

Imparting Knowledge and Skills at the Forefront of Interdisciplinary Research – a case study on course design at the intersection of music and engineering

Elaine Chew

Viterbi Early Career Assistant Professor, University of Southern California Viterbi School of Engineering,
Integrated Media Systems Center, and Epstein Department of Industrial and Systems Engineering
Los Angeles, CA 90089-0193 echew@usc.edu

Abstract – This paper presents a constructivist approach to advanced course development in an interdisciplinary field by focusing on a case study at the intersection of music and engineering. Engineering approaches to music modeling and cognition is a fast growing and rapidly evolving focus of study due to the explosion of digital music information. The challenges of designing a course to impart the expertise necessary to innovate at the forefront of the field include the lack of a formal body of scholarly knowledge in the form of a text, the lack of formal structures to support the course in the traditional disciplinary framework, the relatively small number of students with suitably strong backgrounds in both computing and music, and misconceptions about the nature of music research. The paper describes how these obstacles were surmounted. An outcome of the course was dissemination of organized knowledge and example projects through the web as a resource for the community.

Index Terms – Interdisciplinary curriculum design, Music and Engineering, Constructivist approach, Hands-on learning.

INTRODUCTION

Engineering is increasingly interfacing with traditionally disparate fields to push the frontiers of human knowledge. The spotlight on interdisciplinary research has necessitated the development of new courses that bridge disciplines and encourage lateral thinking. Training the next generation of engineers to acquire the breadth of expertise required to innovate at the frontiers of interdisciplinary knowledge poses unique challenges for the design of course curricula. This paper presents one solution to advanced course development in a rapidly growing, and continually evolving, interdisciplinary domain by focusing on an example at the intersection of music and engineering – the computer modeling of music processes.

The building of computational models for music processes can be traced back to the theoretical chemist and cognitive scientist, Christopher Longuet-Higgins (in his Letters to a Musical Friend in the 60s [1][2], and computational models in the 70s and 80s, see for example [3]). Longuet-Higgins was also a gifted amateur musician [4] – he was a pianist, conductor and composer – and he took the

seminal steps towards advancing the scientific understanding of music, its structures, and performance. The study of music cognition became an academic specialty in the 80s. The new millennium saw a surge in interest in the field and newcomers from the engineering disciplines, due in part to the maturing of computing technology, and the explosion of digital music information. There exists today renewed and rapidly growing interest in automating the processing and machine understanding of digital music, not only for artificial intelligence, but also for music retrieval by content. The exponential growth in interest in the field can be measured by that in refereed publications generated by newly created conferences such as the International Conferences on Music Information Retrieval (ISMIR, [5], founded 2000), Sound and Music Computing (SMC, [6], founded 2004), and Computer Music Modeling and Retrieval (CMMR, [7], founded 2003).

This paper describes the creation, development, and evolution of a course titled *Topics in Engineering Approaches to Music Cognition*. The goals of the course are: to introduce graduate and advanced undergraduate students to cutting-edge knowledge and research in the computational modeling of music processes; and, to lay the foundations for advanced scientific research at the forefront of the field (including thinking and articulation skills). In the course, the students actively review and present current literature on computational research in music cognition (learn by example), and gain hands-on experience through the design and implementation of a music and computing project (learn by doing). Selected reports and presentations, and all final projects, are posted online as open courseware (www-scf.usc.edu/~ise575), serving as a resource to the community. The course, created in Spring 2003, appeared in its fourth incarnation in Spring 2006. Due to the exponential growth in research in the field, the course is now subdivided into semester-long, focused treatments of special topics. The paper will detail the design goals and principles in the crafting of the course content, including synergistic activities organized to augment the course.

The challenges to mounting and sustaining such a course include the lack of a formal body of knowledge in the form of a text, the lack of formal academic structures to support the course, the relatively small number of students with suitably strong backgrounds in both computing and music, and misconceptions about the nature of music research. The bulk

of the growth in literature in the field has taken place only in the past five years, presenting challenges to the organization of material for a class. As a fast growing field that is still in its infancy, there is often not a natural home in the traditional departmental structures of a university through which one might offer a course on the subject. Mounting a new and interdisciplinary course that does not quite fit in any one department requires an awareness of academic structures beyond that necessary for the teaching of a traditional class. Extra effort also needs to be taken to ensure that students from multiple disciplines who might be interested in, and benefit from, the course know about it.

The interdisciplinary approach espouses a knowledge view that applies tools and language from more than one discipline to examine a central theme or topic (paraphrase of quote by Jacobs, [8] in [9]). An engineering course that focuses on an application area, especially one in the arts, rather than predominantly scientific methods, can often be viewed with suspicion within the traditional and abstract methodology-driven mold. Hence, open communication with one's colleagues about the content and goals of the course can be as important as educating the students in the class. The computational modeling of music cognition draws upon methodologies and tools from music theory, cognitive science, artificial intelligence, experimental psychology, mathematics, signal processing, acoustics, psychoacoustics, and neuroscience. Few, if any, students enrolled in the class are equipped with the requisite knowledge to apprehend all the material. Cross-disciplinary knowledge acquisition is facilitated by opening the course to students not only in engineering, but also in music and psychology. The course adopts a constructivist approach to facilitate learning (see Piaget's theory from [10], summarized in [11]). Care is taken in the design of the course to ensure that each student can gain confidence by building upon domain knowledge that s/he possesses, and acquiring new knowledge through collaborative exchange in a case-based and hands-on learning environment.

The remainder of the paper is organized as follows: the next section expands on the challenges and solutions to the design and implementation of a course on computational modeling of music processes. Examples are provided using the last two instantiations of the course (one in the survey format, and one topical example on computational modeling of expressive performance) as case studies. Synergistic and supporting activities are described. The paper concludes with a summary of the findings.

CHALLENGES AND SOLUTIONS

This section addresses each of the challenges to creating, developing, and sustaining *Topics in Engineering Approaches to Music Cognition*, the interdisciplinary course on computational modeling of music processes, and some concrete solutions to the problems.

I. The Structuring of Content

Since the computational modeling of music processes is still a burgeoning field whose body of knowledge is yet amorphous

and undergoing metamorphosis, the design of a course in this area of study requires first a structuring of the available content. Unlike more mature fields such as Mechanics or Algebra, it may not be immediately obvious how one might organize the growing body of information in the field.

The structure of the original survey course (offered 2003-2005) was formed after consulting numerous course syllabi in the related areas of music theory and music cognition at the Eastman School of Music, at Ohio State University, and at Indiana University, Bloomington, and Temperley's structuring of content in his book *The Cognition of Basic Musical Structures* [12]. Admittedly, these courses based in music schools and departments focused more on the organization of music theoretic knowledge rather than technical applications; however, they draw inspiration from music pedagogy, with a tradition of organizing music domain knowledge that has evolved over centuries. Some guiding principles gleaned from this background research resulted in the structuring of the music knowledge part of the course into the analysis of pitch vs. time structures, and within the pitch domain, the analysis of horizontal (linear) vs. vertical (simultaneous) grouping structures. See, for example, the organization of material in Weeks 4-11 (labeled 'Musical Structures') in the Spring 2005 survey course shown in Figure 1, and available online at www-scf.usc.edu/~ise575/a/syllabus. Earlier examples can be located through www-scf.usc.edu/~ise575.

The *Topics in Engineering Approaches to Music Cognition* course aims to fulfill more than the goal of imparting music cognition and theoretic knowledge. It also strives to associate these concepts with engineering approaches that could provide automated (computational) solutions to the modeling of these processes and structures. The design goals of the class are threefold:

- To provide an introduction to the computational modeling of music processes;
- To lay the foundations for advanced scientific research at the forefront of the field; and,
- To build a structured repository of knowledge in the field as a resource for the broader community.

As a result of participating in the course, by the end of the semester, each student should:

- Understand basic musical structures, and possess the vocabulary with which to describe them;
- Be capable of manipulating digital music in different formats;
- Be able to create computational/automatic means of analyzing, generating, and visualizing music; and,
- Be able to formulate a scientific question and build the computational tools to answer it.

The structure discussed thus far would only address the first action goal. To address the practical goal of being able to handle and manipulate digital music information, a module was added to survey digital music representation in 2005 (see Figure 1, Week 3), and to survey tools for analyzing expression in 2006 (see Figure 2, Week 10). To allow students

to learn by example, several weeks of surveying existing interactive applications in the field were added (see Figure 1, Weeks 12-15, and Figure 2, Weeks 13-14). The third design and fourth action goals will be addressed in the next section.

Home	Syllabus	Week-by-week Guide	Homework FAQ	Archives	Projects
<p>ISE575 CSE575 EE675</p> <p>Engineering Approaches to Music Perception and Cognition</p> <p>Spring 2005 Daniel J. Epstein Department of Industrial and Systems Engineering University of Southern California Andrew and Erna Viterbi School of Engineering</p>					
Introduction					
Week 1					
<p>Introduction [PDF] [HTML] [PDF] Signup information sheet [PDF] Reading plan [PDF]</p> <p>Course syllabus [PDF] Class list [HTML] [PDF] Signup information sheet [PDF] Reading plan [PDF]</p>					
Week 2					
<p>Linking Music</p> <p>Impromptu by Jeanne Bamberg and Armando Hernandez, available at www.bamberg.com</p> <p>Jeanne Bamberg (2004). "The development of intuitive musical understanding: a natural experiment." <i>Psychology of Music</i> 30(1):7-30. [PDF] presentation: [link]</p> <p>Hyperscore by Mary Farhood and Egon Pätzold, available at www.hyperscore.com</p> <p>Kevin Jennings (2004). "Compositional Music: Emulation of a Core Process." <i>Special Issue of ICMPC - Learning with ICT for C21st</i>. [PDF] presentation: [link]</p> <p>Poulenc's Concerto for Two Pianos performed by the Tostoff Piano Duo, Vely Stoyanova and Ila Tostoff, and the USC Thornton Symphony. More info about the duo at www.tostoffpiano.com. Concert info: [link] [reports] [links] [reports]</p> <p>Sample student homeworks and compositions created using Impromptu and Hyperscore: [link] [reports] [links] [reports]</p> <p>Other cool compositions: [link]</p>					
Week 3					
<p>Musical Structures</p> <p>Musical Structures by CAPRI, available at www.musicalstructures.com</p> <p>Walter B. Hewitt (1997). "The MuseData: Representation of Musical Information." In Beyond MIDI: Eleanor Selfidge-Field (ed). The Handbook of Musical Codes. [HTML] presentation: [link]</p> <p>The Humdrum Toolkit by David Huron, available at cml.southern.edu/humdrum</p> <p>David Huron (1999). "Humdrum: A Brief Tour." Ohio State University. [HTML] presentation: [link]</p> <p>GUIDO Music Notation by the SALLER Project group, available at www.musicalstructures.com/guido</p> <p>Hoger Hoon, K.A Hamel, Kai Renc and Juergen Kilian (1998). "The GUIDO Music Notation Format - A Novel Approach for Adequately Representing Score-level Music." ICMI-95-454. [PDF] presentation: [link]</p> <p>Kai Renc and Hoger Hoon (1998). "A Web-based Approach to Music Notation Using GUIDO." ICMI-95-458. [PDF] presentation: [link]</p> <p>Homework 2: [PDF]</p>					
Week 4					
<p>Articulation</p> <p>Articulation Beat</p> <p>Simon Dixon (2001). "Automatic Extraction of Tempo and Beat From Expressive Performances." <i>JNMR</i> 30(1). [PDF] presentation: [link] [reports] [links] [reports]</p> <p>Masataka Goto (2001). "An audio-based real-time beat tracking system for music with or without drum-sounds." <i>JNMR</i>, 30(2):159-171, 2001. [PDF] presentation: [link]</p> <p>Eric Schreier (1998). "Tempo and beat analysis of acoustic musical signals." <i>J. Acoust. Soc. Am.</i> 103(1): 588-601. [PDF] presentation: [link] [reports] [links] [reports]</p> <p>George Zemanek, George Esli, and Pary Cook (2002). "Human perception and computer extraction of beat strength." <i>DAFX</i>. [PDF] presentation: [link] [reports] [links] [reports]</p>					
Week 5					
<p>Rhythm and Meter (conducted by Anja Volk)</p> <p>Anja Volk. "The Generating of Metrical Hierarchies using Inner Metric Analysis" manuscript. reports: [link] [reports] [links] [reports]</p> <p>Douglas Eck (2001). "A Probabilistic Evidence Model for Rhythmic Beat Induction." <i>JNMR</i> 30, pp. 187-200. [PDF] presentation: [link] [reports] [links] [reports]</p> <p>Fabian Gouyon and Simon Dixon (2003). "A review of automatic rhythm detection systems." <i>CMU</i> 29 [PDF] presentation: [link] [reports] [links] [reports]</p> <p>Anssi Klupar, Antti Eronen and Jaakko Astola (2004). "Analysis of the Meter of Acoustic Musical Signals." <i>IEEE Trans. Speech and Audio Processing</i>. [PDF] presentation: [link] [reports] [links] [reports]</p>					
Week 6					
<p>Expressive Performance</p> <p>Samer Abdallah and Mark Plumbley (2004). "Polyphonic Music Transcription by non-negative sparse coding of power spectra." <i>ISMIR</i> [PDF] presentation: [link] [reports] [links] [reports]</p> <p>Alan de Chevigne and Heikki Kaushar (2002). "YIN, a fundamental frequency estimator for speech and music." <i>J. Acoust. Soc. Am.</i> 111:1971-1930. [PDF] presentation: [link] [reports] [links] [reports]</p> <p>Anssi Klupar, Tuomas Virtanen, Antti Eronen, Jarmo Seppänen (2001). "Automatic Transcription of Musical Recordings." Workshop on Consistent and Reliable Acoustic Cues for Speech Analysis. [PDF] presentation: [link] [reports] [links] [reports]</p> <p>Malcolm Slaney and Richard J. Lyon (1990). "A perceptual pitch detector." <i>ICASSP</i> [1:357-360]. [PDF] presentation: [link] [reports] [links] [reports]</p>					
Week 7					
<p>Formal Models</p> <p>Emilia Caimponazzo (1996). "A Formal Theory for the Discovery of Local Boundaries in a Melodic Surface." Proceedings of the III Journées d'Informatique Musicale. [PDF] presentation: [link] [reports] [links] [reports]</p> <p>Elaine Chew and Xiaodan Wu (2004). "Separating Voices in Polyphonic Music: A Contig Mapping Approach." <i>CMMR</i>. Springer Verlag LNCS #310. [PDF] presentation: [link] [reports] [links] [reports]</p> <p>Conklin and Anagnostopoulou (to appear). "Segmental Pattern Discovery in Music." <i>INFORMS JOC</i>. presentation: [link] [reports] [links] [reports]</p> <p>Elzabeta Margala (to appear). "A Model of Melodic Expectation." <i>Music Perception</i>. presentation: [link] [reports] [links] [reports]</p>					
Week 8					
<p>Complexity</p> <p>Steedman M J (1984). "A Generative Grammar for Jazz Chord Sequences." <i>Music Perception</i> 2:1-27. presentation: [link] [reports] [links] [reports]</p> <p>Steedman M J (1996). "The Blues and the Abstract: Truth: Music and Mental Models." In: Garnham A., Oakhill J. (eds.), <i>Mental Models in Cognitive Science</i>. Psychology Press, pp.305-311. [PDF] presentation: [link] [reports] [links] [reports]</p> <p>Nitta Spro (2002). "Combining Grammar-Based and Memory-Based Models of Perception of Time Signature and Phase." In C. Anagnostopoulou, M. Fernand, and A. Small, eds. <i>Music and Artificial Intelligence</i>. Berlin: Springer, pp. 183-194. [PDF] presentation: [link] [reports] [links] [reports]</p>					
Week 9					
<p>Complexity</p> <p>Elaine Chew (to appear). "Slicing it all ways: mathematical models for tonal induction, approximation and segmentation using the Spiral Array." <i>INFORMS JOC</i> [link] [reports] [links] [reports]</p> <p>Emilia Gouyon (to appear). "Tonal Description of Polyphonic Audio by Music Content Processing." <i>INFORMS JOC</i>. presentation: [link] [reports] [links] [reports]</p> <p>Carol Krumhansl (1990). "The Cognitive Foundations of Musical Pitch." <i>Cognitive Foundations of Musical Pitch</i>. Oxford University Press. presentation: [link] [reports] [links] [reports]</p> <p>Ozgur Ilgin and Semih Bilgen (1998). "A Model for Tonal Context Time Course Calculation from Acoustical Input." <i>JNMR</i> 25(5). presentation: [link] [reports] [links] [reports]</p>					
Week 10					
<p>SPRING BREAK</p>					
Week 11					
<p>LINKS TO PAPER PRESENTATIONS, REPORTS</p> <p>David Temperley (2002). "Cognition of Basic Musical Structures." MIT Press, Chapters 1 (intro/skms), 2 (meter) and 6 (harmony).</p> <p>Melisma Music Analyzer by David Temperley and Daniel Sleator, available at www.litk.cs.cmu.edu/music-analysis</p> <p>Homework #1: [PDF]</p>					
Interactive Systems					
Week 12					
<p>Generating Music</p> <p>Gerard Assayag, Shimoju Dubrov (2004). "Using Factor Oracles for Machine Improvisation." <i>IEEE Int. Computing, Special Issue on Formal Systems and Music</i>. [PDF] presentation: [link] [reports] [links] [reports]</p> <p>Daniel Conlon. "Music: Generation from Statistical Models." <i>AISB 2003 Symposium on AI and Creativity in the Arts and Sciences</i>, pp. 30-35. [PDF] presentation: [link] [reports] [links] [reports]</p> <p>Francis Pachet (2003). "The Continuator: Musical Interaction with Style." <i>JNMR</i> 32(3):333-344. [PDF] presentation: [link] [reports] [links] [reports]</p> <p>The Continuator by Francis Pachet, available at www.sony.com/fr/pachet</p> <p>Belinda Thorn (2003). "Interactive Improvisational Music Companionship: A User-Modeling Approach." <i>User Modeling and User-Adapted Interaction Journal</i>. Special Issue on User Modeling and Intelligent Agents. [PDF] presentation: [link] [reports] [links] [reports]</p>					
Week 13					
<p>Creating Expression</p> <p>Josep Lluís Arcos, Ramon Lopez de Mantaras and Xavier Serra (1998). "SaxEx: A Case-Based Reasoning System for Generating Expressive Musical Performances." <i>JNMR</i> 27(1). [PDF] presentation: [link] [reports] [links] [reports]</p> <p>Roberto Bresin (1998). "Artificial Neural Networks Based Models for Automatic Performance of Musical Scores." <i>JNMR</i> 27(1). [link] [reports] [links] [reports]</p> <p>Sunberg J., Friberg A., and Bresin R (2003). Attempts to reproduce a pianist's expressive timing with Director Music's performance rules. <i>JNMR</i> 32(3):317-326. [PDF] presentation: [link] [reports] [links] [reports]</p> <p>Werner Goebel, Elias Pampalk and Gerhard Widmer (2004). "Exploring Expressive Performance Trajectories: Six Famous Pianists Play Chopin Pieces." <i>ICMPC</i>. [PDF] presentation: [link] [reports] [links] [reports]</p>					
Week 14					
<p>Complexity</p> <p>OpenMusic by IRCAM, available at sourceforge.net/projects/ircam-openmusic</p> <p>an open-ended virtual programming environment for music composition based on CommonLisp/CLOS</p> <p>Midi Toolset for Matlab from the U of Jyväskylä, available at www.midi-toolset.com</p> <p>Nyquist by Roger Dannenberg, available at www2.cs.cmu.edu/~dmc/musicml/software.html</p> <p>Improvising by Greg Slep, available at www.usc.edu/~djp</p> <p>C++ environment for writing programs that enable musician/computer interaction using MIDI instruments.</p>					
Week 15					
<p>Complexity</p> <p>Christopher Raphael (2004). "Aligning Musical Scores with Audio using Hybrid Graphical Models." <i>AAAI</i> 2004. [PDF] [reports] [links] [reports]</p> <p>Christopher Raphael and Joshua Sissold (2003). "Harmonic Analysis with Probabilistic Graphical Models." <i>ISMIR</i> 2003. [PDF] [reports] [links] [reports]</p> <p>Christopher Raphael (2003). "Orchestral Music Accompaniment from Synthesized Audio." <i>ICMC</i>. [PDF] [reports] [links] [reports]</p> <p>Christopher Raphael (2001). "A Bayesian Network for Real-Time Musical Accompaniment." <i>NIPS</i>. [PDF] [reports] [links] [reports]</p>					
Week 16					
<p>PROJECTS [Final Projects] [Pictures] [Students]</p>					

FIGURE 1

ONLINE SYLLABUS (2005) + LINKS TO PRESENTATIONS, REPORTS.

To keep pace with rapid developments in the field, the course underwent further revision and expansion to assume a topical form in Spring 2006. Each one-semester part now covers some topic that can be classified generally into the traditional musical categories of analysis, performance, and composition. The Spring 2006 course focused on computational modeling of expressive performance; its syllabus is shown in Figure, 3 and available online at www.scf.usc.edu/~ise575/b/syllabus. This is the first time a course with this topic is being offered, and the material has been organized into the descriptors of performance (tempo, dynamics, articulation), a survey of

analysis and synthesis tools, and a review of some interactive systems, book-ended by survey articles and musicians' views. Other aspects of the course such as the presentation format, reporting and project requirements remain the same.

Home	Syllabus	Week-by-week	Homework FAQ	Archives	Final Projects
<p>ISE575 CSE575 EE675</p> <p>Engineering Approaches by Music Cognition</p> <p>Computational Modeling of Expressive Performance</p> <p>Spring 2006 Daniel J. Epstein Department of Industrial and Systems Engineering University of Southern California Andrew and Erna Viterbi School of Engineering</p>					
Introduction					
Week 1					
<p>Course syllabus [PDF] Class list [HTML] [PDF] Signup information sheet [HTML] Reading plan [PDF]</p> <p>Checklist</p> <p>Palmer, C. (1997). <i>Music Performance. Annual Review of Psychology</i>, 48:115-38. [PDF] presentation: [link]</p> <p>Gabrielsson, A. (1999). <i>Music Performance</i>. In D. Deutsch (Ed.), <i>The Psychology of Music</i>, Vol.2, pp.501-602. San Diego, CA: Academic Press.</p> <p>Read and write a report on one of the following: set A (p.501-523) OR set B (p.523-557, 577-579). presentation (A): [link] reports (A): [link] [link] [link] [link]</p>					
Week 2					
<p>Tempo, Dynamics, Articulation</p> <p>Week 3</p> <p>Todd N. P. (1992). The dynamics of dynamics: A model of musical expression. <i>Journal of the Acoustical Society of America</i>, June 1992, Vol. 91, issue 6, pp.3540-3550. [PDF] presentation: [link] [reports] [links] [reports]</p> <p>Repp, B. H. (1998). A microsecond of musical expression: I. Quantitative analysis of pianists' timing in the initial measures of Chopin's Etude in E major. <i>Journal of the Acoustical Society of America</i>, Aug 1998, Vol. 104, issue 2, pp.1085-1100. [PDF] presentation: [link] [reports] [links] [reports]</p> <p>Repp, B. H. (1999). A millisecond of musical expression: II. Quantitative analysis of pianists' dynamics in the initial measures of Chopin's Etude in E major. <i>Journal of the Acoustical Society of America</i>, Mar 1999, Vol. 105, issue 3, pp.1972-1988. [PDF] presentation: [link] [reports] [links] [reports]</p> <p>Playing Vivaldi's Sonata No. 4 by Eric Cheng. presentation: [link]</p>					
Week 4					
<p>Todd N. P. M. (1998). The kinematics of musical expression. <i>Journal of the Acoustical Society of America</i>, March 1998, Vol. 97, issue 3, pp.1940-1949. [PDF] presentation: [link] [reports] [links] [reports]</p> <p>Friberg A., Sunberg J., and Frykholm, L. (2000). Music from motion: Sound level envelopes of tones expressing human locomotion. <i>Journal of New Music Research</i>, 29 (2000), No. 3, pp. 199-210 [PDF] presentation: [link] [reports] [links] [reports]</p>					
Week 5					
<p>Analysis and Prosody</p> <p>Bresin, R. & Ballew, C. U. (2000). Articulation strategies in expressive piano performance. <i>Journal of New Music Research</i>, Vol.29, No. 3, Sep. 2000, pp.211-224. [PDF] presentation: [link] [reports] [links] [reports]</p> <p>Palmer, C. (2003). The Contribution of Invariant and Performed Aspects to Emotional Expression in Short Tone Sequences. <i>Journal of New Music Research</i>, Vol.32, No.3, Sep.2003, p.289-302. [PDF] presentation: [link] [reports] [links] [reports]</p> <p>Lindner, C. & Hildrich, S. (2005). What is musical prosody? In B. H. Ross (Ed.), <i>Psychology of Learning and Motivation</i> 44. Elsevier Press. presentation: [link] [reports] [links] [reports]</p>					
Week 6					
<p>Musical Structures</p> <p>Palmer, C. (1998). On the assignment of structure in music performance. <i>Music Perception</i>, 14: 23-56. [PDF] presentation: [link] [reports] [links] [reports]</p> <p>Mazzola, G. & Goller, S. (2002). Performance and Interpretation. <i>Journal of New Music Research</i>, Vol.31, No.3 (September 2002) pp.221-232. [PDF] presentation: [link] [reports] [links] [reports]</p> <p>Friberg A., Sunberg J., and Frykholm, L. (2000). Music from motion: Sound level envelopes of tones expressing human locomotion. <i>Journal of New Music Research</i>, Vol.32, No.3, Sep.2003, p.289-302. [PDF] presentation: [link] [reports] [links] [reports]</p> <p>De Paik G. (2004). Musical Expression for Expressiveness Modeling and for Music Performance. <i>Journal of New Music Research</i>, Vol.33 No.3, Sep.2004, pp.189-202. [PDF] presentation: [link] [reports] [links] [reports]</p>					
Analytical and Generative Tools					
Week 7					
<p>Visualization</p> <p>Langner J. & Goebel W. (2002). Visualizing Expressive Performance in Tempo-Loudness Space. <i>Computer Music Journal</i>, Vol.27 No.4, Winter, 2003, pp.69-83. [PDF] presentation: [link] [reports] [links] [reports]</p> <p>Goebel W., Pampalk, E. & Widmer, G. (2004). "Exploring Expressive Performance Trajectories: Six Famous Pianists Play Chopin Pieces." In Proceedings of the 8th International Conference on Music Perception and Cognition. [PDF] presentation: [link] [reports] [links] [reports]</p>					
Week 8					
<p>Analysis</p> <p>Widmer, G. & Goebel W. (2004). Computational Models of Expressive Music Performance: The State of the Art. <i>Journal of New Music Research</i>, Vol.33, No.3, Sep.2004. [PDF] presentation: [link] [reports] [links] [reports]</p> <p>Langner J., Kopeck, R., Stoffel, C., Witt, M. (2003). Real-time Analysis of Dynamic Shaping in Proceedings of the International Conference on Music Perception and Cognition, 5-10 Aug 2003, Keele, England. [PDF] presentation: [link] [reports] [links] [reports]</p>					
Week 9					
<p>Rule-based Models</p> <p>Friberg A. (1991). Generative Rules for Music Performance: A Formal Description of a Rule System. <i>Computer Music Journal</i>, Vol.15 No.2, pp.56-71. presentation: [link] [reports] [links] [reports]</p> <p>Friberg A., Colombo, V., Frykholm, L., Sunberg, J. (2000). Generating Musical Performances with Director Music's Models. <i>Computer Music Journal</i>, 2000. presentation: [link] [reports] [links] [reports]</p> <p>Sunberg, A., Friberg, A., & Bresin, R. (2003). Attempts to Reproduce a Pianist's Expressive Timing with Director Music's Performance Rules. <i>Journal of New Music Research</i>, 2003. [PDF] presentation: [link] [reports] [links] [reports]</p>					
Week 10					
<p>Complexity</p> <p>Arcos, J. L. de Mantaras, R. L. & Serra, X. (1998). SaxEx: A Case-Based Reasoning System for Generating Expressive Musical Performances. <i>Journal of New Music Research</i>, Vol.27 No.3 (Sept) presenter: [link] [reports] [links] [reports]</p> <p>Arcos, J. L. de Mantaras, R. L. (2001). An Interactive Case-Based Reasoning Approach for Generating Expressive Music. <i>Applied Intelligence</i>, Vol.14 No.1, Jan.2001, pp.115-129. [PDF] presentation: [link] [reports] [links] [reports]</p> <p>TOOLS for alignment collected by Miles, Cheng, and Uhai</p> <p>TOOLS for features/parameters collected by Lin, Chiang, and Jer</p> <p>TOOLS for tempo tracking collected by Lin, Chiang, and Chung</p> <p>TOOLS for visualization collected by Chuan, Maddoxian, Moost, and Vaziri</p>					
Week 11					
<p>Machine Learning</p> <p>Bresin, R. (1998). Artificial neural networks based models for automatic performance of musical scores. <i>Journal of New Music Research</i>, Vol.27 No.3, pp.239-270. presentation: [link] [reports] [links] [reports]</p> <p>Imonen, T. & Takala, T. (1999). Conductor following with artificial neural networks. In Proceedings of the International Computer Music Conference, 1999. [PDF] presentation: [link] [reports] [links] [reports]</p> <p>Raphael, C. (2001). Can the Computer Learn to Play Music Expressively? <i>Artificial Intelligence and Statistics</i>, 2001. [PDF] presentation: [link] [reports] [links] [reports]</p> <p>Widmer, G. (2002). Machine Discoveries: A Few Simple, Robust Local Expression Principles. <i>Journal of New Music Research</i>, Vol.31, No.1, Mar.2002, pp.37-50. [PDF] presentation: [link] [reports] [links] [reports]</p> <p>Widmer, G. & Todorov, A. (2003). Playing Mozart by Analogy: Learning Multi-level Timing and Dynamics Strategies. <i>Journal of New Music Research</i>, Vol.32 No.3, Sep.2003, pp.259-266. [PDF] presentation: [link] [reports] [links] [reports]</p>					
Week 12					
<p>Interactive Systems</p> <p>The Radio Baton</p> <p>Mathews, M. V. (1991a). The Conductor Program and Mechanical Baton. In M. V. Mathews and J. R. Pierce, eds. <i>Current Directions in Computer Music Research</i>. Cambridge, Massachusetts: MIT Press, pp.263-281. presentation: [link] [reports] [links] [reports]</p> <p>Mathews, M. V. (1991b). The Radio Baton and Conductor Program, or, Pitch: The Most Important and Least Expressive Part of Music. <i>Computer Music Journal</i>, Vol.15 No.4, pp.37-46. presentation: [link] [reports] [links] [reports]</p> <p>Sunberg, J., Friberg A., Mathews, M., Bennett, G. (2003). Experiences in Combining the Radio Baton and the Director Music's Performance Grammar. Workshop on Current Research Directions in Computer (MOSART) Barcelona, Spain, Nov.15-17, 2001. [PDF] presentation: [link] [reports] [links] [reports]</p> <p>Mathews, M., Friberg, A., Bennett, G., Sapp, C., Sunberg, J. (2003). A marriage of the Radio Baton program and the conductor program: the proceedings of the Stockholm Music Acoustics Conference, August 6-9, 2003 (SMAC'03), Stockholm, Sweden, reports: [link] [reports] [links] [reports]</p>					
Week 13					
<p>Other Interfaces</p> <p>Dixon, S., Goebel, W., Widmer, G. (2005). The "Air Worm": an interface for Real-Time Manipulation of Expressive Musical Performance. Proceedings of the International Computer Music Conference. [PDF] presentation: [link] [reports] [links] [reports]</p> <p>Moe, E., François, A., R. Liu, J. & G. (2005). ESP: A Driving Interface for Musical Expression Synthesis. Conference on New Interfaces for Musical Expression (NIME), Vancouver, B.C., Canada, May 28-28, 2005. [PDF] presentation: [link] [reports] [links] [reports]</p> <p>Wendley M.M. & Ohl, D. (2002). Evaluation of a Driving Interface for Musical Expression: Borrowing Tools from HCI. <i>Computer Music Journal</i>, Vol.26, No.3, Sep. 2002, pp.67-76. [PDF] presentation: [link] [reports] [links] [reports]</p>					
Performers' Perspectives					
Week 15					
<p>Three Pianists' Opinions</p> <p>Bresin, R. (1976). <i>Musical Thoughts and Allen-Thoughts</i>. London: Robson Books. Selected essays. presenter: [link] [reports] [links] [reports]</p> <p>Woff, R. (1979). Schrabal's Interpretation of Piano Music. London: Faber Music in association with Faber and Faber. Chapter 10: Interpretation of Character and Structure, p.120-153. presenter: [link] [reports] [links] [reports]</p> <p>Woff, R. (1979). Schrabal's Interpretation of Piano Music. London: Faber Music in association with Faber and Faber. Chapter 11: The Means of Regulating Sound, Chapter 12: Technique, and Appendix, p.154-182. presenter: [link] [reports] [links] [reports]</p> <p>Chaffin, R. G. Smith, A. F. Leman, and G. Chen (2003). "Being the Big Picture": Piano Practice as Expert Problem Solving. <i>Music Perception</i>, Vol.20, No.4 (Summer 2003), p. 465-490. reports: [link] [reports]</p>					
Week 16					
<p>PROJECTS [Final Projects] [Pictures] [Students]</p>					

FIGURE 2

ONLINE SYLLABUS (2006) + LINKS TO PRESENTATIONS, REPORTS.

II. Building Knowledge in an Interdisciplinary Domain

In order to properly address the central theme of computational modeling of music cognition, one needs to draw upon methodologies and tools from several disciplines, including music (theory, composition, and performance), the cognitive sciences (theory, composition, and performance), the cognitive sciences (artificial intelligence, experimental psychology, psychoacoustics, and neuroscience), and applied mathematics (signal processing, acoustics, algorithm design and implementation). Few students, if any, are equipped with the disciplinary know-how from more than one or two of these

October 28 - 31, 2006, San Diego, CA

specialties. Collaborative learning with cohorts from multiple disciplines is facilitated by opening the class enrolment to students from a variety of engineering disciplines, including electrical engineering (EE), computer science (CS), and industrial and systems engineering (ISE), as well as in music and psychology. The varied backgrounds the students bring to the classroom, each of which is pertinent to some aspect of the course material, and the interdisciplinary nature of computational music cognition, points to the need for a constructivist approach to course instruction. The musician possesses domain-specific knowledge, the computer scientist knowledge of algorithms and data structures, the electrical engineer essentials of digital signal processing, and the psychologist experimental design and statistical know-how.

The course aims to create an environment for case-based and hands-on learning. As a quick induction into the process of music-making, Week 2 in Spring 2005 focused on the creating of short musical pieces using intuitive interfaces that do not require note-reading ability, such as Bamberger and Hernández's *Impromptu* [13] and Farbood and Pasztor's *Hyperscore* [14]. Through experimentation, even students who have not yet engaged in active music making can experience firsthand the basic concepts of musical design and structure. As much as possible, each student is responsible for her/his own learning, and is assigned a few papers to present over the course of the semester. The student is expected not only to summarize and explain the content of the paper, but to seek out and demonstrate additional examples or associated software the author may have distributed online. For example, the third presentation in Week 12 of Spring 2005 included a demonstration of Pachet's Continuator [15] for human-machine improvisation, and the second presentation in Week 9 of Spring 2006 included a demonstration of KTH's Director Musices software [16] for rule-based synthesis of musical expression. The presentation format and the required weekly reports are often foreign to the typical engineering student, and present valuable opportunities for the student to learn to articulate ideas verbally and in writing. No one student is expected to know every detail of a paper, and the other students in the class are encouraged to chime in with relevant information that they may bring to the table from their own backgrounds. For example, the music student is able to explain the concept of meter (higher level groupings of periodic pulses) and demonstrate conducting gestures, while the engineering student can help explain the dynamic programming approach to music alignment.

Finally, to test the students' grasp of the material, and provide an opportunity to learn by doing, they are required to formulate a music-engineering project related to the content of the course, implement and document it, and present a working prototype to the class. Figure 3 shows a part of the final projects website for Spring 2005 (the full list is available at www.scf.usc.edu/~ise575/a/projects), and the 'Percolate' site by Stein and Sundar; Figure 4 shows the projects index for Spring 2006 (www.scf.usc.edu/~ise575/b/projects). Each student has provided web documentation of his/her project, including executable software, whenever possible, for open

download and testing. For example, the 'Percolate' site contains links to open source code and documentation, and example sound files.

FIGURE 3 STUDENT PROJECTS INDEX (PARTIAL, 2005) + SAMPLE PROJECT SITE.

FIGURE 4 STUDENT PROJECTS INDEX (2006) + LINKS TO PROJECT SITES.

The nine projects from 2005 ranged from the generating of music using a gestural interface, to computer assisted composition, to structural analysis of music through signal processing. Two of the projects in the last category provided the impetus for the students' doctoral theses. The ten projects from 2006 ranged from comparisons of performances by renowned violinists, to modeling and color visualization of emotional response to music, to the characterization of scene classification in silent film music, to software tools for extracting beat and tempo curves. The creators of the first two examples will be expanding and developing their projects into Masters theses, and the project on silent film music will be incorporated into ongoing doctoral research.

In the spirit of MIT's OpenCourseware movement [17], all course material (including information and linked content shown in Figures 1 through 4) are openly available on the web through the course website (www-scf.usc.edu/~ise575). Whenever possible, a link to an online copy of the paper and accompanying software is provided. Each paper reviewed is accompanied by the presentation file, as well as exemplary summary reports on the paper, all in PDF format (see annotations added to Week 12 in Figure 1). Project material include a website containing a description of the project, and executable code for open distribution (as shown in Figure 3). This open courseware serves as a resource to the community at large. It provides reviews of research in computational modeling of music cognition, organized by music structures analyzed, and interactive applications; and, includes details and source files from sample projects by students in the class. Students, researchers and educators in the field, nationally and internationally have found the site to be a valuable resource, and an excellent example for creating courseware for an interdisciplinary field, and have provided positive feedback.

III. Synergistic Course-Related Activities

Although the number of researchers in computational music cognition is growing rapidly, the ones in the academe are still few and far between, compared to the more traditional disciplines. To expand the students' exposure to researchers in the field other than their instructor, specialists in the field are invited to give guest lectures to the class. The students were also encouraged to listen to actual performances as much as possible, so as to augment their learning with real-life examples, such as the performance of the concerto for two pianos highlighted in Week 2 of Spring 2005 (Figure 1).

The interaction with the invited speakers makes the lessons learned in class tangible and real, reinforcing the fact that there exists a community of researchers 'out there' who are interested in, and devoting their careers to, the very issues addressed in the classroom. In Spring 2005 (see Figure 1), the two visitors were David Temperley, a music theorist from the Eastman School of Music specializing in music cognition, and Christopher Raphael, a statistician from Indiana University at Bloomington, who developed *Music-Plus-One*, an automated accompaniment system using Bayesian networks. In addition, a postdoctoral researcher, Anja Volk, conducted the session on her area of specialty – Rhythm and Meter – in Week 5. The

presence of actual researchers in the field speaking about their work made a real difference in the way the students viewed their assigned presentations. The papers were no longer simply disassociated writings of some unknown persons, many of them represented the life's work of very real individuals who were deeply passionate about their chosen profession.

Due to the broad appeal of the invited lectures, the presentations by specialists in the field have branched off into a separate lecture series, shown in Figure 5 and documented online at www-rcf.usc.edu/~mucoaco/events/vecc0506.html. These talks by experts in the field continue to inspire and motivate the students in, and augment the knowledge goals of, the computational music cognition class. Several of the talks address directly the theme of the course. For example, Jeanne Bamberger, upon the instructor's request, gave a presentation titled *Expressing the Difference: Comparing Great Performances*, that gave a hands-on approach to listening to, understanding, and describing the differences between performances by renowned artists.



university of southern california viterbi school of engineering
music computation and cognition laboratory
 home :: research :: education :: events :: people :: publications



Viterbi Early Career Chair 2005-2006 Lecture Series

This lecture series is made possible by the Viterbi Early Career Chair funds, with additional support from the Integrated Media Systems Center and the Epstein Department of Industrial and Systems Engineering at the University of Southern California Viterbi School of Engineering.

.....
Bryan Pardo, Computer Science, Northwestern University
Audio Presence: new tools for musical information retrieval and interaction
THU, Nov 3, 2005, 3:30PM-4:30PM @ EEB248 [[PDF](#)]

Carol Krumhansl, Psychology, Cornell University
Statistics and Geometries of Music
FRI, Nov 18, 2005, 2:30PM-3:30PM @ EEB248 [[PDF](#)]

.....
Chris Chafe, CCRMA & Music Department, Stanford
Tapping into the Internet as an Acoustical Medium
FRI, Feb 10, 2006, 2:30PM-3:30PM @ EEB248 [[PDF](#)]

Jeanne Bamberger, MIT Music and Urban Education / UC Berkeley School of Education
Situated Inquiry: Moving between Action and Representation
WED, Feb 15, 2006, 7:00PM-8:50PM @ WPH304G (MUED 520 guest lecture)
Expressing the Difference: Comparing Great Performances
THU, Feb 16, 2006, 2:30PM-3:30PM @ EEB248 (ISE 599/575 guest lecture) [[EDF](#)]

Curtis Roads, Media Arts and Technology, UC Santa Barbara
Synthesis, analysis, and visualization of sound based on Gabor's atomic model
FRI, Mar 3, 2006, 2:30PM-3:30PM @ EEB248 [[PDF](#)]

.....
Gérard Assayag, IRCAM Music Representations Research Group, Paris, France
demo: Improving with the Computer using OMax, a Statistical Learning Environment featuring Dennis Thurmond [+keyboard] and Gérard Assayag [+OMax]
TUE, Apr 4, 2006, 12:30PM-1:30PM @ MUS106 (MacDonald Recital Hall)
general lecture: Computer Assisted Composition at IRCAM: the OpenMusic environment
WED, Apr 5, 2006, 2:30PM-4:00PM @ GER309
workshop: OpenMusic and OMax under the cover
THU, Apr 6, 2006, 2:30PM-4:00PM @ GER309
 [[website](#) [poster](#) [PDF](#)]

Mark Sandler, Electronic Engineering, Queen Mary, University of London, UK
general lecture: Interacting with Digital Music
WED, Apr 19, 2006, 2:30PM-4:00PM @ GER309
technical lecture: Constant Q Analysis and its use in Music Informatics
THU, Apr 20, 2006, 2:30PM-4:00PM @ GER309 [[PDF](#)]

FIGURE 5
 OFFSHOOT LECTURE SERIES (FALL 2005, SPRING 2006).

IV. Finding Paths Through Existing Academic Structures

Some bureaucracy is inevitable in the everyday functioning of any large organization. In order to successfully introduce and maintain a new course offering in a system requires some understanding of the formal structures through which such initiatives operate. Any course proposal must pass the scrutiny of the curriculum review committees. And, a course that cannot sustain a reasonable enrolment on a yearly basis will die a quick death. These challenges are further

exacerbated in the case of interdisciplinary course offerings that either do not have a natural home in the traditional department structures, or may not receive as strong support from a given department because it is not deemed to be central to the discipline.

Special steps were taken to counter these threats in the case of *Topics in Engineering Approaches to Music Cognition*. Computational modeling of music processes is not typically the topic of choice for traditional industrial engineering students in the author's, and hence the course's, home department. To ensure steady enrolment, cross-listing in multiple engineering departments were sought so that students not only from ISE, but CS, and EE could readily take the course for credit towards their respective degrees. Cross listing in Psychology (in a different school) was more recently approved. Steps were taken to ensure that the course was listed as an approved elective in relevant Masters degree programs such as the Masters degrees in Multimedia and Creative Technologies (in CS and EE), in Integrated Media Systems, and in Game Development (within CS). Furthermore, special care was taken to ensure that students who might be interested in the course in all three schools (engineering, music, and the college) receive the announcements.

A course centered on music processes in an engineering environment cannot help but have to face up to some erroneous assumptions about the nature of music research. For example, engineering colleagues who favor the more traditional methodology-focused approach to course design may view an interdisciplinary course of this type as being soft, and not sufficiently rigorous. They may also fear that an interdisciplinary course offering may dilute engineering degree programs. These fears may escalate when they learn that students from music and psychology can enroll in the class, and getting credit alongside engineering students. It may take repeated and patient explanations to clarify that an interdisciplinary course is no less rigorous or difficult than a traditional one, and that students who enroll in the course hoping for an 'easy A' will be disappointed, as all students will be held against the highest standards of excellence. Seminars and guest lectures by well known invited speakers in the field can also help allay the fears of the non-rigorousness of interdisciplinary work. Seeking approval for the expansion of the course to a topical sequence that can be re-taken for credit has been a challenging process for the reasons mentioned above. After a year's debate and several appeals, the proposed changes have been approved.

DISCUSSION AND CONCLUSIONS

In conclusion, this paper has presented a constructivist and constructionist approach to interdisciplinary learning at the intersection of music and engineering, with the goal of introducing young scholars to, and preparing them for, research at the forefront of the field. The students in the class learn by example, through the active survey of current literature, and learn by doing, through various hands-on activities incorporated into the class, including the final implementation project. The success of the class can be

measured through the achievements of the students, as documented in their reports, presentations, and final projects, and the fact that several of them are continuing their projects for their theses. The open sharing of all course material serves as a resource to the community. Other examples of courses in music cognition that incorporate aspects of computing such as representation and algorithms now exist; however, few are as copiously documented as the ones described in this paper, nor is the material as readily re-usable by the community. It is difficult to make formal comparisons at present due to this lack of information. The comprehensive and pragmatic approach to the design and implementation of an advanced course in a new area has been shown to be successful over several repeat offerings. Practical strategies were presented, from structuring of course content, to positioning of the course in existing programs, to organization of synergistic activities, for ensuring that the course not only achieves its goals, but thrives and flourishes in the face of various challenges.

ACKNOWLEDGMENT

This material is based upon work supported by the National Science Foundation (NSF) under Grant No. 0347988, and Cooperative Agreement No. EEC-9529152. The work made use of the Integrated Media Systems Center Shared Facilities. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the author, and do not necessarily reflect the views of NSF.

REFERENCES

- [1] Longuet-Higgins, H., C., "Letter to a Musical Friend", *The Musical Review*, Vol 23, 1962, pp 244-8.
- [2] Longuet-Higgins, H., C., "Second Letter to a Musical Friend", *The Musical Review*, Vol 23, 1962, pp 271-80.
- [3] Longuet-Higgins, H., C., & Steedman, M., J., "On Interpreting Bach", In Meltzer, B., & Michie, D. (editors): *Machine Intelligence*, Vol 6, 1971, pp 221-241, Edinburgh University Press, Edinburgh, UK.
- [4] Darwin, C., "Christopher Longuet-Higgins – cognitive scientist with a flair for Chemistry", *The Guardian* – obituaries, June 10, 2004, www.guardian.co.uk/obituaries/story/0,3604,1235115,00.html
- [5] Intl Conf on Music Information Retrieval (ISMIR) – www.ismir.net
- [6] Intl Conf on Sound and Music Computing (SMC) – smc04.ircam.fr
- [7] Intl Conf on Computer Music Modeling and Retrieval (CMMR) – www.lma.cnrs-mrs.fr/~cmmr2005
- [8] Jacobs, H., H., "Interdisciplinary Curriculum: Design and Implementation", September, 1989, Association for Supervision and Curriculum.
- [9] Concept to Classroom: Workshop on Interdisciplinary Learning in Your Classroom - [www.thirteen.org/edonline/concept2class/interdisciplinary - Educational Broadcasting Corporation, 2004](http://www.thirteen.org/edonline/concept2class/interdisciplinary-Educational Broadcasting Corporation, 2004).
- [10] Piaget, J., *The Psychology of Intelligence*, 1950, New York, Routledge.
- [11] Constructivism (learning theory) - en.wikipedia.org/wiki/Constructivism_%28learning_theory%29
- [12] Temperley, D., *The Cognition of Basic Musical Structures*, 2001, Cambridge, MA: MIT Press.
- [13] Bamberger, J., *Developing Musical Intuitions*, 2000, Oxford University Press, with accompanying software, *Impromptu*, by Bamberger, J., and Hernández, A., www.tuneblocks.com
- [14] Farbood, M., & Pasztor, E., *Hyperscore* – www.hyperscore.com
- [15] Pachet, F., *Continuator* – www.csl.sony.fr/~pachet/Continuator
- [16] KTH's Director Musices Software – www.speech.kth.se/music/performance/download
- [17] MIT OpenCourseware – ocw.mit.edu