justokay on Jan 15, 2010 | parent | favorite | on: Learning About Statistical Learning

Mike Jordan at Berkeley sent me his list on what people should learn for ML. The list is definitely on the more rigorous side (ie aimed at more researchers than practitioners), but going through these books (along with the requisite programming experience) is a useful, if not painful, exercise.

I personally think that everyone in machine learning should be (completely) familiar with essentially all of the material in the following intermediate-level statistics book:

1.) Casella, G. and Berger, R.L. (2001). "Statistical Inference" Duxbury Press.

For a slightly more advanced book that's quite clear on mathematical techniques, the following book is quite good:

2.) Ferguson, T. (1996). "A Course in Large Sample Theory" Chapman & Hall/CRC.

You'll need to learn something about asymptotics at some point, and a good starting place is:

3.) Lehmann, E. (2004). "Elements of Large-Sample Theory" Springer.

Those are all frequentist books. You should also read something Bayesian:

4.) Gelman, A. et al. (2003). "Bayesian Data Analysis" Chapman & Hall/CRC.

and you should start to read about Bayesian computation:

5.) Robert, C. and Casella, G. (2005). "Monte Carlo Statistical Methods" Springer.

On the probability front, a good intermediate text is:

6.) Grimmett, G. and Stirzaker, D. (2001). "Probability and Random Processes" Oxford.

At a more advanced level, a very good text is the following:

7.) Pollard, D. (2001). "A User's Guide to Measure Theoretic Probability" Cambridge.

The standard advanced textbook is Durrett, R. (2005). "Probability: Theory and Examples" Duxbury.

Machine learning research also reposes on optimization theory. A good starting book on linear optimization that will prepare you for convex optimization:

8.) Bertsimas, D. and Tsitsiklis, J. (1997). "Introduction to Linear Optimization" Athena.

And then you can graduate to:

9.) Boyd, S. and Vandenberghe, L. (2004). "Convex Optimization" Cambridge.

Getting a full understanding of algorithmic linear algebra is also important. At some point you should feel familiar with most of the material in

10.) Golub, G., and Van Loan, C. (1996). "Matrix Computations" Johns Hopkins.

It's good to know some information theory. The classic is:

11.) Cover, T. and Thomas, J. "Elements of Information Theory" Wiley.

Finally, if you want to start to learn some more abstract math, you might want to start to learn some functional analysis (if you haven't already). Functional analysis is essentially linear algebra in infinite dimensions, and it's necessary for kernel methods, for nonparametric Bayesian methods, and for various other topics. Here's a book that I find very readable:

12.) Kreyszig, E. (1989). "Introductory Functional Analysis with Applications" Wiley.

ptuzla on Jan 15, 2010 [-]

Can one fit this study list in a life time? Seriously, this has been a problem for me for a long time. Any one of the mentioned books would take me months to study. It'd take me a month to just read a textbook, without any toil. Am I too slow, or are there some study/reading techniques that I'm not aware of?

vdm on Jan 15, 2010 [-]

Just so you know, you're not the only one who feels this way. There's no way you could carry that load AND a day job/family.

Maybe tools like incanter will bridge the gap and let application software practitioners put this research to work.

fgimenez on Jan 15, 2010 [-]

Keep in mind that Mike Jordan is a superhuman math machine. I remember his undergraduate research assistants at Cal were telling me that it would take grad students days to understand 5 minute proofs he would do on the fly.

global-variable on Jan 16, 2010 [-]

It's almost like he's the Michael Jordan of math.

physcab on Jan 15, 2010 [-]

Woot woot for Casella! (UF prof)

My go-to book for Machine Learning is Christopher Bishops Pattern Recognition and Machine Learning. I've read that book cover-to-cover and its got an excellent foundation and covers all those other books in some capacity.

bradfordcross on Jan 15, 2010 [-]

Awesome, I was hoping to hear from some academics who know far more than I do...this an elaborate scheme to fill my wishlist pipeline...muahahahhaha