

Belinda Thom

Curriculum Vitae

Computer Science Department
Carnegie Mellon University
5000 Forbes Avenue

(412) 268-2580
fax: (412) 268-5576
bthom@cs.cmu.edu

Pittsburgh, Pennsylvania 15213

<http://www.cs.cmu.edu/~bthom>

Research Interests:

My research is aimed at creating intelligent agents that collaborate—as first-class citizens—with humans in creative art-based domains. My interests are interdisciplinary. In my graduate career, I have specialized in applying Artificial Intelligence (AI) and Machine Learning (ML) techniques to the customization of interactive, computer music systems. The motivation has been to enable the computer’s equal participation in live, improvised musical exchanges—to bring the computer out of the box, so to speak.

Music Interests: improvisation; the hierarchical perception, generation, and representation of harmonic and melodic structure at different time-scales; automatic transcription and expression, principled methods for evaluating system performance and artistic merit; real-time interaction.

Machine Learning Interests: feature selection, Bayesian and semi-parametric statistical methods, unsupervised learning, Monte Carlo simulation, temporal sequence learning.

Education:

Ph.D., Computer Science, Carnegie Mellon University expected date: May 2001
Thesis: A Customized, Interactive Melodic Improvisation Companion
Advisor: Professor Manuela Veloso
Committee: Doctor Roger Dannenberg, Professor Tom Mitchell, and Professor David Wessel

M.S., Computer Science, Carnegie Mellon University 1997
Focus: Machine Learning and Interactive Computer Music Systems

B.S., Mechanical Engineering, University of California at Berkeley 1988
Graduated with honors

Professional Experience:

Carnegie Mellon University 1994 - present
Graduate Research Assistant, Computer Science Department

Lawrence Berkeley Laboratory 1992 - 1994
Associate Development Engineer, Astrophysics Group

Engineering Services Incorporated 1990 - 1992
Control Systems Engineer

Berkeley Process Control 1989 - 1990
Control Systems Engineer

Lawrence Berkeley Laboratory 1989
Research Assistant, Building Energy Analysis Group

Honors:

National Science Foundation Graduate Fellowship Carnegie Mellon University	1994 - 1997
National Physical Sciences Consortium Graduate Fellowship Award Alternate	1994
Chevron Academic Scholarship Mechanical Engineering Department, University of California at Berkeley	1988
Rotary Club Excellence in Physics Scholarship College of the Desert, Palm Desert	1985

Publications:

Peer-Reviewed Journal Articles

Interactive Improvisational Music Companions, Thom, B., To appear in *User Modeling and User-Adapted Interaction, Special Issue on User Modeling and Intelligent Agents*.

Refereed Conferences

Unsupervised Learning and Interactive Jazz/Blues Improvisation, Thom, B., *Proceedings of the Seventeenth Conference on Artificial Intelligence (AAAI-2000)*, 2000.

BoB: an Interactive Improvisational Companion, Thom, B., *Fourth International Conference on Autonomous Agents (Agents-2000)*, 2000.

Learning Melodic Models for Interactive Melodic Improvisation, Thom, B., *International Computer Music Conference (ICMC-99)*, 1999.

A Machine Learning Approach to Style Recognition, Dannenberg, R.B., Thom, B., and Watson, D., *International Computer Music Conference (ICMC-97)*, 1997.

Predicting Chords in Jazz, Thom, B. and Dannenberg, R.B., *International Computer Music Conference (ICMC-95)*, 1995.

Refereed Workshops

Artificial Intelligence and Real-time Interactive Improvisation, Thom, B., *AAAI-2000 Workshop on Artificial Intelligence and Music*, Seventeenth Conference on Artificial Intelligence, 2000.

Machine Learning and Musical Improvisation, Thom, B., *Agents-2000 / ECML-2000 Joint Workshop on Learning Agents*, Autonomous Agents Conference, 2000.

BoB: an Interactive Improvisational Companion, Thom, B., *Workshop on Interactive Robotics and Entertainment (WIRE-2000)*, 2000.

Predicting Chords in Jazz: the Good, the Bad, and the Ugly, Thom, B., *IJCAI Workshop on AI and Music*, International Joint Conference on Artificial Intelligence, 1995.

Service:

Cofounder / Coorganizer

Carnegie Mellon University, 1998 - 2000

School of Computer Science (SCS) Student Seminar Series:

An informal research seminar held by and for SCS graduate students, the purpose being to set up an explicit forum for presenting one's research to a general audience, in the hopes that interesting cross-disciplinary collaborations will result.

Founder / Organizer:

Carnegie Mellon University, 1995 - 1997

Alternative Machine Learning Reading Group

An informal reading seminar held by Machine Learning and Artificial Intelligence graduate students, the purpose being to provide a forum for reading about specific Machine Learning techniques and their application in non-traditional domains.

Teaching:

Introduction to Artificial Intelligence (15-381, Carnegie Mellon University):

Fall 1998

Guest lectures, designed projects, graded assignments.

Machine Learning (15-681, Carnegie Mellon University):

Fall 1997

Guest lectures, designed projects, graded assignments.

Research Projects:

Band-OUT-of-a-Box (BoB)

Carnegie Mellon University, 1997 – present

I am developing a real-time, interactive architecture that allows a user's (musician's) improvised examples to be used by various Machine Learning algorithms (unsupervised clustering, temporal sequence prediction, etc.). These algorithms are used to customize the musical behavior of an improvisational agent whose goal is to trade solos with the user in real-time (MIDI I/O is recorded and sequenced in milliseconds by an independent, high priority multimedia thread). Configuration (learning) amounts to fitting a probabilistic computational model of solo perception and generation to the user's improvisations. A tree encodes each bar of a solo's rhythm; summary statistics of the tree's leaves describe its tonality, continuity, and melodic contour. The model's abstract perception is learned by clustering these per-bar summaries into "musically similar" partitions. A key aspect of this learning paradigm is that it directly plugs into customized, interactive stochastic generation – the parameterizations learned for a given cluster can be used to control the walk through a graph of note nodes, allowing novel solo responses to be generated for a particular cluster. This research was conducted with Professors Manuela Veloso and Roger Dannenberg.

Probabilistic Histogram Clustering

Carnegie Mellon University, 1998 – 1999

I developed a new EM-based algorithm for clustering histograms. A probabilistic mixture-of-multinomials model was introduced and extended to handle histograms with varying sample-sizes (i.e., varying total numbers of counts). Data was simulated to quantitatively investigate the performance of the model under sparse conditions (histogram dimensions larger than sample-size). Qualitatively, I demonstrated that important musical abstractions emerge when this algorithm was trained on a bar-by-bar basis to the improvisations of Bebop saxophonist Charlie Parker. This research was conducted with Professor Manuela Veloso.

Local Gaussian Bayesian Classifiers

Carnegie Mellon University, Spring 1997

We developed a new algorithm for pattern classification that combined parametric and non-parametric techniques. Locally parametric classifiers were developed so that larger kernel widths could be used. This research was conducted with Professor Andrew Moore.

Learning Musical Style

Carnegie Mellon University, Summer and Fall 1996

Feature vectors were built from local windows of improvisations and labeled according to the musical style that a soloist was instructed to convey. This data set was used to train a classifier to predict playing style in real time. I developed a new classification algorithm that integrated Naïve Bayes with Principal Component Analysis in an attempt to reduce the effective complexity of the model. This research was conducted with Professors Roger Dannenberg and Andrew Moore.

Integrating GUIs and Real-time Systems

Carnegie Mellon University, Summer 1995

I integrated two toolkits developed at Carnegie Mellon University, the *Amulet* user interface environment and the *W* multi-tasking real-time environment. This research was conducted with Professors Roger Dannenberg and Brad Myers.

Predicting Harmony

Carnegie Mellon University, Fall 1995

I developed an n-gram model for predicting the next chord in a song given some previously played chords (history) and a training set. Training was based upon fifty randomly chosen jazz standards. The best performance was achieved with a hybrid model that used both on and off-line learning, and different history sizes. This research was conducted with Roger Dannenberg.

Engineering Projects:**Automated Supernova Telescope**

Lawrence Berkeley Laboratory

1992 - 1994

Engineered a turn-key automated supernova imaging telescope. Performed two simultaneous roles: Project Management and Engineering Development. Project Management included budget and manpower planning and procurement, coordination of scientific and multidisciplinary engineering efforts, and supervision of electrical, mechanical and optical support teams. I was the sole control systems engineer, designing, developing, and integrating all software and hardware for the telescope's observatory control systems, including SPARC-based multitasking supervisory host control, BAM-based real-time multi-axis motion control, PLC-based real-time data acquisition, and real-time fail-safe error handling.

Turn-key Control Systems Integration

Engineering Services, Incorporated

1990 - 1992

Designed and developed hardware and software for turn-key control systems in petrochemical and pharmaceutical industries. Also gained experience in related areas, including instrumentation engineering, control panel design, project procurement and management, system start-up, and field troubleshooting.

Automated Production Test Facility

Berkeley Process Control

1989 - 1990

Engineered and maintained an automated production testing facility for verifying the integrity of the multi-axis motion controllers assembled on-site. Designed and implemented all software and hardware for checking over 150 functions, including motion control, high-speed latching, and temperature and power cycling.

References:**Professor Manuela Veloso**

Computer Science Department

Carnegie Mellon University

Pittsburgh, PA 15213

412-268-1474

mmv@cs.cmu.edu

Doctor Roger Dannenberg
Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213
412-268-3827
rbd@cs.cmu.edu

Professor David Wessel
Director of the Center for New Music and Technology
Department of Music
University of California at Berkeley
Berkeley, California 94709
510-643-9990 x302
wessel@cnmat.berkeley.edu

Doctor Tucker Balch
Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15213
412-268-1780
trb@ri.cmu.edu

Doctor Richard Treffers
Department of Astronomy
University of California at Berkeley
Berkeley, California 94709
925-284-9403
rtreffers@astro.berkeley.edu

Professor Pamela Eibeck
Chair, Department of Mechanical Engineering
University of Northern Arizona
Flagstaff, Arizona 94709
520-523-8175
Pamela.Eibeck@nau.edu