

SYLLABUS

Books

1. AIMA - **Artificial Intelligence: A Modern Approach**, by Stuart Russell, 4th edition, 2021 and also [here](#).. This book is not free and is required.
2. GERON: **"Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow"**, 3rd Edition, 2022. This is free for both NJIT and NYU students, and very useful for those new to Python. It is TF2/Keras based and for those that want an equivalent for Pytorch you can consult the open-source **Dive2Deep Learning** book.
3. DL - **Deep Learning**. This book goes into the necessary depth required for this course especially on all sections that require the development of a statistical learning machine.

Planned Schedule [↗](#)

Part I: Perception and Machine Learning

Lecture 1: We start with an introduction to AI and present a systems approach towards it. We develop a map that will guide us through the rest of the course as we deep dive into each component embedded into *AI agents*. Reading: AIMA Chapters 1 & 2.

Lecture 2: The perception subsystem is the first stage of many AI systems including our brain. Its function is to process and fuse multimodal sensory inputs. Perception is implemented via a number of reflexive agents that map directly perceived state to an primitive action such as regressing on the frame coordinates of an object in the scene. We present *the supervised learning problem* both for classification and regression, starting with classical ML algorithms. Reading: AIMA Chapter 19.

Lecture 3: We expand into *Deep neural networks*. DNNs are developed bottom up from the Perceptron algorithm. MLPs learn via optimization approaches such as Stochastic Gradient Descent. We deep-dive into back-propagation - a fundamental algorithm that efficiently trains DNNs. Reading: AIMA Chapter 21 and DL Chapter 6

Lecture 4: We dive into the most dominant DNN architecture today - *Convolutional Neural Networks (CNNs)*. Reading: DL Chapter 9 & 10 and AIMA Chapter 25 (in part).

Lecture 5: When agents move in the environment they need to abilities such as *scene understanding*. We will go through few key perception building blocks such as Object Detection, Semantic and Instance Segmentation. Reading: AIMA Chapter 25 (in part).

Lecture 6: In this lecture we introduce probabilistic models that process the perceptive predictions over time and understand how the agent will track / update its time-varying belief about the state of the environment. This is achieved with recursive state estimation algorithms acting on dynamic bayesian networks. This lecture introduces Bayesian filters in discrete and continuous state spaces (Kalman filters). All robots one way or another employ such filters. Reading: AIMA Chapters 12, 13 & 14.

Part II: Natural Language Processing

Lecture 7: NLP is the pinnacle of applied AI in every day life - we are all using natural language as the prime means of communicate between us and increasingly between us and robots. In this lecture we pose the NLP problem, understand its components and their mechanics. AIMA Chapter 23.

Lecture 8: We then talk extensively about *language modeling* and start with an approach based on the [RNN / LSTM architecture. The later is used far beyond language modeling and expands into every application that involves sequences. We introduce the concept of *attention* and go through the Transformer framework - perhaps the most successful architecture in NLP today. AIMA Chapter 24 and DL Chapter 10.

Part III: Reasoning and Planning without Interactions

Lecture 9: Our ability to track and and predict the state of the environment is now supplemented by symbolic representations. Knowledge representation and reasoning is the core of the *symbolic subsystem* of AI agents. We will go through the basic building blocks of knowledge representation and reasoning such as propositional logic that allow the agent to evaluate using theorem proving the truth value of logical sentences using an incrementally growing Knowledge Base. Reading: AIMA Chapters 7.

Lecture 10: After the last lecture, the agent has a clear view of the environment state such as what and where the objects that surround it are, its able to track them as they potentially move. It needs to plan the best sequence of actions to reach its goal state and the approach we take here is that of *problem solving*. In fact planning and problem solving are inherently connected as concepts. If the goal state is feasible then the problem to solve becomes that of *search*. For instructive purposes we start from simple environmental conditions that are fully observed, known and deterministic. This is where the A* algorithm comes in. We then relax some of the assumptions and treat environments that are deterministic but the agent takes stochastic actions or when both the environment and agent actions are stochastic. We also investigate what happens when we do not just care about reaching our goal state, but when we, in addition, need to do so with optimality. Optimal planning

under uncertainty is perhaps the cornerstone application today in robotics and other fields.

Readings: Reading: AIMA Chapters 3 & 4 (problem solving) and 11.

Part IV: Planning with Interactions - Reinforcement Learning

Lecture 11: We now make a considerable extension to our assumptions: the utility of the agent now depends on a sequence of decisions and, further, the stochastic environment offers a feedback signal to the agent called *reward*. We review how the agent's policy, the sequence of actions, can be calculated when it fully observes its current state (MDP) and also when it can only partially do so (POMDP). We conclude with the basic taxonomy of the algorithmic space for RL problems.

Readings: AIMA Chapter 16 & 17.

Lecture 12: The algorithms that learn optimal policies in such settings are known as Reinforcement Learning (RL). In this lecture we establish the connection between MDP and RL, by introducing the Bellman expectation backup and Bellman optimality equations. We then use these equations to derive the policy iteration algorithm that is behind the policy-based REINFORCE algorithm that is empowered by approximating the policy function using the Deep Neural Networks that we met in the perception subsystem. AIMA Chapter 22.

Lecture 13:: Review - last lecture before the final.