

AN ANALYSIS OF ISMIR PROCEEDINGS: PATTERNS OF AUTHORSHIP, TOPIC, AND CITATION

Jin Ha Lee

University of Washington
The Information School
jinhalee@uw.edu

M. Cameron Jones

University of Illinois at
Urbana-Champaign
mjones2@illinois.edu

J. Stephen Downie

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

ABSTRACT

This paper presents analyses of peer-reviewed papers and posters published in the past nine years of ISMIR proceedings: examining publication and authorship practices, topics and titles of research, as well as the citation patterns among the ISMIR proceedings. The main objective is to provide an overview of the progress made over the past nine years in the ISMIR community and to obtain some insights into where the community should be heading in the coming years. Overall, the ISMIR community has grown considerably over the past nine years, both in the number of papers and posters published each year, as well as the number of authors contributing. Furthermore, the amount of collaboration among authors, as reflected in co-authorship, has increased. Main areas of research are revealed by an analysis of most commonly used title terms. Also, major authors and research groups are identified by analyzing the co-authorship and citation patterns in ISMIR proceedings.

1. INTRODUCTION

This year, 2009, marks the tenth iteration of the International Symposium on Music Information Retrieval conference series (ISMIR). ISMIR was organized with the hope that the “resulting information interchange will enable scholars to move more quickly towards viable solutions to many problems” [1] in the field of Music Information Retrieval (MIR).

Futrelle & Downie [2] defined MIR as “a rapidly growing interdisciplinary research area encompassing computer science and information retrieval, musicology and music theory, audio engineering and digital signal processing, cognitive science, library science, publishing, and law. Its agenda, roughly, is to develop ways of managing collections of musical material for preservation, access, research, and other uses”. Necessarily, MIR spans both audio and symbolic representations of music [3], but also includes musical metadata, usage data, and other ex-

tra-musical information [4], including user-studies and human-computer interaction studies of music systems. To date, most research in MIR has been content-based [5].

In 2000, MIR was still a fairly new field with a great deal of potential that was gaining the interest of researchers from many different domains. Although ISMIR started as a small-scale symposium, it has grown immensely over the past nine years as more people have recognized the importance of MIR research and have been drawn in to the field. The community has grown to the point of establishing the International Society for Music Information Retrieval, which will help orient, organize, and disseminate the community’s future research.

We performed various informetric analyses on the ISMIR proceedings from 2000 to 2008 in order to discover how the patterns of publications have changed over the past nine years. Through these analyses, we hope to obtain insights into what the ISMIR community has and has not been able to accomplish and which directions it could be heading towards in the coming years.

In the following, we provide descriptive statistics showing the change in the number of publications and authorship patterns over the past nine years. We also provide the results of our analysis of the title terms, looking at the most commonly used single terms as well as bigrams. In addition, we performed analyses on the citation patterns among the publications and authors who have published in the ISMIR proceedings.

2. GROWTH OF THE ISMIR COMMUNITY

The first ISMIR conference had just 10 refereed papers and 16 posters representing 55 authors, with several other scholars presenting invited talks. To date, 881 authors have contributed peer-reviewed papers and posters to the ISMIR proceedings, not to mention the numerous participants in the annual Music Information Retrieval Evaluation eXchange (MIREX), conference workshops, demonstrations, tutorials, and invited talks. The rapid growth in participation has been paralleled by a similar increase in the number of papers and posters accepted to ISMIR. In total, over 700 peer-reviewed papers and posters have been published, comprising a substantial literature on a breadth of topics ranging from signal-processing techniques to user studies of MIR systems.

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.

2.1 Change in the Number of Publications per Year

The number of publications in ISMIR proceedings has been steadily increasing over the past nine years. Exact numbers are presented in Table 1 along with the number of unique authors published in each year.

Year	00	01	02	03	04	05	06	07	08
Papers	10	21	31	23	60	59	59	62	105
Posters	16	16	22	24	44	57	28	65	-
Total	26	37	53	47	104	116	87	127	105
Unique Authors	55	74	113	108	213	232	185	267	262

Table 1. The number of publications and unique authors per year.

We can better observe the changes in the proportion of papers and posters for each year, as well as the changes in the number of authors. The number of publications almost doubled in 2004, jumping from 47 in 2003 to 104. In 2008, there was a change in the submission format and all paper submissions were to have accompanying posters as well. Looking at the number of authors, we can see that there were two sharp increases in 2002 and 2004, and a major drop in 2006. However, the overall number of authors represented at the conference each year has generally grown over the past nine years.

Figure 1 shows the authorship trends within the ISMIR proceedings, tracking the proportion of papers with one, two, three, four, and five-or-more authors. As can be seen, the number of single-authored papers has decreased year-over-year. The number of papers with two co-authors peaked in 2002, and has steadily declined since.

However, the number of papers with three authors has steadily increased year-over-year. The average number of co-authors on papers and posters published each year has increased over the past nine years, from an average of 2.27 authors per publication in 2000 to 2.93 authors per publication in 2008. Clearly ISMIR is becoming a much more collaborative community as the number of authors per paper increases, and the proportion of single, and double-authored papers diminishes in favor of papers with three or more authors.




Figure 1. Co-authorship trends tracking the percentage of papers with 1, 2, 3, 4, and 5+ authors from 2000 to 2008.

2.2 Co-authorship Analysis

We performed an analysis to identify the patterns of co-authorship among all the authors who published in ISMIR proceedings and determine which authors appear as the central hubs in the co-authorship graph. Figure 2

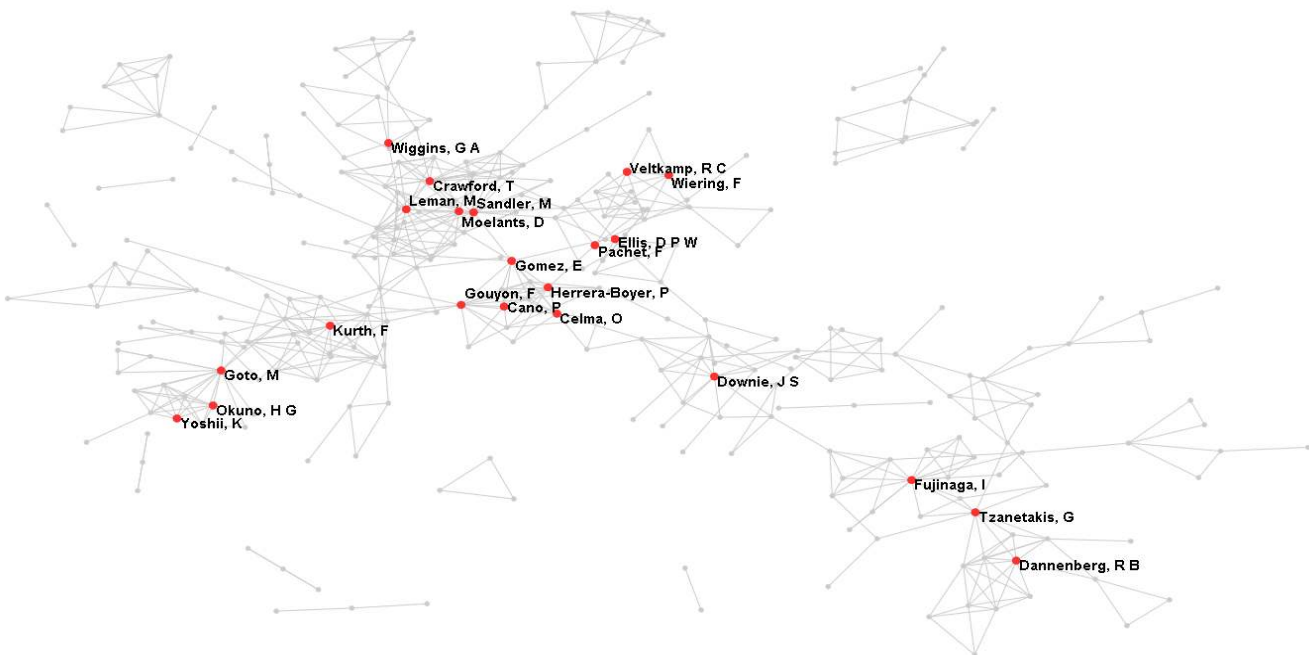


Figure 2. Co-authorship network among ISMIR authors who have published two or more articles. The 22 authors with the largest co-authorship networks have been highlighted.

was generated by using Pajek which is a social networking analysis and visualization tool [6]. In this figure, only the authors who published 2 or more papers/posters are included in order to simplify the network.

Several main clusters of people can be visually identified in the figure showing the close connections among some authors. The top 22 authors with the largest number of distinct co-authors (12+) are labeled in the figure. Often these authors represent an active research group such as National Institute of Advanced Industrial Science and Technology (AIST) headed by M. Goto in Japan, The International Music Information Retrieval Systems Evaluation Laboratory (IMIRSEL) headed by J. S. Downie in Illinois, Distributed Digital Music Archives and Libraries headed by I. Fujinaga in Canada, the Center for Digital Music headed by Mark Sandler in London, and so on.

What is evident from these analyses is the growing role of research labs in the ISMIR community, and how they engender collaboration and increase participation in research. Many European labs and research groups are tightly interconnected, and are difficult to distinguish one from the other based on the co-authorship patterns. Furthermore, not all evidence of collaboration is represented in the co-authorship network; for example, the IMIRSEL lab appears relatively isolated, despite their central role in organizing MIREX. Large, intercontinental, multi-institutional grant projects, such as the Networked Environment for Music Analysis (NEMA) project [7], may start to change the shape of collaboration within the ISMIR community.

3. RESEARCH TOPICS IN ISMIR

The topics explored in the first ISMIR conference laid the foundation for the future growth and evolution of the field. While ISMIR has grown, it has remained true to the original vision laid out in the early conference programs.

In this section, we present an analysis of terms extracted from the titles and abstracts of ISMIR papers. Only title and abstract terms were used as these represent concise summaries of the papers' content.

3.1 The Most Commonly Used Title Terms

In order to get an idea as to which research areas have been of interest over the past nine years, we analyzed the title terms of all peer-reviewed papers and posters in the ISMIR proceedings. All the terms from the titles of the papers and posters were extracted. The words were stemmed using a Perl-based implementation of the Porter stemming algorithm [8], and stop-words were removed using a combination of a standard list of common-usage English-language words, with the stop-word "music", as this term appears in almost all titles in the ISMIR proceedings. Table 2 shows the top terms that appeared in the publication titles for each year. New terms entering the top-ranked lists are highlighted in bold-face.

From the table, we can observe that the most often used title terms were relatively similar for each year; however, it is possible to identify certain trends. For instance, there was a strong interest in query by singing/humming systems in 2002 and 2003 shown by the title term query ("queri") appearing only in the lists of these two years. Research interest in musical genres increased in 2005 and 2006, and interest in music similarity research peaked in 2006. Interest in classification and modeling has been consistent over the past nine years. Additionally, the consistently high rank of the term "audio" suggests that ISMIR researchers have been focused primarily on audio rather than symbolic representations.

What is also evident from the title terms, is how closely the field has stuck to the original framing of MIR as represented in the 2000 ISMIR program. The core concepts have remained prevalent throughout the past dec-

2000	2001	2002	2003	2004	2005	2006	2007	2008
Retriev	Retriev	Retriev	Retriev	Audio	Audio	Audio	Audio	Audio
Inform	Inform	Audio	Automat	Retriev	Retriev	Similar	Retriev	Featur
Model	System	Inform	Model	Automat	Classif	Classif	Similar	Retriev
System	Audio	Queri	Similar	System	Featur	Model	Model	Model
Audio	Approach	System	Database	Classif	Inform	Genr	System	Analysi
Classif	Model	Automat	Audio	Polyphon	Model	Automat	Recognit	Automat
Polyphon	Analysi	Model	Inform	Pattern	Polyphon	Feature	Polyphon	Song
Segment	Similar	Polyphon	Queri	Inform	Extract	Approach	Featur	Inform
Instrument	Match	Similar	System	Extract	Similar	Perform	Analysi	Similar
Techniqu	MIR	Analysi	Classif	Featur	Algorithm	Retriev	Classif	Chord
Languag	Spot	Content		Sound	Genr	Evalu	Automat	Content
		Pattern		Tempo		Key	Approach	
		Voic					Evalu	
							Transcript	
							Algorithm	

Table 2. Top 10 ranked title terms of each year (w/ ties); new terms are highlighted in bold-face font.

ade, yet have accommodated expansion into new areas.

3.2 Title and Abstract Bi-grams

Single-term-concepts present a limited view of research concepts and topics, especially after subtle differences in terms are merged by stemming (e.g., ‘using’ and ‘users’ have the same stem, ‘us’, yet carry different connotations in usage). Furthermore, the limited text available in titles, only provides a glimpse at the complexity of concepts and ideas being researched and published. In order to get at more specific concepts which have taken the interest of ISMIR researchers, we extracted stemmed bi-grams (i.e., 2-word phrases) from the titles and abstracts of all papers and posters. Initially, we examined the bi-grams on a year-by-year basis, much as we did for single term concepts. However, as expected, the number of bi-grams exceeds the number of uni-grams, and the frequency with which any one bi-gram occurs is much lower. No meaningful or interesting patterns arose in the yearly analysis; however, when considered in aggregate, there is stronger evidence of dominant research topics within the field. Table 3 shows the top 20 most commonly used bi-grams in ISMIR proceedings over the last nine years.

Bi-gram (stemmed)	Count
inform_retriev	25
content_base	24
genr_classif	14
web_base	9
hidden_markov	9
queri_hum	9
polyphon_audio	8
real_time	7
system_base	7
optic_recognit	7
audio_featur	7
ground_truth	7
base_similar	6
featur_extract	6
playlist_gener	6
audio_fingerprint	6
sing_voic	6
retriev_system	6
automat_transcript	5
melod_similar	5
similar_measur	5
automat_genr	5

Table 3. Top 20 most commonly used bi-grams from titles and abstracts, reflecting the main research foci, methods, and approaches of the ISMIR community.

The most common bi-gram is “information retrieval”, followed by “content based”, “genre classification”, and so on. Beyond these, we can see the prevalence of the web, and web-based systems, which has paralleled the emergence of “web 2.0” and greater access to music and music systems online within the commercial sector. Although the frequencies of occurrence of some individual concepts are low, overall we find the topics represented

by the bi-gram analysis to be fairly representative of the major research interests in the field: such as “music similarity”, “feature extraction”, and so on.

4. CITATION PATTERNS

Moving beyond terms and bi-grams as representations of research interests, the papers themselves published in the ISMIR proceedings serve as representations of research topics and areas, and references to them serve as a way of highlighting the prevailing research interests of the community. Weinstock [9] outlines 15 motivations for why academics cite each other in scholarly writing including paying homage to pioneers, giving credit for related work, and so on. We examined the references lists of all peer-reviewed ISMIR papers and posters, and looked for references to other peer-reviewed ISMIR papers and posters. We did not consider references to demos, invited talks, tutorials, MIREX abstracts, or workshop papers. We also did not attempt to measure references to publications outside the ISMIR proceedings, nor did we attempt to gauge the number of citations to ISMIR papers from outside.

First, we shall outline and describe the general citation behavior of the ISMIR community. Figure 3 shows the frequency distribution of publications by the number of references to other ISMIR publications they contain. Most ISMIR papers and posters (nearly 50%) do not reference any other ISMIR publications. The average number of ISMIR references per paper/poster was 1.278 with the standard deviation of 2.05 and the maximum of 27.




Figure 3. Number of ISMIR references in ISMIR papers

The reasons for low internal referencing within the ISMIR community may be due to the fact that some authors preferentially cite journals, books, and theses which are extensions of, or refinements of ideas initially published in ISMIR over the ISMIR publications. Other possible explanations include the fact that ISMIR proceedings are not indexed in commonly used digital library portals, such as the ACM Digital Library or CiteSeer, and are inconsistently indexed by Google Scholar. The fantastic resource on <http://www.ismir.net/>, which has been developed and maintained by Michael Fingerhut, contains a near-complete set of the full-text versions of all papers and posters published in the ISMIR proceed-

ings; however, it lacks full-text search capabilities itself, and the site does not provide complete, standardized metadata records which may improve the visibility of ISMIR papers in search engines, and digital library portals.

Author/Title	# Refs
Goto, M., et al. (2002). <i>RWC Music Database: Popular, Classical and Jazz Music Databases</i>	21
Bello, J. & Pickens, J. (2005). <i>A Robust Mid-Level Representation for Harmonic Content in Music Signals</i>	13
Tzanetakis, G., Essl, G. & Cook, P. (2001). <i>Automatic Musical Genre Classification of Audio Signals</i>	13
Aucouturier, J. & Pachet, F. (2002). <i>Music Similarity Measures: What's the use?</i>	13
Sheh, A. & Ellis, D. (2003). <i>Chord segmentation and recognition using EM-trained hidden markov models</i>	12
Pampalk, E., Dixon, S. & Widmer, G. (2003). <i>Exploring music collections by browsing different views</i>	11
Paulus, J. & Kalpuri, A. (2002). <i>Measuring the similarity of Rhythmic Patterns</i>	11
Goto, M., et al. (2003). <i>RWC Music Database: Music genre database and musical instrument sound database</i>	10
Clausen, M., et al. (2000). <i>PROMS: A Web-based Tool for Searching in Polyphonic Music</i>	9
Ellis, D., et al. (2002). <i>The Quest for Ground Truth in Musical Artist Similarity</i>	8
Logan, B. (2000). <i>Mel Frequency Cepstral Coefficients for Music Modeling</i>	8
Birmingham, W., et al. (2001). <i>MUSART: Music Retrieval Via Aural Queries</i>	8
Logan, B. (2004). <i>Music Recommendation from Song Sets</i>	8
Abdallah, S. & Plumbley, M. (2004). <i>Polyphonic transcription by non-negative sparse coding of power spectra</i>	7
Foote, J., Cooper, M. & Nam, U. (2002). <i>Audio Retrieval by Rhythmic Similarity</i>	7
Mazzoni, D. & Dannenberg, R. (2001). <i>Melody Matching Directly from Audio</i>	7
Vinet, H., Herrera-Boyer, P. & Pachet, F. (2002). <i>The CUIDADO Project</i>	7
Soulez, F., Rodet, X. & Scharwz, D. (2003). <i>Improving polyphonic and poly-instrumental music to score alignment</i>	7
Whitman, B. & Ellis, D. (2004). <i>Automatic Record Reviews</i>	7
Whitman, B. & Smaragdis, P. (2002). <i>Combining Musical and Cultural Features for Intelligent Style Detection</i>	7

Table 4. Top cited papers and posters (excluding self-citation).

Working with the references we were able to extract, we filtered self-citations, which we defined as a reference to a paper in which an author of the citing paper is an author on the referenced paper. Table 4 shows the top cited papers and posters in the ISMIR proceedings, ranked by the number of references we were able to find to each.

Among the top cited papers and posters, there is a diversity of topics and publications, from which we may infer a range of motivations. The most cited publication in the ISMIR proceedings is Goto, et al.'s 2002 poster

introducing the RWC database, garnering 21 references. Following Weinstock's taxonomy of citer motivation, the referencing of a data set is most like motivation three: identifying methodology, equipment, etc. The lack of standardized data sets with ground truth is a recurring problem in the MIR community and the RWC database has served as a valuable resource for MIR researchers, as it acts as a de facto standardized collection on which to build and evaluate systems. In fact, the presence of Goto, et al., 2003, and Ellis, et al., 2002 on this list reiterate the importance of standardized data sets with ground truth within MIR research.

There are several other methodological references, including references to Logan (2000), Tzanetakis, et al. (2001), Sheh & Ellis (2003), Goto, et al. (2003). There are also elements of "paying homage" in the references to several papers, especially the seminal work of Beth Logan, who introduced MFCCs to the MIR community.

Author	Ref. Count	Co-author Count	Paper/Poster Count
Goto, M	43	24	21
Ellis, D P W	41	12	12
Hashiguchi, H	34	5	3
Nishimura, T	34	5	3
Oka, R	34	5	3
Widmer, G	34	11	19
Dannenberg, R B	29	15	10
Logan, B	29	4	5
Whitman, B	28	6	5
Downie, J S	26	15	25
Pampalk, E	26	11	12
Tzanetakis, G	24	27	15
Birmingham, W P	23	11	7
Pachet, F	22	12	13
Dixon, S	22	9	9
Meek, C	22	10	5
Pickens, J	21	7	6
Pauws, S	19	6	8
Cook, P	19	7	6
Fujinaga, I	19	31	28

Table 5. Top 20 cited authors (excluding self references).

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. Regardless, the most referenced works comprise a diversity of topics and areas which span the breadth of research within MIR, including references to signal-processing algorithms and methods as well as techniques for handling symbolic representations of music. There are papers covering music transcription, and rhythm analysis, as well as high-level tasks such as genre-classification, search and recommendation algorithms, and approaches to understanding audio similarity.

Table 5 shows the top 20 cited authors excluding self references. The second column shows the count of co-authors each of these authors have in ISMIR proceedings and the third column shows the count of papers/posters each author published. The most heavily cited author was

Masataka Goto with 43 references by other ISMIR authors. Among these top-cited authors, we can see there are those who have many references, in part because they have published many papers (e.g., Goto; Widmer), and there are authors who are highly cited, but have only a few publications (e.g., Logan; Whitman). There is, however, no correlation ($r=0.021$) between reference count and paper count, indicating that the referencing of authors is not merely a product of their productivity within the community. It is worth noting that among the top-cited authors, there is a strong correlation between the number of co-authors an author has, and the number of papers he/she has written ($r=0.815$). This correlation is not that surprising given our findings from section 2 where we discussed the trend towards collaboration and co-authorship among ISMIR authors.

5. CONCLUSION

The ISMIR community has grown significantly, and through the contributions of nearly 900 researchers, the field of Music Information Retrieval has been well-defined and established. The community is a tightly-knit one, with a high-degree of collaboration and co-authorship, focused around a core set of research topics and areas.

The main insights of our analyses can be summarized as follows:

- 1) The ISMIR community is becoming more collaborative as shown by increasing co-authorship;
- 2) The role of research labs is growing in the ISMIR community as they promote collaboration and increased participation in research;
- 3) The focus of research has mainly been on audio so far as revealed by the most commonly used title and abstract terms;
- 4) The most cited works in the ISMIR proceedings comprise a variety of topics, but primarily point to datasets, techniques, and methods;

In their early ISMIR paper discussing the interdisciplinary communities and research issues, Futrelle and Downie [2] lists several key research areas in MIR. Among these, our analyses show that areas such as feature detection and classification/machine learning have been the major topics represented to date in the ISMIR proceedings, whereas topics such as user studies, metadata, work on symbolic representations, and epistemology/ontology have not been as well represented as others. Our advice for the sustained, future growth of the ISMIR community is to encourage greater activity in these areas, as they are relatively uncrowded, open topics of research in which great advances can be made.

We would like to continue our informetric analysis of MIR research, and there are several aspects that can be further analyzed to obtain a broader picture of MIR. One area in which we could improve our understanding of the domain, is to include external sources and references in

our citation analysis, and track the number of ISMIR references found in other related journals and proceedings, references that are not from ISMIR proceedings and so on. Additionally, we explored several clustering analyses in researching this paper, and none provided immediately compelling results. We would like to continue to explore how papers, authors, and research topics cluster based on semantic similarity, co-authorship patterns, citation patterns, and bibliographic coupling.

6. ACKNOWLEDGEMENTS

We would like to acknowledge the work of Michael Fingerhut in maintaining the ISMIR proceedings database (<http://ismir.net/>), which was our primary resource for locating all of the publications, and was vital to verifying our data integrity; we are very thankful for his diligence. We also thank the Andrew W. Mellon Foundation for its financial support.

7. REFERENCES

- [1] ISMIR 2000: International Symposium on Music Information Retrieval. Available at: <http://ciir.cs.umass.edu/music2000/>
- [2] J. Futrelle, J. S. Downie: "Interdisciplinary Communities and Research Issues in Music Information Retrieval," *Proceedings of the International Symposium on Music Information Retrieval*, pp. 215-221, 2002.
- [3] D. Byrd, T. Crawford: "Problems of Music Information Retrieval in the Real World," *Information Processing and Management* 38, pp. 249-272, 2001.
- [4] J. H. Lee: *Analysis of Information Features in Natural Language Queries for Music Information Retrieval: Use Patterns and Accuracy*. University of Illinois, Ph.D. thesis, 2008.
- [5] N. Orio: "Music Retrieval," *Foundations and Trends in Information Retrieval* 1(1), pp. 1-90, 2006.
- [6] V. Batagelj, A. Mrvar: *Pajek - Analysis and Visualization of Large Networks* in Jünger, M., Mutzel, P., (Eds.) *Graph Drawing Software*. Springer, Berlin, pp. 77-103, 2003.
- [7] Networked Environment for Music Analysis (NEMA). Available at: <http://nema.lis.uiuc.edu/>
- [8] M. Porter: An algorithm for suffix stripping in Sparck-Jones, K.; Willett, P. (Eds.): *Readings in Information Retrieval*. Morgan Kaufmann Publishers, pp. 313-316, 1997.
- [9] M. Weinstock: "Citation Indexes," *Encyclopedia of Library and Information Science* 5. Marcel Dekker, New York, 1971.