TISMIR

EDITORIAL

Introduction to the Special Collection "20th Anniversary of ISMIR"

Arthur Flexer*, Florence Levé†, Geoffroy Peeters‡ and Julián Urbano§

1. A Brief Look at 20 Years of ISMIR Evolution

From the International Symposium on Music Information Retrieval (ISMIR) held in Plymouth, Massachusetts USA, in October 2000 to the International Society for Music Information Retrieval Conference held in Delft, The Netherlands, in November 2019, twenty annual meetings took place that have allowed MIR researchers to meet, present and exchange their ideas leading today to an MIR community and identity for more than 2000 researchers (according to the number of subscriptions to the mailing list). Not only the acronym meaning changed from symposium to conference, but also the size of its yearly conference from less than 100 researchers to more than 500. Due to ISMIR's single track conference nature the number of accepted papers has been constant at around 100 since 2004 (solid line in **Figure 1**), but the number of submitted papers is now at a level of about 250 (dashed line in **Figure 1**). This gives an acceptance rate of about 40% at the moment.

Also the topics being addressed by the community, or at least how they are addressed, have changed. With respect to that, it is interesting to consider that research topics such as "optical music recognition" or "symbolic music representation", which had gradually disappeared to the advantage of audio-based MIR around 2010, are now again a full part of the MIR landscape (boosted by deep learning approaches). Some research topics (as presented in the call for papers) disappeared ("music digital libraries", "music indexing and metadata", "query languages for music IR", "intellectual property rights issues"), some have been renamed ("building up music databases" in "MIR datasets and annotation protocols"), and some new topics appeared ("music generation", "lyrics", "web mining", "multimodality", "evaluation metrics", "user-centered MIR").
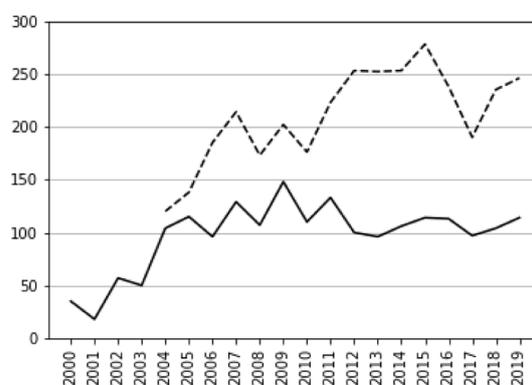
Figure 1: Number of submitted (dashed line) and accepted papers (solid line) on y-axis versus year of ISMIR conference on x-axis. Number of submitted papers before 2004 could not be obtained.

MIR is also younger today than ever (60% of the papers accepted in 2019 were written by students) and more diverse (although MIR diversity has now been considered as a major issue for the community). The importance of the industry in these shifts should not be neglected. During these 20 years, the industrial landscape surrounding MIR has also drastically changed. While the music industry (Philips, Sony) was considered as a potential outlet for MIR research in the past, MIR has today found a fruitful academic-industrial relationship with IT companies (Spotify, Google, Pandora, Deezer, etc.), who help it develop at a larger scale.

2. Putting together this Special Issue

The International Society for Music Information Retrieval Conference (ISMIR) is now the world's leading research forum on processing, searching, organising and accessing music-related data. Commemorating the anniversary of 20 years of ISMIR, this special issue reflects on the development of MIR as a research field by charting its progress over the last two decades and elaborating on its future. This collection of papers consists of submissions to an open call and of invited extended versions of anniversary papers already presented at the 20th ISMIR Conference in Delft, the Netherlands, 2019. We asked for

* Institute of Computational Perception, Johannes Kepler University Linz, AT

† Université de Picardie Jules Verne, MIS, FR

‡ Télécom-Paris, Institut Polytechnique de Paris, FR

§ Multimedia Computation Group, TU Delft, NL

Corresponding author: Arthur Flexer (arthur.flexer@jku.at)

thought-provoking and insightful papers reflecting on the evolution of the MIR field as a whole or particular topics therein. All accepted papers provide a critical state-of-the-art overview of a broad MIR problem area and at the same time discuss its midterm future.

Among the six accepted papers of this special issue, two papers are extended versions of anniversary papers (A. Lerch et al. and P. Knees et al.) presented at ISMIR 2019, one is extended from a keynote of ISMIR 2019 (G. Born), three are new contributions (O. Nieto et al., C. E. Cella, and J. Devaney).

3. Overview of the Special Issue

The six articles in this special issue cover a wide range of topics: an analysis of diversity in the field of MIR and a personal history of the relationship of MIR and contemporary classical music, both coming from colleagues outside of the immediate MIR community but with a close affinity to it; a review of music performance analysis and a proposal to facilitate interdisciplinary research on human engagement with music; two state-of-the-art reports on the more classical tasks of music structure analysis and user interfaces for music discovery. We will now provide brief overviews of these six papers to give our readers a first impression of what to expect from the special issue.

The article **"Diversifying MIR: Knowledge and Real-World Challenges, and New Interdisciplinary Futures"** by Georgina Born, written from her music anthropological view, asks for a stronger diversification of MIR research. According to the author, an intense interdisciplinary dialogue with fields like musicology, music anthropology and sociology should bring about a diversification of who takes part in MIR research, of the music with which MIR engages, and of the kinds of results and knowledge MIR produces. The author invites MIR researchers to join her and others from relevant disciplines to form a "think tank" which is able to bring about "responsible innovation" in MIR beyond computational music genre recognition or recommendation in its current profit-oriented form, setting human musical flourishing as the goal.

The following contribution to this special issue already is very much in line with Georgina Born's plea for diversification of MIR. The article **"Music Information Retrieval and Contemporary Classical Music: A Successful Failure"** by Carmine-Emanuele Cella tells a very personal story of using MIR results and tools in the context of contemporary classical music creation. The author's engagement with MIR over the last 15 years has not been without conflict and his story is able to point towards new research directions for MIR. The author argues that the quest for deep "end-to-end machines that magically solve all our problems brought us on the wrong path" and that machine learning in MIR should work "inside a larger context framed by our knowledge". He also advocates that concentrating on the complex world of contemporary classical music could help re-focus MIR's research goals.

The article **"An interdisciplinary review of music performance analysis"** by Alexander Lerch et al. explores how the variations in the performance of a musical composition are related to the intentions of the performer and their impact on the listener's perception. Aiming at "reducing the gap between MIR and Music Performance Analysis" (MPA), it introduces MPA and its challenges from an MIR perspective. Among the main challenges, it states the importance of a better differentiation between score and performance parameters, allowing to better analyze their respective contributions during evaluation. It highlights the importance of taking into account human perception in MPA research, as for the measurement of the listener's response. It also lists many data-related challenges, in particular concerning the annotation of data.

In the same direction, the article **"Using note-level music encodings to facilitate interdisciplinary research on human engagement with music"** by Johanna Devaney addresses this problem of annotation of data. The author considers that "the acts of analysis, performance, and listening are related but distinct types of human response to music". By surveying the developments in music encoding research, the article highlights that many MIR data annotations "focus on project-specific encoding formats and therefore cannot be utilized to their fullest potential". It proposes concrete suggestions to increase interoperability between datasets by linking data on analysis, performance, and listening with specific musical elements such as notes or chords, and encourages the development of new tools "to assist researchers in encoding their annotations with rich music data."

The article **"Audio-Based Music Structure Analysis: Current Trends, Open Challenges, and Applications"** by Oriol Nieto et al. presents an overview of the work achieved over the last decade for a typical MIR task: music structure discovery from audio signals. This task shares problems common to much of MIR, in that after 20 years of research it still faces problems of task definition, lack of accessible large annotated corpora, and unsolved definition of informative performance measures. Like many MIR tasks it has also faced the deep learning revolution, moving from hand-crafted systems (combining audio features with hand-made decision algorithms) to fully supervised machine-learning systems where knowledge is acquired from annotated data (which reinforces the need for annotated corpora hence for clear task definition). As an opening to a new decade of research, the authors propose a set of challenges to help the field in its further development: how to deal with subjectivity (multiple possible annotations) and ambiguity (multi-dimensionality of structure, novelty vs repetition) of the task or with the inherent hierarchy of music structure. It also highlights a new set of potential applications in music creation, production, generation, recommendation, live performance and visualization.

The article **"Intelligent User Interfaces for Music Discovery"** by Peter Knees et al. reflects on the evolution of user interfaces for browsing and discovery, and how they may have influenced MIR research directions for the past twenty years. As such, this paper is a well-motivated

introduction to our field to share with outsiders. Knees et al. argue for three major developments that shaped user interfaces and listening practices. Phase 1 covers personal music collections or small commercial catalogs, and entails the development of audio description algorithms to facilitate the automatic organization and search of music based on sound qualities. Phase 2 covers collective music web platforms, and entails collaborative and automatic semantic description of music, thus getting closer to users and exploiting user-generated data. Phase 3 covers streaming services, and represents the dominance of recommender systems built upon very large online music collections and user interactions. The authors also discuss how future technology will be more context aware and personalized, with examples from smart speakers, to conversational interaction or automatic music generation.

In summary, the papers of this special collection are what we were hoping for: thought provoking and insightful contributions that open new directions for our research community while also providing solid overviews on topics at the center of MIR. We wish our readers an interesting and informative time with this collection of articles!

Competing Interests

All four authors are members of the editorial board of the Transactions of the International Society for Music Information Retrieval. All four authors were removed completely from all editorial processing. There are no other competing interests to declare.