INFLUENCES OF ISMIR AND MIREX RESEARCH ON TECHNOLOGY PATENTS

Sally Jo Cunningham
Department of Computer Science
University of Waikato
sallyjo@waikato.ac.nz

Jin Ha Lee
Information School
University of Washington
jinhalee@uw.edu

ABSTRACT

Much of the current Music Information Retrieval (MIR) research aims to contribute to the field by creating practical music applications or algorithms that can be used as part of such applications. Understanding how academic research results influence and translate to commercial products can be useful for MIR researchers, especially when we try to measure the impact of our research. This study aims to improve our understanding of the commercial influence of academic MIR research by analyzing the patents citing publications from ISMIR (International Society for Music Information Retrieval) Conference proceedings and its associated MIREX (Music Information Retrieval Evaluation eXchange) MIR algorithm trials. In this paper, we provide our preliminary analyses of the relevant patents as well as the ISMIR publications that are referenced in those patents.

1. INTRODUCTION

The ISMIR (International Society for Music Information Retrieval) conference started as a small-scale symposium in 2000 and continued to grow over the past decade as the field of Music Information Retrieval (MIR) matured. ISMIR has been one of the most important MIR conferences since the early establishment of the field, and serves as a key venue for dissemination of MIR research. The associated MIREX (Music Information Retrieval Evaluation eXchange) event, first run in 2005, has similarly grown to be a focus for MIR system and algorithm evaluation. One of the key objectives of MIR research is to make practical contributions toward the development of commercial music applications and services to improve users' interaction and experience with music. There is anecdotal evidence that our research results inform the development of new music applications and services, especially since ISMIR conferences and MIREX trials have continued to attract participants from the commercial sector. However, to date no research has been conducted to systematically investigate the extent of the practical impact of academic MIR research published through the ISMIR conferences and MIREX trials.

To this end, we first identify patents that reference publications from ISMIR and MIREX (Section 3). We then perform an informetric analysis over these patents and the referenced publications drawn from ISMIR and MIREX, to discover patterns of influence of ISMIR and MIREX on patented MIR technology (Section 4).

2. PREVIOUS WORK

Previous investigations of the characteristics of MIR research have focused exclusively on the field as viewed through academic publication. The research methods used were primarily bibliometric—that is, quantitative measures such as citation analysis, based on data drawn from the metadata and text of the ISMIR and MIREX proceedings. These techniques have been used to paint rich pictures of the state of ISMIR academic research at various stages in the history of ISMIR and MIREX [1], [2], [3], with the emphasis on scholarly publishing.

This present paper applies these bibliometric techniques to a set of patent filings rather than academic papers. 'Patent bibliometrics,' the natural extension of bibliometric techniques to collections of patent metadata and texts, has seen widespread use in technology-related fields since its introduction in 1994 [8]. The introduction of online, free-to-search patent databases has further encouraged patent bibliometric investigations [6]; the most comprehensive and widely used databases are provided by the Unites States Patent and Trademark Office (USPTO) and the European Patent Organisation (EPO).

One natural topic of interest has been the relationship between academic publications and patents in a given field. Data regarding this relationship should be straightforward to draw from a patent database, as each patent filing includes the equivalent of the bibliographic citations in the form of references to prior art, and the prior art can include both earlier patents and relevant conference and journal papers.

Unfortunately, patent databases index only the 'front page' prior art citations, and these are almost exclusively limited to patents. References to prior art in the form of relevant academic publications are typically found in the body of the patent—which is not indexed by the databases—and, though they may appear on the patent's front page as "Other Publications," they are not indexed in the USPTO and EPO patent search engines [7].

Given these hurdles to identifying linkages between patents and academic research directly from the patent databases, researchers have had to draw in evidence from additional sources, or to use crude proxy measures for the linkages. For example, a study of academics in Norway compared the publishing behavior of matched sets of academic patent inventors and non-inventors—but since the databases often did not include inventor affiliations, that had to be determined through expensive, error-prone, and time-consuming surveys of the institutions themselves [5]. Meyer [6] estimates the interactions between a country's academic and patented outputs through the coarse mechanism of comparing patent and publication rates for that country in a small number of narrowly focused fields.

The introduction of search over the full text of a patent by Google Patent Search (GPS) [1] supports patent bibliometric investigations to a far greater degree than has been possible previously, in particular supporting citation analysis of references to published scientific literature in patents. To our knowledge, this present paper is the first to use this facility to directly explore the influence of academic research on the patents.

3. DATA COLLECTION

We used Google Patent Search in order to find all the patents containing references to ISMIR and MIREX. Google Patent Search is a specialized Google search engine that indexes patents and patent applications drawn from public domain patent databases of the USPTO and the EPO. While both the USPTO and EPO offer search facilities for their individual collections, the Google Patent Search presents a unified interface to both databases. More significantly for this present investigation, Google Patent Search indexes the entire patent for search, where the USPTO and EPO support search only over the basic patent metadata (title, classification code, publication date, inventor, etc.). Through GPS, we can link back to the scientific literature supporting a patent, via the references cited in the textual description of the background to the invention and of the invention itself.

Searches for the initial patent datasets were conducted over Google Patent Search in April 2013, using the terms 'ISMIR music' and 'MIREX music'. Where multiple patent filings under the same inventor name and title were found, we retained the earliest filing, and the granted patent record over the application. As a consequence of preferring the record with the earlier filing date, in most cases the US patent record was retained rather than European (a common pattern in technology patents is to file first in the US). As a result, we found a total of 141 patents citing ISMIR papers, and 13 patents citing MIREX.

For each of these patents, we identified: the year of filing; whether this patent instance was an application or had been granted; if granted, the year; the inventor(s) and assignee(s); and the number of references to ISMIR and MIREX in each patent. For each reference to ISMIR or MIREX, we collected basic bibliographic data (title, authors, and year of publication).

4. DATA AND DISCUSSION

This section presents an analysis of both sets of patents—those referencing ISMIR publications, and those referencing MIREX. Given the disparity in the size of these two datasets (141 and 13, respectively), we present the results separately.

4.1 Analysis of the Patents

In this section, we summarize the broad characteristics of the 141 patent documents that reference ISMIR publications: distribution of patents by year of application, unique inventors and assignees, ISMIR references included in the patents, and patent topics.

4.1.1 Year of Application

Of the 141 patents identified as referencing ISMIR publications, 102 have been issued as of the date of our dataset gathering (April 2013), and 39 exist as applications. Figure 1 shows the number of patents by application year.

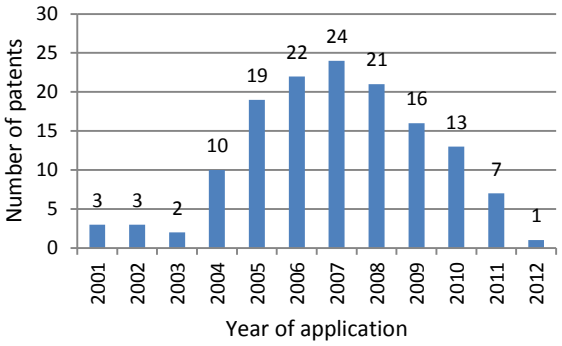

Figure 1. Number of patents by year of application

Examining the dates of application, we see a peak of patent applications referencing ISMIR publications in 2007, with a sharp drop-off in 2011 and 2012. The presence of only a single ISMIR-referencing patent filing in 2012 may be partly explained by a time delay in updates to the Google Patent Search and the underlying USPTO and EPO databases; according to the USPTO website, most patent applications filed on or after November 29, 2000, are published 18 months after the filing date of the application[2]. Also the EPO website states that the patent application is published 18 months after the date of filing

[1] https://www.google.com/patents

[2] http://www.uspto.gov/faq/patents.jsp

or the priority date[1]. Still, the steady decline remains; this may be due to the fact that technologies in some areas are saturated and new developments are appearing in a slower pace than they used to in the earlier stage of the MIR field. We consider possible explanations for this trend as we examine the characteristics of the ISMIR references themselves (Section 4.2).

4.1.2 Analysis of Inventors and Patent Distribution

There are 241 unique inventors in the ISMIR dataset, with an average number of 1.39 patents per inventor. All inventors are natural persons (US patents were preferred in selecting between multiple filings, and under US law corporate entities cannot be registered as inventors). The distribution of patents over inventors (Figure 2) indicates a strong skew towards single-filing inventors; here, approximately 80% of inventors are associated with one filing. We found this a bit surprising as one would assume that typically a team of researchers is involved in developing new algorithms/technologies (as evidenced by a growing trend toward co-authorship in ISMIR [1]).

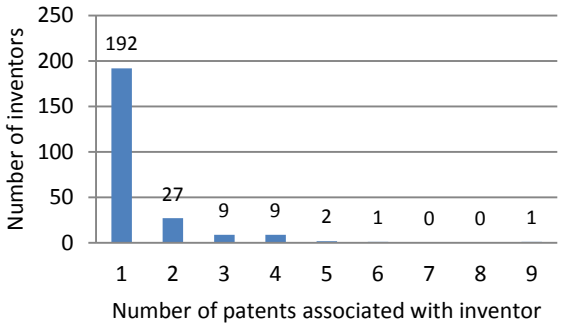

Figure 2. Number of patents sorted by number of inventors

The four inventors associated with the largest number of patents are presented in Table 1. Of the four, Masataka Goto and Brian A. Whitman have published in ISMIR proceedings.

No. of patents	Inventors
9	Louis B. Rosenberg
6	Thomas Kemp
5	Masataka Goto Brian A. Whitman

Table 1. Inventors with the largest numbers of patents referencing ISMIR

4.1.3 Analysis of assignees and Patent Distribution

Looking at assignees, there are 120 unique assignees, with an average of 1.18 patents per assignee. Of the 120 assignees, 21 are individuals, and 99 are corporate entities (primarily commercial organizations and universities).

Figure 3 shows the number of patents per assignee, which ranged from 1 to 12; there were 120 unique assignees in total. Most of the assignees were associated with a single patent (58.3%). Looking at the top 10 patent-holding assignees (Table 2), we see a mix of large IT corporations (Google, Apple, Microsoft, Yahoo!), two electronics corporations (Philips, Sony), a music and video metadata specialist company (Gracenote), and two organizations specializing in patent acquisitions (Colwood Technology, Outland Research).

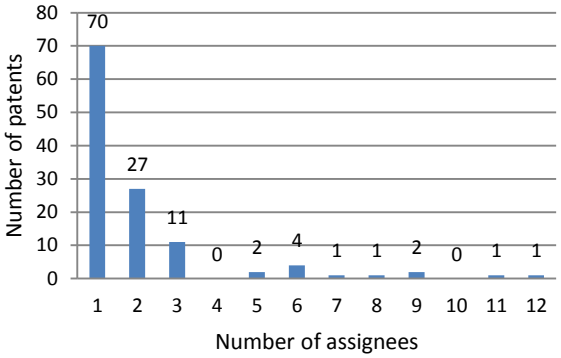

Figure 3. Number of assignees per patent referencing ISMIR

No. of patents	Assignees
12	Google Inc.
11	Apple Inc.
9	Colwood Technology, Llc
9	Outland Research, Llc
8	Strands, Inc.
7	Microsoft Corp.
6	Gracenote, Inc.
6	Koniklijke Philips Electronics N.V.
6	Sony Corp.
6	Yahoo! Inc.

Table 2. Top 10 assignees by number of patents

4.1.4 Number of ISMIR References per Patent

There were a total of 213 references to ISMIR publications in the patents we analyzed. The average number of ISMIR references per patent was 1.5. Unfortunately we cannot estimate the proportion of ISMIR references to all references included in the patents; the lack of standardization in patent descriptions and background formats, compounded by errors in the Google process for identifying patent metadata (including references and patent citations) preclude this type of analysis.

[1] http://www.epo.org/applying/basics.html

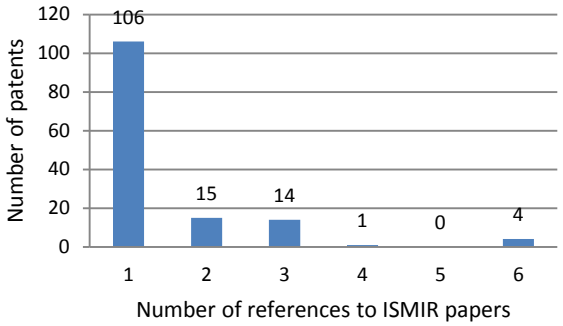

Figure 4. Number of ISMIR reference per patent

4.1.5 Topics of the Patents

In order to get an overview of the patent topics, we manually identified and categorized the main topics represented by each patent. Table 3 shows the top 10 most common topics represented in the patents analyzed.

Topics	Number of patents
Audio fingerprinting	14
Identification of similar songs	12
Music recommendation	9
Automatic playlist generation	9
Audio/Music analysis	6
Music player/interface	6
Music search and display	5
Music visualization	5
Music classification	5
Music retrieval methods	4

Table 3. Top 10 patent topics

The most common topic was audio fingerprinting; 14 patents dealt with various technologies related to audio fingerprinting or thumbprinting. This is followed by 12 patents about using audio similarity algorithms to identify songs similar to a sample song from a music collection. Music recommendation and automatic playlist generation were also popular topics. This is probably related to the increase in the popularity of streaming services and emergence of new types of services such as music identification as all of these technologies are commonly used in popular music services [1], for instance, music streaming services such as Pandora and Spotify, and music identification services such as Shazam or Soundhound. Topics such as music analysis, search, display, classification, and retrieval methods that are important components of music digital libraries/applications also appeared multiple times in the patents. Different techniques for music display and visualization were also found multiple times. Some examples of other topics that appeared two times include: music metadata, composition, audio encoding/decoding, associating music and geographic information, and social ratings.

4.2 Analysis of ISMIR Publications Cited in Patents

In this section, we examine the ISMIR publications cited in the patents: specifically, the distribution of citations from the ISMIR conference series, the ISMIR publications most frequently referenced in the patent dataset, and the overlap between ISMIR authors and inventors.

4.2.1 Year of ISMIR Publications

Figure 5 shows the year of publication of ISMIR papers referenced in the patent dataset. The most striking aspect of this figure is the peak at 2002; just over a (34.9%) of the ISMIR papers referenced were drawn from the 2002 ISMIR conference proceedings. This may be due to the fact that topics related to content-based retrieval (e.g., audio music similarity, automatic generation of playlists) started to become quite popular around that time.

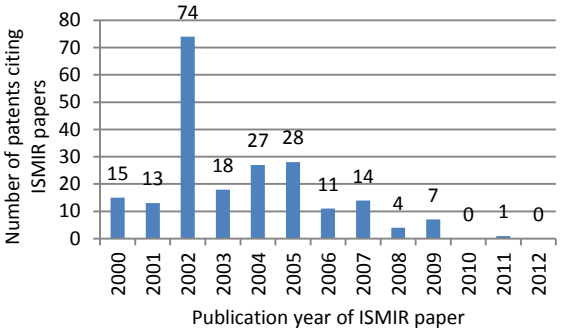

Figure 5. Number of patents citing ISMIR publications, by publication year of ISMIR paper

4.2.2 Most Highly Cited ISMIR Publications in Patents

Table 4 lists the top 10 most highly cited ISMIR publications in the patents we analyzed, sorted by the number of times cited. These publications date primarily from the early years of the ISMIR conference series; this skew is to be expected, given that the older publications have more time to accumulate citations and given the distribution of ISMIR-referencing patent filings (Figure 1).

Authors (pub. year)	Title	Freq
Haitsma J, Kalker T (2003)	A highly robust audio fingerprinting system	21
Cano P, Kaltenbrunner M, Gouyon F, Batlle E (2002)	On the use of FastMap for audio retrieval and browsing	10
Logan B (2002)	Content-based playlist generation: exploratory experiments	9

McKinney MF, Moeland D (2004)	Extracting the perceptual tempo from music	9
Pauws S, Eggen B (2002)	PATS: realization and user evaluation of an automatic playlist generator	9
Aucouturier J-J, Pachet F (2002)	Music similarity measures: What's the use?	7
Pampalk E, Flexer A, Widmer G (2005)	Improvements of audio-based music similarity and genre classification	5
Logan B (2000)	Mel frequency cepstral coefficients for music modeling	5
Berenzweig A, Logan B, Ellis DPW, Whitman B (2004)	A large-scale evaluation of acoustic and subjective music similarity measures	4
Liu D, Lu L (2003)	Automatic mood detection from acoustic music data	4
Tzanetakis G, Essl G, Cook P (2002)	Automatic musical genre classification of audio signals	4
West K, Cox S (2005)	Finding an optimal segmentation for audio genre classification	4
Paulus J, Klapuri A (2002)	Measuring the similarity of rhythmic patterns	4
Oliver N, Kreger-Stickles L (2006)	PAPA: Physiology And Purpose-Aware automatic playlist generation	4

Table 4. Ten most highly cited ISMIR publications by patents

Four of the papers in Table 4 are also listed among the most highly cited ISMIR papers in a 2009 informetric analysis of the ISMIR conference series [1]: specifically, Aucouturier & Pachet (2002), *Music similarity measures: What's the use?*; Logan (2000), *Mel frequency cepstral coefficients for music modeling*; Tzanetakis et al (2002), *Automatic musical genre classification of audio signals*; and Paulus & Klapuri (2002), *Measuring the similarity of rhythmic patterns*.

4.2.3 Inventors Who Publish in ISMIR

A total of 49 inventors have published in ISMIR proceedings; Table 5 presents a list of most prolific ISMIR authors among the the inventors. Masataka Goto (National Institute of Advanced Industrial Science and Technology) emerged as one of the key figures for connecting the academic research and commercial developments, followed by Brain Whitman (The Echo Nest) and Malcolm Slaney (Microsoft).

Author	No. of ISMIR publications	No. of patents invented
Masataka Goto	29	5
Daniel P. W. Ellis	28	1
Elias Pampalk	15	1
Francois Pachet	15	1
Tim Pohle	14	2
Hiroshi Okuno	13	2
Kazuyoshi Yoshii	12	2
Dominik Schnitzer	10	1
Douglas Eck	9	1
Mitsunori Ogihara	9	1
Paul B. Lamere	8	1
Brian A. Whitman	7	5
Malcolm Slaney	7	3
Kristopher C. West	7	2
Josep-Lluis Arcos	6	1

Table 5. Top 15 inventors who are also ISMIR authors (sorted by the number of ISMIR publications)

4.3 Analysis of MIREX References in Patents

Thirteen patents were identified that referenced MIREX: 9 had been granted, and 4 were applications. Table 6 summarizes the filing dates and, for the granted patents, the year in which they were issued. Of these thirteen, six also referenced ISMIR papers. A significant degree of overlap could be expected, given that the research of MIREX participants is also frequently published in the associated ISMIR conference. Indeed, the relationship between the MIREX and ISMIR events may explain the relatively small number of patents including MIREX references: entries to the MIREX trials are frequently accompanied by more detailed submissions regarding the algorithms to the associated ISMIR. Further, MIREX proceedings are informally published and can be difficult to locate and cite [1]—additional reasons why an ISMIR paper might be cited in preference to a similar MIREX publication.

Year	No. of Applications	No. Issued
2006	3	
2007	5	
2008	2	1
2009	1	
2010	2	5
2011		1
2012		2

Table 6. Summary by year of MIREX patents

These thirteen patents included 24 references to MIREX: 21 references to 11 unique papers in the MIREX proceedings, 2 references to an ISMIR paper providing an overview of MIREX results, and 3 more general references to the MIREX trials as a whole. These latter point to the significance of algorithms by relevance to their representation and relative performance in the MIREX trials (e.g., "…systems using this technique regularly rank in

the very top places in the yearly MIREX Automatic Music Recommendation evaluations…" [1]).

The thirteen patents had 22 unique inventors (one inventor was named in two patents). Of these 22 inventors, 11 had published in at least one ISMIR conference (See Figure 6, MIREX inventors publishing in ISMIR by year)—yielding an average of 2.91 inventors who cited MIREX in attendance at each ISMIR, 2002 – 2012.

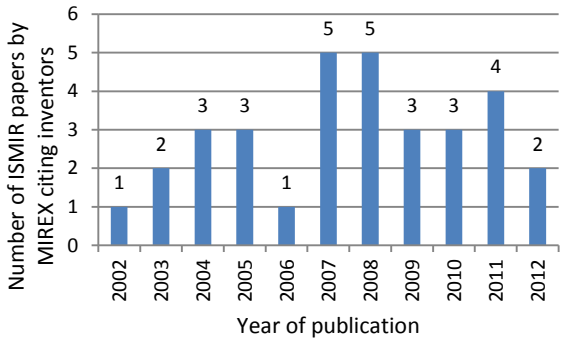

Figure 6. MIREX inventors publishing in ISMIR by year

5. CONCLUSION AND FUTURE WORK

In this paper, we have examined the influence of academic MIR research from the ISMIR conference series and MIREX on patents, as viewed through citation links from patents to the academic publications. We identified over hundred references to ISMIR and MIREX research in the patents, most prominently in early 2000s, and concerning various content-based retrieval technologies such as audio fingerprinting, recommendations, automatic playlist generation, and so on. The investigation indicates the presence of strong and ongoing personal links between academic and commercial MIR research, as evidenced by the number of individuals who produce both academic publications and patents in MIR (Section 4.2.3). It is also encouraging to see significant proportions of the ISMIR/MIREX references in patents not filed by inventors with direct connections to ISMIR (Sections 4.2.1, 4.2.3).

Ironically, it would not be straightforward to investigate the opposite flow of influence—ISMIR/MIREX publications citing patents—because the ISMIR and MIREX events are not associated with the primary computer science and engineering professional societies and consequently are not included in the societies' digital library portals (i.e., the ACM Digital Library and IEEE Xplore). Further, the MIREX papers are incompletely represented in Google Scholar [1], and both MIREX and ISMIR are inconsistently indexed [1], [3]—to the extent that it is not possible to be assured of complete coverage of ISMIR and MIREX publications through Google Scholar searches. Two digital libraries have been developed and maintained by individuals within the MIR community to provide improved access to ISMIR and MIREX publications: Michael Fingerhut's ISMIR net (http://www.ismir.net/) for ISMIR publications, and David Bainbridge's library (http://music-ir.org/mirex-dl/library) for MIREX. However, it lacks full-text search capabilities, and the latter does provide complete, standardized metadata—and so their utility for informetric investigations is reduced. More seriously, the visibility of the ISMIR and MIREX series as a whole is diminished.

In our future work, we plan to conduct a topic analysis of ISMIR publications, to identify shifts in focus over the conference series and for comparison to topics identified in the patents (Section 4.1.3). Also further investigation is required to tease out the factors contributing to the decline in the number of patents referencing ISMIR. One possibility is that the research interests of the academic and commercial communities have indeed diverged (though the overlap between academic publishers and inventors argues against this). Ultimately, we hope to expand our search and identify the patents citing any publications related to MIR in multiple publication venues, not limited to ISMIR conference proceedings. This may help reveal a direction for new research that can make strong impact in the everyday life of music users.

6. REFERENCES

[1] S. J. Cunningham, D. Bainbridge, and J. S. Downie: "The impact of MIREX on scholarly research," *Proc. of the ISMIR*, pp. 259-264, 2012.

[2] J. Futrelle, J. S. Downie: "Interdisciplinary Communities and Research Issues in Music Information Retrieval," *Proc. of the ISMIR*, pp. 215-221, 2002.

[3] J. H. Lee, M. C. Jones, J. S. Downie: "An analysis of ISMIR proceedings: Patterns of authorship, topic, and citation," *Proc. of the ISMIR*, pp. 57-62, 2009.

[4] J. H. Lee and N. M. Waterman: "Understanding user requirements for music information services," *Proc. of the ISMIR*, pp. 253–258, 2012.

[5] A. Klitkou and M. Gulbrandsen: "The relationship between academic patenting and scientific publishing in Norway", *Scientometrics*, Vol. 82, pp. 93-108, 2010.

[6] M. Meyer: "Patent citations in a novel field of technology—what can they tell about interactions between emerging communities of science and technology?" *Scientometrics,* Vol. 48, No. 2, pp. 151-178, 2000.

[7] M. Meyer: "What is special about patent citations? Differences between scientific and patent citations," *Scientometrics*, Vol. 49, No. 1, pp. 93-123, 2000.

[8] F. Narin: "Patent bibliometrics," *Scientometrics*, Vol. 30, pp. 147–155, 1994.

[9] D. Schnitzer: "A method and a system for identifying similar audio tracks." European Patent No. EP 2273384. 12 Jan. 2011.