Machine Epistemology

Munich Center for Mathematical Philosophy
Ludwig-Maximilians-Universität München
Winter 2016/17

Instructor:	Dr. Gregory Wheeler		
Office:	Ludwigstr. 31, 131		
Office Hours:	*By Appointment*		
E-Mail:	gregory.wheeler@lrz.uni-muenchen.de		

Lecture:	Ludwigstr. 31, 021	Lab:	Ludwigstr. 31, 028
Time:	Tuesday, 10:00 – 12:00 c.t.	Time:	Wednesday, 10:00 – 12:00 c.t.
Coursesite:	https://www.coursesites.com/s/_MachineEpistemology		

Course Description

The aim of machine learning is to make a computer learn from data without explicitly programming it to do so. The fruits of machine learning are all around us: email spam filters classify your messages using machine learning techniques, and postal services around the world "read" and sort hand-written addresses using machine learning techniques. The rise of machine learning upends traditional assumptions about hard problems in epistemology – Hume's problem, Duhem's problem, and much of traditional analytic epistemology – as well as traditional presumptions of expertise in statistics concerning the selection of models for data and the setting of a model's parameters. But the success of these methods also raises serious ethical and public policy questions.

This course is an introduction to this philosophically rich and radical method of inquiry. As logic was to 20th century philosophy, so computational methods shall be to this century. So, the course is designed to be a hands-on introduction to machine learning with an emphasis on its philosophical underpinnings and methodological implications.

Requirements

You will be asked to implement a series of machine learning algorithms using the Octave programming environment. Students are recommended to attend tutorials with a working copy of Octave running on their own personal laptop. Students will also write one short philosophical essay ($\sim$ 2000 words) on a topic covered in the course.

Online Registration

To access course materials and submit programming assignments, students **must** register for *Machine Epistemology* online at gregorywheeler.coursesites.com.

To register, **email me to receive an invitation to join the course**.

Software

We will be using the Octave programming environment, an open source variant of the MATLAB. (Students with a MATLAB license are free to use MATLAB for the assignments.) Installation instructions are available online

Grading

Your final grade will be determined by six assignments:

5 Programming Assignments: 75%

1 Philosophical Essay ($\sim$ 2000 words): 25%

Details for each assignment, including specific instructions for labeling files and turning them in, will be made available on the course website.

Course Schedule

The course consists of three parts. The first ten weeks are a hands-on introduction to contemporary regression-based techniques in machine learning, with a focus on supervised learning and classification. The second part is a two week overview of some general methodological issues that arise in selecting a learning algorithm and how to evaluate its performance. The last three weeks of the course considers a range of philosophical topics through the lens of machine learning.

		Lecture Topic		Tutorial Topic	○ Background Reading ● Required Reading ⋆ Assignments
1	18 Oct	Machine Epistemology Intro	19 Oct	Introduction to Octave I	● (Domingos 2012) ● (Wheeler 2016) ○ Review of Linear Algebra (Online)
2	25 Oct	Univariate Linear Regression, Cost Function, Gradient Descent for One Feature	26 Oct	Introduction to Octave II	● (Murphy 2012, Ch. 1)
3	1 Nov	*All Saint's Day* «no lecture»	2 Nov	Gradient Descent	⋆ A1 **due**: Basic Octave Operations
4	8 Nov	Multivariate Linear Regression, Gradient Descent for Multiple Features, Non-linear Regression	15 Nov	Batch Gradient Descent	○ (Petrova and Solov'ev 1997)
5	15 Nov	Classification, Logistic Regression, Maximum Likelihood Estimator	16 Nov	Logistic Regression Cost Function	⋆ A2 **due**: Linear Regression
6	22 Nov	Multiclass Classification, One-vs-all	23 Nov	Overview of Optimization Algorithms: Conjugate Gradient, BFGS, L-BFGS	○ (Cramer 2004)
7	29 Nov	Overfitting and Regularization	30 Nov	Regularizing Linear and Logistic Regression	⋆ A3 **due:** Logistic Regression
8	6 Dec	Neural Networks I	7 Dec	Neural Network Cost Function and Back-propagation	○ (McCulloch and Pitts 1943)
9	13 Dec	Neural Networks II	14 Dec	Non-linear Classification, modeling logical operators	○ (Basheer and Hajmeer 2000)
10	20 Dec	Support Vector Machines	21 Dec	Kernels	⋆ A4 **due:** Neural Networks ○ (Fisher 1936) ○ (Aronszajn 1950)
11	10 Jan	How to Evaluate Learning Algorithms	11 Jan	Bias vs Variance Trade-off	○ Geman et al. (1992) ● (Breiman 2001)
12	17 Jan	Feature Selection	18 Jan	Accuracy vs Complexity	⋆ A5 **due:** SVMs ● (Koller and Sahami 1996) ● (Guyon and Elisseeff 2003)
13	24 Jan	No Free Lunch Theorems	25 Jan	Problem of Induction	● (Wolpert and Macready 1997) ● (Vickers 2016, §2) ● (Okasha 2001) ● (Salmon 1991)
14	31 Jan	Formal Learning Theory	1 Feb	Discovery and Justification	● (Schulte 2014) ● (Glymour 2004) ● (Glymour and Kelly 1992)
15	7 Feb	Causal Inference	8 Feb	The Computational Turn	● Napoletani et al. (2010) ● (Spirtes 2010) ● (Williamson 2004) ● (Dreyfus 2007)
			22 Feb		⋆ A6 **due**: Philosophical Essay

References

Aronszajn, N. (1950). Theory of reproducing kernels. *Transactions of the American Mathematical Society 68*(3), 337–404.

Basheer, I. and M. Hajmeer (2000). Artificial neural networks: Fundamentals, computing, design, and application. *Journal of Microbiological Methods 43*, 3–31.

Breiman, L. (2001). Statistical modeling: The two cultures. *Statistical Science 16*(3), 199–231.

Cramer, J. S. (2004). The early origins of the logit model. *Studies in History and Philosophy of Biological and Biomedical Sciences 35*, 613–626.

Domingos, P. (2012). A few useful things to know about machine learning.

Dreyfus, H. L. (2007). Why Heideggerian AI failed and how fixing it would require making it more Heideggerian. *Artificial Intelligence 171*, 1137–1160.

Fisher, R. A. (1936). The use of multiple measurements in taxonomic problems. *Annals of Human Genetics 7*(2), 179–188.

Geman, S., E. Bienenstock, and R. Doursat (1992). Neural networks and the bias/variance dilemma. *Neural Computation 4*(1), 1–58.

Glymour, C. (2004). The automation of discovery. *Dædelus 133*(1), 69–77.

Glymour, C. and K. T. Kelly (1992). Thoroughly modern meno. In *Inference, Explanation, and Other Frustrations: Essays in the Philosophy of Science*, pp. 3–22. University of California Press.

Guyon, I. and A. Elisseeff (2003). An introduction to variable and feature selection. *Journal of Machine Learning Research 3*, 1157–1182.

Koller, D. and M. Sahami (1996). Toward optimal feature selection. In L. Saitta (Ed.), *Proceedings of the Thirteenth International Conference on Machine Learning (ICML)*, pp. 284–292.

McCulloch, W. S. and W. Pitts (1943). A logical calculus of the ideas immanent in nervous activity. *Bulletin of Mathematical Biophysics 5*, 115–133.

Murphy, K. P. (2012). *Machine Learning: A Probabilistic Perspective*. Cambridge, MA: MIT Press.

Napoletani, D., M. Panza, and D. C. Struppa (2010). Toward a philosophy of data analysis. *Foundations of Science 16*(1), 1–20.

Okasha, S. (2001). What did Hume really show about induction? *The Philosophical Quarterly 51*(204), 307–327.

Petrova, S. S. and A. D. Solov'ev (1997). The origin of the method of steepest descent. *Historia Mathematica*, 361–375.

Salmon, W. C. (1991). Hans Reichenbach's vindication of induction. *Erkenntnis 35*, 99–122.

Schulte, O. (2014). Formal learning theory. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2014 ed.). CSLI Publications.

Spirtes, P. (2010). Introduction to causal inference. *Journal of Machine Learning Research 11*, 1643–1662.

Vickers, J. (2016). The problem of induction. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy*. CSLI Publications.

Wheeler, G. (2016). Machine epistemology and big data. In L. McIntyre and A. Rosenberg (Eds.), *The Routledge Companion to Philosophy of Social Science*. Routledge.

Williamson, J. (2004). A dynamic interaction between machine learning and the philosophy of science. *Minds and Machines 14*, 539–549.

Wolpert, D. H. and W. G. Macready (1997). No free lunch theorems for optimization. *IEEE Transactions on Evolutionary Computation 1*(1), 67–82.