AI-MACHINE LEARNING & ANALYTICS

Bachelor in Data and Business Analytics BDBA SEP-2023
AIML-DBA.3.M.A

Area Others

Number of sessions: 30

Academic year: 23-24

Degree course: THIRD

Number of credits: 6.0

Semester: 1º

Category: COMPULSORY

Language: English

Professor: **LAURA SANCHEZ GARCIA**

E-mail: lsanchezg@faculty.ie.edu

Laura has a mixed academic and industry background. On the academic side, Laura has a PhD, Master and a Bachelor in Physics from Universidad Autónoma de Madrid. She did her PhD in the field of Nanotechnology with a strong focus on experimental data acquisition and analysis. She has published more than 10 peer-reviewed scientific publications and has been speaker at numerous international conferences.On the industry part, she is working as Advanced Analytics Lead in Sandoz Farmacéutica and previously worked for McKinsey & Company as a Senior Data Scientist. She has experience in a broad variety of Data Science use cases and has worked in several industries across the globe.

lsanchezg@faculty.ie.edu

SUBJECT DESCRIPTION

Artificial Intelligence (AI) started in the aftermath of World War II, during the Cold War, as one of the newest fields in science and engineering. Shortly afterwards, AI transformed into a hot topic as its applications spread across numerous industries and businesses, including automation, natural language processing, recommended systems, fraud detection and disease diagnosis in medicine. This path was not straightforward though, AI had to suffer the so called “AI Winters” which shaped the field as the one we have today.

Based on the Turing Test approach to AI, a machine is considered “intelligent” if it possesses the following capabilities:

- **natural language processing** to enable it to communicate successfully;
- **knowledge representation** to store what it knows or hears;
- **automated reasoning** to use the stored information to answer questions and to draw new conclusions;

- **machine learning** to adapt to new circumstances and to detect and extrapolate patterns.
- **computer vision** to perceive objects, and
- **robotics** to manipulate objects and move about.

Paradoxically, although these principles were stated in the 60's, they form most modern applications of AI. These applications have been known for years, but the real boost of AI began in 1989 after Yann LeCun and collaborators applied the standard back-propagation algorithm, developed in the 70's, to train an Artificial Neural Network (ANN) to learn how to recognize handwritten digits. Since then, ANNs have grown in size and complexity, giving rise to the field of Deep Learning, over-performing classical machine learning algorithms in most tasks.

The general objective of this course is to understand the basis of modern AI applications based on Machine Learning and Deep Learning. Using a "from theory to practice" approach; we will apply theory and knowledge to real practical examples across different industries.

LEARNING OBJECTIVES

In this course, you will learn state-of-the-art methodologies in AI used for different applications. The subject is organized in three main groups following a sequential order, namely, an introductory part and review of mathematical concepts required, a second part focused on Classical Machine Learning, and a last part regarding Deep Learning with focus in image and sequence processing. After an introduction to AI, we will review fundamental concepts in Machine Learning that have been already seen in previous courses and will get deeper into Classical Machine Learning. We then address more advanced topics such as Deep Learning.

We expect that at the end of the course, students will obtain a solid understanding of the theoretical background and practical applications of modern AI methodologies regarding both Classical Machine Learning and Deep Learning.

The course has a fair balance between theory and practice. In order to fix the theoretical concepts, students will need to develop their own models for practical applications proposed along the course, using popular Python libraries (e.g., Tensorflow, Keras and many others). Therefore, it is expected that students have a medium to advanced level of programming in Python.

In order to fill the possible gap in mathematical knowledge, the course has a section devoted to reviewing important mathematical concepts that will be used along the course such as matrix algebra or calculus.

TEACHING METHODOLOGY

Learning Activity	**Weighting**	**Estimated time a student should dedicate to prepare for and participate in**
Lectures	30.0 %	45.0 hours
Discussions	10.0 %	15.0 hours
Exercises in class, Asynchronous sessions, Field Work	30.0 %	45.0 hours
Group work	30.0 %	45.0 hours
Individual studying	0.0 %	0.0 hours
TOTAL	100.0 %	150.0 hours

PROGRAM

The program addresses the following topics:

- Introduction to AI (main concepts and ingredients of intelligence systems).
- Review of math for deep learning.
- Classical ML (supervised & unsupervised).
- DL for computer vision, sequences and text analytics.

SUMMARY

Disclaimer: The following description of the material covered is tentative. An attempt will be made to cover all listed topics. However; the pace in the classes will depend on the group performance.

SESSION 1 (LIVE IN-PERSON)

Course introduction and overview of Artificial Intelligence.

Review of programming environments setup and best practices (PyCharm, GitHub, VirtualEnvs, pip or conda installs, needed libraries, GPU configuration or Google Collab).

Book Chapters: Artificial Intelligence: A Modern Approach, Chapters 1 and 2 (See Bibliography)

Book Chapters: Hands-On Machine Learning, Chapter 1 (See Bibliography)

Book Chapters: Python Machine Learning, Chapter 1 (See Bibliography)

Other / Complementary Documentation: 30 python best practices (Will be Provided)

Article: Advanced Jupyter notebooks. A tutorial (Towards Data Science; Jan 3, 2019)

Article: Python Virtual Environments: A Primer (Real Phyton; Apr 13, 2022)

Other / Complementary Documentation: Manage Conda environments (conda.io)

Article: Pipenv vs virtualenv vs conda environments (Medium; Apr 19, 2019)

SESSION 2 (LIVE IN-PERSON)

Review of fundamentals of Mathematics: matrix operations & calculus for ML

Review of matrix algebra operations.

- Mathematics for Machine Learning. Chapter 2.

Review of calculus and optimization.

- Mathematics for Machine Learning. Chapter 5 and 7.

Book Chapters: Mathematics for machine learning, Chapter 2, 5, and 7 (See Bibliography)

SESSIONS 3 - 5 (LIVE IN-PERSON)

Review of supervised learning.

- Python Machine Learning, Chapter 2, 3 and 10. (Required)
- The Elements of Statistical Learning, Chapters 4, 9 and 12. (Required)
- Hands-On Machine Learning, Chapters 3, 4, 5 and 6. (Required)
- Pattern Recognition and Machine Learning, Chapter 3 and 4. (Recommended)

Ensemble Learning: Forests, Boosting, Bagging and Stacking.

- Hands-On Machine Learning, Chapter 7. (Required)
- Python Machine Learning, Chapter 7. (Required)
- The Elements of Statistical Learning, Chapters 9 and 10 (Required)
- Pattern Recognition and Machine Learning, Chapter 14. (Recommended)

Hyper-parameter Optimization.

- Python Machine Learning, Chapter 6. (Required)
- Hyper-parameter optimization with grid search and randomized search. (Recommended)
- *Scikit-optimize: Bayesian optimization framework compatible with scikit-learn. (Recommended)*
- *Optuna: hyper-parameter optimization framework. (Recommended)*
- A Comprehensive List of Hyperparameter Optimization & Tuning Solution. (Recommended)

Book Chapters: Python Machine Learning; Chapters 2, 3, 6, 7, and 10 (See Bibliography)
Book Chapters: The Elements of Statistical Learning; Chapters 4, 9, 10, and 12 (See Bibliography)
Book Chapters: Hands-On Machine Learning; Chapters 3 - 7 (See Bibliography)
Book Chapters: Pattern Recognition and Machine Learning; Chapter 3, 4, and 14 (See Bibliography)(Optional)
Article: Hyper-parameter optimization with grid search and randomized search (Machine Learning Mastery; September 14, 2020) (Optional)
Other / Complementary Documentation: Bayesian optimization framework compatible with scikit-learn (Scikit-optimize)(Optional)
Other / Complementary Documentation: Hyper-parameter optimization framework (Optuna)(Optional)
Article: A Comprehensive List of Hyperparameter Optimization & Tuning Solution (Mikkokotila; May 13, 2018) (Optional)

SESSION 6 (LIVE IN-PERSON)

Interpretability & explainability methods for Classical Machine Learning

- Interpretable Machine Learning.

Book Chapters: Interpretable Machine Learning: A Guide for Making Black Box Models Explainable; Munich : Christoph Molnar, 2022 (See Bibliography)

SESSION 7 (LIVE IN-PERSON)

Model selection & validation. Cross validation

- The Elements of Statistical Learning: Chapter 7 (Required)
- Python Machine Learning, Chapters 4, 6. (Required)
- Hands-On Machine Learning, Chapter 2. (Required)
- Model Selection with Scikit-learn. (Scikit Learn) (Recommended)

Book Chapters: Python Machine Learning; Chapters 4 and 6 (See Bibliography)
Book Chapters: The Elements of Statistical Learning; Chapter 7 (See Bibliography)
Book Chapters: Hands-On Machine Learning; Chapter 2 (See Bibliography)
Technical note: Model Selection with Scikit-learn (Scikit Learn) (Optional)

SESSIONS 8 - 9 (LIVE IN-PERSON)

Unsupervised learning

Clustering algorithms, anomaly detection and dimensional reduction.

Book Chapters: Hands-On Unsupervised Learning with Python; Chapters 2, 3, 4, and 5 (See Bibliography)
Book Chapters: The Elements of Statistical Learning; Chapter 14 (See Bibliography)

SESSION 10 (LIVE IN-PERSON)

Review/Questions/Inquiries.

SESSION 11 (LIVE IN-PERSON)

Introduction to Artificial Neural Networks. Basic Concepts.

Book Chapters: Hands-On Machine Learning; Chapter 10 (See Bibliography)

Book Chapters: Deep Learning with Python; Chapters 1, 2 and 3 (See Bibliography)

Book Chapters: The Elements of Statistical Learning; Chapter 11 (See Bibliography)

SESSION 12 (LIVE IN-PERSON)

Shallow neural networks. Multilayer Perceptron (MLP).

Book Chapters: Hands-On Machine Learning; Chapter 10 (See Bibliography)

Book Chapters: Deep Learning with Python; Chapters 1, 2, and 3 (See Bibliography)

SESSION 13 (LIVE IN-PERSON)

Building and training (Feed Forward) Deep Neural Networks I

The gradient descent and back-propagation algorithms.

Book Chapters: Hands-On Machine Learning; Chapter 10 and 11 (See Bibliography)

SESSION 14 (LIVE IN-PERSON)

Building and training (Feed Forward) Deep Neural Networks II

Initialization, optimization and regularization methods in DNN. Application: image classification with Feed Forward Deep Neural Network.

Book Chapters: Hands-On Machine Learning; Chapter 12 (See Bibliography)

SESSION 15 (LIVE IN-PERSON)

Review/Questions/Inquiries.

SESSION 16 (LIVE IN-PERSON)

Midterm exam.

SESSIONS 17 - 18 (LIVE IN-PERSON)

Convolutional Neural Networks (CNN)

Foundations. Application to image classification. Using pre-trained architectures.

Book Chapters: Hands-On Machine Learning; Chapter 14 (See Bibliography)

Book Chapters: Deep Learning with Python; Chapter 5 (See Bibliography)

SESSION 19 (LIVE IN-PERSON)

Students Projects follow up.

SESSIONS 20 - 21 (LIVE IN-PERSON)

Sequence Models

- *Recurrent Neural Networks (RNN).*
- *Long-Short Term Memory (LSTM).*
- *Application: financial time series analysis.*

Book Chapters: Deep Learning with Python; Chapter 6 (See Bibliography)

SESSION 22 (LIVE IN-PERSON)

Student's Group Projects follow-up.

SESSION 23 (LIVE IN-PERSON)

Generative Adversarial Neural Networks (GANs): Introduction.

Book Chapters: Hands-On Machine Learning; Chapter 17 (See Bibliography)
Book Chapters: Deep Learning with Python; Chapter 8 (See Bibliography)

SESSION 24 (LIVE IN-PERSON)

Generative Adversarial Neural Networks (GANs): Applications.

Book Chapters: Hands-On Machine Learning; Chapter 17 (See Bibliography)
Book Chapters: Deep Learning with Python; Chapter 8 (See Bibliography)

SESSIONS 25 - 26 (LIVE IN-PERSON)

Students projects follow up.

SESSION 27 (LIVE IN-PERSON)

Review/Questions/Inquiries.

SESSION 28 (LIVE IN-PERSON)

Industry/business use cases for Machine Learning algorithms & architectures.

SESSION 29 (LIVE IN-PERSON)

Students' Group Project presentations.

SESSION 30 (LIVE IN-PERSON)

Final exam

EVALUATION CRITERIA

During the course students will be required to read material prior to the sessions, and to participate in discussions during class. Students will be given assignments and quizzes to be completed individually. These will be solved and discussed during class sessions where students are expected to participate actively. The group project will consist of solving a given business or research problem using some of the AI methods learned. The group will present the solution to the problem. The overall grading will be based on the following criteria:

criteria	percentage	Learning Objectives	Comments
Final Exam	20 %		final exam
Individual Work	20 %		includes assignments, voluntary exercises, research.
Midterm Exam	20 %		midterm exam
Group Presentation	30 %		includes quality of group project and delivery
Class Participation	10 %		includes active engagement in class asking/answering questions

RE-SIT / RE-TAKE POLICY

BIBLIOGRAPHY

Recommended

- Stuart J. Russell and Peter Norvig. (2009). *Artificial Intelligence A Modern Approach.* Third. Prentice Hall. ISBN 0136042597 (Digital)
- Marc Peter Deisenroth, A. Aldo Faisal, and Cheng Soon Ong. (2020). *Mathematics for Machine Learning.* First. Published by Cambridge University Press. ISBN 9781108455145 (Digital)
- Christopher M. Bishop. (2006). *Pattern Recognition and Machine Learning.* Springer. ISBN 9780387310732 (Digital)
- Aurélien Géron. (2019). *Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow.* Second. O'Reilly Media, Inc.. ISBN 9781492032649 (Digital)
- Sebastian Raschka, Vahid Mirjalili. (2019). *Python Machine Learning.* Third. Packt Publishing Ltd.. ISBN 9781789955750 (Digital)
- Trevor Hastie, Robert Tibshirani, Jerome Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction.* Springer. ISBN 0387848576 (Digital)
- Giuseppe Bonaccorso. (2019). *Hands-On Unsupervised Learning with Python.* First. Packt. ISBN 9781789348279 (Digital)
- Benjamin Johnston, Aaron Jones, Christopher Kruger. (2019). *Applied Unsupervised Learning with Python.* First. Packt Publishing. ISBN 9781789952292 (Digital)

- Ian Goodfellow and Yoshua Bengio and Aaron Courville. (2016). *Deep Learning.* First. MIT Press. ISBN 9780262035613 (Digital)
- Francois Chollet. (2017). *Deep Learning with Python.* First. O'Reilly Media, Inc.. ISBN 9781617294433 (Digital)
- Rowel Atienza. (2018). *Advanced Deep Learning with Keras.* First. Packt. ISBN 9781788629416 (Digital)
- Christoph Molnar. *Interpretable Machine Learning: A Guide For Making Black Box Models Explainable.* Independently published. ISBN 9798411463330 (Digital)

BEHAVIOR RULES

Please, check the University's Code of Conduct here. The Program Director may provide further indications.

ATTENDANCE POLICY

Please, check the University's Attendance Policy here. The Program Director may provide further indications.

ETHICAL POLICY

Please, check the University's Ethics Code here. The Program Director may provide further indications.