COURSE DESCRIPTION CARD - SYLLABUS

Course name
A THEORETICAL PERSPECTIVE ON MACHINE LEARNING [S5ITIT>TAUM]

Course

Proposed by Discipline
–

Year/Semester
3/5

Level of study
Doctoral School

Course offered in
English

Form of study
full-time

Requirements
elective

Number of hours

Lecture
4

Laboratory classes
0

Other
0

Tutorials
0

Projects/seminars
0

Number of credit points

1,00

Coordinators

dr hab. inż. Wojciech Kotłowski prof. PP
wojciech.kotlowski@put.poznan.pl

Lecturers

dr hab. inż. Wojciech Kotłowski prof. PP
wojciech.kotlowski@put.poznan.pl

Prerequisites

Basic knowledge of university-level mathematics, particularly calculus, linear algebra, probability, and statistics. Ability to work with mathematical notation and to read and understand scientific literature in English.

Course objective

The aim of the course is to introduce doctoral students to the theoretical foundations of machine learning, with particular emphasis on supervised learning, statistical decision theory, and the principles of learning from data. The lecture explains how relatively simple learning methods can be used to solve complex problems, while also addressing the capabilities and fundamental limitations of machine learning. The acquired knowledge should enable doctoral students to apply machine learning methods in scientific research in a more informed and critical manner.

Course-related learning outcomes

Knowledge:

SzD_W01 (P8S_WG): The doctoral student knows and understands the theoretical foundations of machine learning.

SzD_W02 (P8S_WG): The doctoral student is familiar with the main applications and development trends in machine learning and understands its importance in modern artificial intelligence and scientific research.

Skills:

SzD_U01 (P8S_UW): The doctoral student is able to identify research problems that can benefit from supervised machine learning and formulate them as prediction tasks.

SzD_U02 (P8S_UW): The doctoral student is able to critically analyze results obtained using machine learning methods, taking into account overfitting, generalization ability, and limitations arising from available data and underlying assumptions.

Social competencies:

SzD_K03 (P8S_KK): The doctoral student recognizes the importance of theoretical knowledge in solving cognitive and practical problems and is prepared to critically reflect on the capabilities and limitations of machine learning in scientific research.

Methods for verifying learning outcomes and assessment criteria

Learning outcomes presented above are verified as follows:

The course is assessed based on a short essay (1–2 pages) discussing the use of artificial intelligence and machine learning in the student's PhD research or the potential benefits of applying these methods to their research.

The assessment considers the relevance of the discussion to the student's research topic, technical correctness, quality of argumentation, and critical reflection on the capabilities and limitations of the discussed methods.

Successful completion of the course requires preparing and submitting the essay via the eKursy platform by the specified deadline.

Programme content

Introduction to the theoretical foundations of machine learning, with particular emphasis on supervised learning. Statistical formulation of prediction problems, classification and regression, loss functions, and predictive risk. Statistical decision theory and optimal prediction rules for different loss functions. Learning from finite samples, limitations of inductive inference, and the "No Free Lunch" theorem. Empirical risk minimization, the role of inductive assumptions, and the choice of hypothesis classes. Overfitting and generalization, theoretical bounds on prediction error, and the relationship between sample size and model complexity. Decomposition of prediction error into approximation and estimation components.

Course topics

Introduction to machine learning: basic concepts, learning paradigms, and examples of applications; supervised learning as a prediction problem.

Statistical foundations of learning: formalization of classification and regression problems, loss functions, statistical risk, and evaluation of prediction quality.

Optimal prediction decisions: the optimal Bayes classifier, optimal predictions in regression, conditional expectation and median, excess risk.

The problem of learning from data: inductive inference, fundamental limitations of learning, the "No Free Lunch" theorem, and the importance of regularity and assumptions about data.

Generalization guarantees: learning from finite samples, dependence of predictive risk on sample size and hypothesis class complexity.

Empirical risk minimization (ERM): empirical and true risk, hypothesis class selection, inductive bias, and

its role in learning.

Overfitting and model complexity: differences between training error and performance on unseen data, approximation and estimation error, and the trade-off between model complexity and generalization ability.

Teaching methods

Lecture - slide presentation

Bibliography

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

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

3. Barber, D., Bayesian Reasoning and Machine Learning, Cambridge University Press, 2012.

4. Abu-Mostafa, Y. S., Magdon-Ismail, M., Lin, H.-T., Learning From Data, AMLBook, 2012.

Breakdown of average student's workload

	Hours	ECTS
Total workload	25	1,00
Classes requiring direct contact with the teacher	4	0,00
Doctoral student's own work (literature studies, preparation for laboratory classes/tutorials, preparation for tests/exam, project preparation)	21	1,00