

The background of the slide is a detailed, close-up photograph of a complex mechanical device, likely a historical automaton or a large clockwork. It features numerous interlocking gears of various sizes, some with decorative patterns. The mechanism is made of dark, polished metal. In the lower right foreground, an open book with aged, yellowed pages lies flat on a wooden surface, possibly part of the machine's base. The lighting is dramatic, with strong highlights on the metallic surfaces and the book's pages, set against a dark, blurred background.

Large Language Models

Course
introduction

Flavio Giobergia
Riccardo Coppola

Welcome to *LLM4SE*!

- *Large Language Models for Software Engineering*
- *6 CFU* course
- Optional course, mainly for:
 - MSc in Computer Engineering
 - MSc in Data Science & Engineering

Course division

- (*p1*) 3 CFU
 - Theoretical introduction to LLM
 - Intro to deep learning, Transformers, Large Language Models
 - ft. Flavio Giobergia
 - Approx. from today to early November
- (*p2*) 3 CFU
 - *Practical applications of LLM to Software Engineering*
 - Intro to software engineering, using LLM in SE processes
 - ft. Riccardo Coppola
 - Approx. from early November to end of Semester

Learning objectives

- (*p1*) Theoretical foundations of LLM
 - Transformers architecture
 - LLM taxonomy and evolution
 - LLM training, fine-tuning, evaluation
- (*p2*) Practical applications of LLM to SE
 - Understanding of SE
 - Prompt engineering
 - Ethical & Societal impacts

Prerequisites

- *Core concepts of Machine Learning*
 - Today we will have an intro to deep learning
 - But, we assume that you have an understanding of basic ML concepts
- *Solid knowledge of Python and main libraries*
 - We will use PyTorch, but as long as you're familiar with NumPy, you will learn PyTorch just fine!

Timeslots



Mondays



8:30 – 11:30 (3h)



Room 2P



Wednesdays



11:30 – 13:00 (1.5h)



Room 3P

Classes' structure

- Classes will be mainly divided into:
 - *Lectures*
 - covering the main topics of the course in traditional lecture-style classes
 - *Workshops*
 - Hands-on sessions where you will either work on assignments, or follow the instructor in solving practical exercises

Final exam

- *Written test*

- 18 points
- Covers theoretical concepts of the course
- Closed- and open-ended questions
- 75 minutes

- *Group project*

- 15 points (mandatory)
- Application of LLM skills (e.g., to address SE-related activities)
- Written report + oral presentation
- Project assigned once, delivered once and valid for the entire academic year

Office hours

- No predefined office hours
- Send us an email to schedule an appointment!
 - flavio.giobergia@polito.it
 - riccardo.coppola@polito.it

Course materials

- Slides
 - Available for both parts of the course
 - Course website:
 - https://dbdmg.polito.it/dbdmg_web/2025/large-language-models-for-software-engineering-2025-26/
 - Uploaded before each lecture
 - Constitute the main reference for the course
- Lecture recordings
 - All lectures will be recorded & available online