



Trends in content-based recommendation

Preface to the special issue on Recommender systems based on rich item descriptions

Pasquale Lops¹ · Dietmar Jannach² · Cataldo Musto¹ · Toine Bogers³ · Marijn Koolen⁴

© Springer Nature B.V. 2019

1 Introduction

Automated recommendations have become a pervasive feature of our online user experience, and due to their practical importance, recommender systems also represent an active area of scientific research. Historically, the two main approaches of building recommender systems are collaborative filtering (CF) and content-based filtering (CBF). In its pure form, collaborative filtering relies on the identification of preference patterns in a user community. Content-based filtering approaches, in contrast, only consider the past preferences of an individual user and try to learn a preference model based on a feature-based representation of the content of recommendable items.

Due to the potential limitations of considering only the past preferences of an individual user, a number of proposals were made over the past two decades to combine both algorithmic approaches in hybrid methods. In more recent years, we further-

✉ Pasquale Lops
pasquale.lops@uniba.it

Dietmar Jannach
dietmar.jannach@aau.at

Cataldo Musto
cataldo.musto@uniba.it

Toine Bogers
toine@hum.aau.dk

Marijn Koolen
marijn.koolen@di.huc.knaw.nl

¹ University of Bari, Bari, Italy

² University of Klagenfurt, Klagenfurt, Austria

³ Aalborg University Copenhagen, Copenhagen, Denmark

⁴ Royal Netherlands Academy of Arts and Sciences, Amsterdam, The Netherlands

more observe various attempts to incorporate additional side information and external knowledge sources into the recommendation process. In the case of content-based methods, this side information predominantly contains additional knowledge about the recommendable items, e.g., in terms of their features, metadata, category assignments, relations to other items, user-provided tags and comments, or related textual or multimedia content.

While pure collaborative filtering methods dominate the research landscape (Jannach et al. 2012), considering information about the “content” (i.e., the features of the items) is highly important in many of today’s practical applications of recommenders, and will be increasingly relevant in the future for a number of reasons:

- First, a number of research works show that content-based and hybrid approaches, despite being sometimes less accurate in offline experiments, are in some application domains more effective in *real-world applications* than their pure CF-based counterparts (Jannach and Hegelich 2009; Kirshenbaum et al. 2012; Liu et al. 2010).
- Second, certain aspects of a recommendation system can only be implemented in an effective way when one considers knowledge about the features of the items. These aspects for example include the diversification of the resulting recommendation lists, the generation of explanations, or the creation of understandable user models.
- Third, we continue to see new knowledge sources that provide information about the items and which can be used in the recommendation process. These sources include both structured ones like DBPedia, semi-structured ones like Wikipedia, and unstructured ones. This latter class includes all sorts of user-generated content like reviews and tags, ultimately leading to a much richer item representations than those that were available in the past.

This special issue contains a selection of recent work on building recommender systems using rich item descriptions. We outline these papers in Sect. 3. Before that, in Sect. 2, we briefly reflect on the history of the field, discuss recent trends, and sketch potential future developments.

2 History, trends, and future of content-based recommenders

2.1 Historical developments

Some of the underlying ideas of content-based filtering go back to the 1960s and to early ideas of what was called “Selective Dissemination of Information” (Hensley 1963). Here, the goal was to distribute information based on matching newly arriving information items with the assumed interests of the recipients that are stored in *user profiles*. Another root of CB-filtering systems lies in the field of Information Retrieval (IR). Content-based approaches for example often rely on document encodings that were developed in this field (Salton and McGill 1986). In the Web era, content-based techniques were later successfully applied in different domains, e.g., to make personalized recommendations of interesting web pages (Pazzani et al. 1996).

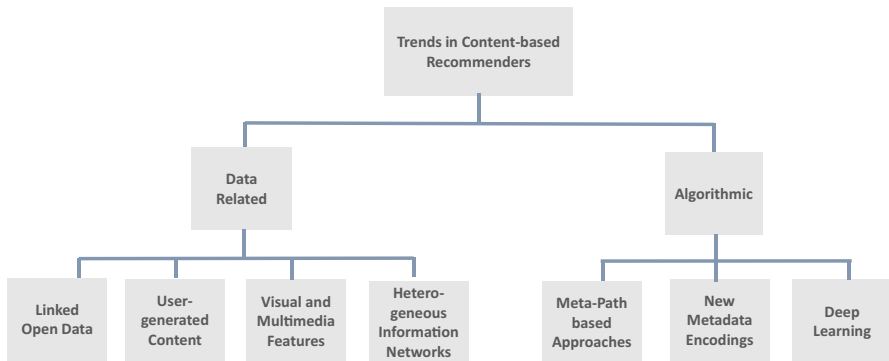


Fig. 1 Recent trends in content-based filtering recommender systems

Soon, however, it turned out that pure content-based filtering approaches can have several limitations in many application scenarios, in particular when compared to collaborative filtering systems. One main problem is that CBF systems mostly do not consider the *quality* of the items in the recommendation process. For example, a content-based recommendation in the movie domain might result in an obscure, low-quality movie similar to those the user liked in the past. CF methods usually take such information into account, either by considering explicit item ratings or by indirectly factoring in item popularity in implicit feedback scenarios.

With a decades-long focus of academic research on finding the best model in terms of prediction accuracy, *pure* CBF methods almost seemed obsolete and they were instead mostly used to boost the performance of collaborative filtering approaches, e.g., in cold-start situations (Melville et al. 2002). However, along with the availability of new knowledge sources—including both structured and unstructured ones, such as user-generated content—comes a steady stream of new proposals for hybrid systems that leverage such information to make better predictions. Furthermore, in some new application scenarios, it is actually important that the recommended items are actually content-wise similar to a reference item, e.g., for similar item recommendation (Yao and Harper 2018). Also, in an automated music radio station, content features can be used to ensure that the radio does not drift away too much from the seed tracks. Finally, content information also enables the generation of better *explanations* (Gedikli et al. 2014), which is becoming increasingly important in the context of fair and transparent recommender systems.

2.2 Recent trends

In the last decade, we could therefore observe a number of works that either use *new types of side information* (e.g., metadata or user-generated content) for recommendation or propose new *algorithmic approaches* to process the available information. Figure 1 shows an overview of these new trends.

2.2.1 Data-related trends

Linked open data Traditionally, content-based recommender systems used metadata descriptions of items, or full-text indexing of textual items (Belkin and Croft 1992). Today, the introduction of Linked Open Data (LOD) offers new ways to extend item descriptions with external knowledge sources (Di Noia et al. 2012; Musto et al. 2017b). Early work by Passant (2010), for example, used DBpedia as an external knowledge source for music recommendation. Since then, Linked Open Data was used to enrich item descriptions for many application domains, including movie and book recommendation (Musto et al. 2017a), and music and sound recommendation (Oramas et al. 2017). One advantage of the explicit semantics of Linked Data is that they can also be used to optimize specific measures like diversity (Musto et al. 2017a). Furthermore, Linked Data can serve as a basis to generate explanations and to thereby increase recommendation transparency (Musto et al. 2019). Overall, data from the LOD cloud were successfully applied for recommendations in several domains, and various opportunities exist to connect LOD information with other types of side information, as shown in Oramas et al. (2017) for music recommendation.

User-generated content (UGC) With the emergence of the participatory web (Web 2.0), various types of UGC became available, such as product reviews, user tags and forum discussions. Most of the early works that tried to leverage this information in the recommendation context focused on user-provided *tags*, which were used in a variety of ways, e.g., to enhance user and item profiles, or to explain recommendations to users (Jäschke et al. 2007; Vig et al. 2009).¹ In another stream of works, researchers focused on *user reviews* and tried to extract various types of information from them, including semantics or user opinions and sentiments, that can be used in the recommendation process. Like for the tag-based approaches, recommendation methods based on UGC not only leverage this side information to improve recommendation quality, but also as a means to explain the recommendations to users. An in-depth review of such approaches can be found in Chen et al. (2015).

Visual and multimedia features As CBF-methods were traditionally text-based, non-textual objects were commonly represented by metadata descriptions. Advances in image and video analysis however made it possible to represent multimedia object by features that were extracted from the objects themselves. Many of these features are in fact difficult to represent in text, e.g., textures or stylistic aspects. McAuley et al. (2015), for example, trained a Convolutional Neural Network on product images to learn how different visual feature dimensions relate to each other across different types of products. The resulting “style space” can then be used to recommend, e.g., trousers that go with a particular pair of shoes. Another example in that context is the work by Elahi et al. (2017), who extract low-level features such as colors, textures, motion and lighting from movies to build a hybrid recommendation system. In a related study Deldjoo et al. (2018), the authors found that these stylistic properties can lead to better recommendations than semantic (object) information that is extracted from the movies through a pre-trained neural network. For the future, we expect that visual

¹ See Bogers (2018) for an overview of tag-based recommendation approaches.

and multimedia features will be used for recommendation in many other domains and scenarios where such data is available, such as advertisement, e-commerce and games.

Heterogeneous information networks Finally, a number of recent works aim to connect various types of side information in parallel. In such approaches, the available data is represented as a Heterogeneous Information Network (HIN), where the semantic relationships between objects are represented through different types of relations between the network nodes. Early work on using social networks for recommendation focused on modelling trust (Golbeck 2006) between connected users. An example of a more recent work is that by Kleinerman et al. (2018) who leveraged such a network to generate explanations in the context of reciprocal recommendation scenarios like online dating. Generally, with the increasing number of knowledge sources that become available, HIN-based approaches will continue to gain importance, e.g., to create multi-dimensional user profiles (Jannach et al. 2017).

2.2.2 Algorithm-related trends

Meta-path based approaches From an algorithmic perspective, the availability of typed links between different types of objects in HINs inspired the development of *meta-path based* algorithms, where a meta-path connects two objects through a sequence of labeled relations between object types (Sun et al. 2011). Shi et al. (2015) for example demonstrated that meta-path-based algorithms are useful for HINs in many recommendation domains, e.g. for movie recommendation with types such as users, movies, actors, and directors or between customers, products, businesses and locations. In another approach, Yu et al. (2014) exploited meta-path-based latent features to represent the connectivity between users and items along different types of paths. From the implicit feedback on items, a user-specific weighting of the heterogeneous relationships between the items and other entities is learned to offer personalized recommendations. As HINs become richer and more ubiquitous, we expect that meta-path based approaches will receive more attention in the future.

New metadata encodings While TF-IDF or latent space representations are still common to encode textual item descriptions in CBF approaches, *embeddings* emerged in recent years as a promising new encoding method that can be used to identify latent connections between the terms in a document. Meta-Prod2Vec (Vasile et al. 2016), for instance, is an approach that computes low-dimensional embeddings of item metadata for sequence-based item recommendation in a hybrid model that uses music playlists and listening data as interactions for CF. Kula (2015), on the other hand, combined user and item metadata to train embeddings for generating recommendations in the fashion domain. In their application scenarios, the embeddings were specifically used to enable transfer learning to improve content-based recommendations for cold-start items and users. Given the flexibility of embeddings, we expect to see more complex models that combine metadata and textual and non-textual content features.

Deep learning Deep Learning (DL) approaches offer a flexible framework to discover existing structures in item and user data, including external knowledge resources (Musto et al. 2018) and non-textual features (Deldjoo et al. 2018), to learn better feature

representations. Another advantage of deep neural networks is their ability to model temporal and sequential aspects, which makes them suitable for session-based and sequence-aware recommendation tasks in domains such as e-commerce (Hidasi et al. 2016) and news (Song et al. 2016). In addition, they can be trained to generate rich explanations from user reviews (Lu et al. 2018). Zhang et al. (2017) recently demonstrated the potential of joint representation learning based on the increasing number of heterogeneous sources of textual and multimedia content features, external knowledge resources as well as user context and interaction data. Another recent research direction is to improve the interpretability of DL models for transparent and explainable recommendations (Seo et al. 2017b).

2.3 Future directions

We foresee a number of research directions for exploring the role of content, metadata and side information in recommender systems.

Transparent recommendations In recent years and with the use of increasingly complex machine learning models, we can observe a rising interest in the topic of *Explainable AI*. This is in particular the case in scenarios of complex human decision-making, where accountability and explainability are important properties, such as health. Explanations have been explored for several years in the context of recommender systems (Nunes and Jannach 2017; Tintarev and Masthoff 2008), and different types of side information have been used in existing explanation approaches. While existing approaches often use side information to highlight features of the recommended items (Gedikli et al. 2014; Vig et al. 2009), one area of future work could be to use side information to *explain the internal mechanisms* that guide otherwise black-box-like algorithms. This could, for instance, be done by “labeling” internal components of the derived models (e.g., a specific layer of a neural network or a latent factor of a factorized matrix), or by using descriptive features to explain the overall behavior of a recommendation algorithm. In this context, one open research question relates to the choice and possible combination of the features that are used in the explanation process, such as descriptive features (e.g., item properties and curated metadata), user-generated content (e.g., textual reviews and tags), collaborative information, and other sources.

New algorithms and tasks Content-based recommendation has, because of its reliance on complex textual content, traditionally been inspired by the developments in the fields of computational linguistics and natural language processing (NLP). For example, (word) embeddings are beginning to see widespread application in recommendation after its successful application in many NLP tasks. We expect this transfer of methods and techniques to continue in the future. For instance, contextual embeddings have recently been shown to outperform the state-of-the-art on certain NLP tasks (Seo et al. 2017a), and they could show promise for content-based recommendation algorithms as well. Besides such algorithmic improvements, another fruitful research direction would be to investigate the role of content information in alternative types of *recommendation tasks* (Jannach and Adomavicius 2016) besides standard item rank-

ing or ratings prediction. For instance, generating recommendations for related or similar items (Yao and Harper 2018), or recommending composites or sequences of items.

New types of content Some domains are rich in non-textual content, which is currently still difficult to utilize effectively, such as multimedia content. So far, multimedia recommendation mostly relies on implicit and explicit user feedback, but we expect that the rapid developments in DL algorithms will enable more structured and effective ways of extracting and using semantic and stylistic features in the future, as done in Messina et al. (2018) in this issue. Open research questions here include how and which features can be extracted from multimedia content, and how they should be integrated with user-generated content and user preference data to provide the most relevant recommendations. A more thorough understanding of cross-domain recommendation—how content from one domain can help recommendation in another—would also be important here; see also Hernández-Rubio et al. (2018) (this issue).

3 Papers in this issue

The papers in this issue continue many of the recent trends that were discussed above. They use various types of side information including User-Generated Content and Linked Open Data as well as features that are automatically derived from the multimedia objects themselves. From an algorithmic perspective, these papers often rely on deep learning approaches for feature extraction or recommendation. Various application domains are considered as well, including traditional ones like music or news recommendation, as well as novel ones like artwork recommendation.

The paper *Content-Based Artwork Recommendation: Integrating Painting Metadata with Neural and Manually-Engineered Visual Feature* (Messina et al. 2018) introduces a novel approach for artwork recommendation. The authors combine a variety of information sources—including item metadata as well as low-level visual features that are extracted through a deep neural network—and show the advantage of their combined approach based on an experiment with real-world data provided by an online artwork store.

The paper *Movie Genome: Alleviating New Item Cold Start in Movie Recommendation* (Deldjoo et al. 2019) adopts a similar ideas and combines a variety of features, both meta-data as well low-level visual and audio features, in a recommender system for movies. They furthermore combine their content-based model with collaborative information in a hybrid approach. Both an offline analysis as well as a preliminary user study confirm the usefulness of their model for new item cold start situations.

The work *Affective recommender systems in online news industry: how emotions influence reading choices* (Mizgajski and Morzy 2018) studies the role of emotions in the recommendation process. Based on a set of *affective item features*, a multi-dimensional model of emotions for news item recommendation is proposed. Experiments revealed the potential of considering emotional reactions and in their case, the enhanced recommendation method led to higher click-through-rates.

The paper *A comparative analysis of recommender systems based on item aspect opinions extracted from user reviews* (Hernández-Rubio et al. 2018) focuses on user-generated content. It provides an in-depth survey of recommender systems that exploit information extracted from user reviews and aims to identify the best techniques for each step when recommending based on user opinions. Furthermore, the paper outlines a number of future directions for review-based recommenders and contributes a number of different helpful resources for researchers, including domain-specific aspect vocabularies and lexicons.

Data-related aspects and the use of Linked Open Data is the focus of the paper *Addressing the user cold start with cross-domain collaborative filtering: Exploiting item metadata in matrix factorization* (Fernández-Tobías et al. 2019). In their work, the authors specifically investigate the use of LOD-based metadata to improve matrix factorization models for cross-domain collaborative filtering. Their experiments emphasize the potential value of considering such knowledge sources, not only to make more accurate recommendations particularly in cold-start situations, but also to balance the trade-off between accuracy and diversity.

The paper *Connectionist Recommendation in the Wild: On the utility, scrutability, and scalability of neural networks for personalized course guidance* (Pardos et al. 2019) is an example of a work that relies on embeddings to encode metadata. The authors propose Course2Vec, an adaptation of Word2Vec, to build a vector space representation of university courses. Their ultimate goal is to recommend a sequence of courses to enroll to students. Their experiments not only showed the effectiveness of the approach but also demonstrated that the recommendations are “scrutable”, which can be a main prerequisite in real-world applications.

Finally, the paper *Feature-combination hybrid recommender systems for automated music playlist continuation* (Vall et al. 2019) contributes to the area of music recommendation. The authors introduce two feature-combination hybrid recommender systems that combine collaborative information from curated music playlists with song features. Their experiments show that the proposed methodology leads to competitive results when compared to CF-based systems, and that it can outperform CF methods when training data is scarce. The work therefore provides additional evidence of the value of content information in cold-start situations.

References

- Belkin, N.J., Croft, W.B.: Information filtering and information retrieval: two sides of the same coin? *Commun. ACM* **35**(12), 29–38 (1992)
- Bogers, T.: Tag-based recommendation. In: *Social Information Access—Systems and Technologies*, Lecture Notes in Computer Science, vol. 10100, pp. 441–479. Springer (2018)
- Chen, L., Chen, G., Wang, F.: Recommender systems based on user reviews: the state of the art. *User Model. User-Adapt. Interact.* **25**(2), 99–154 (2015)
- Deldjoo, Y., Elahi, M., Quadrana, M., Cremonesi, P.: Using visual features based on MPEG-7 and deep learning for movie recommendation. *Int. J. Multimedia Inf. Retr.* **7**(4), 207–219 (2018)
- Deldjoo, Y., Ferrari Dacrema, M., Constantin, M.G., Eghbal-zadeh, H., Cereda, S., Schedl, M., Ionescu, B., Cremonesi, P.: Movie genome: alleviating new item cold start in movie recommendation. *User Model. User-Adapt. Interact.* (2019). <https://doi.org/10.1007/s11257-019-09221-y>

- Di Noia, T., Mirizzi, R., Ostuni, V.C., Romito, D., Zanker, M.: Linked open data to support content-based recommender systems. In: Proceedings of the 8th International Conference on Semantic Systems, pp. 1–8. ACM (2012)
- Elahi, M., Deldjoo, Y., Bakhshandegan Moghaddam, F., Cella, L., Cereda, S., Cremonesi, P.: Exploring the semantic gap for movie recommendations. In: Proceedings of the Eleventh ACM Conference on Recommender Systems, pp. 326–330 (2017)
- Fernández-Tobías, I., Cantador, I., Tomeo, P., Anelli, V.W., Di Noia, T.: Addressing the user cold start with cross-domain collaborative filtering: exploiting item metadata in matrix factorization. *User Model. User-Adapt. Interact.* (2019). <https://doi.org/10.1007/s11257-018-9217-6>
- Gedikli, F., Jannach, D., Ge, M.: How should I explain? A comparison of different explanation types for recommender systems. *Int. J. Hum. Comput. Stud.* **72**(4), 367–382 (2014)
- Golbeck, J.: Generating predictive movie recommendations from trust in social networks. In: International Conference on Trust Management, pp. 93–104. Springer (2006)
- Hensley, C.B.: Selective Dissemination of Information (SDI): State of the Art in May, 1963. In: Proceedings of the May 21–23, 1963, Spring Joint Computer Conference, AFIPS '63 (Spring), pp. 257–262 (1963)
- Hernández-Rubio, M., Cantador, I., Bellogín, A.: A comparative analysis of recommender systems based on item aspect opinions extracted from user reviews. *User Model. User-Adapt. Interact.* (2018). <https://doi.org/10.1007/s11257-018-9214-9>
- Hidasi, B., Karatzoglou, A., Baltrunas, L., Tikk, D.: Session-based recommendations with recurrent neural networks. In: 4th International Conference on Learning Representations (ICLR) (2016)
- Jannach, D., Adomavicius, G.: Recommendations with a purpose. In: Proceedings of the 10th ACM Conference on Recommender Systems (RecSys 2016), pp. 7–10. Boston, Massachusetts, USA (2016). <https://doi.org/10.1145/2959100.2959186>
- Jannach, D., Hegelich, K.: A case study on the effectiveness of recommendations in the mobile internet. In: Proceedings of the 3rd ACM Conference on Recommender Systems, pp. 205–208 (2009)
- Jannach, D., Zanker, M., Ge, M., Gröning, M.: Recommender systems in computer science and information systems—a landscape of research. In: 13th International Conference on Electronic Commerce and Web Technologies, pp. 76–87. Vienna, Austria (2012)
- Jannach, D., Kamehkhosh, I., Lerche, L.: Leveraging multi-dimensional user models for personalized next-track music recommendation. In: Proceedings of the Symposium on Applied Computing, pp. 1635–1642 (2017)
- Jäschke, R., Marinho, L., Hotho, A., Schmidt-Thieme, L., Stumme, G.: Tag recommendations in folksonomies. In: European Conference on Principles of Data Mining and Knowledge Discovery, pp. 506–514. Springer (2007)
- Kirshenbaum, E., Forman, G., Dugan, M.: A live comparison of methods for personalized article recommendation at Forbes.com. In: Machine Learning and Knowledge Discovery in Databases, pp. 51–66. Springer, Berlin (2012)
- Kleinerman, A., Rosenfeld, A., Kraus, S.: Providing explanations for recommendations in reciprocal environments. In: Proceedings of the 12th ACM Conference on Recommender Systems, pp. 22–30. ACM (2018)
- Kula, M.: Metadata embeddings for user and item cold-start recommendations. In: Proceedings of the 2nd ACM RecSys Workshop on New Trends in Content-Based Recommender Systems, *CEUR Workshop Proceedings*, vol. 1448, pp. 14–21. CEUR-WS.org (2015)
- Liu, J., Dolan, P., Pedersen, E.R.: Personalized news recommendation based on click behavior. In: Proceedings of the 15th International Conference on Intelligent User Interfaces, IUI '10, pp. 31–40 (2010)
- Lu, Y., Dong, R., Smyth, B.: Why I like it: multi-task learning for recommendation and explanation. In: Proceedings of the 12th ACM Conference on Recommender Systems, pp. 4–12 (2018)
- McAuley, J., Targett, C., Shi, Q., Van Den Hengel, A.: Image-based recommendations on styles and substitutes. In: Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 43–52 (2015)
- Melville, P., Mooney, R.J., Nagarajan, R.: Content-boosted collaborative filtering for improved recommendations. In: Eighteenth National Conference on Artificial Intelligence, pp. 187–192 (2002)
- Messina, P., Dominguez, V., Parra, D., Trattner, C., Soto, A.: Content-based artwork recommendation: integrating painting metadata with neural and manually-engineered visual features. *User Model. User-Adapt. Interact.* (2018). <https://doi.org/10.1007/s11257-018-9206-9>

- Mizgajski, J., Morzy, M.: Affective recommender systems in online news industry: how emotions influence reading choices. *User Model. User-Adapt. Interact.* (2018). <https://doi.org/10.1007/s11257-018-9213-x>
- Musto, C., Basile, P., Lops, P., de Gemmis, M., Semeraro, G.: Introducing linked open data in graph-based recommender systems. *Inf. Process. Manag.* **53**(2), 405–435 (2017a)
- Musto, C., Lops, P., de Gemmis, M., Semeraro, G.: Semantics-aware recommender systems exploiting linked open data and graph-based features. *Knowl.-Based Syst.* **136**, 1–14 (2017b)
- Musto, C., Franza, T., Semeraro, G., de Gemmis, M., Lops, P.: Deep content-based recommender systems exploiting recurrent neural networks and linked open data. In: *Adjunct Publication of the 26th Conference on User Modeling, Adaptation and Personalization*, pp. 239–244. ACM (2018)
- Musto, C., Narducci, F., Lops, P., de Gemmis, M., Semeraro, G.: Linked open data-based explanations for transparent recommender systems. *Int. J. Hum. Comput. Stud.* **121**, 93–107 (2019)
- Nunes, I., Jannach, D.: A systematic review and taxonomy of explanations in decision support and recommender systems. *User-Model User-Adapt. Interact.* **27**(3–5), 393–444 (2017)
- Oramas, S., Ostuni, V.C., Noia, T.D., Serra, X., Sciascio, E.D.: Sound and music recommendation with knowledge graphs. *ACM Trans. Intell. Syst. Technol. (TIST)* **8**(2), 21 (2017)
- Pardos, Z.A., Fan, Z., Jiang, W.: Connectionist recommendation in the wild: on the utility and scrutability of neural networks for personalized course guidance. *User Model. User-Adapt. Interact.* (2019). <https://doi.org/10.1007/s11257-019-09218-7>
- Passant, A.: dbrec—music recommendations using DBpedia. In: *International Semantic Web Conference*, pp. 209–224. Springer (2010)
- Pazzani, M., Muramatsu, J., Billsus, D.: Syskill & Webert: Identifying interesting web sites. In: *Proceedings of the Thirteenth National Conference on Artificial Intelligence, AAAI'96*, pp. 54–61 (1996)
- Salton, G., McGill, M.J.: *Introduction to Modern Information Retrieval*. McGraw-Hill Inc, New York (1986)
- Seo, M., Kembhavi, A., Farhadi, A., Hajishirzi, H.: Bidirectional attention flow for machine comprehension. In: *Proceedings of the 2017 International Conference on Learning Representations* (2017a)
- Seo, S., Huang, J., Yang, H., Liu, Y.: Interpretable convolutional neural networks with dual local and global attention for review rating prediction. In: *Proceedings of the Eleventh ACM Conference on Recommender Systems*, pp. 297–305. ACM (2017b)
- Shi, C., Zhang, Z., Luo, P., Yu, P.S., Yue, Y., Wu, B.: Semantic path based personalized recommendation on weighted heterogeneous information networks. In: *Proceedings of the 24th ACM International Conference on Information and Knowledge Management*, pp. 453–462 (2015)
- Song, Y., Elkahky, A.M., He, X.: Multi-rate deep learning for temporal recommendation. In: *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval*, pp. 909–912. ACM (2016)
- Sun, Y., Han, J., Yan, X., Yu, P.S., Wu, T.: PathSim: Meta path-based top-K similarity search in heterogeneous information networks. *Proc. VLDB Endow.* **4**(11), 992–1003 (2011)
- Tintarev, N., Masthoff, J.: The effectiveness of personalized movie explanations: an experiment using commercial meta-data. In: *International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems*, pp. 204–213. Springer (2008)
- Vall, A., Dorfer, M., Eghbal-zadeh, H., Schedl, M., Burjorjee, K., Widmer, G.: Feature-combination hybrid recommender systems for automated music playlist continuation. *User Model. User-Adapt. Interact.* (2019). <https://doi.org/10.1007/s11257-018-9215-8>
- Vasile, F., Smirnova, E., Conneau, A.: Meta-prod2vec: Product embeddings using side-information for recommendation. In: *Proceedings of the 10th ACM Conference on Recommender Systems*, pp. 225–232 (2016)
- Fig, J., Sen, S., Riedl, J.: Tagsplanations: explaining recommendations using tags. In: *Proceedings of the 14th International Conference on Intelligent User Interfaces*, pp. 47–56 (2009)
- Yao, Y., Harper, F.M.: Judging similarity: a user-centric study of related item recommendations. In: *Proceedings of the 12th ACM Conference on Recommender Systems*, pp. 288–296 (2018)
- Yu, X., Ren, X., Sun, Y., Gu, Q., Sturt, B., Khandelwal, U., Norick, B., Han, J.: Personalized entity recommendation: A heterogeneous information network approach. In: *Proceedings of the 7th ACM international conference on Web Search and Data Mining*, pp. 283–292 (2014)
- Zhang, Y., Ai, Q., Chen, X., Croft, W.B.: Joint representation learning for top-n recommendation with heterogeneous information sources. In: *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*, pp. 1449–1458. ACM (2017)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Pasquale Lops is an Assistant Professor at the University of Bari, Italy, with a main research focus on content-based recommender systems.

Dietmar Jannach is a Full Professor at AAU Klagenfurt, Austria. His research focus is on e-commerce recommender systems.

Cataldo Musto is an Assistant Professor University of Bari, Italy. He focuses on user modeling and content-based recommender systems.

Toine Bogers is an Associate Professor at the Aalborg University Copenhagen, Denmark. His main research focus is on content-based recommendation and dealing with the complexity of user's recommendation and retrieval needs.

Marijn Koolen is a researcher at the Royal Netherlands Academy of Arts and Sciences. He focuses on search and recommendation with complex item information.