

750753, Machine Learning

3 hours per week, 3 credit hours, prerequisite: **none**

Teaching Method: 37 hours Lectures (2-3 hours per week), 8 hours Seminars (1 per 2 weeks)

Aims: The development of "intelligent" software systems is an area that is becoming increasingly important in applied Computer Science. One view of such systems postulates that truly intelligent machines must include the ability to learn from experience and observation. The purpose of this module is to introduce students to some prominent "Machine learning" techniques. These allow the development of software systems that are capable of adaptive behaviors in complex real-world environments. Some examples are programs that assist users or adapt to their behaviors; identify interesting patterns in databases- like those generated by the Human Genome Project; and predict stock-market movements. In this module students will learn symbolic learning (decision and regression trees), learning using neural-networks, and probabilistic modeling techniques.

Learning Outcomes:

On completion of this module, the student should be able to:

- Understand some prominent "Machine learning" techniques.
- Understand symbolic learning and learning using neural-networks.
- Develop "intelligent" software systems that are capable of adaptive behaviors in complex realworld environments.
- Develop systems that learn from experience and observation.
- Be involved in the research, development and exploitation of machine learning and adaptive computing.

Textbooks and Supporting Materials:

1. Thomas M. Mitchell, Machine Learning, McGraw Hill, 1997.
2. I. Witten and E. Frank, Data Mining: Practical Machine Learning Tools and Techniques with Java Implementations, Morgan Kaufman, 1999.
3. Richard S. Sutton, Aderw G. Barto, Reinforcement Learning, 1998.

4. Judea Pearl, Probabilistic Reasoning in Intelligent Systems, 1995

Synopsis:

1- Machine learning is concerned with the design of algorithms that, rather than encoding explicit instructions or programs for the solution of specific problems, encode inductive mechanisms whereby solutions to broad classes of problems may be derived from examples. Whether the approach is inspired by biology (e.g., artificial neural networks) or by cognitive psychology (e.g., traditional AI), machine learning aims at building computer systems capable of adapting to new tasks and learning from their experience.

2- Applications of machine learning and adaptive computing techniques in business and commerce. This is matched by research, which is directed at the provision of a new class of intelligent computer systems. The development of this area is dependent on the availability of a pool of personnel who are able both to understand and appreciate research findings and to use these to implement appropriate solutions to practical problems.

3- Topics in Evolutionary Computing, Neural Networks, Knowledge Systems, and Inductive Methods, and a series of smaller topics in Reinforcement Learning, Artificial Life, Agent-Based Systems, and Organizational and Business Aspects of IT. All topics cover the techniques associated with each area together with their applications.

4- A research project in this subject may provide a valuable opportunity for students to gain work experience and to consider their future direction.

25

Assessment: Two 1-hour midterm exams (15% each); Assignments (10%); Seminars (10%); 2-hours

Final Exam (50%)