



LLM PROGRAMME AI IN LEGAL PRACTICE AND ITS REGULATION

Curricular unit

AI and Law Module A

Responsible Academic staff and respective workload in the curricular

Il Memming Park, João Dias, Eric Lacosse, Francesco Trapani, Carolina Gonçalves

Syllabus

Lecture I (Francesco Trapani & Eric Lacosse): History of AI

We introduce the idea of intelligence, both natural and artificial, and discuss how the field of AI has evolved throughout history. We start with historical AI systems and move through biologically-inspired connectionist models, until today's data-driven deep learning systems. The goal is to understand how AI systems work in principle, how they evolved, and which technical properties—opacity, unpredictability, etc.— can create challenges for law and regulation.

Lecture II (Francesco Trapani): Neural Networks

We discuss the basic blocks of AI systems: artificial neurons. We start from the simplest learning system (1 neuron) and understand how it can learn using a feedback error signal. From there, we build up to small networks, and understand what it means to approximate a function. The goal is to understand the basics of the learning process, also keeping in mind which are the intrinsic limitations, constraints and biases.

Some of the concepts will be illustrated with concrete calculations using a Python notebook, for which I would advise to bring a laptop if you have one.

Lecture III (Eric Lacosse): Language Models and Cognitive Architectures for Language Agents

Language serves as a fundamental cognitive tool, transforming the complex "messy" continuous external world into the "neat" discrete internal world of human cognition. Agents are autonomous computational systems designed to perceive and act within their environment. The ultimate goal is to develop synthetic agents with human-level intelligence, capable of learning, remembering, representing knowledge, and exhibiting meaningful behavior across diverse tasks. Cognitive Architectures provide a model for the mind, aiming to build robustly helpful, honest, and harmless (HHH) AI systems for both individuals (micro-alignment) and societies (macro-alignment).

Lecture IV (Francesco Trapani & Eric Lacosse) Reinforcement Learning and Agentic AI

We explore AI systems that interact with the world through actions. We first discuss Reinforcement Learning (RL), the core paradigm for such systems. We discuss its basic components, the classical learning approaches, and the distinction between model-free and model-based learning.

Tutorial I (João Dias): Going down the slope to optimize

Tutorial II (Il Memming Park & Carolina Gonçalves): The basics of LLMs - words as vectors

Breaking down how Large Language Models work, what is word-masking and how words are represented in this space as embeddings.

Tutorial III (João Dias): How to train your AI

Exploring the practical concepts behind the processing of training and evaluating AI models, such as how model parameters affect outputs, what is regularization and understanding a loss curve.



Tutorial IV (Il Memming Park & Carolina Gonçalves): Agents in law

A session on how to leverage current state-of-the-art LLMs and agentic systems in the AI-augmented law environment.

Teaching methodologies (including evaluation)

The teaching methodology will be composed of a mixture of lectures and interactive tutorials. The course will be evaluated based on the results of the tutorials, sets of exercises or multiple choice questions.

Main Bibliography

An Introduction to Reinforcement Learning, Sutton & Barto (2018)

Sumers, et al., 2024 Cognitive Architectures for Language Agents <https://arxiv.org/abs/2309.02427>

Goodfellow, I., Bengio, Y., & Courville, A. (2016). Deep learning. MIT Press

Speech and Language Processing (3rd ed. draft), Dan Jurafsky and James H. Martin (2022,

<https://web.stanford.edu/~jurafsky/slp3/>)