

THE ISMIR EXPLORER – A VISUAL INTERFACE FOR EXPLORING 20 YEARS OF ISMIR PUBLICATIONS

Thomas Low¹ Christian Hentschel² Sayantan Polley¹ Anustup Das¹
Harald Sack³ Andreas Nürnberger¹ Sebastian Stober¹

¹ Faculty of Computer Science, Otto von Guericke University Magdeburg, Germany

² Hasso Plattner Institute for IT Systems Engineering, University of Potsdam, Germany

³ FIZ Karlsruhe – Leibniz Institute for Information Infrastructure, Karlsruhe, Germany

stober@ovgu.de

ABSTRACT

Ever since the first International Symposium on Music Information Retrieval in 2000, the proceedings have been made publicly available to interested researchers. After 20 years of annual conferences and workshops, this number has grown to an impressive amount of almost 2,000 papers. When restricted to linear search and retrieval in a document collection of this size, it becomes inherently hard to identify topics, related work and trends in scientific research. Therefore, this paper presents and evaluates a map-based user interface for exploring 20 years of ISMIR publications. The interface visualizes k -nearest neighbor subsets of semantically similar papers. Users may jump from one neighborhood to the next by selecting another paper from the current subset. Through animated transitions between local k -nn maps, the interface creates the impression of panning a large global map. Evaluation results of a small user study suggest that users are able to discover interesting links between papers. Due to its generic approach, the interface is easily applicable to other document collections as well. The search interface and its source code are made publicly available.¹

1. INTRODUCTION

Music Information Retrieval (MIR) has been a steadily growing field of research as documented by the official statistics from the International Society of Music Information Retrieval (ISMIR).² For the 20th anniversary of the annual ISMIR conference, the number of published papers at this event is expected to exceed 2,000. This work aims to make this great resource more accessible for researchers who would like to familiarize themselves with the field of MIR.

The classical search interface based on keyword search and result lists is well suited for users who know what they are

¹ <https://doi.org/10.6084/m9.figshare.8342594>

² <https://www.ismir.net/stats.html>



© Thomas Low, Christian Hentschel, Sayantan Polley, Anustup Das, Harald Sack, Andreas Nürnberger and Sebastian Stober. Licensed under a Creative Commons Attribution 4.0 International License (CC BY 4.0). **Attribution:** Thomas Low, Christian Hentschel, Sayantan Polley, Anustup Das, Harald Sack, Andreas Nürnberger and Sebastian Stober. “The ISMIR Explorer – A Visual Interface for Exploring 20 Years of ISMIR Publications”, 20th International Society for Music Information Retrieval Conference, Delft, The Netherlands, 2019.

looking for. However, in other scenarios, this might not be the ideal search approach and the amount and diversity of papers and topics can be overwhelming. Exploratory information retrieval systems support users with vague or evolving information needs by providing special user interfaces and interaction models. The ISMIR Cloud Browser³ [4], last updated in 2013, allows browsing through links of related terms. Our new web-based search interface, the ISMIR Paper Explorer, takes a different – more visual – approach to exploration using similarity-based two-dimensional maps. It can be accessed at <https://ismir-explorer.ai.ovgu.de>.

Section 2 covers related and prior research that led to or inspired our approach. Details on the underlying dataset and preprocessing are provided in Section 3. We then describe the user interface in Section 4 and present the results from a usability study of our first prototype in Section 5. Furthermore, we also provide insights gained from analyzing the cumulative ISMIR proceedings to detect trends like shifts in topics and methods, and demonstrate whether and how these findings show in the visual search interface (Section 6). Finally, Section 7 offers conclusions.

2. RELATED WORK

In contrast to keyword-based search or recommendation, the goal of exploration is not only to present relevant or related items, but to allow the user to learn about the information space and its properties. This is often desirable when dealing with large document collections and when an overview over a collection is preferred over a pinpoint search.

A common problem when designing interfaces for exploratory search is the typically very high-dimensional feature space used to describe the contents of the respective documents. Considering the task of exploring text documents such as scientific papers, the document contents is commonly expressed as a (weighted) vector of term frequencies (tf-idf). Depending on the selected vocabulary, these often gather up to thousands of dimensions, which are hard to be displayed on a common computer display. Thus, map-based visualizations of inter-document similarities requires projection of these features into a two-dimensional display space. The neighborhood relations in high-dimensional space, however, should be

³ <http://dc.ofai.at/browser>

preserved throughout the projection so that similarity can be easily perceived by the user.

In order to reduce the dimensionality of the feature spaces, different dimension reduction techniques have been proposed (a comprehensive survey is given in [2]). Popular approaches for dimensionality reduction are Self-Organizing Maps (SOMs) [6], Principle Component Analysis (PCA) [5] and Multidimensional Scaling (MDS) techniques [8].

Inevitably, any reduction of dimensionality will also reduce the information conveyed. None of the aforementioned methods is capable of fully preserving the structure of the collection (unless the collection already has a 2-D structure). A projection will cause errors in a sense that documents, whose feature vectors are very close in high-dimensional space might be projected at large distances from each other and documents that are very dissimilar are placed next to each other respectively. Most dimension reduction techniques try to minimize this error (e.g. stress optimization in MDS), however, a complete resolution is impossible due to the nature of the error. In [13] the authors reviewed and compared different dimensionality reduction algorithms for the visualization of music collections. Based on a user study, MDS was favored as best layout algorithm when the collection undergoes changes due to newly added items.

Although applied in the context of image retrieval, Rubner et al. [11] were among the first to propose multidimensional scaling for iterated search in large document collections. The authors suggest to compute a *local* MDS on the top ten best matching images for a given query. Follow-up queries are emitted by using a representative image of a specific area in the MDS as the next query. A *global* MDS is suggested by the authors to provide a broad overview of the entire collection. While this global approach gave good results for rather small sets of images (the authors tested with a 2-D map of 500 images) it can be easily seen that very large collections (such as the collection of all ISMIR papers) cannot be reasonably presented at a time. Moreover, the number of projection errors will increase with increasing number of documents to be displayed.

A different approach to local MDS is given in [14]. The authors discuss the problems of dimension reduction and introduce the measures of *trustworthiness* (closest neighbors of a document in display space are also neighbors in original space) and *continuity* (all proximities in original space are visualized in display space). The proposed method to local multidimensional scaling is different from the one suggested by [11] in that it allows to parameterize and adjust the compromise between trustworthiness and continuity using an adapted cost function. The MDS is computed on varying neighborhood sizes and trustworthiness as well as continuity are reported to decrease with increasing neighborhood size. In a subsequent work the authors present an unsupervised and a supervised approach to optimize both measures [15].

As discussed, naturally, none of these approaches is capable of actually solving the problems of dimension reduction. In this paper, we therefore follow the idea of reducing the projection error itself rather than trying to attenuate the errors that occur from dimension reduction. Similar to [11] we compute a local MDS on a small subset of the entire collection. Our approach followed in this paper takes these ideas one

step further by computing a significantly larger virtual map using MDS of which only a small subset is actually displayed. By selecting one of the displayed documents as follow-up query document, the user can navigate through the collection. Thus, we increase the overlap of any two consecutive maps and thereby increase continuity throughout the navigation process. Furthermore, as suggested in [13] we use Procrustes analysis [3] to better align newly generated maps with their respective predecessors, which leads to a reduced confusion of users when navigating with aligned maps.

3. DATASET PREPARATION

We collected all accessible papers from the years 2000 to 2018⁴ by scraping the respective conference websites, the DBLP index⁵ and the recently established Zenodo repository.⁶ From a total of 1822 URLs, 28 were broken. We were able to retrieve 16 of these papers through a manual web search. 12 papers are still missing. Meta-data such as paper title and authors, publication year and abstract was retrieved from DBLP. In order to obtain paper citation counts, Google Scholar⁷ was queried.

All papers were pre-processed in order to extract the plain text from the originating PDFs (where possible). For several older papers, we had to apply OCR techniques because they only contained the text as image or used obscure font encodings – possibly to hamper copying text from the paper. We experimented with several text extraction tools and eventually settled for *pdftotext* as provided by the Xpdf PDF viewer and toolkit⁸ which consistently provided good extraction results. 3 PDFs turned out to be corrupted, leaving 1807 paper to be indexed.

As the meta-data is missing abstracts for most of the papers, we automatically extracted abstracts from the texts using simple heuristics. Furthermore, a preview image of the first page of each paper was extracted using *pdftopng* provided by Xpdf. For imports of future proceedings, scripts have been created to automate the process.

4. PROPOSED USER-INTERFACE FOR VISUAL BERRY-PICKING

In classic information retrieval, *berrypicking* refers to the process of finding relevant information through a series of modified search queries [1]. With each query the user inspects a new set of search results and gains a better understanding of the underlying document space as well as his or her information need. Due to this incremental learning process the user is able to formulate new search terms and continue searching.

The proposed user interface adapts this pattern from a list-based presentation of search results to a map-based visualization. Instead of sorting results by relevance with respect to search terms, they are arranged in a similarity preserving two-dimensional projection. Additionally, the burden of having to formulate search terms is substituted by a query-by-example

⁴ We aim to update the explorer until the 2019'th ISMIR conference in order to provide access to all 20 years of ISMIR publications.

⁵ <https://dblp.uni-trier.de/db/conf/ismir>

⁶ <https://zenodo.org/communities/ismir>

⁷ <https://scholar.google.com/>

⁸ <https://www.xpdfreader.com>

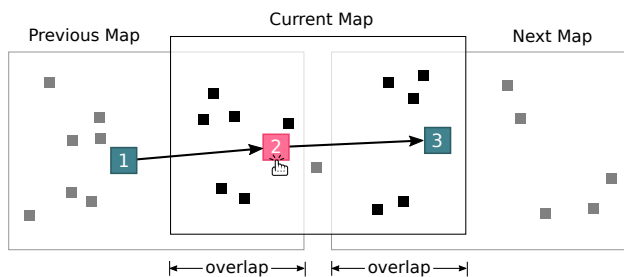


Figure 1. Similarity-based projection of nearest neighbors (squares) using three seed items (colored squares 1,2,3). Common neighbors (black squares) overlap between consecutive maps and are used for alignment when navigating from one item to the next.

search strategy. Such a user interface is hypothesized to be more in line with the goal of learning about the document space.

4.1 General Approach

Instead of trying to visualize the whole proceedings collection in a single similarity-based two-dimensional map, we visualize only the set of k -nearest neighbors for a given seed document in a small map [9]. This circumvents the challenge of the degrading projection quality (c.f. Section 2) and the increasing computational costs for on-the-fly map generation. By selecting a new seed document amongst the k -nearest neighbors, the user is able to hop from one neighborhood map to another. Consecutive maps are aligned as illustrated in Figure 1 to create a consistent transition that is, ideally, perceived as panning a large (global) map. Users are able to transfer knowledge about the content and the relevance of documents accumulated during the exploration process from one visualization to the next. This allows the user to navigate step-by-step through the whole collection.

4.2 Methods

We apply Multi-dimensional Scaling (MDS) [8] for dimensionality reduction. By limiting the number of documents used to compute the projection, we reduce the impact of projection errors and thus visualizations become more reasonable. Transitions between consecutive maps are animated. In order to make these transitions as consistent as possible, we align consecutive maps on their common neighbors. We use Procrustes analysis [3] to reduce the sum of the squared distances between the two sets of images that remain visible during the transition.

As preparation within the Python-based backend, all extracted texts are indexed using the Whoosh library.⁹ We also pre-compute pairwise document similarities (stored in a distance matrix) using the tf-idf vector space model [12] considering the 1,000 most frequent terms after applying a Porter Stemmer [10]. This pre-computation allows to quickly retrieve k -nearest neighbors and the corresponding distance matrix for the MDS projection.

The front-end is implemented in HTML5 based on the jQuery¹⁰ javascript library. Cookies are used to store pinned and highlighted papers as well as the user interface settings

(cf. Section 4.3). They also allow to track individual sessions for evaluation purposes.

4.3 User Interface

The user interface as shown in Figure 2 consists of three major components: a search bar on top that provides a keyword search interface and access to the settings for the visualization, the map visualization on the left and a details pane that displays the meta-data and the abstract or a preview image for the paper that is currently in focus.

A search can be started by clicking on one of the initially displayed papers or by typing query terms into the search bar. This will generate a map of the k most relevant ISMIR papers for the query. The matching query terms used to populate the initial view are highlighted in the presented information on the map as well as in the details pane. The parameter k is initially set to 30 but can be altered by the user. By clicking on any paper, it will gain focus (indicated by a blue border) and the map will change to display the k most similar ISMIR papers using this paper as new query. A paper can also gain temporary focus (without map update) by hovering the mouse pointer over it. As alternative to the default map visualization, the papers can be arranged in a grid layout that preserves similarities and, in addition, avoids overlaps, see Figure 6 (right).

On the map, each paper is visualized as a rectangle with the title and the last names of the first and last author inside. The fill color saturation visualizes the citation count whereas the border color highlights the current paper in focus (in strong blue) as well as the recent history of focused papers (in light blue). On hovering over a paper, functions for pinning and highlighting become available which are visualized by a landmark pin and a star respectively as explained in Figure 3. Both become permanently displayed in yellow if selected. Pinning keeps a paper on the map, even if it no longer belongs to the k -nearest neighbors of the paper in focus. This can be especially helpful when identifying related work and using the map to visualize similarity relations between the found papers. Highlighted papers are easy to spot because of the star and the additional yellow border color. This function can, for instance, be used to help orientation on the map. Furthermore, papers can be bookmarked using the bookmarking function built into the web browser. The browser's back and forward buttons allow to step through the search history.

The details pane on the right displays the meta-data of the paper in focus comprising the title, the full list of authors, the publication year, and the citation count at the top. Below are two tabs containing the abstract and a preview of the first page of the PDF respectively. Clicking on the title opens the PDF. There is further a link to the DBLP BibTeX record. The content of the details pane can be temporarily changed by hovering over a paper on the map. This allows to inspect paper details without having to focus the paper which would imply a map update.

A first user interface prototype with the functions described above was evaluated in a usability study as described in Section 5. Afterwards, several additional functions were introduced to support a historical analysis of the cumulative ISMIR proceedings for the 20th anniversary. To visualize the year of publication for papers displayed on the map, the fill color hue

⁹ <https://bitbucket.org/mchaput/whoosh/>

¹⁰ <https://jquery.com>

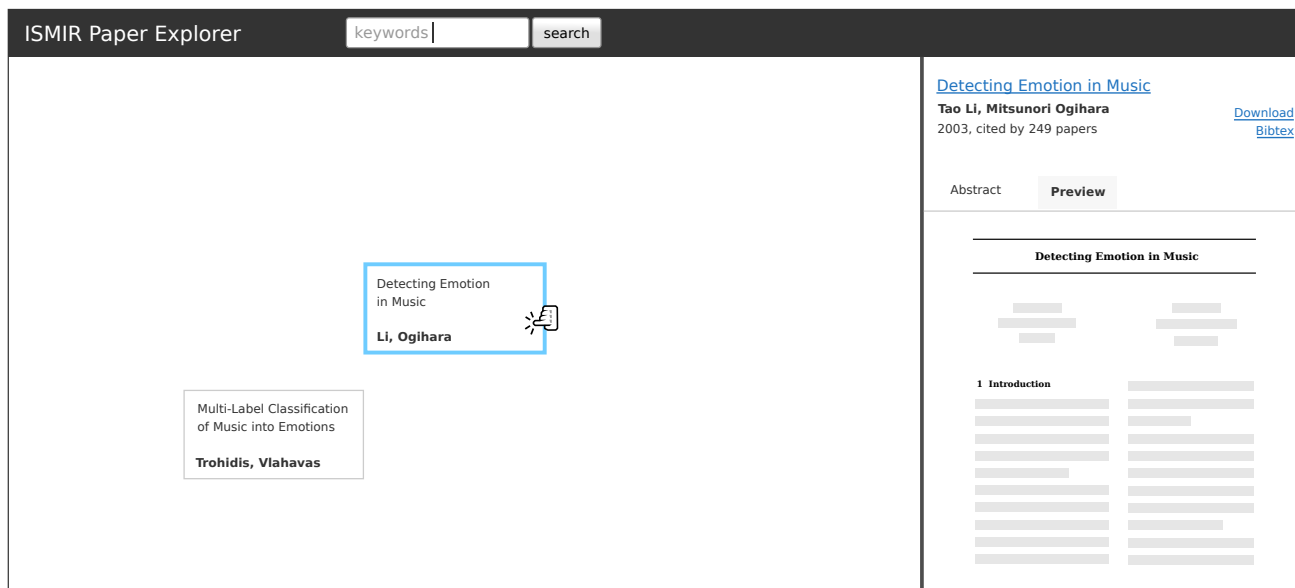


Figure 2. User interface schematic.




Figure 3. Explanation of the visual elements used in the map (accessible through a help button in the search bar at the top of the user interface).

was adapted. To this end, a continuous color map was divided into 20 steps representing the years. This is compatible with the original coloring to indicate the citation count as this only relies on the color saturation. Furthermore, instead of starting with a random paper in focus and its k -nearest neighbors, we now start with the k most cited papers which naturally cover a wide range of topics addressed in MIR. This makes for a nice starting point to explore the different sub-fields of MIR.

5. EVALUATION

Based on the prototypical implementation of the exploration interface described above, we conducted a small anonymous user study with the aim to evaluate user satisfaction with the presented design. Study participants were asked to compare the explorer interface with available traditional linear paper browsing interfaces for the ISMIR proceedings. Specifically, we asked about experiences using DBLP⁵, the Cumulative ISMIR Proceedings browser¹¹ as well as the ISMIR Cloud Browser³ [4]. Furthermore, participants were also free to mention tools apart from these. We were interested to see whether the prototype helped to identify related literature and discover additional related topics, which were the fundamental goals when designing the interface. We also wanted to know whether the prototype helped to efficiently find what the user was looking for.

Furthermore, we asked the participants about their overall satisfaction with regard to the proposed interface. Here, we

were mainly interested to see whether a fundamentally different approach to paper exploration is still considered as *easy to use*. A key aspect of the interface is the exploitation of the 2-D map to spatially arrange similar papers. We therefore wanted to know whether the spatial grouping helped to *infer relevance* for unknown papers and whether the user were able to effectively *refine the topic focus* during navigation. Finally, as mentioned above, applying the MDS to a restricted neighborhood may result in some confusion when changing the views. We were interested to understand whether the arrangement of papers during navigation was still perceived as *plausible*.

Finally, the users were able to enter additional comments and proposals for improvements in free text.

5.1 Evaluation Results

We received a total of 20 valid responses. Figure 4 presents the aggregated results of all responses. Out of the 20 participants, 13 users were experienced with some of the alternative search tools for finding relevant ISMIR papers. 2 participants mentioned Google or Google Scholar as alternative tools for browsing, which are both linear search and retrieval tools. For users, who were not experienced in any other alternative method for ISMIR paper exploration, we rated their answers regarding the comparison with the proposed prototype as neutral.

In general, the prototype was rated very positively by the study participants. We were pleased to see that a majority of the participants (65%) agreed that the proposed approach helped them to identify literature related to their topic of interest. An overwhelming majority of 90% of the participants furthermore agreed or strongly agreed to the claim that the interface was easy to use, which surprised us due to the fundamentally different browsing experience. Furthermore, according to the results of the study, the suggested spatial alignment helped 70% of the participants to infer the relevance of unknown papers based on the grouping of the neighborhood. As could have been expected, a small number of participants (20%) were not satisfied with the arrangement of the papers

¹¹ <http://www.ismir.net/proceedings/index.php>

Compared to [...], the system better helped me to ...

- ... identify literature related to my topic of interest.
- ... discover additional related topics.
- ... efficiently find what I was looking for.

Overall, ...

- ... the interface was easy to use.
- ... the spatial grouping helped me to infer the relevance for unknown papers.
- ... I was able to effectively change (refine or extend) the topic focus through navigation.
- ... during navigation, changes in the arrangement of the papers were comprehensible/plausible.




Figure 4. Evaluation results presented in diverging stacked bar charts showing the percentage of participants rating the proposed interface according to the respective statements.

during navigation as it was not considered comprehensible.

A number of participants asked us to publicly release the source code for future development. A few other comments were made with suggestions to further improvements of the usability. One participant mentioned that it is currently not possible to retrieve a list of bookmarked papers, which is a valuable feature for a future release. A few users mentioned that a persistently visible tutorial or legend would help a novice to better use the interface. Finally, an issue was raised about the overlapping surrogates in the non-grid view, which make it sometimes hard to reach them with the mouse. While we encourage the use of the grid view in these cases, we are aware that the grid layout also breaks the spatial alignment of the papers.

6. EXPLORATORY SEARCH

To identify interesting trends in the dataset, we first followed a conventional approach. We binned the publications into 5-year spans and performed a basic bi-gram analysis on the paper abstracts. Table 1 shows the most frequent bi-grams for each bin. Next, we applied a simple heuristic proposed in [7] to detect the most popular methods by assuming that the name of the method appears after phrases like “based on” or “using.”

Having identified the most prominent methods, we counted all papers per year mentioning them. For this, we also considered variations like acronyms or slightly different ways of spelling. A plot for the most popular methods is shown in Figure 5 (top). The rise of deep learning since about 2012 cannot be overlooked here. The middle plot in Figure 5 shows the development for select MIR topics like query-by-humming (QBH) which was particularly popular in the early 2000s, optical music recognition (OMR) which has been of constant interest, the development of graphical user interfaces (GUIs) which peaked in 2004 and 2008 mainly driven by approaches based on self-organizing maps (SOMs), and music recommendation which gained massive popularity around the turn of the decade and has leveled off slightly since then. Finally, Figure 5 (bottom) shows the popularity of selected music features over time. Here, the increasing use of chroma features and the slow development of MIR research addressing emotions stand out. At the same time, melody had a very high initial popularity (likely due to a focus on symbolic music representations and QBH in early years) that it lost to some extent over time but may gain back thanks to recently sparked interest in generative models for music. This trend could also explain the emergence of “symbolic music” as a top-5 bi-gram in most recent paper abstracts.

Next, we wanted to see, whether we can make similar

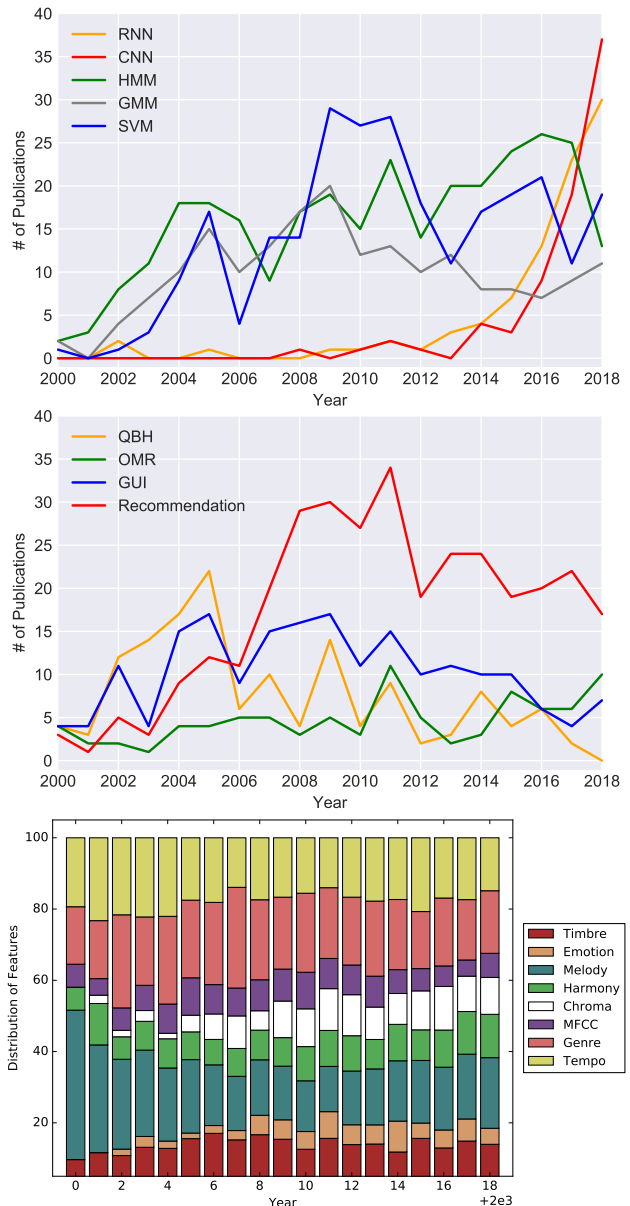


Figure 5. Examples for development over time. Top: popular machine learning techniques. Middle: selected MIR topics. Bottom: selected aspects/features of music.

observations with our exploratory search interface. Here, the same data is used in the back but only a small part of it is shown at a time to not overwhelm the user. Figure 6 (left) visualizes the top-30 papers for the query “query by humming.” Based on the dominating orange and green hues, most of

Table 1. Top-5 bi-grams for each 5-year bin. Numbers correspond to the frequency of each bi-gram within paper abstracts.

2000–2004	2005–2009	2010–2014	2015–2018
content based 30	polyphonic music 30	matrix factorization 27	neural network 52
polyphonic music 26	Gaussian mixture 15	music structure 18	recurrent neural 33
similarity measures 13	music recommendation 10	interactively browse 15	convolutional neural 32
pattern matching 10	collaborative filtering 9	emotion recognition 15	time frequency 22
acoustic features 6	rhythmic information 7	automatic transcription 11	symbolic music 16

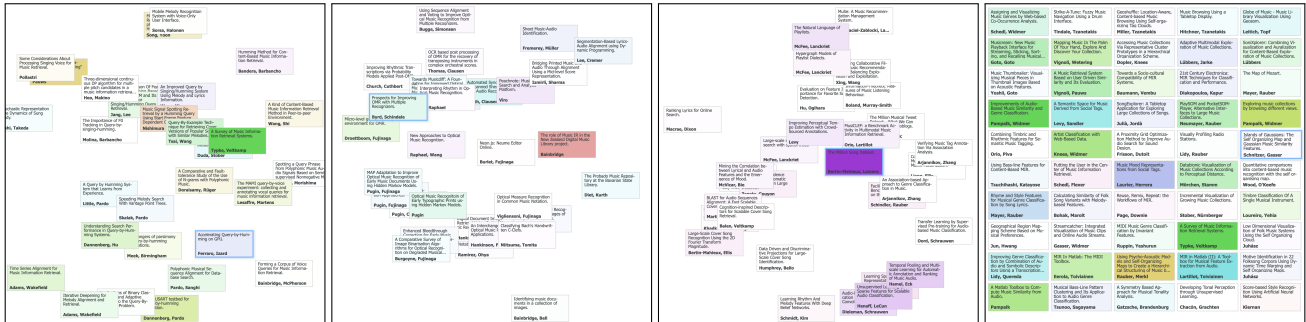


Figure 6. Map visualizations for selected initial search queries – from left to right: “query by humming”, “OMR”, “million song dataset”, “SOM”. All except the rightmost sub-figure show the top-30 results. The rightmost sub-figure shows the top-50 results for “SOM” in grid mode. The color saturation is proportional to the citation counts and the hue corresponds to the year starting with red in 2000 and continuing in equal steps through the spectrum.

the papers come from the early-to-mid 2000s with the by far most-cited paper being a survey of MIR systems from 2005. This matches the respective plot in Figure 5 (middle). The advantage in the view is that clusters can be identified like the aspects of singing in the top left, alignment in the lower left or system performance at the bottom.

For the topic OMR (second map from the left in Figure 6), there are some highly-cited papers around 2010 but there is also plenty of coverage in recent years which is in line with Figure 5 (middle). Recognizable clusters comprise, for instance, alignment (top right), applications in libraries (middle right) and old data (bottom).

The third map from the left in Figure 6 shows papers mentioning the Million Song Dataset. As can be seen easily, these are all fairly recent papers with the original publication standing out clearly due to its high citation count. This visualization nicely illustrates the different ways this dataset has been used in research since its publication. There are, for instance, clusters on cover songs (lower left), training deep neural networks (lower right), and on playlists and recommendation (top).

Finally, the rightmost map in Figure 6 shows the top-50 papers for the query “SOM.” They are displayed in grid mode to avoid clutter which starts to become an issue at this setting of k for typical screen resolutions. Nevertheless, the 50 papers can still be displayed comfortably in grid mode. This visualization covers almost all relevant papers for this topic. Here, the most-cited papers can be attributed to the period between 2005 and 2010 which was a time of high popularity for this technique being used in a large variety of graphical user interfaces. This also correlates with the two peaks around this time for GUI in Figure 5 (middle).

These are just a few examples about what can be discovered

by exploring the cumulative ISMIR proceedings using our proposed user interface.

7. CONCLUSION AND OUTLOOK

In this paper, we presented and evaluated a novel map-based user interface to explore the past 20 years of ISMIR publications. The proposed approach uses an MDS dimensionality reduction on a locally-restricted neighborhood of a query paper in order to project a high-dimensional feature space into 2-D while at the same time keeping projection distortions at a minimum. Consecutive maps are aligned in order to create a consistent impression of panning a global map.

Based on the interface, the user is able to find out a number of interesting topics in the ISMIR collection that would be hard to identify using a linear search interface. As an example, the topic “query by humming” almost disappeared in the early 2010’s after having peaked in the early 2000’s. While this fact would be hard to extract manually from a linear search interface, it is revealed at a single glance using the proposed prototype.

Besides updating the data to include recent ISMIR publication, one possibility for extension is to consider MIR-related papers not published at ISMIR. Respective candidates could be discovered by analyzing the reference sections. This would also allow an additional co-reference analysis which could enrich the interface further. We will also provide support to integrate the ISMIR Explorer into the official ISMIR website.

8. ACKNOWLEDGMENTS

We would like to thank all participants of the user study and the reviewers for their constructive feedback.

9. REFERENCES

- [1] Marcia J. Bates. The design of browsing and berrypicking techniques for the online search interface. *Online Information Review*, 13(5):407–424, 1989.
- [2] Imola K. Fodor. A survey of dimension reduction techniques. Technical Report UCRL-ID-148494, Center for Applied Scientific Computing, Lawrence Livermore National Laboratory, June 2002.
- [3] John C. Gower. Generalized procrustes analysis. *Psychometrika*, 40(1):33–51, 1975.
- [4] Maarten Grachten, Markus Schedl, Tim Pohle, and Gerhard Widmer. The ISMIR cloud: A decade of ISMIR conferences at your fingertips. In *Proceedings of the 10th International Society for Music Information Retrieval Conference, ISMIR'09*, pages 63–68, 2009.
- [5] I. T. Jolliffe. *Principal component analysis*. Springer Verlag, 1986.
- [6] Teuvo Kohonen. Self-organized formation of topologically correct feature maps. *Biological Cybernetics*, 43(1):59–69, 1982.
- [7] Tomoki Kondo, Hidetsugu Nanba, Toshiyuki Takezawa, and Manabu Okumura. Technical Trend Analysis by Analyzing Research Papers Titles. In Zygmunt Vetulani, editor, *Human Language Technology. Challenges for Computer Science and Linguistics*, number 6562 in Lecture Notes in Computer Science, pages 512–521. Springer Berlin Heidelberg, 2011.
- [8] Joseph B. Kruskal. Multidimensional scaling by optimizing goodness of fit to a nonmetric hypothesis. *Psychometrika*, 29:1–27, 1964.
- [9] Thomas Low, Christian Hentschel, Sebastian Stober, Harald Sack, and Andreas Nürnberger. Visual berrypicking in large image collections. In *Proceedings of the 8th Nordic Conference on Human-Computer Interaction: Fun, Fast, Foundational (NordiCHI'14)*, pages 1043–1046, 2014.
- [10] Martin F. Porter. An algorithm for suffix stripping. *Program*, 14(3):130–137, 1980.
- [11] Yossi Rubner, Leonidas J. Guibas, and Carlo Tomasi. The earth mover's distance, multi-dimensional scaling, and color-based image retrieval. *Proceedings of the ARPA Image Understanding Workshop*, pages 661–668, 1997.
- [12] Gerard Salton, Anita Wong, and Chung-Shu Yang. A vector space model for automatic indexing. *Communications of the ACM*, 18(11):613–620, 1975.
- [13] Sebastian Stober, Thomas Low, Tatiana Gossen, and Andreas Nürnberger. Incremental visualization of growing music collections. In *14th Intl. Conf. on Music Information Retrieval, ISMIR '13*, pages 433–438, 2013.
- [14] Jarkko Venna and Samuel Kaski. Local multidimensional scaling. *Neural Networks*, 19:889–899, 2006.
- [15] Jarkko Venna, Jaakko Peltonen, Kristian Nybo, Helena Aidos, and Samuel Kaski. Information retrieval perspective to nonlinear dimensionality reduction for data visualization. *Journal of Machine Learning Research*, 11:451–490, 2010.