

CSCFUA 9472 Artificial Intelligence - Syllabus

ARTIFICIAL INTELLIGENCE

● OVERVIEW

There are many tasks that humans can do easily and almost unconsciously but that have proven extremely difficult to program on a computer. Artificial Intelligence (AI) aims at developing computer systems that can carry out some of these tasks. This body of knowledge and techniques at the interface between mathematics and computer science has made constant progress in recent decades and will be at the origin of many technological innovations with a strong scientific, societal or economic impact. Computer vision or the automatic recognition of speech or handwriting are successful applications that are the most frequently mentioned to illustrate the efficiency of AI solutions. These successes are being presented as harbingers of a revolution that will spare no field of human activity. Expectations are high. AI is supposed to allow for the development of a personalized medicine that will adapt a treatment to the patient's genetic traits. It is to be used to design systems of predictive maintenance for complex infrastructures, such as electricity grids. It would help make aircraft with systems for the early detection of "weak" signals that announce breakdowns, and would thus serve to plan the replacement of components before their probable failure. The vehicles using AI would be safer and fully autonomous, and be in service longer. This is even more so with machine learning, the branch of AI that seeks to develop techniques for automatically analyzing masses of data, usually for predictive purposes. Motivated at the origin by problems related to the recognition of shapes and stimulated by the cognitive sciences, machine learning does not serve just to process data and adjust the parameters of a more or less rigid, predefined model. It is used to design algorithms that "learn" automatically, from examples, which model in an immense catalog of models will be the most efficient. The efficiency of learning algorithms (*e.g.*, "neural networks", "support vector machines" or "random forests") and the development of optimized software for implementing them have popularized the use of AI over the past two decades. Machine learning draws from several fields of applied mathematics : in particular and of course, from Probability, since it provides the right language for describing data variability and formulating predictive problems, Statistics, Optimization and Image/Signal

Processing among others. AI is a very vast interdisciplinary domain and the problems it can be used to tackle are extremely diverse. Undeniably, they cannot be covered through a single course... The main purpose of this class is to familiarize you with some of the fundamental mathematical and algorithmic concepts used to develop intelligent systems. We will focus on several aspects of AI : machine learning, the main paradigm of AI research now, including reinforcement learning (i.e. situations where the machine/agent interacts with its environment) and some of its popular applications such as computer vision, biometrics, recommending systems and natural language processing (NLP), certain crucial notions of symbolic AI as well (e.g. search, logic and automated reasoning) to design knowledge-based systems. We will underline the importance of representations (of data/knowledge) in AI approaches and will explain that the level of “delegation” to be granted to “smart systems” in the near future will heavily depend on how research in AI methodology will reply to questions of ethics (the processing of personal data while respecting privacy, fairness) and reliability/robustness.

● OBJECTIVES

The purpose of this class is to familiarize you with certain crucial notions of AI through the presentation of some major problems (e.g. Pattern recognition, Q-learning, unsupervised learning, collaborative filtering) and popular use cases (e.g. facial recognition, automated medical diagnosis, fraud detection, recommender systems, virtual assistants, the analytics of opinions or feelings). Although focus will be on practice definitely, a certain amount of formalism cannot be avoided however, in particular to gain insight into concepts such as complexity, model under/overfitting, generalization capacity and to describe the problems and the methods in AI. The knowledge of basics in Probability/Statistics, Linear Algebra and Mathematical Analysis is advisable but not mandatory. Regarding the implementation of AI methods, the Python language will be used. However, this is not a coding class and softwares such as <https://scikit-learn.org> will be used extensively.

● AGENDA

May 23, 24 General introduction : scope, basic definitions, historical perspectives, agenda, organization, evaluation - Getting started with scikit-learn – The (monolayer) Perceptron.

May 25, 26, 30 – Pattern recognition : formulation, use cases (e.g. handwritten character recognition, medical diagnosis), first algorithms (linear discriminant analysis, decision trees,

linear SVM, nearest neighbours, neural nets), concepts of statistical learning (training/test, complexity, model selection)

May 31, June 1, 2, 6 – Towards more predictive approaches : advanced learning algorithms (ensemble learning, deep nets, kernel methods) – applications to biometrics

June 7, 8, 9 – Recommending systems – Preference learning – Collaborative Filtering

June 13, 14 15 - Unsupervised learning, 'learning without a teacher' : clustering, blind source separation, novelty/anomaly detection, autoencoders

June 16 Reliable/trustworthy AI : bias, fairness, explainability/interpretability, robustness

June 20, 21, 22, 23 Sequential Learning – Stochastic Bandits – Q learning

June 27, 28, 29, 30 - Natural Language Processing, Symbolic AI : knowledge-based systems (systems of experts), logic reasoning

● ORGANIZATION

The course is composed of two components, taking place remotely on **Tuesday, Wednesday, Thursday and Friday between 11:15am and 1:45pm**

Lectures

The purpose of the lectures is to convey essential concepts in AI, explain problems and methods, and describe use cases. These theory classes include interaction between the students and the instructors.

Lab classes

Understanding the notions presented during lectures will be greatly facilitated by laboratory courses.

Office hours By appointment

All the material will be available on-line.

Evaluation

- Problem sets and programming assignments, prolongating three of the labs (60% of the grade: 20% each)
- Final exam/quiz (40% of the grade)

● PREREQUISITES

V22.0201 (Computer Systems Design I) and V22.0301 (Basic Algorithms)

● REFERENCES

Russel, S. & Norvig, P. (2020). *Artificial Intelligence: a Modern Approach*. Pearson

Vapnik V. (1998). *The Nature of Statistical Learning Theory*, John Wiley.

Friedman J., Hastie T. & Tibshirani, R. (2002). *The Elements of Statistical Learning: Data-mining, Inference and Prediction*, Springer

Sutton, R. & Barto, A. (1998). *Reinforcement Learning: An Introduction*, The MIT Press.

Bishop, C. (2006). *Recognition and Machine Learning*, Springer.

Goodfellow, I., Bengio, Y. & Courville, A. (2016). *Deep Learning Theory*, The MIT Press.