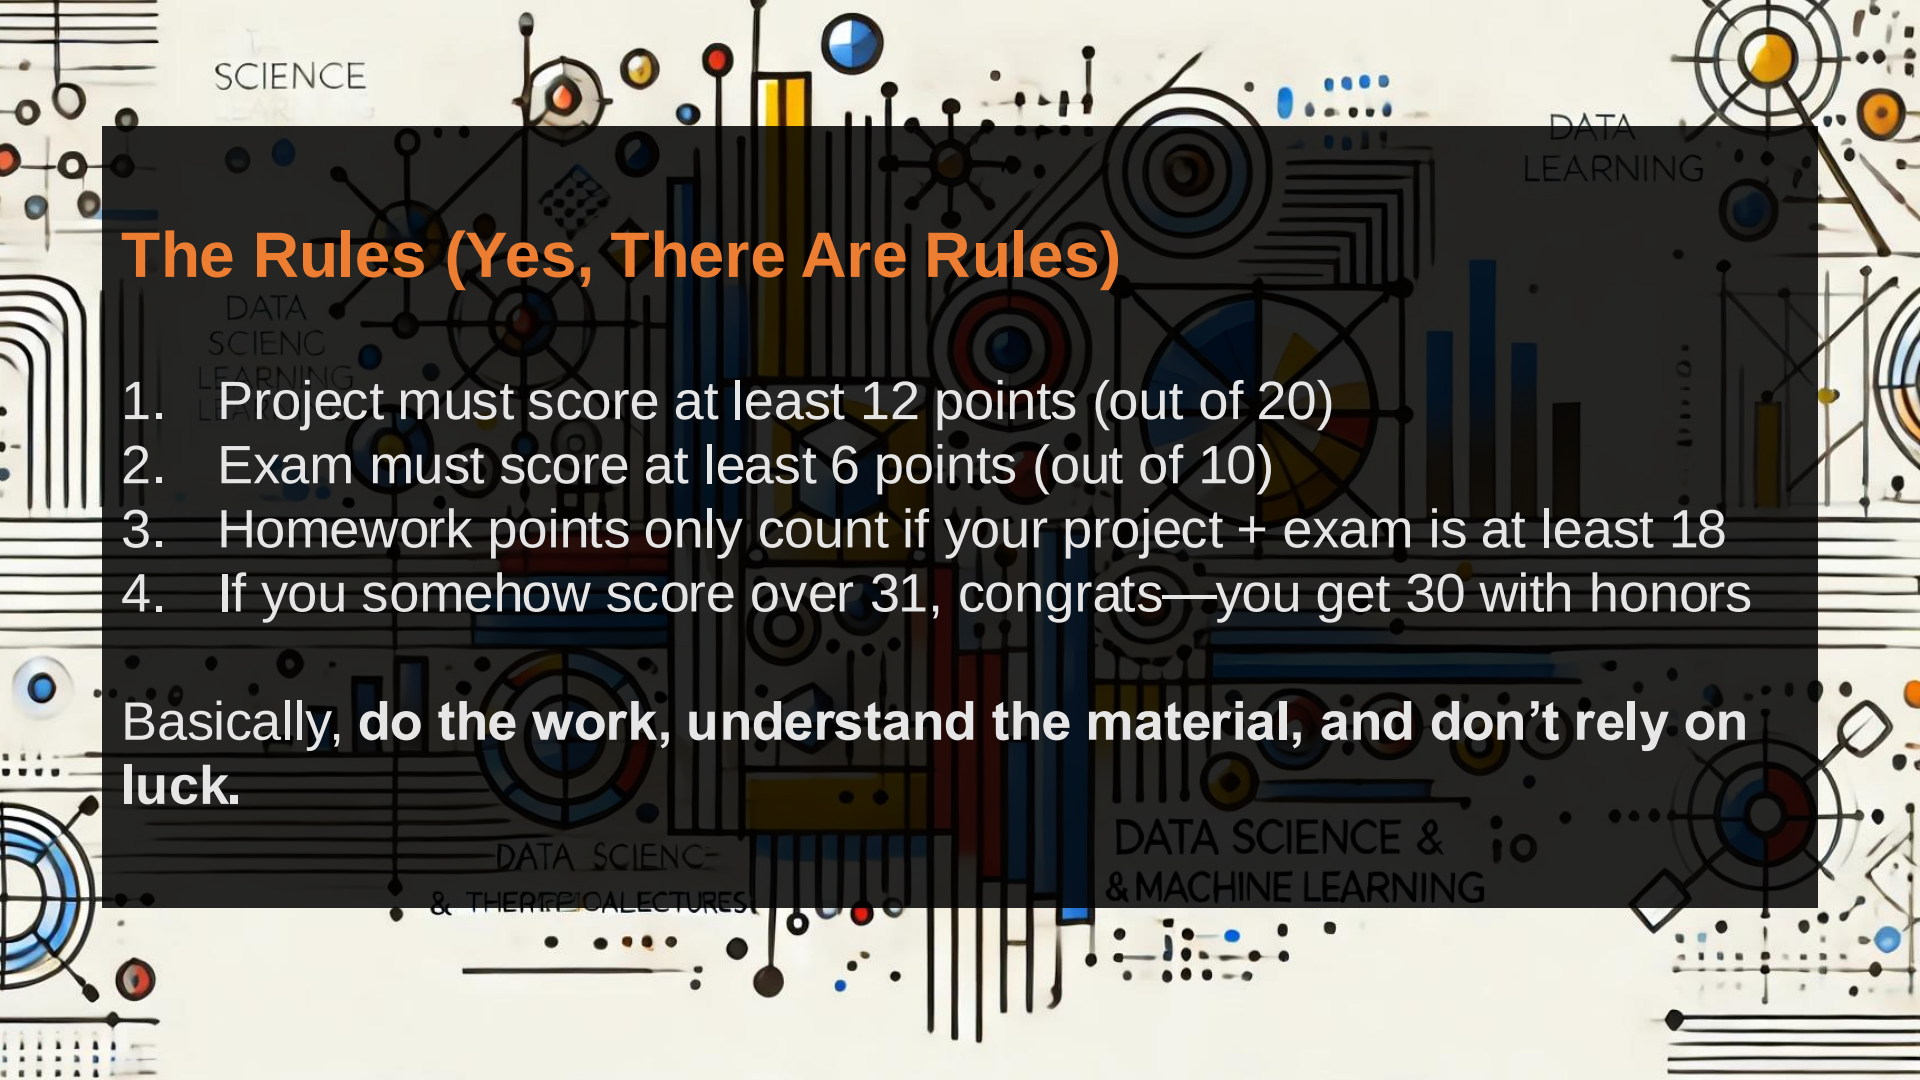

The Rules (Yes, There Are Rules)

1. Project must score at least 12 points (out of 20)
2. Exam must score at least 6 points (out of 10)
3. Homework points only count if your project + exam is at least 18
4. If you somehow score over 31, congrats—you get 30 with honors

Basically, **do the work, understand the material, and don’t rely on luck.**

What's Actually Being Tested?

1. Can you **design and implement** a **data science process**?
2. Do you know how to **evaluate ML algorithms**?
3. Can you use **Python and ML libraries** effectively?
4. Do you understand **key concepts in data mining and machine learning**?

In short: You're being tested on **what you can actually do**, not just what you can memorize.

Final Thought

If you put in the effort, you'll leave this course with **real skills that matter**. If you don't… well, machine learning won't learn itself.

Homework: Yes, You Actually Have to Do It

Homework is your chance to practice **Python, machine learning algorithms, and data science techniques**.

The Breakdown

1. **7 assignments** throughout the course.
   A. **First 6** = 0.5 points each (if submitted on time).
   B. **Final one** = 1 point (because it's tougher and involves **advanced ML on an image dataset**).
2. **Total possible points: 4** (not much, but enough to make a difference).

Homework: Rules (Because We Have to Have Some)

1. Deadlines matter - late submissions = no points.

2. Homework is directly related to lab work - so if you pay attention, it’s easy.

3. Points are only valid until January 2026 (included).

Homework: Why Bother?

1. You get **hands-on experience** with real-world ML problems.
2. You practice using **Python and ML libraries**—which, let's face it, is why you're here.
3. It's a **free 4 points** if you do the work.

Final Thought

Do the homework, and the final project/exam will feel a lot easier. Skip it, and well… good luck.

Group Project: Work Together

This is where you prove you can **apply data science and machine learning** to a real engineering problem - without breaking your group apart in the process ;)

The Assignment

1. **Teams of 6-8 students** (yes, you have to work with others - welcome to real life).
2. **Design and implement a full data science process** using ML algorithms.
3. **Write a report** explaining what you did and why.

Group Project: When & How?

1. Assigned after ~ **7-8 weeks**—so you actually know what you're doing.

2. Evaluated based on **performance, accuracy, completeness, and robustness** (basically, how good, well-documented, and stable your solution is).

3. The **better your model, the better your score**—so maybe don't just copy-paste random code from the internet.

Group Project: Grading (a.k.a. Why This Matters)

1. Worth **20 points**—so it's a big deal.
2. Your model must actually work **with real data**.
3. If your project score is **below 12, you fail**.

Final Thought

If you work well as a team, you'll build something impressive. If you don't… well, hopefully you made friends in the process.

The Written Exam

This is where we check if you actually understand **data science and machine learning**—or if you've just been nodding along the whole time.

What's on It?

1. **Multiple-choice questions.**
2. Covers **data preparation, classification, regression, clustering, and association rule mining**.
3. Includes **machine learning and deep learning for engineering applications**—so don't skip the lectures.

The Written Exam: Rules

1. **60-minute time limit**—so no time for daydreaming.
2. **No textbooks, notes, or electronic devices**—it's just you and your brain.
3. If you score **below 6/10, you fail**.

Final Thought

If you've been paying attention and doing the work, this will be fine. If not… well, at least it's only 60 minutes.


DATA SCIENCE
& MACHE EARNING
PYTHON
DATA SCIENCE
& MACHINE LEARNING
DATA SCIENCE
& MACHE LEARNING
& MACHINE LEARNING
PYHON