CSCI 4155/6505 (2017): Machine Learning

Thomas P. Trappenberg
Dalhousie University

Acknowledgements

These lecture notes have been inspired by several great sources. I would specifically like to acknowledge the influence of Andrew Ng's lecture notes on earlier versions of this manuscript, as well as the teaching material from the NVIDEA Deep Learning Institute.

There are now a variety of excellent books on the theory of machine learning that I would like to recommend for further readings. This includes the by *Introduction to Machine Learning* by Ethem Alpaydin, 2nd edition, MIT Press 2010, and *Pattern Recognition and Machine Learning* by Christopher Bishop, Springer 2006. The standard book on RL is *Reinforcement Learning: An Introduction* by Richard Sutton and Andrew Barto, MIT press, 1998. The standard book for AI, *Artificial Intelligence: A Modern Approach* by Stuart Russell and Peter Norvig, 2nd edition, Prentice Hall, 2003, does also include some chapters on Machine Learning.

The most in-depth book on probabilistic machine learning is likely the book by Kevin Murphy, Machine Learning: A Probabilistic Perspective, MIT Press, 2012. The great book on Deep Learning is the book by Ian Goodfellow, Yoshua Bengio, and Aaron Courville, MIT Press, 2016. These are two books that are invaluable for researchers in machine learning.

Several people have contributed considerably to this lecture notes. In particular I would thank Paul Hollensen, Patrick Connor, and Hossein Parvar.

Contents