Interesting Things from NIPS

Ryan Prescott Adams

Cavendish Laboratory
University of Cambridge
http://www.inference.phy.cam.ac.uk/rpa23/

14 February 2007

Free Lunches: Insights from Behavioral Economics

Dan Ariely, MIT

- How do people really act? Irrationally.
- Spouse Example (It's Valentine's Day!)
- Cheating unrelated to P(getting caught).

Take Home:

Economics is broken. If people don't know what they want, how can they be making optimal decisions?

http://web.mit.edu/ariely/www/

Greedy Layer-Wise Training of Deep Networks

Bengio, et al (Montreal)

- "shallow" versus "deep" learning architectures
- Why does layer-wise training work?
- How to handle continuous cases.

Take Home:

Unsupervised learning of abstraction and representation is key to supervised "deep" learning.

http://books.nips.cc/papers/files/nips19/NIPS2006_0739.pdf

Gaussian and Wishart Hyperkernels

Kondor and Jebara (Columbia)

- “hyperkernels” - kernels on kernels
- Use the kernel directly to learn structure.

Take Home:

An RKHS formalism for representing the relationships between things rather than predictive values, with conventional kernel interpretations.

http://books.nips.cc/papers/files/nips19/NIPS2006_0710.pdf

Bayesian Detection of Infrequent Differences in Sets of Time Series with Shared Structure

Listgarten, et al (Toronto)

- Align and detect differences in time series.
- Multiple classes with shared structure.
- Infer alignment *and* differences simultaneously.

Take Home:

Hierarchical Bayes can solve this with MCMC when classical methods fail.

http://books.nips.cc/papers/files/nips19/NIPS2006_0558.pdf

NASA solenoid data

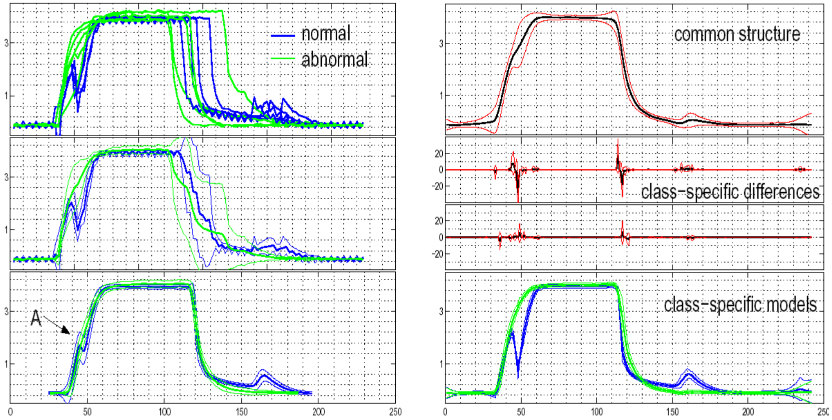

Particle Filtering for Nonparametric Bayesian Matrix Factorization

Wood (Brown) and Griffiths (Berkeley)

- Latent features of data: $\boldsymbol{X} \approx \boldsymbol{Z}\boldsymbol{Y}$
- Don't know the latent dimension.
- Typical approach: Gibbs sample with Dirichlet Process Prior
- Use an Indian Buffet Process instead.

Take Home:

IBP makes particle filtering feasible, and this proves faster than Gibbs sampling.

http://books.nips.cc/papers/files/nips19/NIPS2006_0499.pdf

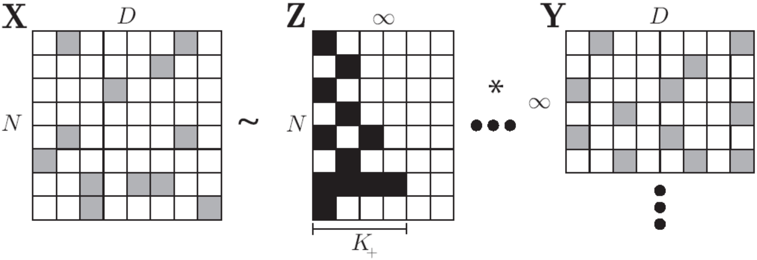
X
D
N
~
Z
∞
N
K_+
*
∞
Y
D

A Scalable Machine Learning Approach to Go

Wu and Baldi (UCI)

- Learn a "scalable" evaluation function. (expected territory)
- "scalability" meaning "not handcrafted"
- Compares to David Stern's stuff (sort of)

Take Home:

Recurrent neural networks might be useful for territory prediction.

`http://books.nips.cc/papers/files/nips19/NIPS2006_0223.pdf`