

What can MIR teach us about music? What can music teach us in MIR?

Cory McKay

Marianopolis College, Canada

Women in Music Information Retrieval (WiMIR) Workshop
October 30, 2021
On-line workshop

Principle themes of this talk

- How MIR might engage more deeply with fundamental, foundational musical inquiry
- How we might enrich such inquiry by better engaging with experts not (yet) involved in MIR

What is MIR?

- **Wikipedia:**
 - “. . . the interdisciplinary science of retrieving information from music”
- **Interesting questions:**
 - What are the implications of a field being fundamentally interdisciplinary?
 - Must MIR be a science?
 - If not, what else could it be?
 - What does “music” comprise?
 - Concrete audio, symbolic or textual representations?
 - Information about music creators, consumers, curators, etc.?
 - Information about larger social or cultural contexts?
 - More?

What is MIR?

- ISMIR website:
 - “. . . a field that aims at developing computational tools for processing, searching, organizing, and accessing music-related data.”
- Interesting questions:
 - Are “tools” the essential goal of MIR?
 - Can asking deeper or broader questions about music also be important?
 - “processing, searching, organizing, and accessing”
 - Could “learning from” or “trying to understand” be added to this list?
 - Or using music-related data to create interesting music?

What are the **goals** of MIR?

- musicinformationretrieval.com:
 - “discover, organize, monetize media collections”
 - “search (‘find me something that sounds like this’) songs, loops, speech, environmental sounds, sound effects”
 - “workflows in consumer products through machine hearing”
 - “automatic control of software and mobile devices”
- **Interesting question:**
 - Are commercial applications associated with “monetizing” and “consumer products” of central importance to MIR?

Who is involved in MIR?

- Wikipedia again:
 - “Those involved in MIR may have a background in musicology, psychoacoustics, psychology, academic music study, signal processing, informatics, machine learning, optical music recognition, computational intelligence or some combination of these.”
- Other possibilities?
 - Music curators (librarians, critics, etc.)?
 - Researchers in sociology, cultural studies, etc.?
 - Music economists, business experts, lawyers, etc.?
 - Health researchers / professionals (e.g. music therapists, neuroscientists, etc.)?
 - Music consumers (both active and passive)?
 - Music creators (composers, performers, etc.)?
 - Music educators?
 - Others?

In the beginning (of ISMIR) . . .

- Collected ISMIR proceedings:
 - <https://ismir.net/conferences/>
- Fascinating historical ISMIR reads:
 - 2000: “Music IR: Past, Present, and Future” by Alexandra L. Uitdenbogerd
 - 2001: “Music Information Retrieval as Music Understanding” by Roger B. Dannenberg
 - “. . . the central problem of Music Information Retrieval is Music Understanding . . .”

In the beginning (of ISMIR) . . .

- Emphasis on both multi-disciplinarity and cross-disciplinarity
 - Single track presentations
- Particularly strong representation of experts in library science, musicology and signal processing
- Less emphasis on machine learning
 - HMMs were especially popular among those who did use it
- Stronger focus on symbolic music
- Musical queries / search / retrieval were of central interest
- Relatively little direct industry involvement
- Relatively informal approach to evaluation

ISMIR since then: Evaluation

- Evaluation has become much more rigorous
- Overfitting and biased datasets are now treated as unacceptable
- Significance / hypothesis testing appears more often
 - Although not always applied correctly
- Approaches are often compared directly, on the same data
- MIREX (starting in 2005) made a huge contribution to improving the quality and scope of evaluation / exchange
 - And the accessibility of shared datasets
 - Also a useful forum for community reflection on evaluation

ISMIR since then: Evaluation

- But . . .
 - Have we become too focused on established tasks where the output of different approaches can be easily compared?
 - With easy-to-access existing datasets
 - With existing bodies of work to which new results can be compared
 - Has our focus on evaluation narrowed or biased the kind of work being done in MIR? Caused us to treat certain problems in ways that are too shallow?
 - e.g. evaluation-centric attitudes can lead to discomfort with ambiguity (and improper treatment of it), but ambiguity is intrinsic to many aspects of music, like harmony
 - Have we become reluctant to ask important research questions whose potential answers cannot be easily evaluated?
 - Do we as reviewers disincentivize such work?

ISMIR since then:

Deep learning and big data

- Enable us do things that would have seemed like magic in 2000
- Techniques can be accessible (and ostensibly effective) even to researchers with little or no musical domain knowledge

ISMIR since then:

Deep learning and big data

- But . . .
 - Can accessibility to researchers with limited musical expertise lead to naïvely posed research questions and interpretation of results?
 - Are those models that have limited or no interpretability able to teach us much about music at a fundamental level?
 - Are assumptions that the size of (usually very) noisy large datasets will enable trained models to smooth over that noise actually valid?

ISMIR since then:

Deep learning and big data

- But . . .
 - Has overconfidence in (and overreliance on) these techniques caused us to move away from the important complementary role that hand-crafted features and high-quality (but usually of necessity small) curated datasets might play?
 - What about interesting problems where only small datasets are available or possible?
 - Are we as reviewers also disincentivizing work in such areas?

ISMIR since then: **Multimodality**

- Many types of musical sources are studied now
 - Symbolic, audio, text, user data, images, videos, etc.
- Permits broad and deep studies of music
 - Each source type can potentially reveal different kinds of insights
 - Both into the music itself and in how music communities, consumers and creators engage with it

ISMIR since then: Multimodality

- But . . .
 - Are we actually using this diverse data to perform these kinds of studies, or are we just using it to get small improvements in our answers to the same possibly overly simplistic questions we were already asking?
 - How multimodal is multimodal research?
 - Most MIR papers still use 1, 2 or, rarely, 3 source types
 - The ISMIR papers by McKay et al. (2010) and McFee & Lanckriet (2012) are still the only ones I know of that used 4 or more source types
 - How accessible are published multimodal datasets in practice?
 - Are all source types actually publicly accessible?
 - How robust are links between source types as time passes?

ISMIR since then:

Industry involvement

- Industry has come to play a core role in both MIR and ISMIR
 - e.g. Shazam (2002) provided early proof of concept for MIR's general utility and scalability
- Provides substantial funding for events (including WiMIR)
- Provides invaluable APIs making music analysis and processing more easily accessible
 - e.g. The Echo Nest and Last.fm APIs provided early APIs that changed the course of MIR
- Provides invaluable datasets
 - e.g. The Million Song Dataset (2011) was built on core contributions from The Echo Nest, Last.fm, musiXmatch and SecondHandSongs
- Creates an expanded public profile for MIR
- Provide the field with a useful grounding
 - In what is practical and possible
- Provides jobs

ISMIR since then:

Industry involvement

- But . . .
 - Commercial APIs and datasets have limited longevity
 - e.g. after The Echonest and Last.fm were acquired by Spotify and CBS Interactive, respectively
 - e.g. aspects of my own jMIR software were terminally broken when Google and others changed their public web services
 - Commercial datasets tend to provide incomplete music (e.g. 30 second excerpts)
 - Easy availability of such data has disincentivized the creation and expansion of datasets comprised of complete open music
 - Risks biasing research against studies of music incorporating longer-scale structural elements, for example

ISMIR since then:

Industry involvement

- But . . .
 - Commercial MIR APIs are rarely open source, or even transparent in how their algorithms work
 - Some algorithms may accidentally be applied inappropriately or naively by unwitting users
 - Makes it hard to move forward the state of public knowledge
 - Industry is interested primarily in aspects of MIR that can be monetized
 - Has clearly lead to many useful, interesting and meaningful contributions
 - But tends to channel research away from less easily monetizable foundational questions about the nature of music and how we as humans interact with it, and away from pure research in general

ISMIR since then:

People and disciplines involved

- Important efforts to make the MIR community more diverse
 - More women
 - More people from outside the traditional MIR centers of Western Europe, East Asia and North America
 - Increased accessibility for those with limited or no funding
- Important efforts to expand MIR's scope to include more musical traditions
 - Popular, traditional and “art” musics from around the world
- Participants from a broader range of fields have joined what was already a richly multidisciplinary field to begin with
 - e.g. neuroscience, composition, etc.
- More experts in computer science and software development
 - Better-informed, more sophisticated use of machine learning
 - More professional and more easily extensible software architectures
 - Better user interfaces

ISMIR since then:

People and disciplines involved

- But . . .
 - How introspective have we really been about fundamental changes in how non-Western musical traditions should be represented, analyzed, thought about and interacted with?
 - Are we just plugging this music into existing frameworks?
 - Has the professionalization of MIR caused us to avoid difficult foundational questions about music?
 - That might be hard to formulate or answer using methodologies and ways of thought that engineers and software developers are trained on?
 - Is MIR treated by some more as a toy problem for machine learning than as an area of fundamental inquiry?

ISMIR since then:

People and disciplines involved

- But . . .
 - Has significant training or experience in music itself become rarer among MIR researchers?
 - With potential consequences relating to improper assumptions, naïve conclusions, etc.?
 - Are important fields like musicology, sociology, music perception/cognition or the library sciences gradually becoming marginalized within MIR?
 - We are undoubtedly a deeply multidisciplinary community, but are we fully taking advantage of this through interdisciplinary collaborations?
 - Or are we operating primarily discretely in sub-disciplinary silos?

Sample problem domain:

Genre classification

- A persistent but much maligned ISMIR topic
 - And the topic of my first ISMIR paper in 2004, before I eventually turned my focus to other things
- A common criticism:
 - Not the most useful way for users to retrieve or organize music
 - Ontologically unconstrained similarity-based approaches often proposed as better than genre for such purposes

Sample problem domain:

Genre classification

- These are arguably accurate points, but does that mean MIR research in genre should be abandoned?
 - Does a genre play a role in the way its fans construct their personal identities?
 - Or perhaps give us insight into the intersections of their broader identities?
 - Can empirical genre study provide insight into genre conventions that are informally understood by some but not (yet) clearly articulated?
 - Useful for education, music theory, musicology, etc.?
 - Can systematic studies of genre changes (individually and in relation to one another) over time provide insight on broader musical or social trends?

Sample problem domain:

Genre classification

- Most MIR genre research to date has incorporated unsupportable simplifying assumptions that compromise the meaningfulness of results
 - e.g. ontologies with oversimplified flat or hierarchical structuring
 - e.g. non-overlapping genre classes
 - e.g. single genre memberships for individual songs
 - e.g. focus primarily on audio content (mostly timbre), to the exclusion (mostly) of higher-level musical traits or cultural context
- Why have such problematic assumptions been made?
 - Perhaps they make a musically and culturally sophisticated problem more accessible?
 - Perhaps they make evaluation and comparison easier?

Sample problem domain:

Genre classification

- So, how might MIR approach studying genre in more meaningful ways?
 - As a musical problem using computation, rather than a computational problem using music?
 - Cross-disciplinary collaborations with genre experts in other fields?
- Genre is certainly an area of interest for experts in fields other than MIR, although they do not (yet?) typically using MIR-related methodologies
 - Musicology and music cognition in particular

Interdisciplinary collaboration

- Speaking generally, perhaps working directly with experts outside our own sub-disciplines can help us expand the scope, informedness, and meaningfulness of MIR research?
 - Others within the ISMIR community
 - People from outside the ISMIR community

Interdisciplinary collaboration

- Do people from other relevant communities know what MIR and ISMIR are?
 - Musicologists? Music theorists? Music psychologists? Composers? Performers? Sociologists? Music librarians?
- Certain people from these fields are already deeply involved in ISMIR
 - e.g. Audrey Laplante, Stephen Downie, Jin Ha Lee, Sally Jo Cunningham, Xiao Hu, David Bainbridge, etc. have made prominent ISMIR contributions relating to library science and user studies
- But are they representative of their larger disciplines?
 - e.g. would a survey of participants at musicology conferences/meetings like IMS, AMS, MedRen, etc. reveal an awareness of ISMIR, or even MIR techniques in general?
- Are MIR researchers really impacting other research fields (musical and otherwise) meaningfully?

Interdisciplinary collaboration

- How can we interest experts in other fields in working with us?
 - By attending conferences in relevant fields we have little expertise in?
 - Doing better in making experts in diverse fields feel welcome and valued at MIR events?
- How can we enrich MIR with their knowledge, methods and perspectives?
 - And how can we enrich their fields in turn?
- How can we communicate effectively with them?
 - Overcoming (benefiting from?) varied goals, priorities, terminologies, methodologies, fundamental assumptions and ways of thinking?

Sample domain:

Renaissance musicology

- Since 2016 my primary research focus has been on Renaissance musicology
- I knew very little about the field at the time, but have benefited hugely from the insights of collaborators in musicology and music theory
 - Julie Cumming, María Elena Cuenca, Esperanza Rodríguez-García, Peter Schubert and Sylvain Margot

Sample domain:

Renaissance musicology

- My original approach was naïve and foolishly MIR-centric
- I assumed goals that were interesting from an MIR perspective were interesting from a musicological perspective
 - And was unaware of certain essential musicological goals
- I was unaware of essential aspects of how Renaissance music is annotated / encoded
 - e.g. unannotated elements like *musica ficta*
 - e.g. non-standardized rhythmic note value multiples
 - Cumming, McKay, Stuchbery and Fujinaga (2018) expand on such issues

Sample domain:

Renaissance musicology

- I assumed it would be appropriate to give musicologists information formatted in ways I would find useful myself
 - e.g. tables of raw feature values, information gain analysis results, direct results of clustering or supervised machine learning experiments, etc.
- I was unaware of less formally structured but essential musicological sources of insight
 - e.g. provenance analysis of manuscripts, historical texts, illuminations, etc.

Sample domain:

Renaissance musicology

- What have I learned to do better?
 - Humility, an openness to different ways of thinking and a collaborative mindset are essential
- Learn from domain experts about what matters to them and why
 - But also offer suggestions about new types of research questions that MIR might be able to help answer
- Learn about methodologies used in other fields
 - Some may be more subjective or anecdotal than scientists might typically be comfortable with, but there can be good reasons why such approaches are used
 - And also offer suggestions about how MIR-oriented techniques might be integrated with their methodologies
- Frans Wiering has offered interesting insights on such issues

Sample domain:

Renaissance musicology

- Learning to communicate effectively can require significant effort for all involved, but it is very much worth it
- Finding ways of graphically representing data and results of statistical analyses in ways that are accessible to expert human interpretation can be very helpful
 - Large-scale broad data analysis is one of the biggest benefits MIR can bring to musicology, but results must be relatively easily interpretable by domain experts
 - Paula Muñoz-Lago and her DIDONE project collaborators are doing some very interesting work in this area

Sample domain:

Renaissance musicology

- Identifying musical salience and causality can be essential
 - e.g. a composer's works may demonstrate a statistical pattern that can be used to reliably identify their work, but is this enough to be musicologically interesting?
 - Maybe yes, if all you are trying to do is credit unattributed works
 - Probably not if the pattern is not musically salient and you wish to gain musically meaningful insights about the composer's style or influences, for example

Sample domain:

Renaissance musicology

- Be willing to adapt or develop new MIR techniques, do not just try to apply existing systems if they are not a good fit
 - e.g. implementing new kinds of features, in my case, such as those associated with contrapuntal imitation
- Everyone involved should be open to both the conventional and then unconventional
 - e.g. MIR techniques can be used to systematically test established musicological or theoretical hypotheses
 - e.g. MIR techniques can reveal unexpected but musically meaningful insights
 - Such as the importance of rhythm, or influences on the madrigal's development, in my case

What next?

- MIR as a field has historically made impressive efforts towards introspective self-examination
- Let's continue that tradition
 - What kinds of research questions do you think MIR should be asking?
 - Both in general and specifically?
 - Are we in fact asking such questions?
 - Have we been successful in answering such questions so far?
 - How can we improve interdisciplinary collaboration?
 - Within the ISMIR community?
 - With research communities not aware of ISMIR?