

TEN YEARS OF ISMIR: REFLECTIONS ON CHALLENGES AND OPPORTUNITIES

J. Stephen Downie

University of Illinois
at Urbana-Champaign
jdownie@illinois.edu

Donald Byrd

Indiana University
at Bloomington
donbyrd@indiana.edu

Tim Crawford

Goldsmiths College
University of London
t.crawford@gold.ac.uk

ABSTRACT

The *International Symposium on Music Information Retrieval* (ISMIR) was born on 13 August 1999. This paper expresses the opinions of three of ISMIR's founders as they reflect upon what has happened during its first decade. The paper provides the background context for the events that led to the establishment of ISMIR. We highlight the first ISMIR, held in Plymouth, MA in October of 2000, and use it to elucidate key trends that have influenced subsequent ISMIRs. Indicators of growth and success drawn from ISMIR publication data are presented. The role that the Music Information Retrieval Evaluation eXchange (MIREX) has played at ISMIR is examined. The factors contributing to ISMIR's growth and success are also enumerated. The paper concludes with a set of challenges and opportunities that the newly formed *International Society for Music Information Retrieval* should embrace to ensure the future vitality of the conference series and the ISMIR community.

1. ORIGINS OF ISMIR

In mid-August 1999, Byrd and Downie were at the Radisson Hotel Berkeley Marina conference center in Berkeley, California: Byrd for the ACM (Association for Computing Machinery) Digital Library Conference (DL '99), Downie for the ACM SIGIR conference, which immediately followed DL '99. We had not met before, but our paths had been converging for some time, and in retrospect, it is hardly surprising that something special came out of our face-to-face encounter. Crawford was in England at the time, but he and Byrd had been collaborating since the early 1990s. Crawford and Byrd had recently received word that their "Online Music Recognition and Searching" (OMRAS) project [1] would be jointly funded by the Joint Information Systems Committee (JISC) of the UK and the National Science Foundation (NSF) of the USA. Steve Griffin,

the project's NSF program officer, had already suggested to Byrd and Crawford independently that a music-IR workshop be organized in conjunction with OMRAS. Furthermore, Crawford was organizing another workshop on music IR, as part of the "Digital Resources for the Humanities" conference to be held in London in September 1999. Finally, Downie, with the assistance of David Huron (Ohio State University) and Craig Nevill-Manning (then of Rutgers University), had organized "The Exploratory Workshop on Music Information Retrieval" at SIGIR '99.¹ Before going to Berkeley, Downie was already thinking of a larger-scale follow-up event as this was an explicit goal of his SIGIR workshop. One of the workshop presenters, Michael Fingerhut of IRCAM, would later play a pivotal role in the success of ISMIR through his establishment and maintenance of vital community resources (see Section 3.2).

With the encouragement of Bruce Croft (University of Massachusetts, Amherst)—a very well-known researcher in the text IR world, and Byrd's boss at the time—Downie and Byrd decided on the spot to join forces to plan a larger-scale event instead of a workshop in the normal sense, and they came up with the name "International Symposium on Music Information Retrieval."

Most of the above has been described in print before [2]. Previously unreported, however, are some informal meetings convened in Berkeley, which variously included Byrd, Downie, Nevill-Manning, David Bainbridge (University of Waikato), Matthew Dovey (University of Oxford), and Massimo Melucci (University of Padua). It is interesting that Byrd's notes of these meetings show a heavy emphasis on music in symbolic form over audio, and quite a bit of discussion of TREC²-like evaluations of music-IR systems.

1.1 What's in a Name?: Evolution of "ISMIR"

The ISMIR acronym, decided upon during the August 1999 meetings, was carefully crafted. First, both Byrd and Downie wanted to strongly encourage the participation of researchers from around the world, so *International* was chosen without hesitation. Second, the word *symposium* has its roots in the Greek verb, *sympotein*,

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page.

¹ See http://nema.lis.uiuc.edu/sigir99_mir_wshop.pdf.

² The Text Retrieval Conference upon which MIREX is based.

which means “to drink together.”¹ As many know, Downie is particularly fond of symposia. Besides its social connotations, the term *Symposium* was agreed upon as it indicated a certain academic middle-ground between a workshop and a full-fledged conference. Before long, however, some participants noted that they were having difficulties obtaining travel funding to attend “a mere symposium,” and in 2002 ISMIR became the “International Conference on Music Information Retrieval.” Over the years, ISMIR organizers explored affiliation opportunities with such organizations as the Association for Computing Machinery (ACM), the Institute of Electrical and Electronics Engineers (IEEE), and the International Computer Music Association (ICMA); none worked out. Undeterred, Ichiro Fujinaga of McGill University led the way to formally establishing ISMIR as an independent society. On 4 July 2008, the “International Society for Music Information Retrieval” was officially born. By the time ISMIR 2009 in Kobe concludes, the music-IR community will have elected its first roster of ISMIR executive officers and held its first Annual General Meeting.

2. ISMIR 2000 AT PLYMOUTH, MA: LANDING OF THE MUSIC-IR PILGRIMS

In accordance with the events of 1999 described above, ISMIR 2000² was held in Plymouth, Massachusetts (the site of the Pilgrims’ 1620 arrival in the New World) from 23 to 25 October 2000. Byrd was general chair and Downie was program chair. The other organizing committee members were Crawford, Croft, and Nevill-Manning. In addition, Jeremy Pickens, then a PhD student working on the OMRAS project, became, by virtue of his good nature, *the* local organizer—i.e., audio-visual person and general helper—during the conference.

In terms of statistics, 88 people attended ISMIR 2000: not bad at all for a first conference in the field, and about twice the attendance at the first computer-music conference (which Byrd had attended in 1974). Furthermore, attendance was already very international: 29 attendees (33%) came from 11 countries outside the United States. ISMIR 2000 was very heavy on invited papers, of which there were nine. An additional 33 papers were submitted; 10 were accepted as papers, 16 as posters.

2.1 ISMIR 2000: Highlights and Commentary

Many of the intellectual themes, challenges, and opportunities that would resonate throughout subsequent convenings of ISMIR were already evident in Plymouth. To illustrate this, a selection of ISMIR 2000 highlights with *editorial comments* follows:

- Marvin Minsky delivered the keynote address. *His talk was uniquely creative and pointed out several connections that are still relevant, e.g., to artificial intelligence, improvisation vs. written-out music, and even to his institution, MIT.*
- Beth Logan gave one of the first papers formally examining the implications of using Mel Frequency Cepstral Coefficients (MFCC) for music; this created a fair amount of controversy. *We wish we had a penny for each MFCC calculated since 2000!*
- There were two papers on music digital library applications: Jon Dunn spoke on the “Variations” system; David Bainbridge talked about the “New Zealand Digital Music Library.” *Downie, as a library science professor, notes with some sadness that the digital library theme has not gained much traction in subsequent ISMIRs.*
- Byrd, Crawford, and Steve Larson led a “Lecture, Recital, Discussion, and Survey” session on music similarity. Centered on Mozart’s piano piece *Variations on Ah! Vous dirai-je, maman* (the melody English speakers call “Twinkle, Twinkle, Little Star”), Larson played the piece, and attendees filled out survey forms to say how similar they felt each of the selected variations was to the theme. *This session led to our choosing three measures from the Mozart variations for the ISMIR logo (Figure 1). The “similarity problem” remains a huge challenge, not least because of the difficulty of establishing “ground-truth” in this subjective area.*




Figure 1. The Mozart-based official ISMIR logo.

- Mary Levering of the U.S. Patent and Trademark Office talked about “Intellectual Property Rights in Musical Works.” *This is a problem that continues to plague many music-related activities, including music-IR research.*
- George Tzanetakis and Perry Cook gave a paper on audio-IR tools. *Tzanetakis’s MARSYAS is now one of the most widely used music-IR toolkits.*
- Jonathan Foote gave a paper on recognizing pieces of orchestral music regardless of performance differences. *Foote’s approach looked solely at low-level (though long-term) audio features. Numerous music-IR papers since 2000 ignore musical knowledge and instead employ low-level features that seem to work; this paper foreshadows the trend.*
- There were papers on musicology applications, transcription from audio, retrieval from audio, Optical Mu-

¹ See <http://en.wikipedia.org/wiki/Symposium>.

² See <http://ismir2000.ismir.net/>.

sic Recognition (OMR), language modeling, and XML representation of music notation. *All of these except language modeling have been the subjects of numerous ISMIR papers since.*

- Eric Allamanche of Fraunhofer gave an informal demo of an audio fingerprinting application designed to identify broadcast music in real time. *Similar systems are now widely available in the commercial sphere.*
- The creators of MusicXML and MEI each gave posters on early versions of their representations. *Since 2000 MusicXML has become the most popular XML form for music content, but MEI has recently been the subject of development for specific applications, particularly in the musicological domain.*
- Suzanne Lodato of the Andrew W. Mellon Foundation took an active role in the plenary planning and discussions sessions of the symposium. *The Mellon Foundation would go on to provide critical funding to prepare for, establish, and run MIREX.*

Of the 19 full papers presented, six were mostly or entirely on audio; nine mostly or entirely on music in symbolic form (including metadata); and four about equally on audio and symbolic music. As is obvious to anyone who has attended the last five or six ISMIRs, the predominance of symbolic music (reflecting the backgrounds of the original organizers) has not persisted; we will say more about this later.

Despite the inexperience of the organizers and the novelty of the subject area, ISMIR 2000 was universally regarded as a resounding success.

3. SUCCESS AND GROWTH OF ISMIR: 2000–2009

The ongoing success and growth of ISMIR since 2000 is both remarkable and encouraging. The vitality of the community is readily apparent from even the most cursory examination of the statistics. For example, Table 1 presents the number of published items (both posters and papers), number of pages published, and the number of unique authors represented in the proceedings of each ISMIR from 2000 to 2009. The table shows a 251% increase in the number of published items per year, from 35 to 123. The number of pages went up even more: 155 to 729 is a 370% increase.

The number of unique authors represented also grew tremendously, by 363% from 2000 (63) to 2009 (292). On average, 183 unique authors made contributions to each of the 10 ISMIRs under consideration. For us, the growth in the number of unique authors is the best statistic of the set, since it indicates that ISMIR has attracted the most important asset of any conference: active, engaged and publishing researchers.

Mailing list statistics also confirm ISMIR’s success. The *music-ir@ircam.fr* list, established in October 2000, is the ISMIR community’s primary communications me-

chanism. This list, as of 22 August 2009, has 1190 registered subscriptions. It has broadcast nearly 3000 messages for an average of 28 per month. These are strong numbers for such a specialized research area as music IR.

YEAR	LOCATION	ITEMS	PAGES	UNIQUE AUTHORS
2000	Plymouth, MA	35	155	63
2001	Bloomington, IN	41	222	86
2002	Paris, FR	57	300	117
2003	Baltimore, MD	50	209	111
2004	Barcelona, ES	105	582	214
2005	London, UK	114	697	233
2006	Victoria, BC	95	397	198
2007	Vienna, AT	127	486	267
2008	Philadelphia, PA	105	630	253
2009	Kobe, JP	123	729	292
TOTALS	----	852	4407	----

Table 1. ISMIR publication and author data 2000–2009.¹

3.1 The Audio Description Contest and MIREX

As mentioned in Section 1, the formal evaluation of music-IR systems has been part of the ISMIR “wish list” since its inception. However, notwithstanding strong community interest, it was surprisingly difficult to institute a formal evaluation framework along the lines of TREC. There were many challenges to overcome, the greatest of which was the lack of high-quality test collection and ground-truth data caused primarily by the very restrictive intellectual property regimes governing music [4]. After a series of exploratory workshops led by Downie and funded by Mellon and NSF [5], the organizers of ISMIR 2004 in Barcelona were able to put together the “Audio Description Contest” (ADC) [6]. Many valuable lessons were learned in the running of ADC and these were subsequently incorporated into MIREX. After receiving substantial long-term funding from both Mellon and NSF, MIREX established itself as a permanent fixture in time for ISMIR 2005 [7].

	2005	2006	2007	2008
Number of Task (and Subtask) “Sets”	10	13	12	18
Number of Individuals	82	50	73	84
Number of Countries	19	14	15	19
Number of Runs	86	92	122	169

Table 2. MIREX descriptive data 2005–2008 [8].

Like the publication data examined previously, the MIREX data are quite encouraging. Table 2 summarizes the key descriptive data for MIREX between 2005 and

¹ 2000–2008 data sourced from the *Preface* of the ISMIR 2008 proceedings [3].

2008. A fuller explication of the MIREX data can be found in [4, 8].

The number of task and subtask sets grew by 80% from 2005 (10) to 2008 (18). This growth can be attributed to growing interest in MIREX and the donation of new high-quality data sets from community members.

In keeping with ISMIR’s international mission, the number of countries represented was a strong, but flat, 19 for both 2005 and 2008 with an average of 17 per year. Most of these numbers come from European countries with Japan, China, and Taiwan also represented. Likewise, the number of individual participants has not appreciably increased between 2005 (82) and 2008 (84). We do note the lack of growth in the country and participant numbers as something that needs addressing.

The most heartening MIREX statistic concerns the number of individual runs performed: this went from 86 to 169, an increase of 96%. Note that the increase is greater than the increases in both participants and tasks: participants are more likely now to submit multiple variations on their algorithms. This fact suggests to us that MIREX has been successful in its message that MIREX exists as an exploratory mechanism designed to try out new ideas and *not* a “contest” to be won or lost.

In total, MIREX has run 469 algorithms. It is interesting to note the distribution of runs over areas of interest:

- 129 (28%) can be categorized as “train-test” machine-learning classification experiments (e.g., Audio Genre Classification, Audio Mood Classification, etc.).
- 139 (30%) can be categorized as “search” experiments (e.g., Audio Cover Song Identification, Audio Music Similarity, etc.)
- 201 (43%) can be categorized as “low-level” feature experiments (e.g., Audio Onset Detection, Audio Beat Tracking, etc.)

We must also note that, of the 22 unique task sets run over 2005 to 2008, only three (14%) have dealt exclusively with symbolic music data (i.e., Symbolic Genre Classification, Symbolic Key Finding, and Symbolic Melodic Similarity). Not one of these symbolic tasks was run in 2008 and not one proposed for MIREX 2009. 16 task sets (73%) have been exclusively audio-based (e.g., Audio Tempo Extraction, Audio Key Finding, etc.), and three tasks have involved a combination of audio and symbolic data (i.e., Query-by-Singing/Humming, Query-by-Tapping, and Score Following). As these data show, MIREX has been quite successful in growing evaluation activity in the audio domain, but not at all successful in helping the symbolic sub-community to flourish: this is perhaps MIREX’s most serious weakness.

3.2 Success and Growth Factors

Many factors have contributed to the success and growth of ISMIR over the years. These factors are both external

and internal to ISMIR. Like many things in life, ISMIR has been successful through a combination of good timing, thoughtfulness, and hard work.

From the beginning, ISMIR’s timing was good; it has benefitted from several important external opportunities and trends that developed in parallel. These developments have provided ISMIR with a larger body of researchers and research themes to draw upon than we could have anticipated, especially in the audio domain. We believe these external factors include:

- The success of the audio compression research community in developing techniques specifically designed for, and tested against, music. It was this success and the subsequent acceptance of these approaches that afforded the opportunity to create, share, and store large collections of music audio.
- The explosive growth in the availability of audio files, mostly MP3’s, via the Internet. This growth resulted to a great extent from the audio-compression research described above, but in turn it created a demand for better search and retrieval mechanisms. Napster, for example, was established in 1999.
- The work of such standards bodies as the MPEG-7 group that brought together important industry players with leading academic research groups. The MPEG-7 first working draft came out in December 1999.¹
- The success of such search engines as Google, Yahoo, etc., that encouraged researchers to seek fame and fortune in the music domain. The great “dot.com bubble” of 1998–2001 was contemporaneous with ISMIR’s early development.

The internal factors that have contributed to ISMIR’s success are founded upon the thoughtful actions, goodwill, and hard work of community members acting either as individuals, in small groups, or collectively. These factors include:

- The establishment of the communication resources housed at IRCAM. The *music-ir@ircam.fr* mailing list, the hosting of the conference websites, and the archiving of the collected ISMIR proceedings are resources without which ISMIR might not exist today. Each of these has contributed inestimably to the openness, continuity, and intellectual life of the ISMIR community. We applaud Michael Fingerhut for his continued service.
- The diversity of backgrounds and disciplines represented on the ISMIR Steering Committee (SC). The SC has worked hard over the years to ensure that the broadest possible range of research interests is present at each ISMIR. Ichiro Fujinaga has been the SC’s coordinator for years, and he is especially commended for his ability to guide the SC through its deliberations.
- The great fortune ISMIR has had in the quality of the chairs and program committee (PC) members for each

¹See <http://www.chiariglione.org/mpeg/standards/mpeg-7/mpeg-7.htm>.

conference. We have nothing but praise for the ISMIR PC teams; each ISMIR has been organized and run with enthusiasm, integrity, and efficiency.

- The implicit policy of inclusiveness that has pervaded the conference programming ethos of each ISMIR. Unlike other technology-related conference series, ISMIR has not measured its intrinsic value through high rejection rates. In fact, the ISMIR PCs are to be applauded for finding mechanisms like expanded poster presentation opportunities to allow for the maximum level of participation yet maintaining academic research quality through strong peer-reviewing. We believe that it is precisely this policy of inclusiveness that has allowed for the all-important growth in unique author participation noted in Section 3. The ISMIR community as a whole is also to be praised for its consistent efforts to make the peer-review process simultaneously as fair, open-minded, and rigorous as possible.
- The ongoing PC and general community support for ADC and MIREX. This support has contributed to ISMIR by fostering a sense of common purpose and exploration among researchers in many of ISMIR's sub-fields. MIREX has also helped to set standards in many sub-fields for what constitutes proper evaluation. Finally, MIREX has provided an extra opportunity for participation in ISMIR for those researchers whose work could not be included in the official proceedings. We must acknowledge here the extra-special efforts made by Kris West, M. Cameron Jones, Andreas F. Ehmann, and Mert Bay in making MIREX run well.

4. CHALLENGES AND OPPORTUNITIES

In its first 10 years, ISMIR has grown into a vibrant and enthusiastic research community. We now need to turn our attention to making ISMIR's next 10 years, its "teen" years, even more rewarding and successful. Like a teenager, ISMIR will undoubtedly stop growing in size at some point; this is only natural. But if ISMIR—both as a conference series and as a society—is to have a successful "adulthood," it will need to address some challenges that it has not fully engaged with before. It must recast these challenges as opportunities and engage them with its growing maturity and its youthful vigor. Five of the most important challenges are:

1. ISMIR needs to more actively encourage the participation of potential users of music-IR systems. Notwithstanding the laudable efforts made by the ISMIR Steering Committee, ISMIR has tended to focus much less on the potential *users* of music-IR technology than on its developers. These users might include, for example, performing musicians, film-makers, musicologists, music librarians, sound archivists, music educators, and music enthusiasts of all types. The knowledge acquired by interacting with users like these can only improve the quality

of the community's research output. It will also go a long way to helping ISMIR researchers create truly useful music-IR systems.

2. ISMIR research projects must dig deeper into the music itself. Notwithstanding some recent—and heartening—developments in such areas as, for example, chord detection, cover song detection, and structural analysis, etc., a large amount of ISMIR research effort, especially in timbre-based audio matching, has gone into attempts to enhance a few basic features and matching algorithms. However, it seems likely that there is a point beyond which improved matching performance using any single feature cannot be achieved [9]. On the other hand, the incorporation of multiple features in what might be thought of as "hybrid" matching tends to be more successful. But such combining of features needs to be done in a way that is understood and principled, and much more research needs to be done in understanding what such combinations *actually represent in musical terms*. The integration of symbolic music data to create hybrid audio + symbolic music-IR systems could help in this regard.

3. Time has come for ISMIR to expand its musical horizons. The vast majority of ISMIR's collective music-IR research has been conducted on Western popular musics of the late-20th and early-21st centuries. This is a serious problem because there is an enormous amount of music in existence that is utterly different from these corpora. There is no reason to assume algorithms that work superbly for the *Beach Boys* will do anything useful with *Tuvan throat singing*, *musique concrète*, or *Indian Raga*.

4. ISMIR must rebalance the portfolio of music information types with which it engages. Music information is inherently multifaceted. Each of its manifestations—audio, symbolic, and metadata—contributes different but equally important features to the experience of music. We celebrate the accomplishments of ISMIR's audio researchers but, as noted before, research exploiting the symbolic aspects of music information has not thrived under ISMIR. We are thrilled to see, however, the growing body of work that strives to unite social metadata and audio information. Rather than "pushing down" on the audio side of ISMIR research, we challenge ISMIR to make special efforts to "pull up" symbolic and metadata research to create a more productive, synergistic, and harmonious balance among the three.

5. ISMIR must encourage the development and deployment of full-featured, multifaceted, robust, and scalable music-IR systems with helpful user-interfaces. During ISMIR's first decade, we have seen a great deal of effort expended on the development of the various sub-components of music-IR systems. Unfortunately, we have not yet seen much in the way of a successful integration of these sub-systems into real-world-useable re-

sources. This state of affairs cannot be sustained for the next decade as the community needs these full-featured systems to exist in order to inspire the development of the next generation of refinements and improvements. In the text IR world, and starting in the 1960s, such systems as “SMART,”¹ “Managing Gigabytes,”² and “Terrier,”³ have fulfilled this important, if not imperative, role.

4.1 The Grand Challenge

We see our “complete system” challenge as “*The Grand Challenge*” for ISMIR’s second decade. By embracing this challenge, the preceding ones will necessarily have to be engaged. We do recognize, however, that meeting this “Grand Challenge” will not be easy. We believe there will be difficulties because academic researchers traditionally have obtained little academic credit for comprehensive system development. Future ISMIR program committees need to find a mechanism through which the developers of such systems can acquire full academic credit for accomplishments. One possibility is to have ISMIR create a rigorous set of peer-reviewing criteria specifically designed to handle this type of work. Along these lines, the demonstration of complete systems should receive the same status now afforded to paper presentations. Special awards should also be considered.

5. CLOSING REMARKS

As we noted in the beginning of this paper, the founders of ISMIR, because of their backgrounds, had conceived of music IR as an intersection of music and symbolic IR techniques. As early as ISMIR 2000, it became readily apparent that this conception was much, much too limiting. ISMIR research papers now cover a wide range of activities and recent “Calls for Papers” have reflected this broadening of scope explicitly. We now challenge the ISMIR community to consider whether the term “music IR” has outlived its usefulness. Is it possible that “information retrieval” is too narrow a concept to fully encapsulate what ISMIR researchers actually do? Byrd has proposed several times making the “R” in “ISMIR” stand for “Research” instead of “Retrieval” which could better describe the breadth of the organization without losing ISMIR’s name recognition. A related idea is to refer to “music informatics” instead of “music information.”

We will leave these questions open in the hope that they will inspire some healthy, self-reflective, debate about the future of ISMIR. It will be through such reflections that ISMIR will continue to be vibrant, energetic, and successful well past its second decade.

6. ACKNOWLEDGEMENTS

We thank Jeremy Pickens, Michael Casey, Ichiro Fujinaga, Masataka Goto, Keiji Hirata, and the Kobe 2009 Organizing Committee, for their assistance in writing this paper. We also thank the NSF and the Mellon Foundation for their financial support of MIREX. We are especially grateful to Steve Griffin of NSF for the financial support he arranged for ISMIR 2000 and 2001.

7. REFERENCES

- [1] D. Byrd and T. Crawford. “Problems of Music Information Retrieval in the Real World,” *Information Processing and Management*, Vol. 38, No. 2, pp. 249–272, 2002.
- [2] D. Byrd and M. Fingerhut. “The History of ISMIR - A Short Happy Tale.” *D-Lib Magazine*, Vol. 8, No. 11, 2002. See <http://www.dlib.org/dlib/november02/11inbrief.html#BYRD>.
- [3] D. P. W. Ellis, Y. Kim, J. P. Bello, and E. Chew. “Preface,” *Proceedings of the International Conference on Music Information Retrieval (ISMIR 2008)*, pp. 9–11, 2008.
- [4] J. S. Downie. “The Music Information Retrieval Evaluation Exchange (2005–2007): A Window into Music Information Retrieval Research,” *Acoustical Science and Technology*, Vol. 29, No. 4, pp. 247–255, 2008. See <http://dx.doi.org/10.1250/ast.29.247>.
- [5] J. S. Downie. “The Scientific Evaluation of Music Information Retrieval Systems: Foundations and Future,” *Computer Music Journal*, Vol. 28, No. 3, pp. 12–23, 2004.
- [6] P. Cano, E. Gomez, F. Gouyon, P. Herrera, M. Koppenberger, B. Ong, X. Serra, S. Streich, and N. Wack. *ISMIR 2004 Audio Description Contest. MTG Technical Report, MTG-TR-2006-02*, Music Technology Group, Barcelona, 2004.
- [7] J. S. Downie, K. West, A. F. Ehmman, and E. Vincent. “The 2005 Music Information Retrieval Evaluation eXchange (MIREX 2005): Preliminary Overview,” *Proceedings of the International Conference on Music Information Retrieval (ISMIR 2005)*, pp. 320–323, 2005.
- [8] J. S. Downie, A. F. Ehmman, M. Bay, and M. C. Jones. “The Music Information Retrieval Evaluation eXchange: Some Observations and Insights,” *Advances in Music Information Retrieval*, Springer, New York, in press.
- [9] J.-J. Aucouturier and F. Pachet, “Improving Timbre Similarity: How High is the Sky?” *Journal of Negative Results in Speech and Audio Sciences*, Vol. 1, 2004. See <http://www.csl.sony.fr/downloads/papers/uploads/aucouturier-04b.pdf>.

¹ See http://en.wikipedia.org/wiki/SMART_Information_Retrieval_System.

² See <http://www.cs.mu.oz.au/mg/>.

³ See <http://ir.dcs.gla.ac.uk/terrier/>.