## Course Overview

**Problem solving** is a goal-directed activity. The field called "**Artificial Intelligence**" (AI) has always considered problem solving this kind of activity, and many problems studied in the cognitive literature have been approached in this way. The list of such problems includes the Traveling Salesman Problem, the Tower of Hanoi, visual navigation, math and physics problems, as well as creative problems such as formulating scientific theories. It turns out that all goal-directed actions require forming mental representations (models) of the problem at hand. At the very minimum, this mental representation must include the start state, the goal state and a list of possible actions that lead to a sequence of steps through intermediate states towards the goal. When the search space is large, as is the case with NP-hard problems, the agent must have a way to monitor the progress towards the goal. This naturally leads to the use of optimization methods that allow the agent to solve the problem by minimizing the estimated distance to the goal. This characterization of problem representation is how Newell and Simon, the fathers of AI, viewed the essential features of their General Problem Solver. Solving problems is likely to benefit from inherent regularities characterizing the problem space, such as similarity of legal moves and similarity of states, as is the case with problems such as the 15-puzzle or Go. These regularities allow the agent to plan global actions before executing local steps in a remote part of the problem. If the regularities of the problem space cannot be inferred from the description of a problem, they must be learned through repeated interactions with the problem. The nature and methods of learning are of fundamental importance in AI, but they have been largely neglected in psychology of problem solving. This course will introduce the theory of machine learning in the context of problem solving, with a special emphasis on reinforcement learning, as well as neural network implementations of learning. In particular, we will review some very recent developments in deep reinforcement learning algorithms, meta-learning and algorithm learning with neural Turing machines.

This course is ideal for students interested in computational models of intelligent behavior. The course is intended to be accessible to students from a broad range of disciplines, with varying background knowledge in the field. However, quantitative reasoning skills, including basic calculus and computer programming are necessary to understand and implement the core concepts.

In case of doubt, interested students are encouraged to e-mail the instructors.

The programs distributed for the hands-on exercises and assignments will be written as IPython Notebooks. If necessary, the instructor will introduce the basic Python programming.

As a Python tutorial, refer to these resources:
http://www.scipy-lectures.org/intro/index.html (sections 1.1 through 1.4).

Students will be evaluated on an individual project of their choice and present it orally on week 10 to the class. Ideally, this project should be related to the students' research and, if applicable, with data already collected.

## References

The course is based on material from the following books and articles, and may provide useful complementary information:

- **Russell and Norvig, Artificial Intelligence, 3rd Edition, 2010**

- Silver et al. 2018, Mastering the game of Go with deep neural networks and tree search (Links to an external site.)
- Graves *et al.* 2016 Hybrid computing using a neural network with dynamic external memory (Links to an external site.)
- Sutton and Barto, Reinforcement learning: An Introduction, 2018 (Links to an external site.)
- (Links to an external site.)Lawler et al. 1985, The Traveling Salesman Problem: A Guided Tour of Combinatorial Optimization 1st Edition (Links to an external site.)

## Tentative Course Schedule:

| Content | Instructor | Week |
|---|---|---|
| Introduction to the course | | |
| • Problem solving as a goal-directed behavior | | |
| • Problem solving as optimization | ZP,EN | Jan 8 |
| • The role of representations | | |
| • Theory of Mind | | |
| Combinatorial Optimization Problems - Intro | | |
| • Computational complexity | ZP | Jan 15 |
| • NP-hard and NP-complete problems | | |
| Human performance in TSP | | |
| • Time and error | ZP | Jan 22 |
| • Clustering, trees, shortest paths | | |
| TSP approximation and near-optimal algorithms - visual mechanisms | ZP | Jan 29 |
| Machine Learning and Optimization | | |
| • Learning from data, supervised, unsupervised, RL | | |
| Reinforcement Learning, | EN | Feb 5 |
| • Markov Decision Processes | | |
| • Bellman Equation | | |

- Applications

| | | |
|---|---|---|
| Reinforcement Learning,<br>• Model-based approaches<br>• Monte Carlo Methods<br>• Alpha Go | EN | Feb 12 |
| Reinforcement learning in Neuroscience:<br>• Reward-based Learning | EN | Feb 19 |
| Neural Networks<br>• Basic Concepts<br>• Differentiable Neural Computers (Neural Turing Machines)<br>• Applications | EN | Feb 26 |
| Insight Problems<br>• Changing mental representations<br>• Theories of insight<br>• Scientific discovery | ZP | Mar 5 |
| Project Presentations | All | Mar 12 |