UNIVERSITÀ
DEGLI STUDI DI MILANO-BICOCCA

SYLLABUS DEL CORSO

Modelli Probabilistici per le Decisioni

2627-1-F1802Q106

Aims

The course will provide the main concepts and operational tools, based on probabilistic and computational methods, for representing uncertainty, updating beliefs from incomplete or noisy data, modelling the evolution of dynamic systems, and supporting decision-making under conditions of uncertainty. Students will acquire the ability to use the concepts and methods learned to solve problems of inference, prediction, and decision-making. In particular, students will develop skills in building and evaluating probabilistic models, using simulation and sampling techniques, analysing Markov processes and temporal models, and formulating sequential decision-making problems.

Knowledge and understanding (DdD 1)
By the end of the course, students will be able to:
• Know and understand the main concepts, theories, and methods related to probabilistic reasoning under uncertainty.
• Know temporal probabilistic models, Markov chains, Hidden Markov Models, Bayesian filters, and sequential decision-making models.
• Understand the role of inference, prediction, simulation, and decision-making in intelligent systems and computing applications.
• Recognize the connections among probability, artificial intelligence, machine learning, sequence analysis, and reinforcement learning.

Applying knowledge and understanding (DdD 2)
Students will be able to:
• Apply the theoretical knowledge acquired to solve specific problems or case studies involving probabilistic and decision models.
• Use analytical tools and techniques to design concrete solutions in the presence of uncertainty.
• Conduct experiments, practical activities, or simulations.
• Implement simple procedures for inference, sampling, filtering, and sequential decision-making.

Making judgements (DdD 3)

Through laboratory activities, exercises, and guided discussions, students will develop:
• Critical ability in interpreting data, evidence, and collected information.
• The ability to evaluate different modelling and decision-making strategies, weighing their advantages, limitations, and assumptions.
• The ability to formulate well-reasoned and well-argued judgements independently with respect to the choice and use of probabilistic models.

Communication skills (DdD 4)
During the course, students will be encouraged to:
• Communicate the results of their analyses effectively, both in written form, through assignments and exercises, and orally, through discussions and presentations.
• Use appropriate and clear technical language.
• Justify modelling choices and critically interpret the results obtained.

Learning skills (DdD 5)
By the end of the course, students will be able to:
• Independently organize and plan their own path of further study and professional development.
• Use bibliographic and digital resources, as well as self-learning tools, to continue studying effectively.
• Develop a critical and reflective approach to new knowledge and technologies related to probabilistic models, decision-making systems, and artificial intelligence.

Contents

Representation of uncertainty in decision-making problems.
Probabilistic inference and Bayes' rule.
Probabilistic simulation, pseudo-random numbers, and Monte Carlo methods.
Markov chains and discrete-time stochastic processes.
Probabilistic reasoning over time and Hidden Markov Models.
Bayesian filtering, Kalman filters, and particle filtering.
Sequential decisions, Markov Decision Processes, the Bellman equation, and links with reinforcement learning.

Detailed program

1. Representation of uncertainty in decision-making problems.
Action under uncertainty. Basic notions of probability. Joint and conditional distributions. Inference using the full joint distribution. Independence and conditional independence. Bayes' rule and applications.
Reference: Russell & Norvig, Artificial Intelligence: A Modern Approach, 4th ed., Chapter 12, Sections 12.1–12.5 (3ª ed. it.: Capitolo 13, sezioni 13.1–13.5).

2. Probabilistic simulation and generation of random variables.
Random and pseudo-random numbers. Pseudo-random numbers, seed, period, reproducibility, and statistical properties of the generated sequences. Generation from a uniform distribution, inverse transform sampling, generation from discrete distributions, and the acceptance-rejection method. Introduction to Monte Carlo simulation for estimating probabilities and expected values.
Reference: material provided by the instructor; Rubinstein & Kroese, Simulation and the Monte Carlo Method, 3rd ed., Chapter 2, Sections 2.2–2.4; Chapter 4, Sections 4.2–4.3.

3. Markov chains.
Discrete-time stochastic processes. Markov property. State space, initial distribution, and transition matrix. Multi-

step probabilities and evolution of the state distribution. Classification of states, absorbing states, recurrence and transience. Stationary distributions and asymptotic behaviour. Random walks and applications.
Reference: material provided by the instructor; Norris, Markov Chains, Chapter 1.

4. Hidden Markov Models and sequential inference.
Extension of Markov chains to the case of states that are not directly observable. Hidden states and noisy observations. Definition of a Hidden Markov Model through an initial distribution, transition model, and emission model. Markov assumption, local observation assumption, and factorization of the joint distribution along the sequence. Filtering, prediction, smoothing, and likelihood of an observation sequence. Forward algorithm, forward-backward algorithm, and Viterbi algorithm. Notes on parameter learning via Baum-Welch/EM.
Reference: Russell & Norvig, Artificial Intelligence: A Modern Approach, 4th ed., Chapter 14, Sections 14.1–14.3 (3ª ed. it.: Capitolo 15, sezioni 15.1–15.3); Rabiner, "A Tutorial on Hidden Markov Models and Selected Applications in Speech Recognition", limited to the definition of HMMs, the three fundamental problems, and the forward, forward-backward, Viterbi, and Baum-Welch algorithms. For notes on parameter learning: Russell & Norvig, Artificial Intelligence: A Modern Approach, 4th Global ed., Chapter 21, Section 21.3 (3ª ed. it.: Capitolo 20, sezione 20.3).

5. Continuous Bayesian filtering.
The Bayes filter as a general framework for belief updating. Continuous states and noisy observations. Linear Gaussian models and the Kalman filter. Prediction, update, mean, covariance, innovation, and Kalman gain. Notes on the Extended Kalman Filter and Particle Filter for nonlinear, non-Gaussian, or multimodal models.
Reference: Russell & Norvig, Artificial Intelligence: A Modern Approach, 4th ed., Chapter 14, Sections 14.1–14.2 and 14.4 (3ª ed. it.: Capitolo 15, sezioni 15.1–15.2 e 15.4); material provided by the instructor; Thrun, Burgard, Fox, Probabilistic Robotics.

6. Sequential decisions and Markov Decision Processes.
Sequential decision-making problems. States, actions, probabilistic transitions, rewards, and policies. Markov Decision Processes. Value functions. Bellman equation and principle of optimality. Value iteration and policy iteration. Notes on Partially Observable Markov Decision Processes and links with reinforcement learning.
Reference: Russell & Norvig, Artificial Intelligence: A Modern Approach, 4th Global ed., Chapter 16, Sections 16.1–16.2 and 16.4 (3ª ed. it.: Capitolo 17, sezioni 17.1–17.4). For possible notes on algorithms for POMDPs: Russell & Norvig, Chapter 16, Section 16.5 (3ª ed. it.: Capitolo 17, sezione 17.4). For the connection with reinforcement learning: Sutton & Barto, Reinforcement Learning: An Introduction, 2nd ed., Chapter 3 and Chapter 4, Sections 4.1–4.4.

Prerequisites

Basic notions of probability, statistics, and linear algebra. Basic knowledge of programming and algorithms is useful.

Teaching form

Lectures, exercises, and laboratory activities.
The course is delivered in Italian.

The course consists of 54 total hours, divided into theoretical lectures, exercises, and in-person laboratory activities, with possible interactive activities according to the arrangements set out in the academic calendar.

Textbook and teaching resource

S. Russell, P. Norvig, “Intelligenza Artificiale: Un Approccio Moderno”, Prentice Hall, III Edizione.
Material provided by the instructor: slides, exercises, notebooks, and selected papers.

Additional references:
• J. R. Norris, Markov Chains.
• L. R. Rabiner, A Tutorial on Hidden Markov Models and Selected Applications in Speech Recognition.
• S. Thrun, W. Burgard, D. Fox, Probabilistic Robotics.
• R. S. Sutton, A. G. Barto, Reinforcement Learning: An Introduction.
• K. P. Murphy, Machine Learning: A Probabilistic Perspective or Probabilistic Machine Learning: An Introduction.

Semester

Second Semester

Assessment method

Written examination + optional oral examination.

Office hours

By appointment

Sustainable Development Goals

INDUSTRY, INNOVATION AND INFRASTRUCTURE