Tutorial proposal:
Machine Learning for Multimedia Applications

By Flavia Sparacino, Sensing Places, and Yuan (Alan) Qi, MIT

This document and the tutorial's support material also at:
http://www.sensingplaces.com/acm/

Abstract

This one day workshop will introduce the audience to Bayesian inference and learning for graphical models with examples and applications drawn from the multimedia field. The workshop will begin with a hands-on introduction to Bayesian reasoning and Bayesian networks, using HUGIN. The speakers will then present more advanced topics on Bayesian approximation techniques. They will give a brief review on Bayesian approximate inference on nonlinear and non-Gaussian dynamic systems, with application to wireless communications. Then they will present efficient inference algorithms on loopy graphical models.

During the second half of the workshop, the speakers will first present Bayesian feature selection methods for micro-array data classification and Bayesian conditional random fields, which are applied to text parsing and recognition of handwritten organization charts. Then the speakers will present a series of case-based studies of applications of Bayesian networks to multimedia, ranging from user modeling for personalized multimedia content delivery, to predicting user behavior during www navigation, probabilistic location estimation, and modeling of computer game agents.

Target Audience

Researchers and practitioners in the multimedia field who are interested in Bayesian machine learning techniques.

Prerequisites

Attendants are expected to have the following background:

- Familiarity with basic probability theory
- Familiarity with basic linear algebra
- Knowledge of basic computer science principles and skills

Syllabus

9:00-9:15	Introductions
9:15-10:15	Probabilistic Reasoning, Bayesian Modeling, Influence Diagrams and working examples (Flavia)
10:15-10:30	Generative vs Conditional Models (Alan)
10:30-11:00	Coffee Break
11:00-11:45	Introduction to Bayesian Inference (Flavia) Inference in Bayesian Networks Learning in Bayesian Networks Dynamic Bayesian Networks
11:45-12:30	Bayesian Approximation Techniques (Alan) Review of Expectation Propagation Generalizing Kalman Smoothing and comparison with Monte Carlo methods (Particle Filters and Smoothers, MCMC) Joint wireless signal detection and channel estimation From Belief Propagation to Tree-Structured Expectation Propagation
12:30-2:00	Lunch Break
2:00-2:45	From Belief Propagation to Tree-Structured Expectation Propagation (Alan) Bayesian Feature Selection Micro-array data classification for genomics
2:45-3:00	Bayesian Conditional Random Fields (Alan) Text retrieval: parsing FAQs files Ink recognition: classifying handwritten organization charts
3:00-3:30	Coffee Break
4:00-4:45	User classification for personalized multimedia presentations (Flavia)
4:45-5:00	Case studies: (Flavia) Location measurements using probabilistic approaches Predicting user's requests on the WWW Using Bayesian networks to model computer game agents
5:00-5:30	Summary and Q&A (Flavia and Alan)

Essential References

Books

- Jensen, F.V. Bayesian Networks and Decision Graphs. Springer-Verlag, New York, 2001.
- MacKay David. Information Theory, Inference, and Learning Algorithms, Cambridge University Press, 2003
- Bishop, Christopher, M. Neural Networks for Pattern Recognition, Oxford University Press, 1995

Papers

- Predictive Automatic Relevance Determination by Expectation Propagation, Yuan Qi, Thomas P. Minka, Rosalind W. Picard, and Zoubin Ghahramani, In Proceedings of Twenty-first International Conference on Machine Learning, July 4-8, 2004, Banff, Alberta, Canada.
- Tree-structured Approximations by Expectation Propagation, Thomas Minka and Yuan Qi, Neural Information Processing Systems, December 2003, British Columbia, Canada.
- Expectation Propagation for Signal Detection in Flat-fading Channels, Yuan Qi and Thomas Minka, MIT Media Lab Technical Report Vismod-TR-555. Also, in the proceedings of IEEE International Symposium on Information Theory, June, 2003, Yokohama, Japan.
- F. Sparacino. Sto(ry)chastics: a Bayesian Network Architecture for User Modeling and Computational Storytelling in Interactive Spaces. UbiComp 2003.
- F. Sparacino. Sto(ry)chastics: a Bayesian network architecture for combined user modeling, sensor fusion, and computational storytelling for interactive spaces. MIT PhD Thesis, 2001
- John Krumm, “Probabilistic Inferencing for Location”, Proceedings of the 2003 Workshop on Location-Aware Computing, October 2003.
- I Zukerman, D W Albrecht and A E Nicholson: Predicting users' requests on the WWW, UM 99.
- Vaerge, Jarlskov, and Kjaerulff. Using Bayesian Networks for Modeling Computer Game Agents, Technical Report, Aalborg University, June 2003

Additional References

Links

- Tom Minka’s tutorials at: http://www.stat.cmu.edu/~minka/papers/learning.html
- Kevin Murphy’s Bayesian Network Toolbox http://www.ai.mit.edu/~murphyk/Software/BNT/bnt.html
- Zoubin Ghahramani's Tutorial on Bayesian Methods for Machine Learning at: http://www.gatsby.ucl.ac.uk/~zoubin/ICML04-tutorial.html
- Intel OpenPNL library: http://www.intel.com/research/mrl/pnl/
- HUGIN: http://www.hugin.com

Books

- Frey B., Graphical Models for Machine Learning and Digital Communication, MIT Press, 1998.
- Jordan M. I., Learning in Graphical Models. MIT Press, 1998.
- E. T. Jaynes, G. Larry Bretthorst. Probability Theory: The Logic of Science, Cambridge University Press, 2003.

Papers

- Section 6 of D. Heckerman, (1995) “Tutorial on learning in Bayesian networks,” Microsoft Technical Report.
- N. Friedman. (1998) The Bayesian Structural EM Algorithm. In Fourteenth Conf. on Uncertainty in Artificial Intelligence (UAI). 1998.
- Eugene Charniak, Bayesian Networks without Tears, AI Magazine 12 (1991), 4, 50-63.
- Jordan M.I. and Weiss Y. Graphical models: probabilistic inference. In Arbib, M. (ed): Handbook of Neural Networks and Brain Theory. 2nd edition. MIT Press. postscript 457k (2002)
- Andersen, S. “HUGIN – a Shell for Building Bayesian Belief Universes for Expert Systems.” In: Proceedings of the Eleventh International Joint Conference on Artificial Intelligence, IJCAI, 1080-1085. Menlo Park, California, 1989.
- Pavlovic V., Rehg J., Cham T-J., Murphy K. “A Dynamic Bayesian Network Approach to Figure Tracking Using Learned Dynamic Models”. Proceedings of ICCV (Int'l Conf. on Computer Vision) 1999.
- G. Zweig and S. Russell. Speech Recognition with Dynamic Bayesian Networks In Proceedings of AAAI 98.

A/V Equipment Requirements

The speakers will need a VGA projector and audio speakers. They will use their own laptop computer for the presentation.

Speakers' bios

Flavia Sparacino

Flavia Sparacino is an interactive space and experience designer. She is the head of Sensing Places, an MIT spinoff, and a digital art and technology firm that builds interactive narrative spaces using ambient and body sensors, statistical mathematical models, interactive multimedia, and 3-D graphics. Her emphasis is on natural man-machine interfaces, which use unencumbering techniques (computer vision and wireless sensors) to understand people's presence, movement, and behavior. Flavia designed the first interactive table prototype for MOMA's Unprivate House exhibition (1999) based on her earlier prototype for Unbuilt Ruins (MIT Compton Gallery 1999), and designed the interactive prototypes for SFMOMA's Points of Departure museum exhibit (2001). More recently she designed the numerous interactive installations for Milan's La Scala "Puccini Set Designer" exhibit, including wearable computers, an immersive cinema, and an interactive video documentary presentation table. Currently she acts as technology curator and technical consultant for various museums in the US and Europe. Flavia holds five academic degrees plus an honorary one, including a PhD from MIT's Media Lab. Italian-born, she was nominated Knight of the Republic of Italy ("Cavaliere della Repubblica"), by the Italian president Carlo Azeglio Ciampi, in the year 2000, for her contributions to innovative communication of art and culture supported by new technologies.

Contacts:
flavia@media.mit.edu http://www.flavia.media.mit.edu/~flavia
flavia@sensingplaces.com http://www.sensingplaces.com

Yuan (Alan) Qi

Yuan (Alan) Qi just completed his Ph.D. studies from the Massachusetts Institute of Technology. He received a Master of Science in Electrical Engineering from University of Maryland at College Park, a Master of Science in Pattern Recognition and Artificial Intelligence from Institute of Automation, Chinese Academy of Sciences, and a Bachelor of Engineer in Automatic Control from Huazhong University of Science and Technology, China. His research interests include machine learning, bioinformatics, statistical signal processing and digital wireless communication. He has published at NIPS, ICML, ICPR, ICASSP, and other venues. His more recent research includes Bayesian conditional random fields for joint classification of relational data, nonparametric density estimation by Gaussian processes. Alan has given several invited talks and lectures on Pattern Recognition and Machine Learning at MIT, CMU, Toshiba Research Labs in Japan, University College London, Microsoft Research Cambridge, and Boston University

Contacts:
yuanqi@media.mit.edu http://web.media.mit.edu/~yuanqi/