

Gabi Teodoru

gabiteodoru@gmail.com – +44 (0)7596 118 268

Location: London, UK. Citizenship: Romanian

PERSONAL STATEMENT

Experienced researcher keen to apply machine learning and data science techniques and code efficient implementations of novel statistical inference and prediction algorithms to optimise companies' performance and returns, and use my strong communication skills to clearly explain techniques to those who will be working with them.

EDUCATION

PhD Candidate in Machine Learning, 2008 – 2013

Gatsby Computational Neuroscience Unit, University College London, London, UK

Recipient of a research studentship from the Gatsby Charitable Foundation

MSc in Music Informatics, 2006 – 2008. GPA: 3.765 (pending)

Indiana University School of Informatics, Bloomington, IN, USA

Recipient of Barbara and David Jacobs Fellowship, and Informatics Grad Support Fellowship

BMus in Music Composition, 2002 – 2006. GPA: 3.788

Indiana University Jacobs School of Music, Bloomington, IN, USA

Recipient of the Dean's Award

SKILLS

- Machine learning/maths: Bayesian learning and inference, time series models, Hidden Markov models, the EM algorithm, belief propagation, sampling and MCMC, Support Vector Machines, Markov Random Fields, Stochastic Context Free Grammars, convex and non-convex optimization, Fourier analysis
- Programming and data analysis: C++, Matlab, Python, R, Maple, Mathematica, Flash (ActionScript 3.0), PHP, MySQL, Java, real-time data processing, multithreading
- Interpersonal skills: collaborations with experts in diverse fields; seeking requirements and implementing feedback from end users to produce software tailored to their needs; giving talks with Q&A sessions on maths, Bayesian methods, and music topics to non-technical audiences

SELECTED RESEARCH EXPERIENCE

PhD thesis topic: Spectral Learning for Latent Variable Models – Collaborations and joint work with Ankur Parikh, Le Song, Mariya Ishteva, Eric Xing, Jeff Beck and Maneesh Sahani.

Gatsby Unit, UCL, 2012-2013

Spectral learning is a novel method for learning latent variable models, which combines the method of moments with ideas from subspace identification and predictive-state representations.

Skills: Graphical models, Time Series models, Stochastic Context-Free Grammars (SCFG), linear algebra and optimization, statistical estimator properties, belief propagation.

My main contribution was the insight that the original spectral learning estimator may be written as an optimization; this allowed for a wider variety of similar estimators that were more effective when applied to real-world data, yielding better results. This insight was made possible by understanding the optimization interpretations of matrix pseudoinverses and singular value decomposition, as well as by a different approach to thinking about graphical models, using inverse conditional probability matrices to obtain an operation very similar to, but fundamentally different from traditional message passing in graphical models.

MSc thesis topic: Conducting a Virtual Orchestra – joint work with Christopher Raphael.

Indiana University, 2007-2008

Pedagogical software that allows conductors to practice conducting an orchestra using a computer and a webcam.

Skills: C++, real-time data processing (audio and video), multithreading, Hidden Markov Models, EM algorithm, working with conducting students and reacting to feedback, object-oriented coding style.

Investigation and modeling of human memorization of discrete sequences – joint work with Adam Sanborn and Maneesh Sahani

Gatsby Unit, UCL, 2011

Research project to understand the human mind's representation of discrete sequences.

Skills: Matlab exploratory data analysis; Flash (ActionScript 3), PHP and MySQL; crowdsourcing with Amazon Turk; outsourcing graphic design with vworker.com. Efficient SCFG inference implementation: 2000x speed gain in rewriting the Matlab prototype into C++ and then C.

Sparse multiview methods for classification of musical genre from magnetoencephalography recordings; joint work with Tom Diethe, Nick Furl and John Shawe-Taylor

Computer Science Department, UCL, 2009

Research project to classify music genres from listeners' brain scans.

Skills: Working with big data (~5 GB of brain imaging data), experience in scientific experimental setup

Freelance work: Piano tuner

I wrote a stand-alone application that employs spectral analysis to aid me in my piano tuning freelancing. "I have tuned pianos for clients including St Martin-in-the-Fields church (Trafalgar Square, London)

TEACHING EXPERIENCE

Machine Learning – Teaching Assistant: October 2009 – January 2010

Gatsby Computational Neuroscience Unit, University College of London

Assisted in teaching 25 postgraduate students for the duration of a course in Machine Learning with a strong Bayesian emphasis. Topics included: Bayesian inference, time series models, the EM algorithm, belief propagation, variational approximations, expectation propagation, convex relaxations, sampling and MCMC.

Math and Logic for Cognitive Science – Teaching Assistant: January 2008 – May 2008

Indiana University, Cognitive Science Programme

Assisted in teaching 30 postgraduate students for the duration of the course. Topics included: Hopfield networks, Bayesian networks, Markov Chains and Hidden Markov Models, Markov Decision Processes and Q-Learning, Entropy, Singular Value Decomposition and Latent Semantic Analysis.

Music Theory – Teaching Assistant: August 2006 – May 2007

Indiana University, Jacobs School of Music

Taught two separate classes (15 students each) to supplement lecture teaching. The Jacobs School of Music is consistently ranked within the Top 5 best music schools in the US for music theory.

PUBLICATIONS

Parikh, A.; Song, L.; Ishteva, M.; Teodoru G., and Xing, E. P.; [A Spectral Algorithm for Latent Junction Trees](#), *Proceedings of the 27th International Conference on Conference on Uncertainty in Artificial Intelligence (UAI 2012)*. Catalina Island, California, USA

Teodoru, G; Sahani, M, [Finding modal states in a discrete, sparsely connected, loopy graph by Iterated Cutset Conditioning](#). *NIPS 2010 Workshop on Discrete Optimization in Machine Learning, Whistler, Canada*

Diethe, T; Teodoru, G; Furl, N; Shawe-Taylor J. [Sparse Multiview Methods for Classification of Musical Genre from Magnetoencephalography Recordings](#). *Proceedings of the 7th Triennial Conference of European Society for the Cognitive Sciences of Music (ESCOM 2009) Jyväskylä, Finland*

Teodoru, G; Raphael, C, [Pitch Spelling with Conditionally Independent Voices](#), *Proceedings of the Sixth International Conference on Music Information Retrieval, ISMIR 2007* pp.201-206, Vienna, 2007.

PERSONAL INTERESTS

I am passionate about board games, and with every new game I try to find the optimal strategy or think about how I would design a good game AI. I am an assistant organizer of London's largest board gaming club, [London on Board](#), as well as a regular host, explaining game rules to new members and making them feel welcome at our club.

I regularly attend gymnastics classes, which have built the confidence necessary to attempt some dangerous moves, such as back handsprings and somersaults, and have taught me to multitask and think quickly when under pressure.

I have been educated as a professional musician. I regularly play piano and organ, and participate in the UCL Chamber Music Society. I am fascinated by fugues, and my approach to maths is influenced by my musical thinking.