ASSOCIATIONS BETWEEN MUSICOLOGY AND MUSIC INFORMATION RETRIEVAL

Kerstin Neubarth
Canterbury Christ Church University
Canterbury, UK
kerstin.neubarth@canterbury.ac.uk

Mathieu Bergeron
CIRMMT
McGill University
Montreal, Canada

Darrell Conklin
Universidad del País Vasco
San Sebastián, Spain
and IKERBASQUE,
Basque Foundation for Science

ABSTRACT

A higher level of interdisciplinary collaboration between music information retrieval (MIR) and musicology has been proposed both in terms of MIR tools for musicology, and musicological motivation and interpretation of MIR research. Applying association mining and content citation analysis methods to musicology references in ISMIR papers, this paper explores which musicological subject areas are of interest to MIR, whether references to specific musicology areas are significantly over-represented in specific MIR areas, and precisely why musicology is cited in MIR.

1. INTRODUCTION

At the tenth anniversary ISMIR 2009 several contributions discussed challenges in the further development of music information retrieval (MIR) as a discipline, including requests for deeper musical motivation and interpretation of MIR questions and results, and envisaging closer interaction with source disciplines such as computer science, cognitive science and musicology [4, 9, 20].

Suggestions for interdisciplinary collaboration have often considered musicology as a *target* discipline, emphasising the usefulness of MIR tools to musicology (e.g. [18]). Occasionally mutual benefits have been explored (e.g. [15]). The current study addresses associations between MIR and musicology as a *source* discipline. It presents a systematic analysis of how MIR, as represented at ISMIR, has drawn on musicology so far, by applying data mining and content citation analysis to musicology references in ISMIR publications.

Related quantitative ISMIR surveys mainly analyse topics and trends [3, 7, 10]. The study by Lee et al. [10] performed a citation analysis of papers within ISMIR, counting references to individual authors and papers. It briefly addressed citer motivations such as identification of data and methods or paying homage, and concluded that: "Without a more in-depth analysis of the individual contexts surrounding each citation, it is difficult to tease out the precise motivations for all the references" (p. 61). Functions of references to one particular study were discussed in the editorial to the JNMR Special Issue on MIR in 2008 [1].

This study extends previous work in several ways: It analyses inter-disciplinary references (musicology cited in MIR) rather than intra-MIR references; also, the references are analysed at the level of MIR and musicology subject categories instead of individual papers. The quantitative analysis goes beyond citation counts; *association mining* is used here to yield interdisciplinary associations. In addition, citation contexts are analysed in depth to reveal functions of musicology citations in ISMIR papers, taking into account both MIR-specific functions and more general referencing purposes to allow comparison with existing studies.

2. DATA SELECTION AND ANALYSIS

This section presents the corpus development and the association mining and citation analysis methods used for extracting and analysing associations between musicology and music information retrieval.

2.1 Sampling

From the cumulative ISMIR proceedings (www.ismir.net) as a sampling frame, first all available full papers from 2000 until 2007, and all oral/plenary session papers for 2008 to 2010, were selected. This resulted in 416 papers. Then the reference lists of those papers were screened and a purposive sample was taken of all papers which contain references to musicology as a source discipline, excluding self-citations and references to other ISMIR papers. The final analysis corpus consisted of 184 papers. These papers are identified by their IDs within the cumulative proceedings (e.g. ID135).

2.2 Encoding

The 184 papers of the analysis corpus were labelled according to their MIR research topic and the musicology areas that they cite, using the following categorisations.

MIR Categories. In a first step of encoding, the 184 papers were classified into MIR research areas (Table 1). As no single standardised and comprehensive taxonomy of MIR topics exists [3,6], an organisation of topics was developed based on ISMIR calls and programs, harmonising categories across conferences. Each ISMIR paper is assigned to exactly one MIR category. Numbers in brackets in Table 1 indicate the number of papers in each category.

Research paradigms (6)
Epistemology, interdisciplinarity
Representation & metrics (24)
Representation, metrics, similarity
Data & metadata (11)
Databases, data collection & organisation, metadata, annotation
Transcription (42)
Segmentation, voice & source separation, alignment, beat tracking & tempo estimation, key estimation, pitch tracking & spelling
OMR (5)
Optical music recognition, optical lyrics extraction
Computational music analysis (21)
Pattern discovery & extraction, summarisation, chord labelling, musical analysis (melody & bassline, harmonic, rhythm and form analysis)
Retrieval (19)
Query-by-example
Classification (32)
Genre classification, geographical classification, artist classification & performer identification, instrument-voice classification (instrument recognition, instrument vs voice distinction, classification of vocal textures), mood & emotion classification
Recommendation (5)
Recommendation methods & systems, playlist generation, recommendation contexts
Music generation (4)
Music prediction, improvisation, interactive instruments
Software systems (10)
Prototypes & toolboxes, user interfaces & usability, visualisation
User studies (5)
User behaviour (music discovery, collection organisation)

Table 1. Thematic categories and examples of topics of MIR research.

History, criticism & philosophy
History of music (8)
Philosophy of music & music semiotics (3)
Textual criticism, archival research & bibliography (22)
Electronic & computer music (7)
Popular & jazz music studies (5)
Film music studies (1)
Theory & analysis
Music theory & analysis (36)
Performance studies (6)
Ethnomusicology
Ethnomusicology (non-Western) (5)
Ethnomusicology (folk music) (9)
Ethnomusicology (other) (1)
Systematic Musicology
Acoustics (11)
Psychology of music (perception & cognition) (93)
Psychology of music (emotion & affect) (8)
Psychology of music (other) (1)
Sociology & sociopsychology (18)

Table 2. Thematic categories of musicology.

Musicology Categories. In a second step, the musicology references in the 184 papers were assigned to musicology areas (Table 2). As the interest of this study is in musicology as a source discipline, the labels used are based on traditional subject organisations (e.g. [11, 14, 16]) rather than more recent developments such as empirical or computational musicology which potentially overlap with music information retrieval (e.g. [18]). Category counts in Table 2 refer to the number of ISMIR papers citing this musicology area one or more times. A paper may reference more than one musicology category.

2.3 Association Mining

Data mining of the analysis corpus is used to reveal associations between papers in specific MIR categories and papers that have citations to specific musicology categories. For every musicology category A and MIR category B, the *support* (number of papers) $s(A)$ and $s(B)$ were computed. Also the support $s(A, B)$ of an association $\langle A, B\rangle$ (number of papers containing references to musicology category A which are also in MIR category B) and the statistical significance of the association were computed.

The null hypothesis is that for an association $\langle A, B\rangle$ the two categories are statistically independent, i.e. that the proportion of papers citing musicology category A that are in MIR category B does not differ significantly from the relative frequency of MIR category B in the general population. Given the small corpus and low counts for many cat-

egories, the appropriate test for statistical independence is Fisher's one-tailed exact test on a 2x2 contingency table [5]. For an association $\langle A, B\rangle$ with support $s(A, B)$, this gives the probability (p-value) of finding $s(A, B)$ or more papers of category B in $s(A)$ samples (without replacement) in $n = 184$ total papers. If the computed p-value is less than the significance level $\alpha = 0.05$, then we reject the null hypothesis that the categories are independent.

Prior to computing the p-values, the counts of all MIR categories B (and hence the p-values of associations) are slightly adjusted upwards to account for the fact that only papers citing musicology were included in the sample of $n = 184$ papers from the larger corpus of 416 papers. Under the null hypothesis of independence, it is assumed that the larger set of papers has the same distribution of MIR categories as the smaller corpus. The adjustment is done by increasing n to 416 and $s(B)$ to $416 \times s(B)/184$.

In line with the view of musicology as a source discipline, significant associations were oriented into rules from musicology to MIR categories. For every significant association $\langle A, B\rangle$, the *confidence* of the oriented rule $A \rightarrow B$ was computed as $s(A, B)/s(A)$, indicating the empirical probability of a paper being in MIR category B given that it cites a paper in musicology category A.

2.4 Content Citation Analysis of Referencing Functions

Papers supporting significant associations were analysed in more detail to reveal functions of musicology references. Studies of citation behaviour have proposed several classifications of citer motivation (e.g. [2, 12, 13, 17]). In our analysis, we are mainly interested in (a) the function of the reference in the citing paper rather than conclusions about the cited work, and (b) referencing purposes that can be suggested from the content of the citing paper and the co-text of the citation. Musicology references were analysed in their context in the ISMIR paper, and recurring referencing functions extracted and linked to existing citation classifications.

3. RESULTS

This section presents the associations and referencing functions uncovered in the corpus.

3.1 Associations

Figure 1 presents the network of all associations with support ≥ 3 extracted by the association mining method described in Section 2.3. The figure highlights that MIR areas generally draw on more than one musicology discipline. But there are differences in the level of co-citation, i.e. occurrences within the same ISMIR papers: For example, eight out of the ten papers on representation and metrics which cite music theory and analysis literature also cite psychological work on perception and cognition. On the other hand, perception and cognition research and acoustics are cited by different subsets of papers on transcription.

Comparing Figure 1 against MIR topics in Table 1, additional links could have been expected e.g. between papers on data and metadata and references to textual criticism, archival research and bibliography or history of music [15] or between classification and performance studies (for performer identification), acoustics (for instrument-voice classification) and popular music studies or history of music (for genre classification). However, these relations are supported by only one or two papers each and thus do not appear in Figure 1. Surprisingly, the category of ethnomusicology (folk music) (Table 2) does not feature in associations with MIR categories above the support threshold.

Of the 21 associations shown in Figure 1, nine are statistically significant (Section 2.3). Table 3 enumerates those associations that have a p-value less than $\alpha = 0.05$. In Figure 1 these particular associations are shown in bold lines, with a directed arrow indicating in brackets the confidence of the oriented rule. Overall the relatively low confidence values confirm that in general musicology areas are cited across MIR categories.

Generally an association will be significant if the association support $s(A, B)$ is high relative to the size of one involved category. Here this applies in particular for small categories like ethnomusicology (non-Western), OMR, user studies or recommendation. Significance becomes harder to achieve for associations between large categories; it is more likely to achieve the observed level of support at random given the individual category distributions in the corpus. For example, perception and cognition research is linked to several MIR categories with high support, but the distribution of those MIR categories across the 93 papers citing perception and cognition does not differ significantly from their distribution across all sampled papers.

3.2 Referencing Functions

For the content citation analysis we selected the ISMIR papers supporting the associations in Table 3, as these papers are examples of musicology and MIR categories that are significantly correlated. Of these 47 papers, 17 papers (5 computational music analysis papers, 8 representation and metrics papers and 4 retrieval papers) also cite perception and cognition research; these references were also considered. The in-depth analysis of citation contexts in these papers demonstrates that musicology is used for a variety of purposes. Figure 2 presents a taxonomy of referencing functions and the references' contribution in the MIR work (boxes), with examples of co-text. Related features from the citation analysis literature [2, 12, 13, 17] are included in italics.

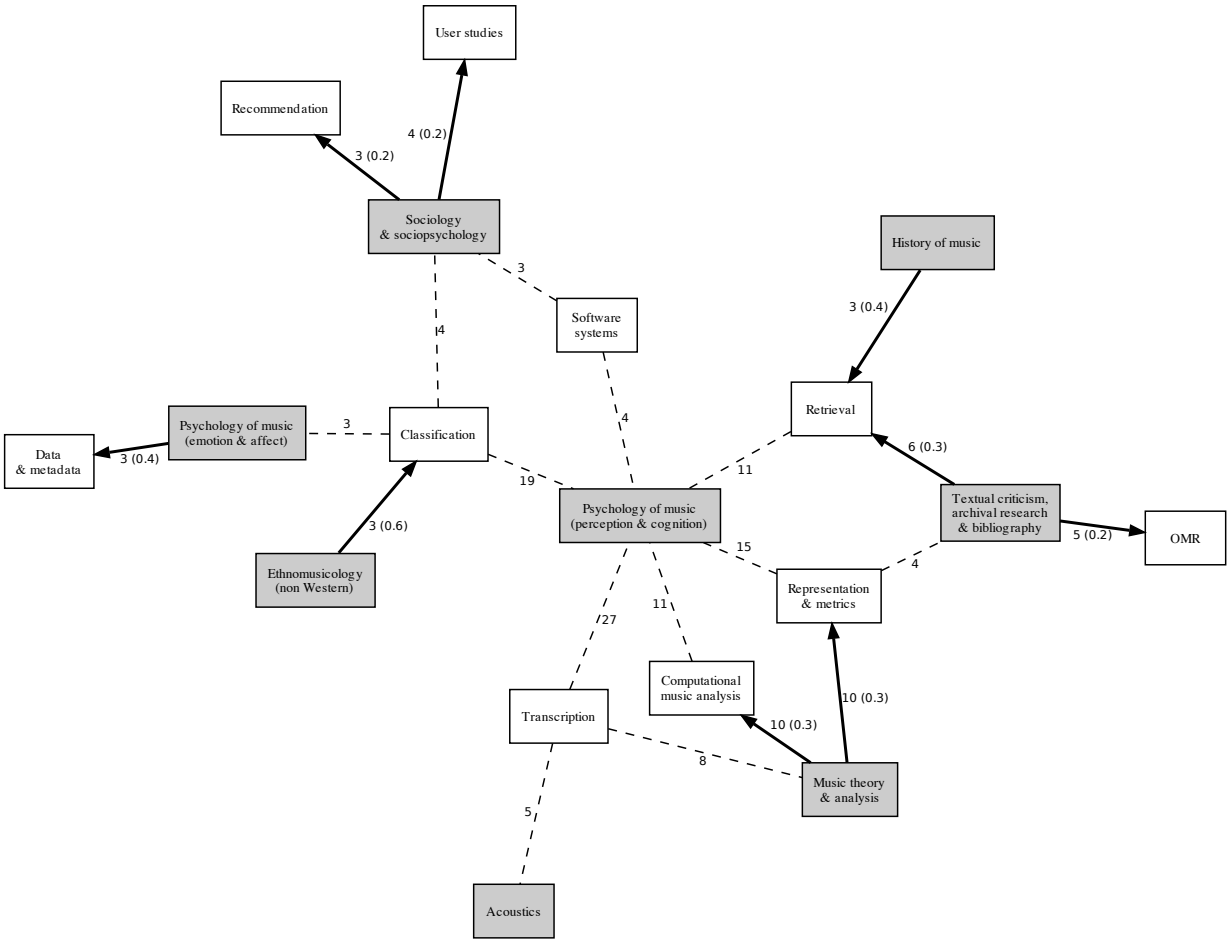

Figure 1. Associations (support ≥ 3) between musicology categories (dark boxes) and MIR categories. Edges are labelled with the support of the association, and significant ($\alpha = 0.05$) associations are indicated with dark oriented edges. Rule confidence is indicated in brackets for significant associations.

A	B	$s(A, B)$	p-value
textual criticism, archival research & bibliography	omr	5	9.7e-05
sociology & sociopsychology	user studies	4	0.00068
music theory & analysis	computational music analysis	10	0.0035
psychology of music (emotion & affect)	data & metadata	3	0.0078
sociology & sociopsychology	recommendation	3	0.0091
music theory & analysis	representation & metrics	10	0.01
textual criticism, archival research & bibliography	retrieval	6	0.016
history of music	retrieval	3	0.037
ethnomusicology (non western)	classification	3	0.038

Table 3. Significant ($\alpha = 0.05$) associations found in the corpus. A: musicology category; B: MIR category; $s(A, B)$: support of the association; p-value of the association.

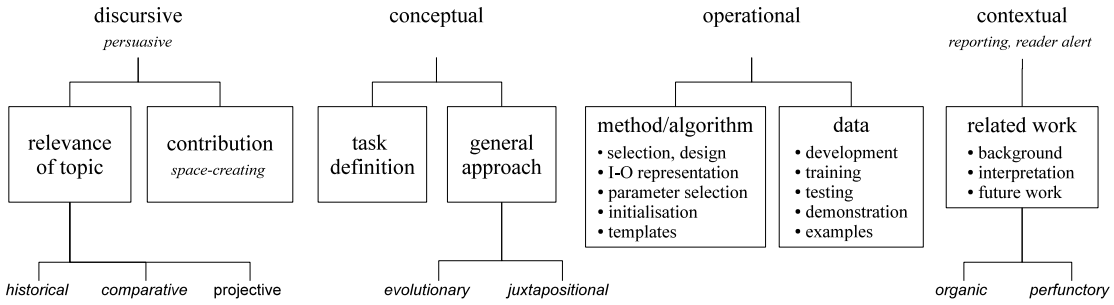

Function	Co-text examples
Relevance	"Repeated patterns [...] represent therefore one of the most salient characteristics of musical works [music theory references]" (ID242)
Contribution	"This paper addresses systematic differences in the performance of final ritardandi by different pianists [...] the kinetic model is arguably too simple [...] In this work [...] [psychology of music references]" (ID159)
Task definition	"As stated in [music theory reference] musical analysis is 'the resolution of a musical structure into relatively simpler constituent elements, and the investigation of the functions of these elements within that structure"' (ID24)
General approach	"The [basic] idea is motivated by the results of musicological studies, such as [...]" (ID859)
Method/algorithm	"HMM initialization [...] The covariance matrix should also reflect our musical knowledge [...], gained both from music theory as well as empirical evidence [psychology of music reference]" (ID30)
Data	"we evaluated both [OMR] systems on the same set of pages to measure their accuracy [...] [textual criticism, archival research & bibliography reference]" (ID729)
Related work	"dimensions of dissimilarity have been interpreted to be e.g. [...] [psychology of music reference]" (ID345)

Figure 2. Taxonomy of referencing functions (top) and selected examples of co-text (bottom).

4. DISCUSSION AND CONCLUSIONS

The findings presented in this study are based on direct and explicit references to musicology. However, not all references are explicit: papers sometimes refer to musicological work reported in earlier MIR publications; incorporate concepts or approaches like music-analytical methods into the main text without including specific references; characterise the considered repertoire such as non-Western traditions without making explicit whether the description is derived from musicological research, common cultural knowledge or the researchers' personal experience; or use music examples without citing a musicological source. Taking into account such references is expected to strengthen rather than change the picture of associations presented here.

For this study the analysis corpus only contained full ISMIR papers (Section 2.1). Future work could extend the corpus to also include posters, in particular those from ISMIR 2008 onward (because these are of equal length and status to full papers); apply multilevel association mining methods [8] to hierarchical subject categorisations; allow a paper to be within multiple MIR categories; and evaluate whether the associations found here persist and whether new significant associations arise. Furthermore, if an encoding of the complete ISMIR proceedings was available, other interesting types of analysis would be possible, e.g. exploring whether certain MIR areas are over- or under-represented in the corpus of papers citing musicology, or comparing use of musicology references against other source disciplines like computer science or cognitive science.

Several observations can be drawn from the results presented in this paper. First, less than half of the full papers in the cumulative ISMIR proceedings (184 out of 416) cite musicology. Given the close interdisciplinary links between MIR and musicology a larger percentage had been expected. Second, the most frequently cited category is music perception and cognition research (93 citing papers across all MIR categories in our corpus). On the other hand, historical musicology and especially history of music appear to be under-represented in our corpus, compared to their traditional weight in musicology [16]. Third, the association mining has revealed significant associations between certain musicology and MIR categories. However, most pairings are not significant, and this may indicate opportunities for category refinement and for specific interdisciplinary collaboration. Fourth, the content citation analysis yields a range of citation purposes, from justifying the MIR topic and specifying the MIR task, through informing methods or providing data, to references which demonstrate awareness of the research context but remain without direct implications for the MIR work.

Following the discussions at ISMIR 2009 [4, 9, 20], in the further development of MIR we would expect that with the increasing interest in ethnic music (e.g. [3, 19]) ethnomusicology will more strongly feed not only into classification but also MIR areas such as representation and metrics, transcription or retrieval; envisage more musicology references, including history of music, in defining MIR research questions and in interpreting MIR results; and encourage more projective references highlighting potential of MIR achievements for musicology, beyond providing technological tools. The association mining and content analysis methods applied in this paper will be invaluable to study the continuing evolution of the field of music information retrieval.

5. REFERENCES

[1] J.-J. Aucouturier and E. Pampalk. Introduction – From genres to tags: A little epistemology of music information retrieval research. *Journal of New Music Research*, 37(2):87–92, 2008.

[2] T. A. Brooks. Private acts and public objects: An investigation of citer motivations. *Journal of the American Society of Information Science*, 36(4):223–229, 1985.

[3] O. Cornelis, M. Lesaffre, D. Moelants, and M. Leman. Access to ethnic music: Advances and perspectives in content-based music information retrieval. *Signal Processing*, 90:1008–1031, 2010.

[4] J. Downie, D. Byrd, and T. Crawford. Ten years of ISMIR: Reflections on challenges and opportunities. In *International Society for Music Information Retrieval Conference*, pages 13–18, 2009.

[5] S. Falcon and R. Gentleman. Hypergeometric testing used for gene set enrichment analysis. In F. Hahne, W. Huber, R. Gentleman, and S. Falcon, editors, *Bioconductor Case Studies*, pages 207–220. Springer, 2008.

[6] J. Futrelle and J. Downie. Interdisciplinary research issues in music information retrieval: ISMIR 2000–2002. *Journal of New Music Research*, 32(2):121–131, 2003.

[7] M. Grachten, M. Schedl, T. Pohle, and G. Widmer. The ISMIR cloud: a decade of ISMIR conferences at your fingertips. In *International Society for Music Information Retrieval Conference*, pages 63–68, 2009.

[8] J. Han and Y. Fu. Discovery of multiple-level association rules from large databases. In *International Conference on Very Large Data Bases*, pages 420–431, 1995.

[9] P. Herrera, J. Serrà, C. Laurier, E. Guaus, E. Gómez, and X. Serra. The discipline formerly known as MIR. In *International Society for Music Information Retrieval Conference, fMIR workshop*, 2009.

[10] J. Lee, M. Jones, and J. Downie. An analysis of ISMIR proceedings: Patterns of authorship, topic, and citation. In *International Society for Music Information Retrieval Conference*, pages 57–62, 2009.

[11] Library of Congress. *Library of Congress Classification Outline: Class M – Music*. http://www.loc.gov/catdir/cpso/lcco/.

[12] M. Liu. Citation functions and related determinants: A study of Chinese physics publications. *Journal of Library and Information Science*, 19(1):1–13, 1993.

[13] M. Moravcsik and P. Murugesan. Some results on the function and quality of citations. *Social Studies of Science*, 5(1):86–92, 1975.

[14] R. Parncutt. Systematic musicology and the history and future of western musical scholarship. *Journal of Interdisciplinary Music Studies*, 1(1):1–32, 2007.

[15] J. Riley and C. Mayer. Ask a librarian: The role of librarians in the music information retrieval community. In *Proceedings of the International Conference on Music Information Retrieval*, 2006. w/o pages.

[16] S. Sadie, editor. *The New Grove Dictionary of Music and Musicians*, chapter "Ethnomusicology", "Music Analysis", "Musicology", "Psychology of Music", and "Theory, Theorists. 15: New Theoretical Paradigms, 1980-2000". Macmillan, London, 2001.

[17] J. Swales. Citation analysis and discourse analysis. *Applied Linguistics*, 7(1):39–56, 1986.

[18] G. Tzanetakis, A. Kapur, W. Schloss, and M. Wright. Computational ethnomusicology. *Journal of Interdisciplinary Music Studies*, 1(2):1–24, 2007.

[19] P. van Kranenburg, J. Garbers, A. Volk, F. Wiering, L. Grijp, and R. C. Veltkamp. Collaborative perspectives for folk song research and music information retrieval: The indispensable role of computational musicology. *Journal of Interdisciplinary Music Studies*, 4(1):17–43, 2010.

[20] F. Wiering. Meaningful music retrieval. In *International Society for Music Information Retrieval Conference, fMIR workshop*, 2009.