“DIGITISE, OPTIMISE AND VISUALISE”: **Applied Machine Learning**

Simon Scheidegger
Department of Finance, HEC, University of Lausanne
simon.scheidegger@unil.ch

May 15, 2020

1 Purpose of the school:

This is a series of introductory lectures to “applied machine learning” in the context of applications in economics and finance, held at USI Lugano. Supervised (classification and regression), unsupervised, and reinforcement learning will be introduced in the context of formulating and solving appropriate optimisation problems—that is, so-called cost-functions. Furthermore, Gaussian Process Regression and Deep Neural Networks will serve us as showcase models to illustrate the potential of machine learning in actual applications. Finally, a large fraction of the course will be dedicated to introduce numerical concepts for dealing with the said taks on a computer.

2 Instructor:

Simon Scheidegger, Professor of advanced data analytics, Department of Finance, HEC, University of Lausanne. Web page: `https://sites.google.com/site/simonscheidegger`.

3 Course description:

The unprecedented availability of (big) data paired with substantial advances in machine learning (ML) theory, numerics, as well as computer hardware and powerful, simple-to-use software has allowed for an unprecedented rise of artificial intelligence and ML everyday life over the past few years. In this series of lectures, we want to demystify ML by looking under its hood in the context of Gaussian Process Regression (GPs) as well as Deep Neural Networks.

Gaussian Process Regression and Deep Neural Networks have—among other ML models—proven tremendous potential in financial and economic applications, ranging from option pricing, hedging, asset pricing in general, uncertainty quantification, or optimal control and optimal stopping problems (see, e.g., [3, 4, 11, 5, 7, 1] for an incomplete list of references).

This series of lectures aims at providing students a self-contained introduction to machine learning in general, and to GPs as well as Deep Neural Networks and the related numerical concepts concretely. In particular, we will cover key concepts such as formulating loss functions and likelihood functions, and numerical optimisation algorithms such as gradient descent and stochastic gradient descent to determine the optimal model parameters when confronted with data. This fundamental knowledge will serve the learners to apply the presented materials appropriately to relevant research questions in finance and economics.

Furthermore, ample time will be dedicated to applying the introduced ideas to actual data sets. During the lectures, the course attendees will have full access to contemporary machine learning software packages such as Scikit Learn (https://scikit-learn.org), tensorflow (https://www.tensorflow.org), and Keras (https://keras.io/) on Nuvolos (http://nuvolos.cloud), an integrated data science and computational platform. In practical sessions, the learners will implement and apply in Python some of the introduced models and optimisation methods, based on real financial data from established research databases.

After the course, learners should be able to:

- Understand the basics principles of the different variants of ML.
- Formulate customised cost functions appropriate to a research question at hand, and deal with them numerically.
- Apply GPs and Deep Neural Networks in their daily research.
- Understand the strengths and limitations of GPs and Deep Neural Networks.
- Use contemporary ML packages.

4 Prerequisites:

Undergraduate calculus and statistics, as well as programming experience in Python.

5 Preparatory materials and readings that accompany the lectures:

Below, we list both preparatory materials as well as resources that accompany the topics covered in the lectures.

5.1 Programming (preparatory materials):

- Introduction to Python: https://python-programming.quantecon.org/index_learning_python.html.
- Introduction to Computation and Programming Using Python [9].

5.2 Mathematics (preparatory materials):

- Mathematics for Machine Learning [6].

5.3 Machine Learning in general (accompaying materials):

- Pattern Recognition and Machine Learning [2].

5.4 Gaussian Process Regression (accompaying materials):

- Gaussian Processes for Machine Learning [10].

5.5 Artificial Neural Networks / Deep Learning (accompaying materials):

- Deep Learning [8].

6 Basic schedule & Teaching philosophy

This series of lectures aims at equipping learners in a workshop-like style with key skills and competences in ML, and specifically in GPs and Deep Neural Networks. This philosophy of teaching should allow the attendees to apply the acquired know-how to selected data-science and optimisation methods to relevant research questions in finance and economics.

The basics schedule of the course is seven days. Days 1-5 will be lectures will be held in the morning from 9:00-12:00. Software workshops will be held from 13:30-15:00, while theoretical and coding exercises in groups will be held from 15:30-17:00. Days 6-7 will be dedicated to working individually and in groups on a proposed project.

7 Preliminary table of content

- Day 1
  - Introduction and motivation to ML.
  - Cost Functions, Likelihood, Hyperparameters.
  - Numerical optimisation (e.g., gradient descent, stochastic gradient descent, Newton's method, BFGS) and numerical integration.
  - Python workshop: a round-trip in scikit-learn.
  - Hands-on exercises to the day's topics.
- Day 2
  - Basics of Gaussian Process Regression.
  - Noise-free kernels.
  - Kernels with Noise.
  - Gaussian Process classification.
  - Python workshop: Hands-on code demos with GPs.
  - Hands-on exercises to the day's topics.
- Day 3
  - Limitations of GPs.
  - GPs and "big data": scalable GPs (sparse kernels).
  - The curse of dimensionality and how to deal with it (e.g., Active Subspaces).
  - Active Bayesian Learning.
  - Frontier-topics of GPs ("Kiss" GPs, Deep Kernel GPs)
  - Python workshop: Hands-on code demos with scalable GPs, Active Subspaces, and frontier-topics.
  - Hands-on exercises to the day's topics.
- Day 4
  - A basic introduction to Artificial Neural Networks.
  - Deep Neural Networks.
  - Python workshop: A basic Neural Network code in Python.

  - Hands-on exercises to the day's topics.
  - **Attendees have to propose a project.**

- Day 5
  - State-of-the art software: Tensorflow and Keras.
  - A glimpse into reinforcement learning.
  - Hands-on exercises to the day's topics.
  - Wrap-up of the course.

8 Grading and school project:

Grading will be based for 30% on active class participation, for 70% on the graded take-home project and the corresponding oral exam. The school project consists of (i) a set of selected theoretical and code exercises and (ii) a small research project to be proposed developed in groups of up to 3 students per group.

References

[1] S. Becker, P. Cheridito, and A. Jentzen. Deep optimal stopping. *arXiv preprint arXiv:1804.05394*, 2018.

[2] C. M. Bishop. *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer-Verlag, Berlin, Heidelberg, 2006.

[3] H. Buehler, L. Gonon, J. Teichmann, and B. Wood. Deep hedging. *Quantitative Finance*, 0(0):1–21, 2019.

[4] H. Buehler, L. Gonon, J. Teichmann, B. Wood, B. Mohan, and J. Kochems. Deep hedging: Hedging derivatives under generic market frictions using reinforcement learning. Available at SSRN 3355706, 2019.

[5] L. Chen, M. Pelger, and J. Zhu. Deep learning in asset pricing. *Available at SSRN 3350138*, 2020.

[6] M. P. Deisenroth, A. A. Faisal, and C. S. Ong. *Mathematics for Machine Learning*. Cambridge University Press, 2020.

[7] A. Didisheim, D. Karyampas, and S. Scheidegger. Implied risk aversion smile. *Available at SSRN 3533089*, 2020.

[8] I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. MIT Press, 2016. `http://www.deeplearningbook.org`.

[9] J. Guttag. *Introduction to Computation and Programming Using Python: With Application to Understanding Data*. The MIT Press. MIT Press, 2016.

[10] C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press, 2005.

[11] S. Scheidegger and I. Bilionis. Machine learning for high-dimensional dynamic stochastic economies. *Journal of Computational Science*, 33:68 – 82, 2019.