A MULTICULTURAL APPROACH IN MUSIC INFORMATION RESEARCH

Xavier Serra
Music Technology Group
Universitat Pompeu Fabra, Barcelona
xavier.serra@upf.edu

ABSTRACT

Our information technologies do not respond to the world's multicultural reality; in fact, we are imposing the paradigms of our market-driven western culture also on IT, thus facilitating the access of a small part of the world's information to a small part of the world's population. The current IT research efforts may even make it worse, and future IT will accentuate this information bias. Most IT research is being carried out with a western centered approach and as a result, most of our data models, cognition models, user models, interaction models, ontologies, etc., are culturally biased. This fact is quite evident in music information research, since, despite the world's richness in terms of musical culture, most research is centered on CDs and metadata of western commercial music. This is the motivation behind a large and ambitious project funded by the European Research Council entitled "CompMusic: Computational Models for the discovery of the world's music." In this paper we present the ideas supporting this project, the challenges that we want to work on, and the proposed approaches to tackle these challenges.

1. INTRODUCTION

In the last decade there has been great progress in the field of Music Information Retrieval, but the rate of improvement in most retrieval tasks, like the ones evaluated within the MIREX[1] initiative, is clearly slowing down. We are starting to see the limits of the current signal processing and machine learning approaches and efforts are being made to find new ways to advance in terms of the currently identified problems. However, maybe more importantly, there are many relevant problems that have not yet been looked at for which current methodologies may not work at all. For these new problems, it is even more important to concentrate on new research approximations for the computational processing of music information. But, which approximations should we use and what new problems should we look into?

The most general and common approaches used in music information processing are based on data modeling; they start from large data repositories and use signal processing and machine learning techniques to construct models. We should combine these approaches with other research methodologies, like the ones coming from Musicology, Cognition or Human Computer Interaction. The approaches coming from Musicology develop models originating from music theory in which a thorough formalization contributes to an understanding of the theory itself, its predictions, and its scope. On the other hand, the approaches coming from Cognition aim at constructing theories of music cognition, formalizing the mental processes involved in listening to and performing music. The approaches than have been developed by the Human-Computer Interaction research community focus on the users and bring in methodologies coming from behavioral sciences. There is a need to expand the information modeling methodologies but we should do it with a cross-fertilization approach.

Making sense of music is much more than decoding and parsing an incoming stream of sound waves into higher-level objects such as onsets, notes, melodies and harmonies. Music is embedded in a rich web of cultural, historical, commercial, technological and social contexts that influence how it is produced, disseminated and perceived. For example, many qualities or categorizations attributed to a piece of music by listeners cannot solely be explained by the content of the audio signal itself. It is thus clear that high-quality automatic music understanding and description can only be achieved by taking into account information sources that are external to the actual music signals.

Despite recent efforts by the MIR research community to open up towards non-western music [1] the major focus has been on the study of a few aspects of the western commercial music of the past few decades. This music repertoire has conditioned the problems that we are working on and thus the solutions obtained. If we study other aspects, and especially other types of music, we will find new interesting problems to be solved that will require new methodologies and new solutions.

For the development of music information models, it is of great advantage to work with musical repertoires, thus

[1] http://www.music-ir.org/mirex

musical cultures, that have a classical tradition and for which there is a relevant body of academic studies available. Despite the trend towards a musical monoculture [2] there are robust music cultures with a classical tradition in places like China, India, Turkey, Indonesia, or in the Arab world, traditions that form a counterpoint to the western music context. A few of these musics have excellent musicological and cultural studies available; they maintain performance practice traditions and they exist within real social contexts. Thus, some of these traditions can be an excellent ground on which to build new information models and the means to challenge the dominant western-centered information paradigms. If we are able to describe and formalize the music of these cultures we might open up the current information models in a way to better capture the richness of our world's music [3].

CompMusic, a research project funded by the European Research Council, was born from these concerns and from the realization that unless we do something, we might lose our world's cultural richness. What the European colonization of many parts of the world did not completely achieve, that is, to westernize their cultures, will now happen by the global use of our western-centric information technologies. These information technologies, and even more, the future technologies emerging from current research, will have a great impact on the way we maintain, access, and use the world's information resources, and thus they will condition the survival of that information and of the cultures that produce it. Unless the information technologies support the diversity of cultural perspectives that exists in the world, we will slowly lose our wonderful world cultures.

It could be argued that there is no need to look outside western music to find new problems with which to improve our MIR techniques. It is completely true: there are many aspects of our classical, folk and other western music traditions for which current MIR methodologies do not work. But it is also true that when we try to analyze some of the non-western classical traditions, our research limitations become even more obvious, and it might be easier to look at the problems with a fresh perspective, without being contaminated by our traditional view on western music. At the same time, the opportunity to help bring new musical traditions into the world of Music Computing is a very rewarding task.

2. PROJECT OBJECTIVES AND APPROACHES

The main objective of the CompMusic project is to promote and develop multicultural perspectives in Music Computing research. We want to identify music problems coming from culture-specific contexts and work on solutions that might result in new computational methodologies of interest for a wide variety of music information processing problems.

In music computing research there is a need to advance in the description and formalization of music with the aim to make it more accessible to computational approaches. In this project we will work on finding ways to reduce the known gap between audio signal descriptions and semantically meaningful music concepts. But we believe that we will only make progress if we approach this issue by combining academic disciplines, such as Information Processing, Computational Musicology, Music Cognition and Human-Computer Interaction, thus using both qualitative/quantitative methodologies and scientific/engineering approaches. We also need to work with a variety of information sources, such as audio features, symbolic scores, text commentaries, or user evaluations. Therefore, we need to open up the research methodologies being used.

CompMusic will take a cultural approach, meaning that information modeling techniques will be developed for specific music repertories with emphasis on their cultural contexts. To model a musical repertoire we need to understand and model the music together with the user community that creates, enjoys and supports that particular music.

An important consideration in a cultural-specific research approach is the issue of who should be involved. We need to involve different cultural perspectives. We believe that is not possible to carry out such a project solely with a research group immersed in a cultural context different to the one to be studied. CompMusic will involve experts, research teams and users immersed in the musical cultures being studied.

Given the availability of musicological/cultural studies and our existing collaborations with researchers and experts within diverse musical cultures, in CompMusic we will study art music traditions in India (Hindustani and Carnatic), Turkey (Ottoman), North Africa (Andalusian), and China (Han). All these traditions offer large and useful information sources from which to develop our project. However, knowing that we will not be able to cover it all with the same depth, we have decided to first focus on the Hindustani, Carnatic and Ottoman traditions, and progressively incorporate specific aspects of the other cultures.

The Hindustani and Carnatic music traditions of the Indian subcontinent offer the possibility to study relevant problems in all aspects of interest to this project. Their instruments, such as the Tambura, the Veena or the Sitar, have been built to emphasize sonic characteristics that are quite different to those of the typical western musical instruments. The concepts of Raga and Tala are completely different to the western concepts used to describe melody and rhythm, thus their understanding and computational description requires new approaches. Their particular improvisatory nature based on the Ragas makes these repertoires very alive and constantly evolving. The musical scores used in both Hindustani and Carnatic traditions serve a different purpose to those of western music and thus have to be studied differently. The tight musical and sonic relationship between the singing voice, the other melodic instruments, and the percussion accompaniment within a piece, requires going beyond the analytical and

modular approaches commonly used in MIR. The special and participatory communication established between performers and audience in concerts, offers great opportunities to study issues of social cognition. The specific differences between Hindustani and Carnatic traditions in all the aspects mentioned permit the study of musical and cultural differentiation issues. Even though we will only be able to tackle some of these problems in this project, we hope to encourage other researchers to study many of the others.

Ottoman art music is very much in the intersection of many music traditions. The Turkish culture is at the junction of European and Asian cultures, and its musics reflect this geo-strategic position. The makam and its performance practice, the basis for the melodic organization of Ottoman music, offer a very rich source of microtonal studies, and there is already extensive theoretical work on its characteristic non-equal-tempered pitch organization. Its complex rhythmic patterns also challenge the audio analysis approaches currently used in MIR. This music culture offers a great context in which to study the open debate in many music traditions on the relevance of theory versus practice in the understanding of a given music repertoire. The adaptation of the western music notation to fit the microtonal aspects of their music, together with other European music influences, also offer a challenging set of musical style characterization problems. It is remarkable that despite all the cultural influences and political pressures, Ottoman art music has maintained a strong personality and an active social context. In fact, it is a music that has influenced many music traditions of both Europe and the Middle East, an aspect that also deserves a thorough study.

Andalusi music can be traced back to the time when the Arabs had a flourishing culture in Al-Andalus (now part of Spain) for many centuries. This music tradition had to move its main center to North Africa when the Arabs were forced to leave Al-Andalus at the end of the 15th century. These historical and geographical developments give us very interesting cultural context problems to be studied. Andalusi music shares many fundamental characteristics with Ottoman tradition, for example the maqams used are very much related to the Ottoman makams. Thus, the comparative study of these two traditions should be very fruitful. Also the fretless Oud, a preeminent instrument upon which most of the theoretical studies of this music tradition have been based, can give us relevant insights into the characteristics of the music. The approach to this classical tradition is more difficult for us than the others given the lack of research activity in the Arab countries of North Africa related to computational musicology. It will be difficult to find collaborators for this project in these countries.

The other music culture that we want to study, the Han music of China, is also an amazing source of information processing problems. The relationship between language and music in a culture that has a tonal language offers interesting research problems. The importance of the traditional philosophies and nature in all Chinese classical art forms gives a completely different context to the Han music. For example, the concept of emotion in this music has a different meaning to the one used in the west, and the social function of the music is also very different. From all the cultures that we want to study, the Han culture is the one that has had the least contact with the western academic world and most musicological studies have been published only in Mandarin and available only in China. Thus, collaboration with researchers and experts immersed in that culture is even more important.

Apart from the issues that can be studied on each music repertoire, it is very relevant to work on the problems that arise from comparative studies. Typically, the available comparative studies focus on the comparison between any music cultures and western classical music. But it is important to also make comparative studies between the different non-western musical cultures. For example, it is interesting to compare the Indian concept of raga with the Ottoman concept of makam, or the improvisatory strategies used in Indian music with the ones used also in Ottoman music. In fact, all the musical cultures selected in this project were part of what is known as the Silk Road. This was a network of trade routes across the Asian continent that connected East, South, and Western Asia with the Mediterranean world, as well as North, East, and Northeast Africa and Europe, for almost 3,000 years. This trade connection had a great impact on the development of all these musical cultures and it is thus fascinating to understand these mutual influences. In this project, we will promote a tight collaboration between the researchers working in the different countries and on the different culture-specific problems in order to understand these cross-cultural issues.

3. PROJECT TASKS

The culture-specific problems on which we will start working are not yet defined; in fact, this is now the main work in progress, together with putting together the initial research team. So far we have defined the main general and transversal tasks and research approaches that we will use to tackle our ambitious aims. These tasks are: (1) to gather and organize audio recordings, metadata, descriptions, scores, plus all the needed contextual information of the selected music repertories; (2) to identify and study the required musicological references in order to understand the chosen repertories within their cultural context; (3) to design the ontologies needed to annotate and analyze the gathered music collections; (4) to work on audio content analysis approaches to help describe the music collections chosen; (5) to work on a social-computing approach to characterize users and communities, modeling their musical preferences and behaviors; and (6) to develop systems that, by integrating the results of this project, can show the relevance of this research approach for the discovery of our world's music.

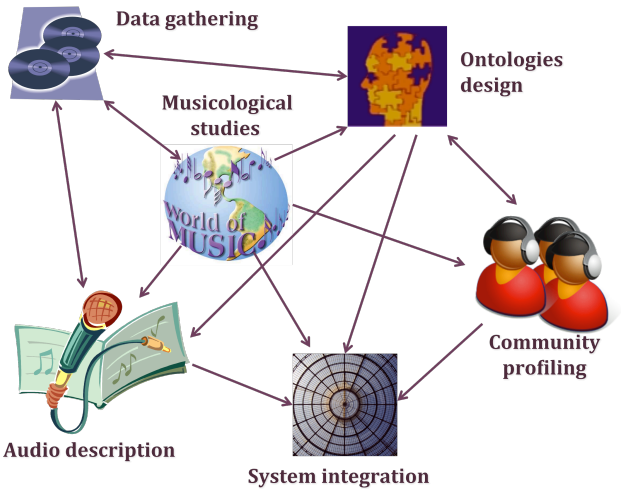

Figure1. General project tasks.

3.1 Data gathering

A major issue in the gathering and organization of the relevant information for this project is the data formats used. Most databases and protocols that we have used in previous information processing projects were designed for housing western pop music. As a result, important information would be lost if we just used the same schemas to organize the types of music that this project works on. We will convert our current framework into a more freely structured database, based on RDF[1] (Resource Description Framework) and graph database technologies. This will enhance the storage system's support for heterogeneous data and allow for the flexible addition of new data as it becomes available. We also want to contribute and take advantage of the existing online data repositories that use open licenses, like Wikipedia (general encyclopedia), MusicBrainz (CDs metadata), Mendeley (bibliographical references), or Freesound (audio clips)[2]. The goal is to be active in the Open Data Movement[3] promoting it within our research community.

We will also have to make some recordings and carry out fieldwork in order to gather specific data with which to tackle some of the identified research problems. All this data will be made available through the existing open repositories, or, if not appropriate, it will be placed on this project's website, also using open licenses. The CompMusic website[4] is used both as this project's dissemination portal and as a collaborative framework in which anyone interested in the objectives of this project can get involved and contribute information.

3.2 Musicological studies

There is no lack of musicological research on the chosen musical cultures, e.g. [4], and we need to study it to understand the musical specificities of each repertoire, the similarities and differences between the selected music cultures, and what can distinguish them from classical western culture. Each classical music tradition has developed its own instruments, compositional styles, performance practices, social uses, and contexts. However, we find very different points of view, sometimes opposing ones, in existing literature, and we will be collaborating with musicologist from the different music cultures in order to understand those points of view and to be able to develop appropriate musicological perspectives for our identified problems.

An important part of the traditional musicological research of western music is based on the study of symbolic representations, scores, and given that our chosen repertoires are mainly based on oral traditions, their existing symbolic representations have been studied much less. There is some work on the extension and use of digital representations like the Humdrum toolkit[5] or MusicXML[6] to analyze non-western music notations that we want to continue. This type of analysis is very complementary to that which is audio-based and can be used to study some compositional and semantic aspects of music.

The development of a musicological framework is fundamental to properly contextualizing most of our research work. In particular, to identify the relevant problems to be studied, to agree on the terminology to use for the different culture-specific musical concepts, to balance the theoretical with the practical perspectives that will be involved in the different problems, and also to interpret the computational results obtained.

3.3 Ontologies design

To process and share information in open environments we need a common understanding of the meaning of the data, thus we need to specify how concepts and terms of a given domain are understood, specifications that are known as ontologies. For example, Music Ontology[7] [5] is being developed by an open community of experts as a formal framework for dealing with music-related information on the Semantic Web, including editorial, cultural and acoustic information. It was started with the goal of capturing the musical production process, but we propose to extend it to express the variety of concepts that will be indentified in the study of the selected musical cultures. We will also explore the approaches coming from what is known as emergent semantics [6] in order to develop ontologies from our gathered data repositories and from mining music-specific

[1] http://www.w3.org/RDF

[2] http://wikipedia.org, http://musicbrainz.org, http://mendeley.com, http://freesound.org

[3] http://en.wikipedia.org/wiki/Open_Data

[4] http://compmusic.upf.edu

[5] http://musicog.ohio-state.edu/Humdrum

[6] http://www.recordare.com/musicxml

[7] http://musicontology.com

websites (e.g., music magazines, deep musical discussion forums, or reputed blogs about music).

As soon as we take a culture-specific approach in the study of music, the issues of terminology and language become very important. We will have to develop multi-language ontologies from which to capture the musical concepts used by the different cultures, a task that will be done by musical experts. But we will have to deal with the diversity of meanings that a single term may have within a single tradition or the fact that different terms can signify the same or very similar meanings within that same culture.

3.4 Audio description

The extraction of musically and culturally meaningful features from audio recordings of different music repertoires requires introducing new approaches, and for sure, extending the existing signal analysis techniques and machine learning methods. The fact that the music traditions chosen in this project have fundamental differences from western music traditions, such as different musical instruments, tuning systems, performance styles, or musical forms, imply that at the level of feature analysis, most of the descriptors and extraction methodologies being used to analyze commercial western music are not appropriate, or at least they have to be developed much further. Even at the level of acoustics there are very interesting and important differences between western musical instruments commonly used and the instruments used in other classical traditions. This basic consideration has profound implications at all musical levels.

Both the fact that most musical audio features are interrelated and that their musical meaning is very much context dependent will have to be emphasized in our project. The recent approaches in audio analysis based on joint estimation of audio features [7] are especially relevant to finding more perceptually meaningful music characterizations. Also the current trend towards knowledge-based and top-down signal processing [8] is very relevant in our case since many of the music features of a given music repertoire cannot even be analyzed without using information from the cultural tradition they are part of.

Within most cultural traditions, music performance is learned by imitation and music appreciation is developed implicitly from listening. One of the consequences of this is that most people, even the experts, have a very difficult time verbalizing their musical opinions, but they are able to express those opinions through examples. Based on this, our audio description work will use informed audio references supplied by experts as our main data sources, then we will focus on extracting musically meaningful features from them and on developing similarity measures between them.

3.5 Community profiling

The understanding and characterization of the musical preferences and behaviors of groups of people sharing a given musical culture is a new and complex problem with very little prior work. Theoretically, this profiling could be developed explicitly by asking music lovers and experts of a given community, but in this project we want to focus on the implicit characterization of online communities, that is, characterizing groups of people that leave a trace on the Web. We would like to capture the musical behavior of a given online user group by exploiting information sources such as web search statistics, musical preferences extracted from social networks, or data inferred from demographical, geographical, and psychographical data (personality, values, attitudes, and interests). The actual musical profile of a group should be expressed through the audio descriptors of the music they interact with, ontologies generated by them, and from the automatically extracted information about their cultural context and behaviors. Provided that we have access to the appropriate data, with this type of information we should be able to characterize any music community active online.

A good starting point for this work is the current research within the field of social computing [9], concerned with the intersection of social behavior and computational systems. Especially relevant is the research on the analysis, modeling and knowledge extraction from communities of users. The work following the ideas of emergent semantics proposes the analysis of online communities for the extraction of ontologies that structure the concepts and terms of a particular domain [6]. Peter Mika [10] proposes a unified model of social networks and semantics where social tagging systems can be modeled as a tripartite graph with Actors, Concepts and Instances (Actor-Concept-Instance model). By analyzing the relations between concepts both on the basis of co-occurrence in instances and common usage by actors (users), lightweight ontologies can emerge from online communities. A completely different approach to community knowledge extraction for the design of ontologies is proposed in [11], where a Web portal with collaborative ontology management capabilities is implemented. However, the field is very young and there is still a conceptual gap between this vision and its possible implementation; only very simple ontologies can emerge.

A computational approach to community profiling requires having access to data from relevant online communities. In our project we are starting with Freesound, an open collaborative sound database developed by our research group where people with different interests share audio clips. Given its characteristics and our involvement in its development, it offers a good platform to carry out studies on social networking, specifically on community profiling, and to explore technological solutions to promote specific musical activities, which is also an important aspect of our project. There is a lot to be done in understanding the general issues of profiling of online communities. Only after we get a handle on the basic issues will we be able to face the specific profiling problems of the music communities of relevance to the CompMusic project. At this time we have

not yet identified online communities that would be appropriate for our research objectives, but given the fast growth of online communities we should be able to find relevant music sites to study our culture-specific issues.

3.6 System integration

Quite a number of musical applications could take advantage of the expected research results of this project. We can focus on application areas such as music education, music creation, music appreciation, music recommendation, etc., or any particular application that might benefit from an information processing engine able to process our world's music, respecting its culture specificity and that of the users. In this project we want to exemplify the benefits of our multicultural approach to music computing by either developing new music systems or by extending the functionality of existing ones—functionalities related to automatic music analysis, categorization, learning or discovery.

The development of a culture-driven system should be much more than the addition of a set of components; our system should be interactive and evolve with the users, with the context, and with the availability of new information. Even the concept of interface design becomes a critical issue when approached from a multicultural perspective. Can we have a single interface for a system to be used by different cultures? We think not. The interface has to adapt to the users' cognitive/cultural structures and has to engage them using the values and attributes of their own culture. All these aspects require specific research before even considering any system development.

Within CompMusic there is no particular aim of developing complete systems, but we want to put together prototypes with which to demonstrate our research results. However since this task will be worked on during the last part of this project, at this stage it does not make much sense to continue to describe the details or even the type of system that will be developed.

4. CONCLUSIONS

CompMusic is a big and ambitious project that aims at having an impact not just within the Music Computing field but also, more generally, on the overall field of information technology. With funding from the European Research Council we will be able to support the research work of quite a number of PhDs and post-doctorates from different parts of the world covering all areas described in this article. In addition, we also aim to bring in a disruptive point of view in IT by promoting, and showing the validity of, a new research approach rooted in a multicultural perspective.

5. ACKNOWLEDGEMENTS

I want to thank the organizers of WISSAP-2010 that took place in IIT-Bombay in January 2010 for inviting me to give a series of tutorial lectures on Music Computing, a trip that gave me the opportunity to realize the need for this project. I would also like to thank the many musicians and researchers, working outside the well-established academic western institutions, from which I have learned alternative ways to approach and understand music. Finally I have to acknowledge the contributions of many researchers from the MTG-UPF in the definition of this project.

The CompMusic project has received funding from the European Research Council under the European Union's Seventh Framework Program (FP7/2007-2013) / ERC grant agreement 267583 and will run for five years starting on July 2011.

6. REFERENCES

[1] G. Tzanetakis et al.: "Computational Ethnomusicology," *Journal of Interdisciplinary Music Studies*, 1(2), pp. 1-24, 2007.

[2] B. Nettl: *The Western Impact on World Music*. Schirmer Books, 1985.

[3] D. Huron: "Issues and Prospects in Studying Cognitive Cultural Diversity," *Proc. of the 8th International Conference on Music Perception & Cognition*, pp. 83-96, 2004.

[4] Routledge (publisher): *The Garland Encyclopedia of World Music* (10 volumes), 1997-2001.

[5] Y. Raimond: *A Distributed Music Information System*, PhD Thesis, 2008.

[6] K. Aberer, et al.: "Emergent Semantics Principles and Issues," *Database Systems for Advanced Applications*, pp. 25-38, 2004.

[7] H. Papadopoulos and G. Peeters: "Joint Estimation of Chords and Downbeats from an Audio Signal," *IEEE Transactions on Audio, Speech, and Language Processing*, 19(1), pp. 138-152, 2011.

[8] A. Klapuri and M. Davy (Eds.): *Signal Processing Methods for Music Transcription*, Springer, 2006.

[9] S. Chai et al. (Eds.): *Advances in Social Computing*, Springer, 2010.

[10] P. Mika: "Ontologies are us: A Unified Model of Social Networks and Semantics," *Journal of Web Semantics*, 5(1), pp. 5-15, 2007.

[11] A. V. Zhdanova: "Community-driven Ontology Construction in Social Networking Portals," *Web Intelligence and Agent Systems: An International Journal*, vol. 6, pp. 93-121, 2008.