Foundations and Applications of Machine Learning

Proposal for MAP-I Curricular Unit 2023/2024

Rita P. Ribeiro[1], Pétia Georgieva[2], Pedro G. Ferreira[1], and Miguel Rocha[3]

[1] Dep. of Computer Science, Faculty of Sciences, University of Porto
[2] Dep. of Electronics Telecommunications and Informatics, University of Aveiro
[3] Dep. of Informatics, University of Minho

A. Programmatic component

1 Context

Big data is overwhelming nearly every field of knowledge, from life sciences, Internet, finances and banking or social networks. Daily, 2.5 quintillion bytes of data are being generated. It is becoming more and more important to be able to make sense of and communicate all the knowledge that it represents. It is often the case that generated datasets have high dimensionality, complexity, and heterogeneity, which, associated with their large volumes, signify a hard analysis problem. Modern Machine Learning (ML) techniques represent a powerful approach to the analysis of such large-scale datasets by deriving novel representations that augment domain knowledge and support informed decision-making processes. ML is considered by many to be the driver of the next wave of innovation and Artificial Intelligence. There are currently worldwide companies such as Facebook, Google or NVidia that are making significant breakthroughs by developing disruptive ML solutions and incorporating them into new products. Other companies are following similar paths making record investments in ML.

The application of ML technologies is certainly not only business oriented. Life sciences disciplines such as physics and biology are also steadily progressing with the applications of ML to the analysis of Terabytes to Petabytes of data generated by modern devices and technologies. One exciting example is the field of biomedicine. Biomedical data, including those generated from large-scale genomic projects, electronic health records or clinical exams, is growing at an unprecedented scale. ML is increasingly a critical tool to extract value from these data on different domains.

The high demand for ML specialists to work on problems, such as self-driving cars, DNA genome analysis or cancer prediction, climate change and many other fields prompts the need to train the next generation of computer scientists with the theoretical and practical knowledge of Machine Learning that allows them to develop projects that use the latest technologies following the best implementation practices. See, for instance, the McKinsey Global Institute report: "A significant constraint on realizing value from big data will be a shortage of talent, particularly of people with deep expertise in statistics and machine learning." -

From Big data: The next frontier for innovation, competition, and productivity, 2011, McKinsey Global Institute.

To tackle this shortage and the high demand for professionals with a solid background in ML and big data analytics, we propose a curricular unit that teaches Machine Learning using state-of-the-art technologies, including the most recent software libraries and platforms.

2 Pedagogical objectives and learning skills

The main objectives are to supply the students with adequate knowledge and skills in the core principles and techniques of Machine Learning. Thus, students completing this unit should:

- learn the fundamentals of ML: regression, classification, clustering, deep learning.
- understand the connection between learning and optimization; build optimal data representation models;
- learn how to implement and apply predictive, classification, clustering, information retrieval and deep learning algorithms to real datasets;
- develop a critical view and be able to choose, apply and evaluate the most adequate problem-solving techniques in ML;
- be able to design, specify, implement, and validate advanced software tools for specific data analysis problems; assess the quality of the models using the relevant error metrics;
- be able to interact with professionals from different knowledge domains in the process of software development and generate adequate reporting.

Other transversal skills are also approached in this unit, such as:

- to use modern software tools to implement reproducibility and portability during development, testing, and deployment;
- to conduct a short research project, being able to formulate a research problem, review significant literature, evaluate existing solutions and propose alternative approaches implementing those in new software tools;
- to write reports and scientific publications explaining the work developed;
- to be able to communicate with other researchers within multidisciplinary teams, in many cases international;
- to present orally the work developed.

3 Program

The overall aim is to provide an introduction to modern ML approaches for several problems, including regression, classification, clustering, and deep learning, among others.

3.1 Outline

- Fundamentals of ML. What is ML, and what are the challenges?
- Supervised versus Unsupervised ML.
- Data Pre-processing and exploration
  - Detection of outliers; Standardization; Transformation; Dimensionality reduction (Principal Component Analysis, Multi-Dimensional Scaling); Split data into training and test sets
- Model evaluation and model selection methods
  - Cost (loss) function. Cost function convergence. Iterative gradient descent algorithm. Learning curves. Performance measures (error, confusion matrix, sensitivity/specificity, ROC curves); Train and test paradigm; Cross-validation; Bias and Variance; Overfitting. Regularization.
- Classification methods
  - Decision Trees; K-NN; Linear and Nonlinear (Kernel); Support Vector Machines; Neural Networks; Logistic Regression, Ensemble approaches: Bagging, Boosting; Random Forests; Gradient Boosted Decision Trees
- Regression
  - Univariate and Multivariable linear regression; Performance measures: RMSE, R-Squared; Results (Coefficients, residuals) interpretation. Batch/ mini-batch/ stochastic gradient descent.
- Unsupervised learning - fundamentals
  - Distance (similarity) measures; K-means clustering; Hierarchical clustering;t-SNE; Data compression
- Deep Learning fundamentals
  - Shallow and Deep Neural networks; Network architectures (feed-forward, convolutional, recurrent, auto-encoders); Hyper-parameter optimization; Input data transformation; Applications to unsupervised, classification and regression problems; Introduction to deep generative models
- Imbalanced Domain Learning and Anomaly Detection
- Information visualization
  - Scatter plots; Boxplots; Heatmaps; Trees and Dendrograms
- Introduction to scikit-learn, numpy, matplotlib, pandas and tensorflow/ keras Python packages. Implementation and testing of ML pipelines. Create notebooks and docker containers for portability and predictability during development, testing, and deployment.

3.2 Learning outcomes

- Identify potential applications of ML in practice.
- Describe the core differences in analyses enabled by regression, classification, and clustering.
- Select the appropriate ML task for a potential application.
- Apply regression, classification, clustering and deep learning.
- Represent your data as features to serve as input to machine learning models.
- Assess the model quality in terms of relevant error metrics for each task.
- Build an end-to-end application that uses ML at its core.
- Implement ML algorithms and data analysis pipelines in Python, in particular taking advantage of scikit-learn, scipy, tensorflow/ keras and pandas packages.

4 Pedagogical strategies

This unit will consist of theoretical-practical (TP) lectures to introduce and examine ML methods and algorithms. Lecture material will consist of slides prepared by the lecturers and references to textbooks and scientific articles. Practical programming exercises and small projects will be discussed. Here, the appropriate programming libraries will be introduced. A data-driven approach to teaching will be followed. The course will also contain a semester project where students will tackle real-life problems and datasets. Students may propose a dataset under their thesis project.

4.1 Introduction to research

This unit will require students to develop skills essential for scientific research, namely the following ones.

– Characterize the state-of-the-art. This will be achieved by reading recommended textbooks, scientific papers and webpages.
– Concise presentation of results. The report for the final project should be written as an extended article. At the end of the semester, a presentation should also be done to present and defend the work.
– Team spirit. Students should work in small teams with a division of tasks.
– Reproducible research and pipelines. Students should apply recent methodologies to allow their results to be easily reproduced by others in different settings and platforms.
– Advancing state-of-the-art. By using competitive platforms, students will get a sense of how their proposed solutions are advancing current best-performing solutions.

4.2 Tentative calendar

– Class 1 [Lecturer: Rita P. Ribeiro, U. Porto] - Fundamentals of ML, Basic supervised ML pipeline, Data Pre-processing and exploration.
– Class 2 [Lecturer: Pétia Georgieva, U. Aveiro] – Classification Algorithms, Training process.
– Class 3 [Lecturer: Pétia Georgieva, U. Aveiro] - Model selection and performance evaluation, Ensemble methods.
– Class 4 [Lecturer: Rita P. Ribeiro, U. Porto] – Regression Algorithms, Performance Measures. Linear Regression, Regularization, Batch/ Mini-batch/ Stochastic gradient descent.
– Class 5 [Lecturer: Pedro Ferreira, U. Porto] - Unsupervised learning (Clustering, PCA, MDS).
– Class 6: [Lecturer: Rita P. Ribeiro, U. Porto] – Imbalanced Domain Learning and Anomaly Detection.
– Class 7 [Lecturer: Miguel Rocha, U. Minho] - Neural Networks; Deep Learning (Introduction).
– Class 8 [Lecturer: Miguel Rocha, U. Minho] - Deep Learning (Recurrent and convolutional neural networks; transfer learning; applications of DL).
– Final Project Presentation.

4.3 Evaluation criteria

The evaluation will consist of a Project applying Machine and/or Deep Learning approaches to a selected problem approved by the teaching staff. The evaluation of the project will consist of a written report - 70% - and a presentation- 30%).

5 References

- Flach, P. (2012). Machine learning: The art and science of algorithms that make sense of data. Cambridge University Press.
- Aurlien Gron. (2017). Hands-On Machine Learning with Scikit-Learn and TensorFlow: Concepts, Tools, and Techniques to Build Intelligent Systems (1st. ed.). O'Reilly Media, Inc.
- Raschka, S. (2015). Python Machine Learning. Packt Publishing Ltd.
- James, G., Witten, D., Hastie, T., and Tibshirani, R. (2013). An Introduction to Statistical Learning: with Applications in R. Springer Science & Business Media.
- Ni, M., Delong, A., Alipanahi, B., and Frey, B. J. (2016). Machine Learning in Genomic Medicine: A Review of Computational Problems and Data Sets. Proceedings of the IEEE, 104(1), 176–197.
- Chollet, F. (2018). Deep Learning with Python. Manning Publications.
- Goodfellow, I., Bengio, Y., & Courville, A. (2016). Deep Learning. MIT Press.
- Mitchel, T. (1997). Machine Learning (Mcgraw-Hill International Edit).

B. Teaching team

1 Presentation of the team

The team includes four lecturers from the three universities involved in the MAP-I program: Rita P. Ribeiro and Pedro G. Ferreira (U. Porto), Pétia Georgieva (U. Aveiro) and Miguel Rocha (U. Minho). Altogether, the team has a long-standing experience in teaching and researching in the area of this proposal: Machine Learning and its applications. The team has proposed and successfully run this curricular unit in MAP-I for the last 4 editions.

2 Coordinator

Rita P. Ribeiro (ORCID: 0000-0002-6852-8077) is an Assistant Professor at the Department of Computer Science of the Faculty of Sciences of the University of Porto (FCUP) and a Researcher at the Artificial Intelligence and Decision Support Lab (LIAAD) of the Institute for Systems and Computer Engineering, Technology and Science (INESCTEC). She holds a PhD in Computer Science from the University of Porto. Her main research interests focus on machine learning problems in imbalanced domains, anomaly detection, evaluation issues in learning tasks and application problems related to social good and environmental impact. She has been involved in several research projects concerning ecological problems, fraud detection and predictive maintenance applications. She has supervised nearly 20 master's students and currently supervises 3 PhD students. She is a member of the program committee of several international conferences and a reviewer for several international journals. She is one of the Action Editors of the Machine Learning Journal and the International Journal of Data Science and Analytics, both from Springer. She has also been active in the organization of various scientific events. Personal webpage: `https://www.dcc.fc.up.pt/~rpribeiro`

3 Team members

Pétia Georgieva (ORCID: 0000-0002-6424-6590)) is an Associate Professor with Habilitation at the Department of Electronics Telecommunications and Informatics of the University of Aveiro, a Researcher with the Institute of Electronics Engineering and Telematics of Aveiro (IEETA) and a collaborator member of the Institute of Telecommunications, Aveiro. Previously she had academic positions with the University of Porto (2001-2003) and research visiting positions with the Rowan University, USA, Carnegie Mellon University, USA, University of Lancaster, UK and the Bulgarian Academy of Sciences, Bulgaria. Her interests are in the area of machine learning, deep learning, data mining, signal processing and control. Her work involves the development of novel techniques for high dimensional problems (including neuro-computing, brain-computer interfaces, and image processing) and a data science approach in bio-chemical processes and, more recently, in optical communications. She supervised 9 PhD students and over 40 master's students. Dr. Georgieva is a Senior member of IEEE and a Senior Member of the International Neural Network Society (INNS). Her research is funded by sponsors such as the EU, the Portuguese Foundation for Science and Technology (FCT) and industry. Personal webpage: `http://wiki.ieeta.pt/wiki/index.php/Petia_Georgieva`

Pedro G. Ferreira (ORCID: 0000-0001-8439-8172) is an Assistant Professor with Habilitation at the Department of Computer Science, Faculty of Sciences of University of Porto and an Affiliated Researcher at i3s/ipatimup and at LIAAD-INESCTEC, the Artificial Intelligence and Decision Support Lab of the University of Porto. He graduated in Systems and Informatics Engineering from the University of Minho in 2002 and completed his PhD in Artificial Intelligence from the same University in 2007. From 2008 to 2012, he was a Postdoctoral Fellow at the Center for Genomic Regulation, Barcelona and from 2012 to 2014, at the Functional Population Genomics and Genetics of Complex Traits group, University of Geneva. He was involved in several major international consortia including ICGC-CLL, ENCODE, GEUVADIS and he is an active member of the GTEx consortium. In 2015, he was awarded an FCT Investigator Starting grant and joined Ipatimup/i3s. He has experience in the genomics start-up environment, where he developed information systems for personal genomics data interpretation. His main research focuses on the development of methods for a diverse set of problems in genomic data science. In particular, he is interested in unravelling the role of genomics on human health and disease. In order to achieve this goal, he applies and develops data-analytical models using machine learning and probabilistic methods to analyze and interpret diverse, complex and large-scale genomic datasets. Personal webpage: `http://pgferreira.net`

Miguel Rocha (ORCID: 0000-0001-8439-8172) is an Associate Professor with Habilitation at the Department of Informatics and a Senior Researcher at the Centre of Biological Engineering at the University of Minho. He is currently the Director of the Master in Bioinformatics and an elected member of the School of Engineering Scientific Council. He has a background in computer science, a PhD thesis in machine learning, and a vast curriculum in Machine Learning and Bioinformatics/ Systems Biology, including over 200 publications in peer-reviewed journals/conferences, 8 projects as a PI, and a patent application. He supervised 15 PhD students and over 60 master's students. He is the responsible docent of several curricular units related to Bioinformatics, Machine Learning/ Data Mining and programming in both first-degree and master courses. Personal webpage: `http://ceb.uminho.pt/People/Profile/mrocha`