

Training Researchers with the MOVING Platform

Iacopo Vagliano¹, Angela Fessl², Franziska Günther³, Thomas Köhler³,
Vasileios Mezaris³, Ahmed Saleh¹, Ansgar Scherp⁵, and Ilija Šimić²

¹ ZBW – Leibniz Information Centre for Economics, Germany
{i.vagliano,a.saleh}@zbw.eu

² Know Center, Austria
{afessl,isimic}@know-center.at

³ Technische Universität Dresden, Germany
{franziska.guenther1,thomas.koehler}@tu-dresden.de

⁴ Centre for Research and Technology Hellas, Greece
bmezaris@iti.gr

⁵ University of Stirling, Scotland, UK
ansgar.scherp@stir.ac.uk

Abstract. The MOVING platform enables its users to improve their information literacy by training how to exploit data and text mining methods in their daily research tasks. In this paper, we show how it can support researchers in various tasks, and we introduce its main features, such as text and video retrieval and processing, advanced visualizations, and the technologies to assist the learning process.

Keywords: Technology enhanced learning · Information retrieval · Text and video analysis · Recommender systems

1 Introduction

Researchers in academia and industry struggle in being aware of publications, patents, products, and also funding opportunities to properly address a new research topic because of the fast-growing amount of relevant content. Similarly, students need to learn to exploit a large and steadily increasing number of previous publications and other online educational resources (video lectures, tutorials, blogs, etc.). While nowadays information is often freely available in various formats, such as text and media, extracting useful knowledge from all the accessible educational resources requires a lot of time, and often researchers lack strategies and tools to effectively do so.

The MOVING platform⁶ aims to address these issues. It integrates various kinds of educational resources such as documents, videos, and social media data. Through the platform, users can search these resources and display the search results in different ways thanks to the advanced visualizations available. Users

⁶ <http://platform.moving-project.eu>

can improve their skills in information literacy through a curriculum designed by experts in learning, and the platform can track their progress in using its features and support their learning in various ways.

The paper is organized as follows. In Section 2, we briefly describe how young researchers can be trained through the MOVING platform. In Section 3, we introduce the platform, while in Section 4 its main features. We conclude the paper in Section 5.

2 Training Young Researchers

The MOVING platform can assist researchers in various ways. Firstly, it can support them in reviewing the state of the art on a research topic. Through its advanced search features, users can search for relevant topics, as well as refine their search. They can filter the results by authors, subject area, content type (videos, books, articles, website, social media posts), etc., and select the open-access publications. The advanced search allows users to search specific fields such as author, title, abstract and define boolean queries.

As the list of results can be very long, users can exploit a graph visualization to start progressively exploring the content, identify key authors and topics, as well as discover relations among topics, authors and documents. Users can personalize the ranking of the results based on their interest. By selecting keywords, they can focus on the topics more relevant to them, and they can assign different weights to the various keywords to give more or less importance to each topic. Tag clouds are available to obtain an overview of the results and charts can show the most frequent topics and sources.

Additionally, the platform helps researchers in finding funding opportunities and suitable partners for research projects. Users can search funding opportunities similarly to other documents. Using the community functionality, they can connect with other users and search them by their research area and skills. Finally, the learning environment enables users to access learning material on information literacy and data mining, such as tutorials and MOOCs.

3 The MOVING platform at a glance

The MOVING platform integrates various kinds of educational resources such as documents, videos, and social media data. Some of these resources are automatically harvested from the Web and social networks, while some datasets are manually integrated from Videlectures.net⁷, EconBiz⁸, and the Social Science Open Access Repository (SSOAR)⁹. Through the platform, users can search these resources and display the search results in different ways thanks to the

⁷ <http://videlectures.net/>

⁸ <https://www.econbiz.de/>

⁹ <https://www.gesis.org/ssoar/home/>

advanced visualizations available. They are also assisted in progressing with the information literacy curriculum and in more effectively using the platform.

The architecture of the platform is depicted in Figure 1. The crawlers automatically ingest data from the Web and store them into the index. Then, data processing techniques, including entity extraction, author disambiguation, data deduplication, and automatic concepts annotation, are applied to the data indexed. The data in the index can then be retrieved through the search engine. WevQuery [1] tracks the users behavior on the platform by capturing UI events which are saved in the user database (User data), while the Adaptive Training Support (ATS) [3] analyses these data to help users to better use the platform and progress in the selected curriculum, depending on their usage patterns. The recommender system builds users' profiles based on their search history and suggests documents according to such profiles. In Section 4, we describe the search engine, the ATS, the video processing techniques used, the advanced visualizations and the recommender system.

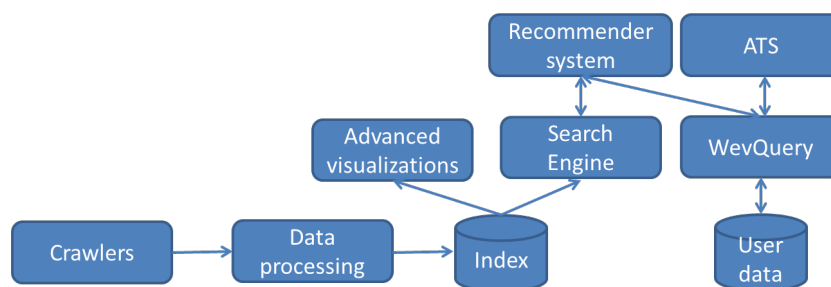


Fig. 1. The MOVING platform's architecture.

4 Main Features of the MOVING Platform

The MOVING Search Engine The MOVING search page (Figure 2) provides multimodal and faceted search and handles multiple document types. To retrieve relevant results, the search engine exploits various metadata such as title, abstract and also the concepts with which documents are annotated. For videos, these concepts are automatically assigned to video fragments. Specific fragments of video can be searched through an ad-hoc facet. In addition to more traditional filters such as author, date, subject area, and document type, users can also filter the open access publications. The search engine exploits HCF-IDF [8], a novel ranking and document profile method, which ranks the search results based on their relevance to the user query relying only on titles. This is key because in the MOVING platform document's titles and other metadata are usually available, but the platform stores a limited number of full texts due to copyright issues.

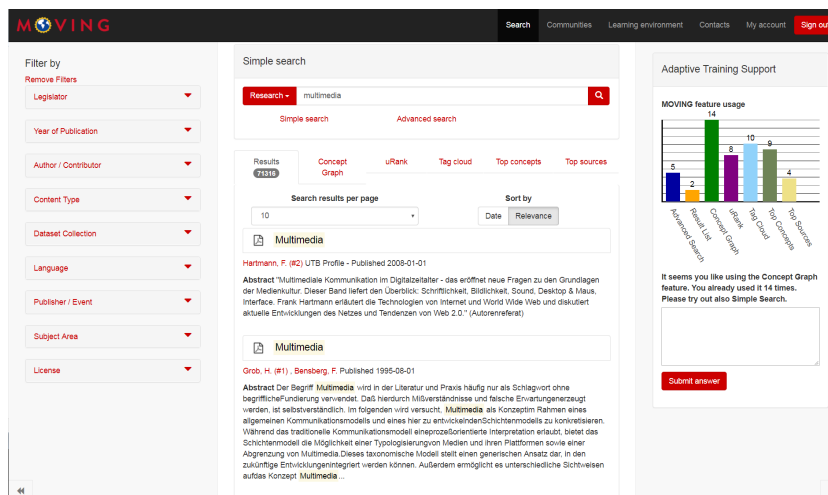


Fig. 2. The search page. On the left, the facets allow users to refine the search; in the middle results are shown; on the right the ATS summarizes the use of the various platform features and provides reflective prompts.

Adaptive Training Support The ATS [3] supports users to improve their use of the MOVING platform by mirroring back the user behavior in combination with reflective prompts. These prompts stimulate users to reflect on their experience in using some features and to try new features. Specifically, it tracks the use of the various features through WevQuery. The ATS also assists users in completing the selected curriculum, summarizing their progress in a progress bar, providing reflective questions related to their progress and suggesting activities to do next.

Video Processing We distinguish between two genres of video in the platform: lecture videos, and non-lecture videos. For both, we use suitable fragmentation and annotation approaches for enabling fragment-level access to the video content. Lecture video fragmentation is performed using a method previously developed [5]. This involves first processing the textual transcripts of the video (generated by an off-the-shelf automatic speech recognition component) to extract *noun phrases*; and then using Word2Vec, a neural-network-based on word embeddings, to represent every word of every noun phrase in a semantic vector space. Sliding temporal windows in the transcripts define larger parts of the text, which are represented in the same semantic space by averaging the aforementioned vector representations of the corresponding words; and the similarity of such windows is assessed using the cosine similarity. Doing so for all pairs of successive sliding windows in a transcript results in a 1D signal. Its local minima (if they also satisfy certain other criteria) are selected as lecture fragment boundaries. Finally, the annotation of each fragment is performed by selecting the most prominent words from the transcripts of each fragment, based on TF-IDF. Alternatively, we can associate each fragment with high-level concepts

from a pre-specified concept pool with an adaptation of a previous work [4]. Non-lecture videos are fragmented to video shots using the method of Apostolidis and Mezaris [2], which detects both abrupt and gradual transitions by appropriately assessing the visual similarity of neighboring frames of the video. Then, a representative keyframe is extracted from each shot and is annotated with concepts from a pre-specified concept pool [7]. This method is based on a deep learning architecture that exploits concept relations at two different levels to learn to more accurately detect the concepts in the video.

Advanced Visualizations In addition to the classical result lists, users can browse the results through alternative and more advanced interfaces. With the concept graph (Fig. 3) users can explore relations between documents, authors, institutions, and topics. Alternatively, uRank (Fig. 4) enables users to personalize the results list by selecting the keywords more relevant to them and assigning to these keywords different weights. Every time a keyword is added or removed as well a weight changed, the results are re-ranked and an information on the positions gained or lost by each result is provided.




Fig. 4. Results displayed in uRank

Fig. 3. An example of a concept graph

Recommender System HCF-IDF [8] is also used to provide recommendations by matching documents with a user profile. The MOVING recommender system builds users' profiles based on users' search history. It obtains the latter from the user data previously logged through WevQuery, and then it suggests documents based on such users profiles. In our case, the user profile refers to the terms users have previously searched for. It is a set of pairs term-weight, $\langle k_i, w_i \rangle$, where k_i is a term and w_i a weight. The weight of a term depends on how often and how recently the term has been searched for. More formally, the weight w_i of a term k_i is defined by $w_i = \alpha_t \cdot \frac{t}{T} + \alpha_h \cdot \frac{h}{H}$, where α_t is the time coefficient and α_h is the hits coefficient. The h term is the number of times a term has been searched by the user and H the total number of searches made by the user; the timestamp (t) of the last search for the term is normalized by the current time (T). We are

also considering to integrate more advanced recommendation techniques based on autoencoders [6] and Linked Data [9].

5 Conclusions

We showed how the MOVING platform may support researchers and its main features. The platform is still being improved and soon new features will be released, such as a more complete learning environment and additional community functionalities.

Acknowledgments This work was supported by the EU's Horizon 2020 programme under grant agreement H2020-693092 MOVING. The Know-Center is funded within the Austrian COMET Program under the auspices of the Austrian Federal Ministry of Transport, Innovation and Technology, the Austrian Federal Ministry of Economy, Family and Youth and by the State of Styria. COMET is managed by the Austrian Research Promotion Agency FFG.

References

1. Apaolaza, A., Vigo, M.: WevQuery: Testing hypotheses about web interaction patterns. Proc. ACM Hum.-Comput. Interact. **1**(EICS), 4:1–4:17 (Jun 2017)
2. Apostolidis, E., Mezaris, V.: Fast shot segmentation combining global and local visual descriptors. In: Proc. IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP) (May 2014)
3. Fessler, A., Wertner, A., Pammer-Schindler, V.: Digging for gold: Motivating users to explore alternative search interfaces. In: Lifelong Technology-Enhanced Learning. pp. 636–639. Springer International Publishing, Cham (2018)
4. Galanopoulos, D., Markatopoulou, F., Mezaris, V., Patras, I.: Concept language models and event-based concept number selection for zero-example event detection. In: Proc. ACM ICMR (2017)
5. Galanopoulos, D., Mezaris, V.: Temporal lecture video fragmentation using word embeddings. In: Proc. 25th Int. Conf. on Multimedia Modeling (MMM). Springer International Publishing (2019)
6. Galke, L., Mai, F., Vagliano, I., Scherp, A.: Multi-modal adversarial autoencoders for recommendations of citations and subject labels. In: Proceedings of the 26th Conference on User Modeling, Adaptation and Personalization. pp. 197–205. UMAP '18, ACM (2018)
7. Markatopoulou, F., Mezaris, V., Patras, I.: Implicit and explicit concept relations in deep neural networks for multi-label video/image annotation. IEEE Transactions on Circuits and Systems for Video Technology pp. 1–1 (2018)
8. Nishioka, C., Scherp, A.: Profiling vs. time vs. content: What does matter for top-k publication recommendation based on twitter profiles? In: Proceedings of the 16th ACM/IEEE-CS on Joint Conference on Digital Libraries. pp. 171–180. JCDL '16, ACM (2016)
9. Vagliano, I., Monti, D., Scherp, A., Morisio, M.: Content recommendation through semantic annotation of user reviews and linked data. In: Proceedings of the Knowledge Capture Conference. pp. 32:1–32:4. K-CAP 2017, ACM (2017)