A BRIEF REVIEW OF CREATIVE MIR

Eric J. Humphrey
Music and Audio Research Lab
New York University
ejhumphrey@nyu.edu

Douglas Turnbull
Ithaca College
dturnbull@ithaca.edu

Tom Collins
Johannes Kepler Univ. Linz
tom.collins@jku.at.edu

1. INTRODUCTION

Originally formed under the auspices of music information retrieval, the field, which might now be better described as music informatics research (MIR), continues to evolve in an effort to better understand and manipulate information related to the phenomena we know as music. Being one of the most readily available data sources, much effort has been invested in the development of systems to extract high-level information from music signals, referred to here as content-based methods. However, once this information is obtained, what else could it be used for beyond retrieval? To these ends, a breakout session was convened to assess the current research climate and state of the art in creative or constructive topics in music informatics research.

In hindsight, it is understandable why the earliest efforts in content-based MIR placed a heavy emphasis on retrieval-centric problems. Within a media ecosystem, there are three fundamental roles an agent may assume —creator, distributor, and consumer— and the degree to which one partakes in each exists on a continuum. Importantly, the roles most affected by the mass embrace of the Internet and the advent of personal media players were those of distributor and consumer. The combination of these two spawned a new challenge never-before faced in the music industry: how *does* one make sense of more music than could ever be listened to in a single lifetime? As a result, the research trajectory of the field was significantly shaped by technological advances and fueled by user expectations of the 20^{th} century.

While creativity-oriented music technology is not a new research topic, there is mounting interest to coalesce a more focused effort on creative topics and applications in MIR. Though individual rationales vary widely, there are several notable reasons why it is critical to start this conversation now. First, society is continuing to progress toward the reality where anyone, anywhere can create multimedia content, referred to at present as *user-generated content*. Not only can music now be recorded at scale, but it is also no longer possible to control, a responsibility formerly upheld by major record labels. Second, "digital natives," the generation born into the age of ubiquitous connectivity and information, do not share or embody the conventional passive-consumer view of media. To paraphrase the oft-spoken words of Clay Shirky, media that fails to be interactive is seen as broken in the eyes of the next generation. It is not merely that digital natives want to engage with their media; they expect it. Additionally, there is a growing interest in the ideas and applications of computational creativity, defined as 'the study and support, through computational means and methods (including simulation), of behaviour exhibited by natural and artificial systems, which would be deemed creative if exhibited by humans' [14]. The final observation, and perhaps the most established, is the extent to which users are not only connected to content but to each other. Users and their experiences are no longer confined to a vacuum, and our research must reflect this reality. As the roles and expectations of users in a media ecosystem evolve, so too must the tools and technologies they use.

2. RECENT RESEARCH EFFORTS

An effort was made by those present to survey the current state of creative music technology by clustering known work into one of a few research areas. Reflecting on our shared knowledge, we were able to arrive at a list of four high-level topics, outlined below. Though this review would surely benefit from a more comprehensive exploration of past work, this will hopefully serve as a reasonable starting point for subsequent inquiries.

2.1 Automatic Mash-Up Systems

Popularized in recent years by artists like Greg Gillis, the musician known as *Girl Talk*, a mash-up is a musical work which takes two or more music recordings and arranges the tracks in novel ways to create an altogether new piece. The act of creating a mash-up consists of the jointly related goals of both finding suitable content and combining the signals in aesthetically satisfying ways. Often a good deal of processing may be necessary to align two waveforms rhythmically via time-stretching or harmonically via pitch-shifting, which implicitly entails the ability of the composer to perform robust beat tracking and chord or key estimation. In this manner, automatic mash-up systems can be seen as the logical progression of classic MIR tasks into the foray of creative applications; extract information from music, and use it to assemble new content. Two such sys-

tems discussed in the session were *Beat-Sync-Mash-Coder* by Griffin et al [1] and *AutoMashUpper* by Davies et al [2]. A particularly interesting consequence of such systems is the capacity to extend the life-span of commercial music, whereby songs that have grown stale can be revitalized by being placed in a new context.

2.2 Recording Tools

Another clear extension of MIR systems and research is found in the development of tools to facilitate or automate the task of recording music audio. Perhaps the most pursued application in this area is the task of automatic mixing, having received a sustained research effort for many years. In this area, the work of Scott et al [4] and Reiss [5] were mentioned during the discussion. Another well researched topic is that of intelligent equalizers, as in the work of Pardo [6] and Cartwright [7], which aim to better map user vocabularies to equalization parameters for improved usability. Noting that modern digital audio workstations are beginning to incorporate more intelligent audio processing, it would appear MIR technologies may play a larger role in recording environments in the near future.

2.3 Music and Games

Due in part to the proliferation of powerful smartphones, the world of gaming has exploded in the last decade, especially in the form of "casual" games. For the purposes of the discussion here, there are two flavors of musical games that relate to creativity and the field of MIR. The first is the potential found in Games With A Purpose (GWAP), such as Law's *Tagatune* [8] and Burgoyne's *Hooked* [9]. As will be discussed, data collection and evaluation can prove challenging in creativity research, and GWAP's may help alleviate some of this burden. Additionally, there is also a growing commercial market for musical, and notably musically creative, games. Companies like Zya [1], Harmonix [2], and Smule [3] have proven there is an audience for engaging, musical experiences, and may be fertile territory for those in the community seeking to put their knowledge into practice.

2.4 Education, Composition and Style

A separate topic in its own right, a discussion of creative technologies naturally overlaps with elements of theory and practice. Sharing some overlap in the realm of recording tools, McKay's jProductionCritic [10] demonstrates how music signal processing can be used to provide feedback to novice users. Score following software, like IRCAM's Antescofo [4] or Arzt's Complete Classical Music Companion [5], can be used to make practice more interactive, or enable improvisation and other creative liberties during performance. Perhaps one of the more untapped prospects in the union of MIR and human learning is the notion that MIR can be leveraged to help illustrate concepts in music theory. For example, it would likely have a greater impact on a student to automatically identify instances of different chords in music that he or she knows intimately; similarly, this same kind of analysis could be used to guide the composition of harmonic progressions. This is closely coupled with Pachet's work on style and form [3], which was covered at length in his keynote talk at ISMIR 2013. Importantly, automated composition systems are free to involve the user in varying degrees, and therefore do not need to solve the challenge of computational creativity definitively to be useful.

Nonetheless, a definitive solution to the challenge of music-computational creativity remains a worthy if ambitious aim. One instance of this challenge is to generate new music that conforms to an existing style [3, 11, 12, 15]. A methodology for evaluating systems that generate new stylistic compositions has been proposed by Pearce and Wiggins [11], deriving from the Consensual Assessment Technique of Amabile [16]. The methodology consists of conducting a listening study in which judges rate the stylistic success of generated passages, which are presented in random order amongst genuine excerpts from the intended style. The system will have succeeded when the difference between ratings for the generated passages and genuine excerpts are not statistically significant [11]. It is possible to include additional questions, e.g. about aesthetic pleasure, or distinguishing the excerpts as computer-generated or human-composed, in order to address questions concerning listener bias [12, 13]. While user studies have always played a role in MIR, it seems likely that they will become more numerous and assume greater importance as researchers confront the challenges of Creative MIR.

3. RELEVANT CONFERENCES

Complementing our modest coverage of research topics, a set of relevant conferences were compiled in the hopes that they too could lead to a more comprehensive survey of the field. Furthermore, it is an unique advantage of looking to the ideas and methodologies of other research communities, so as to build upon that accumulated knowledge and best practices. While many in the MIR community are familiar with *New Interfaces for Musical Expression (NIME)* and the *International Computer Music Conference (ICMC)*, there are other organizations of interest that focus on creativity at large. Three of note are the *International Conference on the Science of Creative Thinking*, organized by the Marconi Institute for Creativity [6], the *International Workshop on Musical Metacreation (MUME)* [7] and the *International Conference on Computational Creativity (ICCC)* [8]. In time, as research efforts in Creative MIR mature, these conferences present possible opportunities for our field to contribute to the greater body of creativity research.

[1] http://www.zyamusic.com/
[2] http://www.harmonixmusic.com/
[3] http://www.smule.com/
[4] http://imtr.ircam.fr/imtr/Score_Following
[5] https://www.youtube.com/watch?v=ZPE_1TwP5DA
[6] http://www.mic-conference.org/
[7] http://www.metacreation.net/mume2013/
[8] www.computationalcreativity.net/iccc2013/?

4. CONCLUSION

Based on the interest at ISMIR and a growing range of topics, creative applications in MIR seems to hold a promising future. It is not without its challenges, however, as it will not be possible to rely solely on a set of predefined metrics in order to evaluate system performance. This style of research may be unfamiliar to some with engineering backgrounds, but the interdisciplinary nature of our community makes it an ideal venue for this course of research. Ultimately, what will likely set Creative MIR apart from other flavors of music technology are those facets that make MIR itself unique: the synthesis of music theory and machine learning, signal processing and psychology, data-driven insights and strong domain knowledge.

5. REFERENCES

[1] Griffin, G., Kim, Y. and Turnbull, D. "Beat-sync-mash-coder: A web application for real-time creation of beat-synchronous music mashups." *Proceedings of the IEEE International Conference on Acoustics Speech and Signal Processing (ICASSP),* 2010.

[2] Davies, M., Hamel, P., Yoshii, K., and Goto, M. "AutoMashUpper: An Automatic Multi-Song Mashup System," *Proceedings of the International Symposium on Music Information Retrieval,* 2013.

[3] Pachet, F., Suzda, J., and Martín, D. "A Comprehensive Online Database of Machine-Readable Leadsheets for Jazz Standards," *Proceedings of the International Symposium on Music Information Retrieval,* 2013.

[4] Scott, J., Prockup, M., Schmidt, E. and Kim, Y. "Automatic multi-track mixing using linear dynamical systems." *Proceedings of the 8th Sound and Music Computing Conference,* 2011.

[5] Reiss, J. "Intelligent systems for mixing multichannel audio." *Proceedings of the International Conference on Digital Signal Processing,* 2011.

[6] Pardo, B., Little, D., and Gergle, D. "Building a personalized audio equalizer interface with transfer learning and active learning." *Proceedings of the 2nd International ACM Workshop on Music Information Retrieval with User-centered and Multimodal Strategies,* 2012.

[7] Cartwright, M. and Pardo, B. "Social-EQ: Crowdsourcing an Equalization Descriptor Map." *Proceedings of the International Symposium on Music Information Retrieval,* 2013.

[8] Law, E., Von Ahn, L., Dannenberg, R., and Crawford, M. "TagATune: A Game for Music and Sound Annotation." *Proceedings of the International Symposium on Music Information Retrieval,* 2007.

[9] Burgoyne, J. A., Bountouridis, D., van Balen, J. M. H., and Honing, H. "Hooked: a Game for Discovering what Makes Music Catchy." *Proceedings of the International Symposium on Music Information Retrieval,* 2013.

[10] McKay, C. "JProductionCritic: An Educational Tool for Detecting Technical Errors in Audio Mixes." *Proceedings of the International Symposium on Music Information Retrieval,* 2013.

[11] Pearce, M.T., and G.A. Wiggins. "Evaluating cognitive models of musical composition," *Proceedings of the Fourth International Joint Workshop on Computational Creativity,* pp. 73–80, 2007.

[12] Collins, T. *Improved methods for pattern discovery in music, with applications in automated stylistic composition,* PhD thesis, Faculty of Mathematics, Computing and Technology, The Open University, 2011.

[13] Moffat, D.C., and M. Kelly. "An investigation into people's bias against computational creativity in music composition," *Proceedings of the Joint Workshop on Computational Creativity,* 2006.

[14] Wiggins, G.A. "A preliminary framework for description, analysis and comparison of creative systems," *Knowledge-Based Systems*, Vol. 19, No. 7, pp. 449–458, 2006.

[15] Gonzalez Thomas, N., Pasquier, P., Eigenfeldt, A., and Maxwell, J.B. "Meta-Melo: A system and methodology for the comparison of melodic generation models." *Proceedings of the International Symposium on Music Information Retrieval,* 2013.

[16] Amabile, T.M. *Creativity in context*, Boulder, CO: Westview Press, 1996.