

## Beliefs and decisions: of minds and machines

CEU SUMMER UNIVERSITY

July 5-9, 2010

### Syllabus

The aim of the course is to demonstrate that some basic principles of decision making can provide a unifying framework for constructing intelligently behaving artefacts on one hand, and for explaining human and animal cognition both in simple as well as in the most complex domains of behaviour on the other hand. To achieve this, lectures will progress via domains of gradually increasing abstraction that machine learning algorithms and humans deal with starting from representing uncertainty, beliefs about unobserved quantities, through learning internal models of the environment, to making adaptive and successful decisions.

The course will be organised around the following three key modules.

1. *Representations of uncertainty.* How is uncertainty represented in machine learning algorithms, what use is it of, what are the main advantages and challenges in practical applications? Does human and animal behaviour reflect the representation of uncertainty? How can networks of neurons represent uncertainty? What are the main sources of evidence for probabilistic computations in neural activity recorded in the brain?
2. *Learning.* How can we make intelligent algorithms that learn without direct supervision? How does learning benefit from probabilistic representations of beliefs? How can principles of learning be formalised mathematically, and how can such formal theories be tested in human or animal behaviour? How do we learn to adjust our movements to our environment, and how does our visual processing become to be adapted to it?
3. *Decision making.* How does uncertainty and learning influence decision making? How can rewards be taken into account in constructing machine learning applica-

tions? How do rewards affect human and animal behaviour? What happens if rewards are delayed rather than immediate, what new challenges does this pose to artificial as well as biological cognitive systems? How can we track the process by which a decision is born in the brain?

The principal format of the course will be 45-minute seminars given by the core faculty. There will be a discussion session after each module where participants will be encouraged to formulate a coherent view based on the lectures. Participants will be expected to critically evaluate competing views represented by a series of papers. In addition, when possible, there will be computer demonstrations of the relevant concepts.

The scope of the course will be broadened by organising round table discussions with representatives of disciplines that are in secondary connection with the main topic of the course. These guest discussants will come from areas such as economics, business, and arts, and they will present their views on how everyday decisions and notions of uncertainty reflect the principles discussed in the main lectures of the course.

Since the schedule will be rather tight, there will be no room for completing a research project during the course. Instead, students will select one of the topics offered by faculty members in advance of the course (each member of the faculty is expected to offer one topic), and during the course they will form small groups to work on a detailed project proposal addressing a relevant question in that topic. During the final day presentations, students will need to argue why the chosen research question is timely in light of the relevant literature, and describe in detail the methodology they propose to follow to address that question. The late afternoons of the preceding days will be spent with the groups preparing these presentations. Each group of students will be supervised by the member of the faculty offering the corresponding topic.