

# AI in Medicine: From Models to Practice

**Biomedical Research Foundation of the Academy of Athens  
MIT Club of Greece**

January 26-28, 2026 | Instructor: Hara Moraitaki, MIT Department of Electrical Engineering and Computer Science

## Program Overview

Artificial Intelligence is rapidly transforming medicine – from improving diagnostic accuracy to enabling early prediction of critical conditions and accelerating drug discovery. Yet, for these technologies to truly benefit patients, clinicians and researchers must understand not only how these models work, but also their limitations, risks, and the conditions under which they fail.

This program combines lectures, hands-on workshops, and project work to give participants both inspiration and practical skills. It is designed not only to showcase what is possible with AI today, but also to spark participants' curiosity and encourage them to pursue further learning opportunities – ensuring that the Greek biomedical community continues to build expertise in responsibly developing and applying AI in healthcare, bringing Greece to the forefront of current innovation.

## What You Will Learn

By the end of this workshop, you will be able to:

- **Understand** how machine learning models work for clinical prediction tasks
- **Run and modify** existing ML pipelines using guided Jupyter notebooks
- **Evaluate** model performance using metrics like AUC, precision, recall, and calibration
- **Critically assess evaluations**: understanding why good benchmark performance doesn't guarantee real-world success
- **Identify fairness concerns** and check model performance across demographic subgroups
- **Use AI coding assistants** to help write and debug code
- **Recognize** when AI tools can (and cannot) be trusted in clinical settings

# What to Bring for the Training & Prerequisites

## Required

- Laptop with a modern web browser (Chrome recommended)
- Google account (for Google Colab access)

## Helpful but Not Required

- Basic familiarity with Python syntax (variables, functions, loops)
- Experience with data analysis in any language (R, MATLAB, etc.)
- Background in biology, medicine, or healthcare research

## About the Instructor

**Hara Moraitaki** is a senior in the Department of Electrical Engineering and Computer Science at the Massachusetts Institute of Technology (MIT) and an AI researcher focused on applications of machine learning in medicine and biomedical research. She has conducted research in AI for health at MIT's Computer Science and Artificial Intelligence Laboratory (CSAIL) and in collaboration with Mass General Brigham, working with real-world clinical data. Her work focuses on building and critically evaluating predictive models, with particular emphasis on fairness, generalization across populations, and responsible deployment in medical settings. In parallel, she has led and participated in applied AI-in-health initiatives, including organizing MIT's first Women's Health AI hackathon, alongside teaching at MIT and at high schools in the Cambridge area, and pursuing an education concentration that includes coursework at the Harvard Graduate School of Education.

# Program Schedule

## January 26th: Opening Lecture and Networking

Time: 15:00-17:00

An introduction to the current state of AI in medicine, with real examples from US hospitals. We will explore both the remarkable successes and the important limitations of these technologies.

The AI Revolution in Healthcare: Where are we today?
Case Studies: Diagnostics, ICU care, drug discovery – successes and failures
Why Models Fail: The evaluation trap – when benchmarks don't match reality
The Greek Opportunity: Why local AI development matters
Q&A and Discussion

## January 27th and 28th: Hands-On Training

Time: 15:00-17:30

A guided, practical introduction to building and evaluating ML models. No advanced coding experience required, as we will work through pre-built notebooks together, using AI assistants to help with code modifications.

<b>Module 1</b>	<b>ML Pipeline and Data Setup</b> Setup Google Colab, load the breast cancer dataset, understand features and labels, train-test splits, run a classifier
<b>Module 2</b>	<b>Model Comparison &amp; Evaluation</b> Compare Logistic Regression, Random Forest, XGBoost. Learn ROC curves, AUC, precision/recall, confusion matrices. Introduction to calibration: when can you trust predicted probabilities?
<b>Module 3</b>	<b>Fairness &amp; Model Interpretability</b> Check performance across subgroups. Introduction to SHAP values: why did the model make this prediction? Discussion: Is your test set representative of the real world?
<b>Module 4</b>	<b>AI-Assisted Coding Demo</b> Live demonstration: using Claude/ChatGPT to modify code, debug errors, and extend analyses. Practical tips for working with AI coding assistants.

*Note: Modules are subject to change.*

## January 28th: Closing Lecture

<b>Closing Lecture: Transfer Learning &amp; The Future of AI in Medicine</b> How AI models trained in one setting can (or cannot) adapt to new hospitals, populations, and countries, emphasizing challenges of generalization, fairness, and domain shift.
<b>Resources &amp; Next Steps</b> Where to continue learning: courses, datasets, communities.
<b>Discussion</b>

