Synthesis: Interdisciplinary Collaboration in Computational Music Research

Hosted by Berkeley Computational Music Research (BCMR)
Saturday, September 28, 2019

8:15-8:55	**Breakfast and Registration**
8:55-9	Opening Remarks
9-10	*Empirical Musicology: A Responsible User's Guide* Claire Arthur, Music Theory and Cognitive Musicology, Georgia Tech
10-10:30	*Electroacoustic Music Analysis: Towards an Interdisciplinary Approach?* Kevin Dahan, Music Technology, De Montfort University / Stanford CCRMA
10:30-11	**Coffee Break**
11-12	*Music Information Retrieval: Goals, Progress, and Three Uncomfortable Statements* Justin Salamon, Audio Research Group, Adobe Research
12-12:30	*Seeing Voices: Using Light to Restore and Preserve Early Recorded Sound* Melanie Hamaguchi, Project IRENE, UC Berkeley
12:30-1:30	**Lunch Buffet**
1:30-2:30	*Reaching Across the Chasm: An Ethnomusicologist's Forays into Computational Analysis* Leslie Tilley, Ethnomusicology, MIT
2:30-3:30	*How Mathematics Can Help in Representing Musical Concepts* Carmine Cella, Composition, UC Berkeley
3:30-4	**Coffee Break**
4-5	*Machine Learning, Music Production, and the Feeling of Creativity* Jon Gillick, School of Information, UC Berkeley
5-5:30	*HitPredict: Using Spotify Data to Predict Billboard Hits* Elena Georgieva, Music Technology, Stanford CCRMA
5:30-5:35	Closing Remarks
6:15	**Dinner for speakers at Berkeley Social Club (invite only)**

BERKELEY STUDENT TECHNOLOGY FUND

Center for New Music & Audio Technologies

Location:

Center for New Music and Audio Technology (CNMAT)
1750 Arch Street
Berkeley, California
94720

Event Information:

Hour-long sessions include a 45-minute talk with 15 minutes of discussion. Half hour sessions include a 20-minute talk with 10 minutes of discussion.

Breakfast, lunch, coffee, and snacks will be provided for all registered participants. Additional restaurants and cafes are located on Euclid Ave. (a five minute walk through the Graduate Theological Union on Ridge Path) or in downtown Berkeley (walk down the hill to Shattuck Ave).

Transportation and Parking:

CNMAT is a 13-minute walk from the Downtown Berkeley BART station. The nearest paid parking garage is the Lower Hearst Garage at 2451 Hearst Ave, which is a five minute walk from CNMAT. From the garage, walk up Scenic Ave and the turn left onto Ridge Path, cutting through the Graduate Theological Union. There is also limited two hour street parking on Arch Street.

BERKELEY STUDENT TECHNOLOGY FUND

CNMAT Center for New Music & Audio Technologies

Abstracts (alphabetical order):

Empirical Musicology: A Responsible User's Guide
Claire Arthur, Music Theory and Cognitive Musicology, Georgia Tech

What kinds of questions can be answered through empirical musicology research? What hasn't been answered yet? And, what may never be answered? This talk will introduce the fields of computational and cognitive musicology as applied to real musicological and cognitive research questions, drawing on my own research as well as the research of others in the field. I will discuss popular themes and topics in cognitive musicology and music informatics (broadly defined), and share my perspective on common methodological pitfalls and limitations. In doing so, I aim to highlight each field's strengths and weaknesses, as well as suggest some "best practices," such as the necessity of interdisciplinary collaborations for extending one's skill sets. Finally, I will present some of my current ongoing research, with an emphasis on the ways in which it synthesizes approaches and methods from music theory, statistics, computer science, cognitive psychology, and digital signal processing.

How Mathematics Can Help in Representing Musical Concepts
Carmine Cella, Composition, UC Berkeley

The talk will present the work done in the last years in searching good signal representations that permit high-level manipulation of musical concepts. After the definition of a geometric approach to signal representation, the theory of sound-types and its application to music will be presented. Finally, recent research on assisted orchestration and augmented instruments will be proposed for musical applications.

Electroacoustic Music Analysis: Towards an Interdisciplinary Approach?
Kevin Dahan, Music Technology, De Montfort University/ Stanford CCRMA

Since the first *Études* of Pierre Schaeffer and the experiments of Lejaren Hiller and Max Mathews with the computer as an instrument, numerous compositions were realized in what is usually referred to as the 'electroacoustic music' genre. While there are many subgenres, works belonging to this overarching category normally share one or several common traits. It is usually easier to characterize them by noting the absence of traditional Western music denominators, such as pitch-based melodies, structured meters, or recurring rhythmic cells. Another usual way of describing them is to stress the importance of *timbre* as a central structural element of the composition. All of these aspects led to a wide variety of compositional 'styles' that are very challenging for the music analyst.
On the other hand, over the past 30 years, tools were developed to describe music compositions in alternative ways (not relying solely on traditional musicological approaches), using concepts, theories and techniques derived from other research fields (e.g. music perception/cognition,

music information retrieval, …). However, probably due to the experimental nature of many approaches and the need for result replication, the usual test case is based on Western traditional music, mostly ignoring the limit case of non-pitch-based, non-metered, timbre-based compositions.

Could current computing tools (such as machine learning approaches) and neuroscientific approaches (such as imagery techniques) be used on the important corpus of electroacoustic music compositions?

We argue that using these approaches on electroacoustic music is needed to: a) update the current toolset and theories used by musicologists and music analysts; b) provide different, multidimensional and complex datasets for data scientists and cognition specialists; c) provide new tools to be used in a compositional and creative context; and d) possibly allow for a better understanding of this musical style, in order to widen its audience.

HitPredict: Using Spotify Data to Predict Billboard Hits

Elena Georgieva, Music Technology, Stanford CCRMA

In this study, we approached the Hit Song Science problem, aiming to predict which songs will become Billboard Hot 100 hits. Music technology companies such as The Echo Nest, ChartMetric, and Next Big Sound have been using data analytics to help artists and labels predict and track a song's success for almost a decade. This problem is referred to as Hit Song Science (HSS) in the Music Information Retrieval (MIR) field. Machine learning is a popular research and industry tool to approach the HSS question.

We collated a dataset of approximately 4,000 hit and non-hit songs. We collected hit songs featured on the Billbard Hot 100 using the Billboard API Library and non-hit songs from the Million Song Dataset. For each song, we extracted audio features from the Spotify Web API. These features include: Danceability, Energy, Speechiness, Acousticness, Instrumentalness, Liveness, Valence, Loudness. To account for artist recognizability, we defined an additional metric: the artist score. Each song was assigned an artist score of 1 if the artist had a previous Billboard Hot 100 hit, and 0 otherwise.

To predict a song's success, we used three different machine- learning algorithms: Logistic Regression (LR), Gaussian Discriminant Analysis (GDA), and a Neural Network (NN). For each supervised learning algorithm, we split the data into training and validation examples using a 75/25 split. LR and GDA yielded a reasonable accuracy of 75.9% and 73.7% against the validation data, with similar accuracy against the training data indicating no overfitting. Our NN with one hidden layer gave a similar accuracy to LR, but generated significantly higher precision. Overall, we feel these results are promising and we are interested in further investigating how machine learning can be used to approach Hit Song Science.

Machine Learning, Music Production, and the Feeling of Creativity
Jon Gillick, School of Information, UC Berkeley (9-10 am)

The word “Creativity” can mean something different to every artist. For music producers, the feeling of creativity can often come from the process of carrying out a high-level artistic vision for a piece of music, even if the concrete steps to execute the vision differ from project to project. New tools and technologies have the potential to dramatically change some of the steps that artists take during the production process. In this talk, I’ll present some of my work on using machine learning to make tools for musicians and producers, and I’ll discuss some of the strengths, weaknesses, and implications of the approach.

Seeing Voices: Using Light to Restore and Preserve Early Recorded Sound
Melanie Hamaguchi, Project IRENE, UC Berkeley

Sound was first recorded and reproduced by Thomas Edison in 1877. Until about 1950, when magnetic tape use became common, most recordings were made on mechanical media such as wax, foil, shellac, lacquer, and plastic. Some of these older recordings contain material of great historical interest, but may be in obsolete formats, and are damaged, decaying, or are now considered too delicate to play. Unlike print and latent image scanning, the playback of mechanical sound carriers has been inherently invasive. Currently, techniques based upon non-contact optical metrology and data analysis, have been applied, non-ivasively, to create and analyze high resolution digital images, and to restore and preserve the audio content, of these materials.

This lecture will discuss the characteristics of early sound recordings and the use of this current technology as applied to a number of notable collections: including field recordings of Native Americans and Canadians from the early 20th Century. Recently studies of degradation in recorded sound media and the use of machine learning techniques in noise reduction and audio feature extraction will be described.

The technology and restoration of historic audio recordings will be illustrated with sounds and images. Additional information can be found at http://irene.lbl.gov/

Music Information Retrieval: Goals, Progress, and Three Uncomfortable Statements
Justin Salamon, Audio Research Group, Adobe Research

Music Information Retrieval, also referred to as Music Informatics Research or MIR, is a research field concerned with the computational analysis, indexing, retrieval, recommendation, separation, transformation and generation of music. Originally a field that relied predominantly on digital signal processing (DSP), similar to computer vision much (but not all) of modern MIR research relies heavily on machine learning and, at the time of writing, deep learning. MIR research has diverse applications ranging from music recommendation on streaming services to computational ethnomusicology. Twenty years since the first International Symposium on Music Information Retrieval (ISMIR), now the International Society for Music Information Retrieval conference, the field has made significant progress, but also faces significant challenges. In this talk I will give a brief overview of MIR goals and progress, followed by a discussion of current challenges motivated by three uncomfortable statements.

Reaching Across the Chasm: An Ethnomusicologist's Forays into Computational Analysis
Leslie Tilley, Ethnomusicology, MIT

The rise of new computer technologies has opened doors for analytical ethnomusicology in recent decades, allowing researchers to organize, search, and understand large collections of musical data in new ways. Yet to offer the deepest insights, the data that such tools generate must often be bolstered by ethnographic approaches and "interpreted by music scholars with an understanding of the specific music(s) involved" (Tsanetakis et.al 2007: 12). Unfortunately, for many of us trained primarily as musicians and music researchers, computational techniques can seem a daunting prospect, perhaps even completely out of reach. How can a non-programming ethnomusicologist begin to think about these kinds of projects, and then begin to work on them? How can an unformalized oral music theory gleaned through field research inform and nuance computational approaches and findings? And how can useful collaborations be fostered that balance and speak to the needs of both fields of knowledge? This paper explores the use – and usefulness – of computational analysis in my own research into the paired improvised Balinese drumming practice *kendang arja*. The work is in its earliest stages, and I present my process as a case study of the mutually supportive potential of ethnography and computational analysis… and the inherent challenges of such an approach. *Arja* drummers are among the most respected musicians in Bali. Unlike drummers of other genres, whose paired patterns are carefully composed to interlock, two *arja* drummers create fast, intricately interlocking patterns through simultaneous improvisation. Musical analysis of their playing suggests unspoken rules and possible model patterns guiding improvisation, and conversation with musicians tells a similar – if less clear-cut – story. Computational analysis has presented an exciting arena for further inquiry, offering fresh perspectives on the conscious and subconscious decisions made by *arja* drummers in the course of performance by revealing statistically relevant patterns and relationships. Yet the process has also been a challenge for me as an ethnomusicologist, and as much as the results, it is the process that interests me here. As such, this paper is as exploratory as it is conclusive. Presenting my analyses and struggles, both, I hope to provide a catalyst for discussing the paths to fruitful exploration and partnership across the disciplinary divide.

Sources Cited

Tzanetakis, George, Ajay Kapur, W.Andrew Schloss, and Matthew Wright. 2007. "Computational Ethnomusicology." *Journal of Interdisciplinary Music Studies* 1(2): 1-24.