

Abstracts

Day 1: Mathematical and Computational approaches	2
Martin Rohrmeier, Digital and Cognitive Musicology Lab École Polytechnique Fédérale de Lausanne.....	2
Isabelle Bloch, Sorbonne University, Paris	3
Dmitri Tymoczko, Princeton University	4
Jason Yust, Boston University	5
Anja Volk, Utrecht University.....	6
Day 2: Machine Learning, Generativity, Interaction.....	7
Cheng-Zhi Anna Huang, University of Montreal, Google DeepMind	7
Sholmo Dubnov, University of California, San Diego	8
Stefania Serafin, Aalborg University Multisensory Experience Lab	9
Dorien Herremans, Singapore University of Technology and Design, AMAAI Lab	10
Nicolas Obin, IRCAM, STMS Lab.....	11
Day 3: Computational Physiology/Medicine	12
Julian F. Thayer, University of California, Irvine & The Ohio State University	12
Ye Wang, National University of Singapore, Sound & Music Computing Lab	13
Kat Agres, National University of Singapore, YST Conservatory of Music, Centre for Music and Health.....	14
Michael Casey, Dartmouth College.....	15
Esa Räsänen, Tampere University.....	16
Day 4: Education, Learning and Creativity.....	17
Andrew Milne, Western Sydney University, Australia.....	17
Rachel S. Y. Chen, Nanyang Technological University, Singapore.....	18
Leah Frederick, University of Michigan	19
Day 5: Student session	20
Paul Lascabettes, Ircam/CNRS STMS Lab, Paris, France.....	20
Gonzalo Romero, Ircam/CNRS STMS Lab, Paris, France	21
Nathanael Koh, Australian National University, Australia.....	22
Yoti Narang, University Pompeu Fabra, Barcelona, Spain.....	23
Mateusz Solinski, King's College London, UK.....	24

Day 1: Mathematical and Computational approaches

Martin Rohrmeier

Digital and Cognitive Musicology Lab

École Polytechnique Fédérale de Lausanne

Syntactic Structures in Music

One of the common points of interest in music theory, music mathematics and music cognition is to understand the properties of musical structure and its sequence of events. For this problem setting, syntactic models and formal grammars provide valuable tools for characterizing the structure of musical progressions and their underlying principles of structure building. They are particularly powerful in expressing forms of abstraction, hierarchical relations, and ways of deriving musical progressions from the recursive generative application of a set of core relations and rules.

In this talk, I will present a synopsis of my research work on musical syntax over the last few years. It will cover motivation and foundations of syntactic models, as well as research on modeling structure building in various musical parameters including harmony, rhythm, voice-leading and form in Western tonal music, as well as melodic structure in non-Western music. The theoretical work presented will be complemented by discussing its implications for computational modeling and cognitive research.

Acknowledgement: This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme under grant agreement No 760081 – PMSB.

[Return to Table of Contents →](#)

Isabelle Bloch

Sorbonne University, Paris

Mathematical morphology: Algebraic setting and applications

In this talk, we will present the algebraic bases of mathematical morphology, in particular in complete lattices. Lattice theory has become a popular mathematical framework in different domains of information processing, and various communities employ its features and properties, e.g. in knowledge representation, in logics, automated reasoning and decision making, in image processing, in information retrieval, in soft computing, in formal concept analysis. Mathematical morphology, in its algebraic and deterministic setting, is based on adjunctions, which endows its operators with strong properties and allows for multiple applications and extensions. Many mathematical structures can be associated with a partial ordering, making them a complete lattice (sets, functions, fuzzy sets, graphs and hypergraphs, logics...) on which mathematical morphology operators can be applied. A few examples will be illustrated. Further examples on music will be described in other talks.

[Return to Table of Contents →](#)

Dmitri Tymoczko

Princeton University

A New Model of Musical Hierarchy

In my talk I will outline ideas from my just-published book, *Tonality: An Owner's Manual*, and explain how they can be used in composition, algorithmic composition, improvisation, and analysis. These ideas include: new theories of chordal inversion, motives, and sequences, and new representations of musical structure that apply voice-leading geometry in recursive harmonic hierarchies. Together these ideas allow us to generalize basic tonal procedures to any chord-and-scale environment (including microtonal environments); this means we can adapt basic tonal strategies to atonal and post-tonal worlds, as well as imagining new forms of tonality. All of these ideas are embodied in a musical programming language called "Arca," which can be used for both analysis and composition.

[Return to Table of Contents →](#)

Jason Yust

Boston University

Coherence of Harmonic and Rhythmic Qualities

This paper explores the musical meaning, in various contexts, of coherence, understood as the phase of a product of Fourier coefficients. Specifically, given three coefficients of a Fourier transform of a pitch-class or time-class set, with indices (frequencies) a , b , and c , such that $a + b = c$, coherence is the phase of the product of the first two coefficients with the conjugate of the third. The coefficients are coherent when the phase of the product is close to zero, indicating a coordination between the three frequencies, and incoherent when it is close to π . I illustrate coherence and incoherence in three different contexts. Coherence values involving coefficients 2, 5, 7, and 9 can be used to define the Balinese Pelog scale and distinguish between regional varieties. Coherence of coefficients 2, 3, and 5 characterize diatonic functional harmony and distinguish it from certain pentatonic and blues-based harmony. In the rhythmic domain, I illustrate the stability of certain coherence values across the rhythmic variations in a master drum performance in the Akan tradition of Adowa.

[Return to Table of Contents →](#)

Anja Volk

Utrecht University

Connecting Computational Models of Musical Structures with Developing Music Technology for Health, Wellbeing and Inclusion

A large part of mathematical and computational research in music is dedicated to the investigation of musical structures, contributing to musicological and music theoretic investigations. Employing models of musical structures for developing music technology in the context of health, wellbeing and inclusion constitutes a growing interdisciplinary research endeavor that provides a new application area for these computational models. In this talk, I present research on investigating musical patterns with computational means to enhance our understanding of music, and on employing these patterns in interaction contexts such as for training and rehabilitation purposes in health, wellbeing, and inclusion. I will discuss the development of e.g., music-based collaborative games for visually impaired children, and applied games for musical attention control training, and the role of musical structures therein. I will reflect on the interdisciplinary challenges and opportunities of developing human-centric interaction technology for health, wellbeing and inclusion.

[Return to Table of Contents →](#)

Day 2: Machine Learning, Generativity, Interaction

Cheng-Zhi Anna Huang

University of Montreal, Google DeepMind

Deep Learning for Musical Creativity

Advances in generative modeling have opened up exciting new possibilities for music making. How can we leverage these models to support human creative processes? First, I'll illustrate how we can design generative models to better support music composition and performance synthesis. Coconet, the ML model behind the Bach Doodle, supports a nonlinear compositional process through an iterative block-Gibbs like generative procedure, while MIDI-DDSP supports intuitive user control in performance synthesis through hierarchical modeling. Second, I'll propose a common framework, Expressive Communication, for evaluating how developments in generative models and steering interfaces are both important for empowering human-ai co-creation, where the goal is to create music that communicates an imagery or mood. Third, I'll introduce the AI Song Contest and discuss some of the technical, creative, and sociocultural challenges musicians face when adapting ML-powered tools into their creative workflows. Looking ahead, I'm excited to co-design with musicians to discover new modes of human-ai collaboration. I'm interested in designing visualizations and interactions that can help musicians understand and steer system behavior, and algorithms that can learn from their feedback in more organic ways. I aim to build systems that musicians can shape, negotiate, and jam with in their creative practice.

[Return to Table of Contents →](#)

Sholmo Dubnov

University of California, San Diego

Modeling Creative Interaction Using Information Theory

In this talk I will present our work on information theoretic modeling of creative interaction between improvising agents. Motivated by Hick's law on constant information rate of human decision making, we use theories of channel coding with a feedback channel to quantify the information loss between interacting agents. Most conditioning schemes modify the layers of the generating network in a way that will produce the desired outcome according to predetermined conditioning vectors. Our approach is different - by maximizing the transfer entropy between multiple improvising agents we seek to build a policy that will adapt their generation to maximize the capacity of their interaction channel. In the case of memoryless sources, the transfer entropy is measured using mutual information neural estimators (MINE). In previous work we showed that maximal mutual information between multiple voices can be obtained by considering a noisy mapping between their music encodings. In the talk I will present an extension of MINE to sequences with memory, known as Directed Information Neural Estimator (DINE), that takes into account causal relations between sequences, thus allowing adaptation of the generative model to novel musical inputs, such as performing in a different style, or interacting in an idiomatic way with previously unheard musician.

[Return to Table of Contents →](#)

Stefania Serafin

Aalborg University Multisensory Experience Lab

Mixed realities in health and culture

Mixed reality technologies are nowadays portable and mobile, in such a way that they can be taken outside laboratory settings and into the clinics and cultural institutions. How can we use these technologies in responsible ways? Can we benefit from AI to provide customized solutions to diverse populations?

In this talk I present several research projects conducted at the Multisensory Experience Lab at Aalborg University in Copenhagen.

In clinical settings, we use virtual reality for example to train children with a hearing impairment to improve their spatial awareness, in collaboration with the Copenhagen Hearing and Balance Center.

In cultural settings, we use virtual reality to reconstruct instruments from the past, in collaboration with the Danish Music Museum.

We discuss several case studies where MR can be used in culture and health, and provide some design guidelines and learning outcomes from our experiences.

[Return to Table of Contents →](#)

Dorien Herremans

Singapore University of Technology and Design, AMAAI Lab

Towards generative music AI models with multimodal controls

The field of music generation has made huge leaps forward. In the last few years we have witnessed an evolution from simple algorithmic models to controllable multimodal generative AI systems. A number of challenges still remain, however, which I will discuss in this talk. For instance, if we can make generative AI models controllable, then we can use them to augment the creative process of music producers and composers. I will discuss a number of such controllable models in detail, including models that allow the generated music to be steered by user-defined emotions and tension, as well as model one that tries to directly match an input video. Finally, I will briefly touch on some of my current work related to modern day interfaces such as text-prompt based music generation (audio) models, as well as ideas for future research.

[Return to Table of Contents →](#)

Nicolas Obin

IRCAM, STMS Lab

Generative modelling of speech: from human-like to expressive speech and beyond - a journey in Wonderland

Speech is the most complex communication system between living organisms on Earth.

Nowadays, it is possible to generate, manipulate or clone a voice in an extremely realistic way and from very small amount of data.

In this presentation, I will give a brief history of modern speech synthesis, from concatenative synthesis to parametric statistical models and neural networks. I will highlight current research trends in neural speech generation, their limitations and their scientific and technological challenges.

I will pay a particular attention to the creative uses of speech synthesis, illustrating this with various examples taken from film and music production. I will conclude by addressing the scientific issues for speech synthesis the concerns, and the ethical and legal issues related to voice generation for artistic creation and in the creative industries.

[Return to Table of Contents →](#)

Day 3: Computational Physiology/Medicine

Julian F. Thayer

University of California, Irvine & The Ohio State University

What Are Emotions and Why Does Music Elicit Them?

Music has long been considered the language of emotions. However, few studies have examined musically induced emotions in the context of a comprehensive model of emotion. We have recently proposed such a comprehensive model of emotion based upon dynamical systems theory and here present an overview of data on musically induced emotions. Much literature supports the idea that emotions can be described by their location in a state space defined by the dimensions of valence and arousal. These dimensions can be conceptualised as motivational, control parameters of affective experience and the discrete emotions as attractors in state space. Importantly, valence maps onto mean pitch level and arousal maps onto tempo.

We will describe the neurovisceral underpinnings of this model including the link between the auditory nerve and the amygdala, and the interaction of medial prefrontal cortex and brainstem outputs to the cardiovascular system via the vagus nerve. This comprehensive model of musically induced emotions is able to accommodate the extant literature and produce testable hypotheses for future research. Suggestions for future investigations using musical interventions are offered based on this integrative account.

[Return to Table of Contents →](#)

Ye Wang

National University of Singapore, Sound & Music Computing Lab

Neuroscience-inspired Sound and Music Computing for Human Health and Potential (SMC4HHP)

Motor synchronization with an auditory beat has been deemed a human cultural universal. Even infants show perceptual sensitivity to and motor coordination with musical rhythms. A temporally stable beat facilitates rhythmic human movement during leisure activities such as exercise. It also serves as the basis for a class of gait rehabilitation therapies known as Rhythmic Auditory Stimulation (RAS) for Parkinson's disease, stroke, and others. Fascinated by the magical power of RAS in 2010, I initiated the development of MusicRx technologies to make the RAS intervention affordable and accessible. In parallel, I led the development of SLIONS technologies as computer-assisted speech interventions for various populations leveraging on the power of singing. MusicRx and SLIONS are the two major initiatives I have explored in the past decade to grow a nascent research field, Neuroscience-inspired Sound and Music Computing for Human Health and Potential (SMC4HHP), into a burgeoning research field today, with increasingly more researchers jumping in. I will reflect on the challenges and opportunities of SMC4HHP, as well as how to make our research program more resilient.

[Return to Table of Contents →](#)

Kat Agres

National University of Singapore, YST Conservatory of Music, Centre for Music and Health

The efficacy of music interventions for mental health and emotion mediation

In recent years, there has been a growing global momentum in the area of music and health, with researchers and musicians around the world investigating why and how music has such a diverse range of impacts on health and well-being. In this talk, Dr. Agres will introduce the Centre for Music and Health, SE Asia's first dedicated research centre investigating the impacts of music on health, which emerged out of this recent global momentum. Her presentation will highlight findings connecting music with mental health and emotion self-regulation, including a music-based brain computer interface (BCI) for real-time emotion regulation, a newly-developed affective music generation system called AffectMachine, and results showing the impact of AffectMachine's generated music on the subjective emotion and physiological states of university students.

[Return to Table of Contents →](#)

Michael Casey

Dartmouth College

Music in Epilepsy Research: Physiological Effects and the Search for their Causes

We studied anepileptic effects of music with surgically implanted (iEEG) patients with epilepsy. We showed that some musical excerpts had significant positive effects on epileptic brain activity. Namely, gamma-frequency (40Hz) auditory tones—previously designed for experiments on mice with Alzheimer’s disease--were also effective for humans with epilepsy, as demonstrated by the suppression of epileptic ‘spikes’ in iEEG signals. Our subsequent results repeated anti-epileptic effects using Mozart’s Sonata for Two Pianos in D Major (K. 448); the effect occurred after 30 seconds of listening to the piece, and sustained for several minutes.

Our follow-up experiments were designed to tease-out the possible causes of the reduction in epileptiform brain activity. One avenue is music anticipation (implication) and reward (realization), which evoke responses in prefrontal and limbic areas of the brain. We demonstrate new biomarkers of anticipation and reward—evoked potentials corresponding to musical phrase closure anticipation and phrase-closure completion. Increased presence of phrase-closure biomarkers in the iEEG signal was correlated with the greatest antiepileptic effects. We also investigated the effects of varying tempo of known anepileptic stimuli. Our work represents a search for clinically-effective musical excerpts that could form a non-invasive, cost-effective listening program that is tailored to persons with epilepsy or other brain disorders.

[Return to Table of Contents →](#)

Esa Räsänen

Tampere University

The Physics of Drumming: Science & Art of Being Slightly Off

It has long been known that musicians do not keep time with the precision of a metronome. The human microtiming deviations from the exact beat typically range from a few milliseconds up to tens of milliseconds. Over time, these deviations are long-range correlated (LRC), i.e., they have a particular mathematical form that resembles a fractal [1]. Astonishingly, these LRC deviations make music sound distinctly human, so that they are preferred by listeners over uncorrelated white noise [2]. Therefore, LRC fluctuations in musical rhythms are not only a natural consequence of human playing, but they also characterize how rhythmic inaccuracies are tolerated or even enjoyed – at least within a few-millisecond range.

Here we present our findings in microtiming deviations of drumming. We demonstrate the nature of rhythmic fluctuations with several examples and discuss the relevance of our results as a subtle element behind musical "groove" – among other elements such as rhythmic syncopation and musical communication. Finally, we present a time-series model [3] able to capture the human feel in drumming and demonstrate it for iconic examples such as Michael Jackson's Billie Jean.

[1] E. Räsänen et al., PLoS ONE 10(6): e0127902 (2015).

[2] H. Hennig et al., PLoS ONE 6(10): e26457 (2011); Phys. Today. 65, 64 (2012).

[3] O. Pulkkinen and E. Räsänen, submitted (2023).

[Return to Table of Contents →](#)

Day 4: Education, Learning and Creativity

Andrew Milne

Western Sydney University, Australia

Enhancing music learning and creativity with maths and AI

Listening to and enjoying music is an almost universal activity so it is surprising that creative engagement with music composition and performance, and understanding of how music works, is remarkably low. One explanation is that music is an already complex task made even more complex through inconsistent ad-hoc 'user-interfaces' (e.g., instrument note-layouts, notational systems, music-theoretical nomenclature). In addition to this inconsistency, most musical user-interfaces do not leverage recent findings of music cognition to reduce cognitive load for conventional music nor allow natural mathematical generalisations to unconventional, unfamiliar, or novel musical forms. There is also a commonly held belief that musical and creative talent is in-born and cannot be improved: a belief that could be modified with the gentle encouragement of a personal tutor, but that is unaffordable for most. In this talk, I will discuss: some strategies and possibilities for how interface design may reduce extraneous cognitive load for beginners and experts and for conventional and unconventional music; and how artificially intelligent assistants may assist users with various musical tasks, and increase motivation through gamification and by encouraging them to adopt incremental beliefs of musical and creative intelligence (i.e., that these can be improved).

[Return to Table of Contents →](#)

Rachel S. Y. Chen

Nanyang Technological University, Singapore

Rhythm in the interactions between Autistic individuals

The repetitive ticks of a metronome could be argued to lack any sense of musical appeal. It is only when time is divided, varied, and organized into a temporal patterns, that any repetitive cycle is transformed into music. Deepening our understanding of rhythm breathes life into the lived experiences of autistic individuals. The manifestation of restricted repetitive behaviors, otherwise known as “stimming”, is phenomenon that clinically defines Autism. Comprising sequences of repetitive movement—swaying, rocking, flapping of hands, tapping—these behaviors have traditionally been conceptualized as anti-social, holding a reputation for being restrictive, rigid. Stimming, however, has been reported by Autistic individuals to instead bring both calmness and joy, having utility as a “natural” and “powerful” means of connecting with others. In this talk, I examine two cases of Autistic individuals stimming with others, one in a drum circle, and another on novel musical mats. I show that when augmented with sound, stimming can present temporal variations that strike a balance between stability and change. Autistic individuals fill their geometric spaces with theme and variation, rhythmic dialogue, polyrhythms, dancing, and varied tempos. Examining how repetitive ticks can become rhythmic flow reveals the “infinite variety” of Autistic musicality, present in the musical fractions of their everyday movements.

[Return to Table of Contents →](#)

Leah Frederick

University of Michigan

Musical Spaces as Instrumental Spaces

Music theorists often represent relationships between musical objects (e.g., notes, scales, and chords) using musical spaces, which Julian Hook (2023) broadly defines as “collection[s] of musically meaningful entities in some visual or geometric arrangement.” Though these musical objects are usually abstract entities, a growing trend in the field of music theory employs these same tools to capture the tangible patterns that underlie the layout of notes on physical musical instruments (e.g., De Souza 2017). This talk explores these instrumental spaces, as well as their relationships to abstract pitch spaces, from a mathematical perspective.

Our primary example concerns the grid-like layout of the violin’s fingerboard. By employing techniques from transformational theory, we show how intervals in this instrumental space map to intervals in abstract pitch space. Drawing on etudes and pedagogical writings for performers, we discuss ways in which such formalizations of instrumental spaces capture relationships that are salient to a performer based on the layout of notes on their instrument.

[Return to Table of Contents →](#)

Day 5: Student session

Paul Lascabettes

Ircam/CNRS STMS Lab, Paris, France

Using Mathematical Morphology to Discover Repeated Patterns in a Multidimensional Representation of Music

How to discover repeated patterns from a multidimensional representation of music? Previous works related to this question rely on the SIA algorithm, which consists in discovering the maximal translatable patterns. Many variants of this algorithm exist and we propose to develop some of the ideas present in these algorithms to improve the discovery of musical patterns from a multidimensional representation of music. In particular, we present a generalization of certain fundamental concepts, such as maximal translatable patterns, and provide answers to some of the open questions of this domain. The originality of these results arises from the use of mathematical morphology, a theory mostly used in image processing with few applications to music. In particular, we demonstrate that the principal operators of binary mathematical morphology fit perfectly into the framework of musical pattern discovery, and that this mathematical theory provides relevant results for this task.

[Return to Table of Contents →](#)

Gonzalo Romero

Ircam/CNRS STMS Lab, Paris, France

Mathematical Morphology as a tool for analyzing symbolic music

Mathematical Morphology has proven to be a very useful tool for analyzing geometric and topological structures, and in particular in the case of image processing. In this talk, I will expose some morphological operators for analyzing music by considering an image-like representation of music (piano rolls). In particular, I will show how it can be useful for harmonic and textural analysis.

[Return to Table of Contents →](#)

Nathanael Koh,
Australian National University, Australia

Exploring New Soundscapes: A Decatonic Framework for 7-limit Just Intonation

The decimal temperament is defined as a 7-limit regular temperament with comma basis $\{25/24, 49/48\}$. By exploring the derivation and usage of the decimal temperament, its applications as an analytical and compositional framework for 7-limit just intonation are discussed. This paper shows evidence of its relationship with the diatonic framework and pentatonic framework, as well as the set of the square roots of all the Pythagorean intervals (examples are $3/2$, $\sqrt{4/3}$, and $\sqrt{3}$). As the diatonic system can be interpreted as a dicot system and the pentatonic system is a semaphore tempered system, one can combine the two frameworks' half-intervals into a new, composite structure that is compatible with the entirety of 7-limit just intonation. Therefore, these relationships can be interpreted as the decimal regular temperament. Utilizing the 10-note (4+6) Moment-of-Symmetry scale of this temperament, novel notational, tonal, and modal systems were created. For musical composition, possible chord constructions using intervals of approximately 250 and 350 cents are detailed. In addition, chord progressions by moving 1 or 2 notes within a chord were established. These new ideas culminated in the composition of "Eternity", a three-movement piece ('of eternal time', 'Life and Knowledge', and 'void').

[Return to Table of Contents →](#)

Yoti Narang

University Pompeu Fabra, Barcelona, Spain

Challenges and Insights in the Analysis of Musical Dynamics in Singing Voice

Musical dynamics, such as 'piano' (soft) and 'forte' (loud), play a crucial role in creating expressive performances, particularly in vocal music. Despite their importance, the analysis and assessment of these dynamics have been relatively overlooked, partly due to the absence of appropriate datasets and a defined framework for evaluation. In my presentation, I will explore the challenges inherent in analyzing these elements, emphasizing the subjective nature of this task. Additionally, I will discuss our findings on the level of objectivity in this analysis, as indicated by the agreement among different annotators on the same set of data points. This research has significant implications for fields such as music education and synthetic voice production. An automated system capable of identifying musical dynamics could provide valuable feedback to learners or enhance the quality of synthetic vocal outputs.

[Return to Table of Contents →](#)

Mateusz Solinski,
King's College London, UK

Computational models of musicians' physiological response to music: from detective work to opportunities

The influence of music on human physiology is indisputable, particularly on musicians' physiology. However, modelling physiological response to music is a complex challenge. To understand the interaction between music and the body, we must find appropriate musical features that can predict the physiological signal changes observed. In the initial stages, the modelling task can resemble detective work – one needs to visualize the signals and, step by step, investigate the possible causes of the changes. There are several “false trails” to deal with; for example, some fluctuations in signals may be due to wandering minds or shifting body postures.

I shall describe a few computational methods for creating descriptive models of performers' physiological responses to playing music. The problems addressed include data collection and preparation, model selection and music feature choices, and variability between subjects. Specific examples will be shown on analyzing performers' physiology in ensemble playing using techniques such as mixed models and simplex plots to explain how beat-to-beat heart dynamics and heart rate variability change with rehearsals. We also use time delay stability to detect when and how players' signals synchronise. This is the first time these computational methods are applied to performers' music and physiology.

Such transparent and explanatory models of physiological response to playing music open new avenues for studying performers' physiology. They provide new tools to better understand how musicians' bodies adapt to the challenges of playing music during practice and react to shape musical interpretations.

[Return to Table of Contents →](#)