



AI Education: Machine Learning Resources

Todd W. Neller (Gettysburg College; tneller@gettysburg.edu)

DOI: [10.1145/3098888.3098893](https://doi.org/10.1145/3098888.3098893)

Introduction

In this column, we focus on resources for learning and teaching three broad categories of machine learning (ML): supervised, unsupervised, and reinforcement learning. In our next column, we will focus specifically on deep neural network learning resources, so if you have any resource recommendations, please email them to the address above.

Machine Learning

In addition to the much-loved and ubiquitous *Artificial Intelligence: a Modern Approach* textbook (Russell & Norvig, 2009), there are a number of excellent introductory texts and tools specific to ML. Christopher Bishop's *Pattern Recognition and Machine Learning* (Bishop, 2006) is among the clearest introductions to ML with emphasis on Bayesian techniques. Kevin Murphy's *Machine Learning: A Probabilistic Perspective* (Murphy, 2012) is more comprehensive and advanced, bringing together the best of many authors into a massive introduction. David Barber's *Bayesian Reasoning and Machine Learning* (Barber, 2012) is widely recommended for advanced undergraduates and graduate students for self-study without considerable prior background. It is also freely available¹. Tom Mitchell's *Machine Learning* (Mitchell, 1997) is the classic introduction to the field, offering a solid foundation and broad perspective.

There are many online resources that are very helpful for the study of ML. I highly recommend Andrew Ng's free online Coursera Machine Learning course². Experiential learning is key, so getting datasets to practice ML techniques with is vital to one's learning. The UC Irvine Machine Learning Reposi-

tory³ is an excellent source for datasets that one can browse datasets by ML task, attribute type, size, application area, etc. Kaggle, a site that supports ML learning through ML competitions, also offers keyword searchable datasets⁴. A variety of ML assignments are available via the Model AI Assignments repository⁵. Many more recommendations of good resources for learning ML may be found at <https://www.quora.com/How-do-I-learn-machine-learning-1>.

Supervised and Unsupervised Learning

For those interested in focusing on supervised and unsupervised learning techniques that span from classical statistical methods with high bias, low variance, and better interpretability (e.g. linear regression) to AI methods with low bias, high variance, and better predictive performance, here are resources I would particularly praise and highlight:

Hastie, Tibshirani, and Friedman's *The Elements of Statistical Learning* (Hastie, Tibshirani, & Friedman, 2009) is a particularly well-written introduction spanning the bias-variance trade-off spectrum, and it is freely available⁶. For a gentler undergraduate-friendly introduction featuring labs using R, I highly recommend Gareth James, Daniela Witten, Trevor Hastie, and Robert Tibshirani's *An Introduction to Statistical Learning with Applications in R* (James, Witten, Hastie, & Tibshirani, 2014), also freely available⁷ and integrating well with the freely available RStudio⁸.

For a free, open-source, Java-based set of Data Mining tools, Weka offers a broad range of classification and regression tools⁹.

³<http://archive.ics.uci.edu/ml/>

⁴<https://www.kaggle.com/datasets>

⁵<https://www.kaggle.com/datasets>

⁶<https://statweb.stanford.edu/tibs/ElemStatLearn/>

⁷<https://www-bcf.usc.edu/gareth/ISL/>

⁸<https://www.rstudio.com/products/rstudio/>

⁹<http://www.cs.waikato.ac.nz/ml/weka/>

Copyright © 2017 by the author(s).

¹web4.cs.ucl.ac.uk/staff/D.Barber/textbook/090310.pdf

²<https://www.coursera.org/learn/machine-learning>

Ian Witten and Eibe Frank's *Data Mining: Practical Machine Learning Tools and Techniques* (Witten & Frank, 2005) is Weka's companion textbook.

Reinforcement Learning

For those interested in reinforcement learning (RL) in particular, I strongly recommend Richard Sutton and Andrew Barto's *Reinforcement Learning: An Introduction* (Sutton & Barto, 1998). This was the first text that presented a clear, unified view of dynamic programming, Monte Carlo, and temporal-difference learning techniques, and remains the best foundational reading for study of RL. It is also freely available¹⁰. More algorithms are briefly described in Csaba Szepesvári's *Algorithms for Reinforcement Learning* (Szepesvári, 2010), also freely available¹¹. Richard Sutton recommends Marco Wiering and Martijn van Otterlo's *Reinforcement Learning: State-of-the-Art* (Wiering & van Otterlo, 2012) as "a valuable resource for students wanting to go beyond the older textbooks and for researchers wanting to easily catch up with recent developments".

Your Favorite Resources?

These are but a few good starting points for learning about ML. If there are other resources you would recommend, we invite you to register with our wiki and add them to our collection at <http://cs.gettysburg.edu/ai-matters/index.php/Resources>.

References

- Barber, D. (2012). *Bayesian reasoning and machine learning*. New York, NY, USA: Cambridge University Press.
- Bishop, C. M. (2006). *Pattern recognition and machine learning (information science and statistics)*. Secaucus, NJ, USA: Springer-Verlag New York, Inc.
- Hastie, T., Tibshirani, R., & Friedman, J. (2009). *The elements of statistical learning: data mining, inference and prediction* (2nd ed.). Springer.

¹⁰<http://incompleteideas.net/sutton/book/the-book.html>

¹¹<https://sites.ualberta.ca/szepesva/RLBook.html>

- James, G., Witten, D., Hastie, T., & Tibshirani, R. (2014). *An introduction to statistical learning: With applications in r*. Springer Publishing Company, Incorporated.
- Mitchell, T. M. (1997). *Machine learning* (1st ed.). New York, NY, USA: McGraw-Hill, Inc.
- Murphy, K. P. (2012). *Machine learning: A probabilistic perspective*. The MIT Press.
- Russell, S., & Norvig, P. (2009). *Artificial intelligence: A modern approach* (3rd ed.). Upper Saddle River, NJ, USA: Prentice Hall Press.
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: an introduction* (1st ed.). Cambridge, MA, USA: MIT Press.
- Szepesvári, C. (2010). *Algorithms for reinforcement learning*. Morgan and Claypool Publishers.
- Wiering, M., & van Otterlo, M. (Eds.). (2012). *Reinforcement learning: State-of-the-art*. Springer.
- Witten, I. H., & Frank, E. (2005). *Data mining: Practical machine learning tools and techniques, second edition (morgan kaufmann series in data management systems)*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.



Todd W. Neller is a Professor of Computer Science at Gettysburg College. A game enthusiast, Neller researches game AI techniques and their uses in undergraduate education.