

# AI in Healthcare

Spring 2024 - last updated: 10/21/2024

**Instructors:** Judith Abécassis ([judith.abecassis@inria.fr](mailto:judith.abecassis@inria.fr))  
**Lectures:** 8 sessions of 3 hours, including 4 labs to submit for grading + project defense  
**Course Website** [TBA](#)

## Course Overview

The healthcare sector has tremendous potential for transformation by using artificial intelligence. During this course, we will delve into the specific applications and challenges of AI in healthcare. We will examine how machine learning algorithms, natural language processing, and computer vision are being used to analyze medical data, improve patient outcomes, and optimize healthcare delivery. Additionally, we will address the ethical and practical considerations that arise when implementing AI in healthcare, such as data privacy, algorithmic bias, and the need for fairness and transparency.

## Course Goals

By the end of this course, you will master

- the different sources of medical data, and the associated constraints for access and use
- the bases for causal inference to obtain prescriptions rather than predictions
- the technical solutions of working with dirty data (noise, missing values, unstructured data)
- the ethical considerations of working with health data.

## Textbooks

The following books are recommended for general field knowledge but not required.

- E. Topol. *Deep Medicine: How Artificial Intelligence Can Make Healthcare Human Again* (2019)
- P. Lee, C. Goldberg, I. Kohane. *The AI Revolution in Medicine: GPT-4 and Beyond*. (2023)
- K. Crawford. *Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence* (2021)

## Evaluation

Students' final grades are based on three components:

- **Labs** (24% of the course grade - 4 labs). The lab report is to be submitted after finishing at home (one week delay to finish).
- **Article presentation or project** (76% of the course grade). A list of possible articles will be given, or you can propose an article of your choice, and work on it after validation from the teaching team

## Computing

This course will involve practical "Lab" sessions with coding in Python. A Jupyter notebook will introduce the topic and the main

## Schedule

### Lecture 1: A journey in the biomedical data world

- legal definition, anonymization and pseudonimization
- Electronic health records
- claim data
- natural text
- medical imaging
- other signals

### Lecture 2: Challenges of AI in healthcare, from prediction to prescription (1/2)

- framing a question with a causal vocabulary
- required assumptions
- overview of estimation approaches
- practical examples and pitfalls

### Lecture 3: Challenges of AI in healthcare, from prediction to prescription (2/2)

Lab 1 Run a causal analysis on health data

### Lecture 4: Working with missing data

- missing data mechanisms
- multiple imputation
- supervised learning with missing values

### Lecture 5: Working with censored data (Survival analysis)

Lab 2 Conduct a survival analysis

### Lecture 6: Working with biomedical texts

Lab 3 Extract relevant information from text

### Lecture 7: The stakes of population imaging (by Bertrand Thirion, Inria)

- variable importance problems
- confounds handling
- analysis pipeline comparison

### Lecture 8: ethical considerations, biases, fair learning

Lab 4 Improve the fairness of your predictions

### Lecture 9: Project defenses