

AI beyond Fit-Predict

Module 3, 2023–2024

Ivan Stelmakh

NES

istelmakh@nes.ru

Course information

Instructor's Office Hours: Friday 19:00 – 20:00. Additional slots are available by appointment

Class Time: M 10:00 – 11:30 & Th 10:00 – 11:30

Room Number: 401

TAs: TBD

Course description

This course complements ML and DL courses to help students develop a broad perspective on AI and its applications to real-world problems. The content of the course is divided into two parts:

Part 1. AI beyond ML

In the first part of the course, we will familiarize ourselves with big ideas in the field of AI and apply them to real-world problems: optimal pathfinding, planning, experimentation, hypothesis testing, and more.

Part 2. ML in Social and Business Context

In the second part of the course, we will focus on the context that defines the application of ML solutions in practice. Where do the data come from? How to ensure user privacy and train models on sensitive data? How to protect models from manipulations? The answers to these questions will prepare students for applying ML in real life.

Course requirements, grading, and attendance policies

Students should be comfortable with Python and basic mathematics. There are no other formal prerequisites for the class. Your final grade will consist of the following:

Component	Description	Weight
Class attendance and participation	There will be quizzes in each class. The main point of quizzes is to check that the content of the lecture is understandable, but participation (and not correctness) will be graded.	10%
Homework assignments	Every week there will be an assignment that will ask you to: <ul style="list-style-type: none">prove some properties or interesting resultscome up with a solution to a practical problem and implement it in code	65%

	These assignments should be submitted through my.nes by the assigned deadline and are subject to the late submission policy (below)	
Final exam	There will be a written in-class final exam	25%

Class attendance and participation. Attendance is required and will be taken in each class through quizzes. You must be in class for your answers to be accepted (i.e., it won't work if your friend just shares a link with you).

Late submission policy. All deadlines are hard, but you have 5 slip days that you may use as you wish (i.e., submit one assignment 5 days late or 5 assignments 1 day late).

Final grade. Grade for each component will be normalized to 0-100 scale and then averaged with given weights. The final grade is determined by the formula:

$$\text{grade} = (\text{score} + 5) // 10$$

Course contents

Block	Topic
Search and planning	Practical search: A^* and Θ^*
	Planning I: Linear Programming, Simplex Method
	Planning II: Integer Linear Programming
	<i>Application: Peer Review & Kidney Exchange</i>
Intro to Reinforcement Learning	Learning from experts: No Regret
	Bandits
	<i>Application: experiments and A/B testing</i>
Data and People	Learning from people I: Aggregation
	Learning from people II: Incentives
	Differential Privacy
	<i>Application: Privacy-preserving Machine Learning</i>
Fairness, Accountability and Transparency	Model interpretability
	Adversarial Attacks
	Robustness
	Fairness
	<i>Application: ML in high-stakes applications</i>

Description of course methodology

The class will use a combination of resources, surrounding AI and its applications to the real-world challenges. The lectures will be accompanied with reading materials and additional links to papers. We will use Python to implement some algorithms.

Sample tasks for course evaluation

There will be both practical and theoretical tasks.

Example of practical tasks:

- Implement the A* algorithm and compare different heuristics
- Solve an ILP instance using CVXPY

(the actual assignments will provide necessary additional details)

Example of theoretical tasks:

- Prove that assignment problem can be solved in polynomial time
- Analyze robustness properties of an ML algorithm

(the actual assignments will provide necessary additional details)

Course materials

Required textbooks and materials

We will not use textbooks. All materials and links to articles will be provided.

Academic integrity policy

Cheating, plagiarism, and any other violations of academic ethics at NES are not tolerated.

Support statement

Academic life may be sometimes stressful, but make sure you prioritize your health and well-being. If you have any concerns or need any help, contact instructors and stay happy 😊