LSE MACHINE LEARNING IN PRACTICE

From chatbots, personalised recommendations on social media, traffic predictions and virtual personal assistants including Siri and Alexa, advances in machine learning are becoming an integral tool that help individuals navigate the modern world.

Increased adoption of such technology across the world has driven massive growth in the volume of data, requiring businesses to harness the power of machine learning to make decisions, learn about and predict customer behaviour to drive strategic advantage.

Combining the fields of engineering, statistics, mathematics and computing, machine learning is one of the leading data science methodologies revolutionising business. This course will cover a wide range of machine learning methods, both model-based and algorithmic. Presenting the theoretical foundations of these methodologies, you will have an opportunity to apply them using real-world examples and datasets.

Computer seminars which enable you to practice your programming skills and a course project give you the opportunity to explore how machine learning can be used innovatively to solve pressing business challenges such as algorithmic trading in the financial industry, predicting customer behaviour, and improving compliance and risk management. By the end of the course, you will have developed the ability to understand how machine learning can be integrated into current business models and the challenges that this poses.

Prerequisites

You should have completed at least one semester of calculus, and at least one semester of probability and statistics to undertake this course. Some minimal experience with computer programming is also required.

Key topics

Introduction to machine learning and its application
Supervised and unsupervised learning
Statistical inference and probability theory
Linear regression and related methods
Classification techniques, logistic regression and discriminant analysis
Tree-based methods
Graphical models and Markov graphs
Neural networks and deep learning
Gaussian processes and support vector machines

Programme structure and assessment

This course is delivered as a combination of lectures, computer seminars, class discussions, problem sets and readings. The computer seminars will give you an opportunity to develop your programming skills and apply the theory using datasets.

The course is assessed through an individual project (50%) and a final examination (50%). Problem sets and computer-based exercises will be provided in class for feedback, allowing you to check your understanding of the course content and to prepare you for the examination. These exercises will not contribute to your final grade.

Course outcomes

Show in-depth knowledge of supervised and unsupervised machine learning algorithms
Learn to perform some of the main techniques and algorithms for regression, classification, tree-based methods and graphical models in R
Discuss the application of clustering and principal component analysis, neural networks and deep learning at an introductory level
Develop an understanding of the process to learn from data and inform decisions on real-world problems
Apply and evaluate suitable methods to various datasets by model selection and predictive performance assessment

Reading materials

The main reading material will consist of lecture slides and related materials, which will be distributed at the beginning of the course. Optional further reading is also recommended from the following textbooks

- Hastie, T., Tibshirani, R. and Friedman, J., The Elements of Statistical Learning: Data Mining, Inference and Prediction. 2nd Edition, Springer, 2009.

- K. Murphy, Machine Learning: A Probabilistic Perspective, MIT Press, 2012.

- C. M. Bishop, Pattern Recognition and Machine Learning, Springer 2006.