ARTICLE TEMPLATE

Recommender Systems: Challenges and Opportunities in the Age of Big Data and Artificial Intelligence

Mehdi Elahi[a], Amin Beheshti[b] and Srinivasa Reddy Goluguri[b]

[a] University of Bergen, Bergen, Norway; [b] Macquarie University, Sydney, Australia

ARTICLE HISTORY
Compiled February 2, 2022

ABSTRACT

Modern Recommendation Systems will require to access and understand the big data built on top of the large data islands. This is important as the growing enhancement in interconnection, storage, as well as data management has made it possible to connect to data deluge from the big data, which in turn, can lead to making intelligent and accurate personalization and recommendations. Despite the effectiveness of the presented methods, in the times of big data and AI, the need for more advanced methods has been a strong force to build intelligent systems for quick, accurate, and personalized recommendations tailored to each customer's needs and preferences. In this book chapter, we will provide an overview of different types of real-world Recommender Systems, along with challenges and opportunities in the age of big data and AI. We will discuss how recent growth in cognitive technology together with advancement in areas such as AI (plus ML, DL and NLP), as well as Knowledge Representation, Interaction and Personalization have resulted in the substantial enhancement in the research of Recommender Systems.

KEYWORDS
recommender systems; algorithms; methods; big data, data science; data analytic,

1. Introduction

In the times of Big Data, choosing the right products to consume is becoming a challenge for consumers due to the massive *Volume, Velocity*, and *Variety* of related data produced online. This can be a reason that the users are getting more and more desperate when making choices among an unlimited set of choices. Recommender Systems are choice support apps that can deal with this challenge by assisting the shoppers when making choices on what to purchase Jannach, Zanker, Felfernig, and Friedrich (2010); Resnick and Varian (1997); Ricci, Rokach, and Shapira (2015). Recommender systems can learn from the *particular* preferences and tastes of users and build personalized suggestions that can tailor to the users' preferences and necessities rather than suggesting based on the mainstream taste Elahi (2011); Elahi, Repsys, and Ricci (2011).

Many recommender softwares and algorithms have been proposed, up to now, by the academic and industrial community. Most of these algorithms are capable of getting, as input data, various data types and exploit them to generate recommendations on top of the data. These data types can describe either the item content (e.g., category, brand, and tags) or the user preferences (e.g., ratings, likes, and clicks). The data is collected and pre-processed, cleaned and then exploited to build a model in which the items are projected as arrays of features. Recommendation list for a specific user is then made by filtering the items that represent alike features to the rest of the item sets that user liked/rated high.

Enhanced capabilities of recommender techniques in understanding the varied categories of user tastes and precisely tackling information burden has enabled them to become an important part of any online shop that requires to tackle the expansion of item catalog Burke (2002); Elahi (2014). Diverse categories of recommender engines has been built in order to generate personalized selection and relevant recommendation of products and services ranging from clothing and outfits to movie and music. Such a personalized selection and suggestion is usually made based on the big data of a huge community of connected users and calculating the patterns and relationships among their preferences Chao, Huiskes, Gritti, and Ciuhu (2009); Elahi (2011); Elahi and Qi (2020); He and McAuley (2016); Nguyen, Almenningen, Havig, Schistad, Kofod-Petersen, Langseth, and Ramampiaro (2014); Quanping (2015); Tu and Dong (2010). The excellency in performance of recommender systems have been validated in the diverse range of e-commerce applications where a choice support mechanism is necessary to handle customers needs and help them when interacting with the online e-commerce. Such an assistance improves the user experiences when shopping or browsing the system catalogue He and McAuley (2016); Tu and Dong (2010).

In this book chapter, we will provide an outline of different types of real-world Recommender Systems, along with challenges and opportunities in the age of big data and AI. We will discuss the progress in cognitive technology in addition to evolutionary development in areas such as AI (with all relevant disciplines such as ML, DL, and NLP), KR, and HCI can empower Recommender Systems in effectively supporting their users.

We discuss that modern recommendation systems will require to access and understand the big data, with all different forms, generated on data islands can be used for building relevant and personalized recommendations tailored to each customer's needs and preferences. We present different application scenarios (including multimedia, fashion, tourism, banking, and education) and review the potential solutions for the recommendation. The remaining parts of the book chapter is organized as the following: section 2 briefly describes the popular methods and algorithms. Section 3 discusses different application scenarios and section 4 reviews real-world challenges and potential solutions. Section 5 extends the previous chapters by providing some advanced topics. Finally, in the section 6 we make a conclusion for this book chapter.

2. Methods

2.1. Classical

Diverse recommendation approaches have been already developed and tested which can be classified within a number of categories. A well-adopted category of methods is named **Content-based** Pazzani and Billsus (2007). Methods within this category suggest items based on their descriptors Balabanović and Shoham (1997). For example, book recommender systems take the terms within the text of a book as descriptors and suggest to user other books that have descriptors close to the book the user liked in the past. Another most popular category is **Collaborative Filtering** Desrosiers and Karypis (2011); Koren and Bell (2011). Collaborative filtering methods predict the preferences (i.e., ratings) of the users by learning the preferences that a set of users provided to items and suggests to users those items with the highest predicted preferences. Methods within **Demographic** Wang, Chan, and Ngai (2012) category generate recommendations by identifying similar users based on the demographics of the users Pazzani (1999). These methods attempt to group the existing users by their personal descriptors and make the relevant suggestions based on their demographic descriptions. **Knowledge-based** Felfernig and Burke (2008) methods are in another category which try to suggest items that are inferred from the needs and constrains entered by usersBurke (2000). Knowledge-based methods are distinguished by their knowledge about how a specific item fulfills a particular user's need Claypool, Gokhale, Miranda, Murnikov, Netes, and Sartin (1999). Hence, they can mine inferences based on the connections within the user's need and the possible recommendation. **Hybrid** Li and Kim (2003) methods combine diverse individual methods among those noted before in order to handle the particular restrictions of an individual method.

2.2. Collaborative Filtering

Collaborative Filtering (CF) is a recommender method used in almost all application domains. This method focuses on effective adoption of the user feedback (e.g., ratings) elicited from the users to be used to make a profile of affinities. Such profiles are used to generate personalized recommendations for them. Hence, collaborative filtering relies on big data of ratings acquired from typically big network of users Desrosiers and Karypis (2011). Using such a data, to collaborative filtering recommends items that a target user has not yet checked while could probably like Koren and Bell (2011). Perhaps a cornerstone for a the systems is to estimate the feedback (the ratings) entered by users for items that they have not produced for any rating yet. Having the predicted ratings, collaborative filtering can sort the items based on the predicted ratings and recommends those with the highest ratings.

Classical methods in collaborative filtering systems are neighbour-based which compute user-to-user or item-to-item similarities based on the co-rating patterns of the users and items. In item-based collaborative filtering, items can be computed as alike if the community of interconnected users have rated those items in a similar way. Analogously, in user-based collaborative filtering, users with similar rating patterns form neighborhoods that are used for rating prediction. Hence, predicting the ratings are performed based on how the item has been co-rated by the other users who were considered as like-minded compared to the target users.

Another category of collaborative filtering systems adopt Latent factor models in order to generate rating prediction. A well adopted category of these methods is *Matrix Factorization* Koren (2008b); Koren and Bell (2011). Matrix factorization build mathematical models on top of rating data and form a set of factors for the users and items. These sets, with the equal length, are learned from every rating elicited from users. Every factor of these sets, assigned to an item, represents the level in which an item projects a particular latent aspect of user preference. In movie domain, as an example, item factors could be interpreted as the genre of the movie while user factors could describe the taste of the users towards such genres.

In order to identify such factors, matrix factorization decomposes the rating matrix into different matrices:

$$R \approx SM^T \tag{1}$$

Where S is a matrix of $|U| \times F$, and M is a matrix of $|I| \times F$.

A well-known implementation of matrix factorization Timely Development (2008) was proposed is called Funk-SVD Funk (2006) and is capable of making predictions using this formula:

$$\hat{r}_{ui} = \sum_{f=1..F} s_{uf} m_{if} \tag{2}$$

where the s_{uf} describes the level of the user u preferences towards the factor f and the m_{if} describes the strength of the factor f is in the item i Koren (2008b).

2.3. Content-Based Recommendation

Content-based methods are also widely adopted in recommender systems. Content-based methods adopt Content-based Filtering (CBF) algorithms in order to build user profiles by associating user preferences to the item content Deldjoo and Atani (2016); Deldjoo, Elahi, Cremonesi, Garzotto, Piazzolla, and Quadrana (2016). As noted before, the user preferences are typically given as ratings to items and item content can be described with diverse forms of features. Content-based recommender systems exploit such content features and make a *Vector Space Model* on top of the content data (Pazzani and Billsus, 2007). This model projects every item into a multi-dimensional space according to the the content features Lops, De Gemmis, and Semeraro (2011). The content-based methods measures a relevancy score associated to the user preferences proportional to the content features.

So far, a diverse spectrum of CBF approaches have been formulated and tested in recommender systems context. A well-adopted method is *K-Nearest Neighbors (KNN)* which exploits the similarities using items content, and build suggestions on top of it. The similarities scores among the item j and all the rest of the items allows us to build a set of nearest neighbors items (i.e., NN_j) containing the items with the maximum similarity scores to the item j. Accordingly, the preferences (e.g., likes/dislikes or the star ratings) that have provided for the items within nearest neighbors set are used to predict the preference $\hat{r}_{ij}$ for user i and item j:

$$\hat{r}_{ij} = \frac{\sum_{j' \in NN_j, r_{ij'} > 0} r_{ij'} s_{sjj'}}{\sum_{j' \in NN_j, r_{ij'} > 0} s_{sjj'}} \tag{3}$$

where $r_{uj} > 0$ reflects the elements of the preferences matrix $\mathcal{R}$, i.e., user ratings included in the matrix of all ratings.

2.4. Hybrid FM

While collaborative filtering method and content-based method are both have been largely adopted by the recommender system community, however, they have a number of restrictions. These restrictions will be explained later on in this book chapter. In

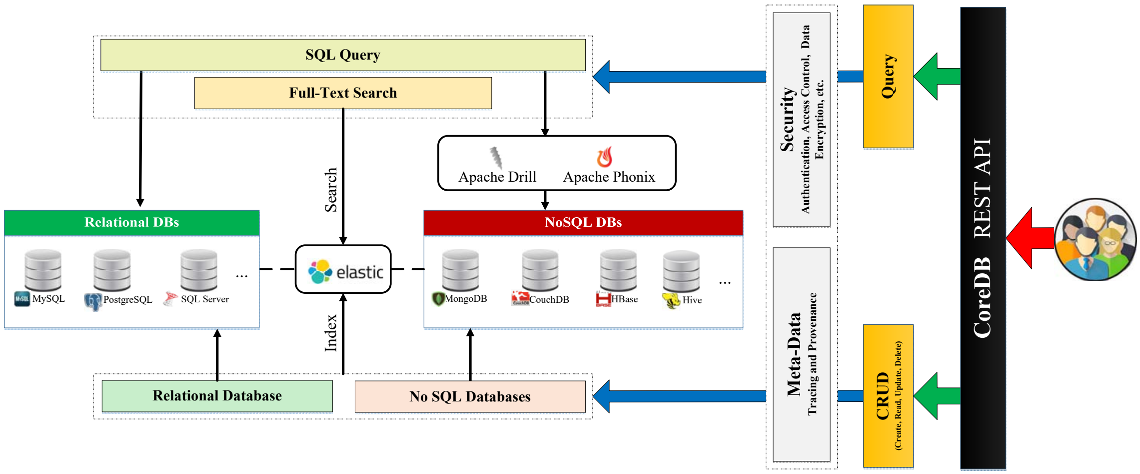

Figure 1. The Data Lake as a Service Architecture (CoreDB Beheshti et al. (2017a)).

order to address such restrictions, *Hybrid* methods have been developed by hybridizing these methods Low, Bickson, Gonzalez, Guestrin, Kyrola, and Hellerstein (2012). While hybrid methods can also have diverse forms, we briefly introduce one of the most recent methods, called *Factorization Machines* (Burke, 2002; Rendle, 2012).

Factorization machines is a recommender method that is formed by extending the classical matrix factorization method TURI (2018). Factorization machines hybridizes matrix factorization by mixing it with a well-none Machine Learning method named *Support Vector Machines (SVM)*. This enables the factorization machines be capable of taking advantages of not only the user preferences (e.g., ratings), but also item descriptions , as well as any addition data attributed by users. This enables factorization machines to adopt a wide range of data, typically referred to as *side information*, e.g., item descriptors (e.g., category, title, or tag) as well as user attributes (e.g., demographics, emotion, mood and personality). Hence, factorization machines build mathematical models on top of user ratings as well as item descriptors or user attributes in order it make preference prediction (Rendle, 2012).

Predicting the user preferences (e.g., like/dislike or rating) is conducted through the next formula:

$$\hat{r}_{ij} = \mu + w_i + w_j + \mathbf{a}^{\mathbf{T}}\mathbf{x_i} + \mathbf{b}^{\mathbf{T}}\mathbf{y_i} + \mathbf{u_i}^{\mathbf{T}}\mathbf{v_j} \tag{4}$$

where μ denotes the bias factor, w_i is the user weight, w_j is the item weight, $\mathbf{x_i}$ and $\mathbf{y_j}$ are feature set for user and item, respectively.

There other advanced models such as (Mooney and Roy, 2000) (Ahn, Brusilovsky, Grady, He, and Syn, 2007) that go beyond traditional methods by building probabilistic models based on the user or item input data. For instance, in Fernández-Tobías and Cantador (2014); Manzato (2013), a model called *gSVD++* has been developed that can takes advantages of content data attributed into MF Koren (2008a).

2.5. Modern Recommender Systems

Despite the effectiveness of the presented methods, in the age of big data and Artificial Intelligence (AI), the need for more advanced methods has been a strong force to build a modern generation of recommender engines. Several improvements in such recommendation engines have enabled them to make quick and accurate recommendations tailored to each customer's needs and preferences. To achieve this goal, modern recommendation systems have focused on three main aspects: data, knowledge, and cognition.

2.6. Data-Driven Recommendations

Modern recommendation systems will require to access and understanding the raw data generated on various data islands including open/private/social data sources Beheshti, Benatallah, Sheng, and Schiliro (2019). This is important as the improvement in data communication/processing enable access to the big data and will enable intelligent and accurate recommendations. In this context, the main challenge in harnessing the big data would be to ingest/organize the big data (from various data islands) onto a centralized repository. The concept of a Data Lake which presents a centralized repository to organize the raw data generated on various data islands. Modern approaches such as CoreDB Beheshti et al. (2017a) propose the notion of Data Lake as a service to facilitate managing and querying the large amount of information (from open, social, IoT, and private data islands) and to enable analysts to deal with the variety of data and non-standard data models. Figure 1 illustrates the CoreDB (Data Lake as a Service) architecture.

To understand the raw data, it would be necessary to leverage AI (Artificial Intelligence) and ML (Machine Learning) technologies to contextualize the raw data will improve the accuracy of recommendations. This will enable the adoption of popular recommender systems, and facilitate the journey from analytical models to deep learning models. The goal here is to generate better predictions by improving correlations between features and attributes. Hence, the concept of Knowledge Lake has been introduced. Accordingly, data curation services can be adopted which will enable automatic transformation of the raw data into curated data. Figure 2 illustrates the architecture of the Knowledge Lake.

As a motivating scenario, we may consider recommendations on social media such as Twitter. Modern recommender systems would need to understand the content and context of Tweets twitted by social users. Considering a Tweet as a raw data, the curation services Beheshti, Tabebordbar, Benatallah, and Nouri (2017b) in the knowledge lake, would be able to extract information (e.g., keyword, phrase, named entity, topic, sentiment, etc.) from the text of the Tweet or a URL in the Tweet, and enrich them using the external knowledge sources and services. A contextualized Tweet (as illustrated in Figure 3) will tell more stories compared to a raw Tweet. For example, if we are able to extract "Barak Obama" from the Text of the Tweet, understand that it is a named entity of type person, and link it to the entity Barak Obama (i.e., the 44th president of the united state) in Wikidata, the recommendation system will understand that this Tweet may be related to the Politics topic. Similarly, if the Tweet contains a keyword related to health or mention of the World Health Organization (WHO) the Tweet would be related to the Health topic.

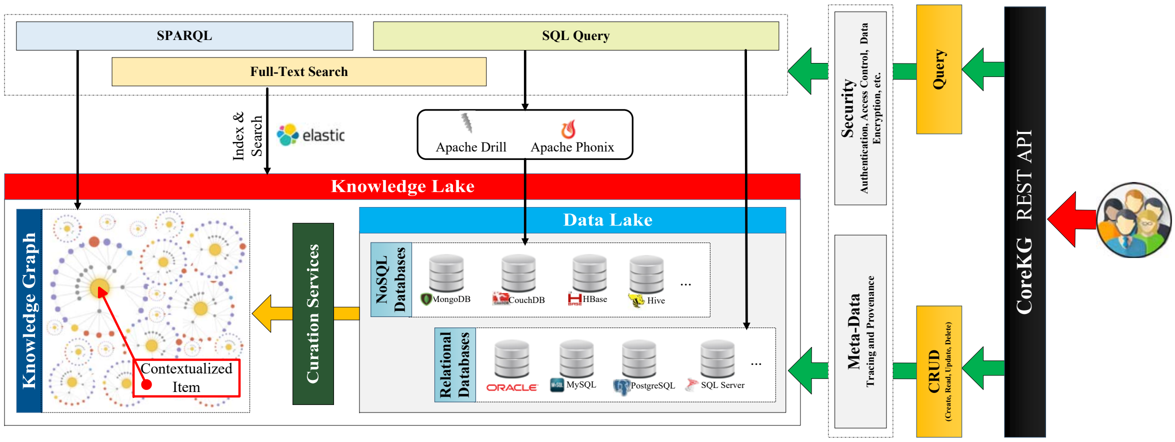

Figure 2. CoreKG: Knowledge Lake as a Service Architecture Beheshti et al. (2018)

.

2.7. Knowledge-Driven Recommendations

Intelligence RSs started to learn from domain experts' experience and knowledge, to understand the domain that the items will be recommended Beheshti, Yakhchi, Mousaeirad, Ghafari, Goluguri, and Edrisi (2020b). For example, a new line of research started Beheshti et al. (2020b) to use crowdsourcing [1] techniques to capture domain experts' knowledge and use them to provide accurate and personalized recommendations. Another line of work, leveraged Intelligent Knowledge Lakes[2] to address the following two challenges: (i) Cold-start problem: Leveraging Intelligent Knowledge Lakes will bring informative data from a crowd of people and use it to generate recommendations. (ii) Bias and Variance: Leveraging Intelligent Knowledge Lakes will be able to guide recommender systems to choose the best next steps by following the best practices learned from domain experts. This is important, as features (used for training recommenders) may be gathered by humans and will enable biases to get into data preparation and training phases. To build an intelligent KL, it is important to mimic domain expert's knowledge. This can be done using techniques such as collecting feedback, organizing interviews, and requesting surveys. To achieve this goal, it would be important to capture important events and entities (and relationships among them) that are happening in real-time in disciplines/fields education, and fintech.

2.8. Cognition-Driven Recommendations

To support accurate and intelligent recommendations, it would be vital, for a recommender system, to identify similar users based on their behaviour, activities, and cognitive thinking. Accordingly, a cognition-driven Recommender System should: (i) facilitate understanding users' personality, emotion, moods, and affinities over time. This task aims to empower the recommender models in exploitation of the cognitive signals and Neural data, including our previous work, Personality2Vec Beheshti, Hashemi, Yakhchi, Motahari-Nezhad, Ghafari, and Yang (2020a), to design mechanisms for per-

[1]crowdsourcing, i.e., the process of obtaining information by engaging a big crowd through platforms such as Amazon Mechanical Turk (mturk.com/) and Figure Eight (figure-eight.com/).

[2]Intelligent Knowledge Lakes Beheshti et al. (2019) has been introduced to facilitate linking Artificial Intelligence (AI) and Data Analytics.

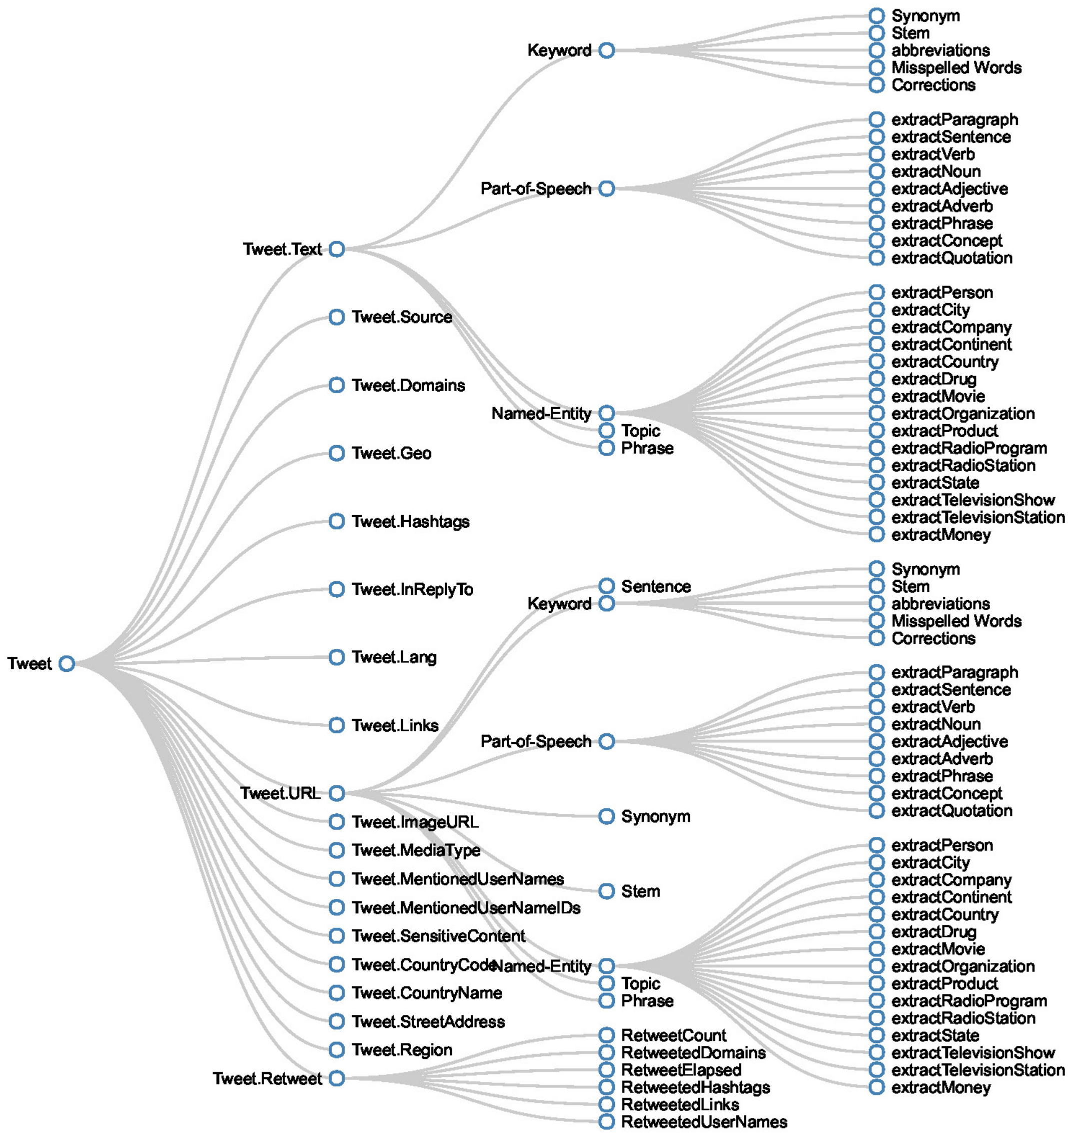

Figure 3. A contextualized Tweet Beheshti et al. (2019).

sonalized task recommendations and to facilitate discovering meaningful patterns from users' social behaviours. A cognitive RS may focus on dimensions such as Explicit behavioural pattern and Implicit behavioural pattern Beheshti et al. (2020b). Explicit patterns may include Text-based methods, Location-based approaches, Action-based, and Feature-based methods. Implicit patterns may focus on Social-based features, Trust-based features, and Action-based features.

Sequential Recommender Systems Wang, Hu, Wang, Cao, Sheng, and Orgun (2019) aim to understand and model user behaviors, however, they do not consider the analysis of users' attitude, behaviour, and Personality over time. Recent work introduced a new type of Recommender Systems, i.e., Cognitive Recommender Systems Beheshti et al. (2020b), which focuses on understanding the users' cognitive aspects.

3. Application

3.1. Classic

3.1.1. Multimedia

Multimedia is probably the most popular application domain in recommender systems. Multimedia recommender systems can exploit different forms of preference data and can use different types of multimedia descriptors when creating recommendations Elahi, Ricci, and Rubens (2012); Hazrati and Elahi. While such features can have different forms, however, we can classify them into a two main categories, i.e., *high-level* and *low-level* forms of descriptors Cantador, Szomszor, Alani, Fernández, and Castells (2008); Hazrati and Elahi.

High-level descriptors illustrate a more of the semantic and syntactic characteristics of multimedia items and can be aggregated from either structured forms of meta-data (e.g., a relational databases or an ontology) Cantador et al. (2008); Mooney and Roy (2000), or from less structured form of data (e.g., user reviews, film plots, and social tags) Ahn et al. (2007); Hazrati and Elahi.

Low-level descriptors, on the other side of the story, are aggregated straightly from multimedia files (e.g., audio or visual files). In the music domain, as an instance, low-level descriptors can represent the acoustic configurations of the songs (e.g., rhythm, energy, and melody) which can be adopted by recommender systems to find similar songs and use it to generate personalized recommendation for a user Bogdanov and Herrera (2011); Bogdanov, Serrà, Wack, Herrera, and Serra (2011); Knees, Pohle, Schedl, and Widmer (2007); Seyerlehner, Schedl, Pohle, and Knees (2010).

In video domain, low-Level descriptors can represent the visual aspects of the videos and hence reflect an artistic *style* Canini, Benini, and Leonardi (2013); Lehinevych, Kokkinis-Ntrenis, Siantikos, Dogruöz, Giannakopoulos, and Konstantopoulos (2014); Yang, Mei, Hua, Yang, Yang, and Li (2007); Zhao, Li, Wang, Yuan, Zha, Li, and Chua (2011).

It is a fact that, recommendation based on low-level features did not draw much attention multimedia recommender systems. On the other hand, such features received massive attention in some related research fields, namely, in Computer Vision Rasheed, Sheikh, and Shah (2005), and Content-Based Video Retrieval. Despite the differences in objectives, these communities share objectives such as formulating the informative descriptors of video and movie items. Hence they reports outcomes and insights that can be beneficial also the context of the multimedia recommender systems Brezeale and Cook (2008); Hu, Xie, Li, Zeng, and Maybank (2011); Rasheed et al. (2005).

3.1.2. Tourism

Another well-studied domain in the research on the recommender systems is tourism. This is a domain where *contextualization* shall play an important role. We can define contextualization as the process of incorporating contextual factors (such as weather condition, travel goals, and means of transportation) in the recommendation generation. The idea is to make personal suggestions by incorporating diverse sources of user data as well as the *condition* represented by contextual factors Adomavicius and Tuzhilin (2011). For example, on bad weather conditions, a group of tourists may be interested to visit the suggested indoor attractions (e.g., museums), although in a nice weather, they may prefer to go for outdoor activities (e.g., hiking). Recommender systems that are capable if using such contextual factors are known as *CARS*.

CARS are well-empowered to exploit mathematical modeling in order to better learn the user preferences in different contextual situations based on diverse sources of data , e.g., the temperature, season, and the geographical position, and even the vehicle type. Due to the popularity of this research domain, a big amount of research has already been conducted in in this domain Baltrunas, Ludwig, Peer, and Ricci (2012); Chen and Chen (2014); Gallego, Woerndl, and Huecas (2013); Hariri, Mobasher, and Burke (2012); Kaminskas, Ricci, and Schedl (2013); Natarajan, Shin, and Dhillon (2013). Majority of these works can exploit the the context experienced by the user in the recommending process.

3.1.3. Food

There are a diverse categories of food recommendation systems that have recently been proposed by the community Trevisiol, Chiarandini, and Baeza-Yates (2014); West, White, and Horvitz (2013). For example, Freyne and Berkovsky Freyne and Berkovsky (2010) built a food recommendation system which, through an effective user interaction model, collects user preferences and generates personalized suggestions. Their system converts the preferences of the users for recipes into preferences for ingredient and then merged these converted preferences to to form user suggestions.

Elahi et al. Elahi, Ge, Ricci, Massimo, and Berkovsky (2014) devised a different approach for food recommendation which can combine the predictions for food along diverse aspects (such as user food preferences, nutrition, ingredients, and expenditure) to measure a score for a potential food (meal). The objective is to take into account measures that shall impact the user's food choices in order to make more beneficial set of recommendations Teng, Lin, and Adamic (2012). In the next paper same authors performed an assessment of the rating prediction method, which used a variant of MF. This method exploits more data than only ratings such as the subjective tags inserted to different recipes by users. It has been discovered that the extra data input on the user preferences allows the noted technique to outperform other baseline methods, e.g., those developed in Freyne and Berkovsky (2013).

Generally speaking, the preferences that are aggregated by a recommender system can have two forms, i.e., long-term affinities or short-term affinities. While obtaining and aggregating both forms of preferences is essential, large body of research on recommender systems do not identify the differences between these two forms. Only limited research works have considered such differences (e.g., Ricci and Nguyen (2007)). The noted example is one of the few works that developed a recommender system , which elicits both generic long-term affinities and specific short-term affinities.

We would like to point out that the traditional line of research on recommender systems typically undermine the importance of human-system interaction model, as an essential component for creating a industrial-grade systems. Hence, they mainly concentrate on rather enhancing the core analytical models, by supposing that the preference acquisition procedure is conducted only in the beginning and then ended.

3.1.4. Fashion

Fashion is traditionally referred to as the prevailing form of clothing and it can be formulated by the concept of *changing*. Fashion includes a diverse characters of self-fashioning such as styles in the street to the other calls of *high* fashion made by designers Bollen, Knijnenburg, Willemsen, and Graus (2010); Person (2019). One of the biggest issues in this types of applications is the growing diversity and expanding number of fashion products. This is an effect that can certainly lead to the choice overload for the fashion consumers. This is not necessarily just negative since the more available options the the higher the likelihood that the consumers finds a desired product. However, such an effect may lead to impossibility to actually choosing a product, i.e., the problem of receiving a many options, particularly when they are very diverse Anderson (2006).

Recommender techniques are powerful tools that can effectively tackle this issue by making relevant suggestions of products tailored to the needs of the users. They can build a filtering mechanism that eliminates uninteresting and irrelevant products from a shortlist of recommendations. They can thoroughly mine the user data in order to learn the particularities among user preferences for each single user. For instance, Amazon can look into the purchase history of users and build predictive models which can ultimately be used to make personalized recommendation for her. Hence, the smart engine behind the recommender can actively understand from the users' behaviors and obtain divers and informative forms of data describing the user tastes, to obtain knowledge on the individual requirements of the every user Rashid, Albert, Cosley, Lam, Mcnee, Konstan, and Riedl (2002); Rubens, Elahi, Sugiyama, and Kaplan (2015); Su and Khoshgoftaar (2009).

3.2. Modern

3.2.1. Financial Technology (Fintech)

Financial technology (Fintech) aims use technology to provide financial services to businesses or consumers. We can discuss that any forms of recommendation method will need to understand three main dimensions to provide intelligent recommendations: (i) banking entities, such as customers and products; (ii) banking domain knowledge, about how different banking segments operate across customers, sales and distribution, products and services, people, process and technology; and (iii) banking processes, to help understanding the best practices learned by knowledge experts in processes such as Fraud detection, Customer segmentation, Managing customer data, Risk modeling for investment banks, and more.

The main shortcoming of existing RSs is that, they do not consider domain experts' knowledge, and hence may not well exploit user side information such as cognitive characteristics of user. These aspects are quite vital to support intelligent and time-aware recommendations.

To support data analytics focusing on customers' cognitive activities, it would be important to understand customers' dimensions both from banking and non-banking perspectives depicted in Figure 4. Modern approaches such as Cognitive Recommender Systems Beheshti et al. (2020b), model the customer behaviour and activities as a graph based data model Beheshti, Benatallah, and Motahari-Nezhad (2016); Hammoud, Rabbou, Nouri, Beheshti, and Sakr (2015) over customers' cognitive graph to personalize the recommendations.

3.2.2. Education

One of the popular application domains in recommender system is Education. Education allows individuals to reach their full potential and aids the development of societies by reducing poverty and decreases social inequalities. Recently, the world has experienced an increasing growth in this domain both on quantity and quality measures.

Figure 4. Users' dimensions in a banking scenario.t Beheshti et al. (2020b).

This in turn has already generated several challenges in the education system, such as: instructors' workload in dealing with assessments and providing recommendations based on students' performance and skills assessment.

In this context, recommender systems can be significantly important tools in personalizing teaching and learning, by understanding and analyzing important indicators such as: Knowledge, Performance (e.g., Cognitive, Affective, and Psychomotor indicators) and skills (e.g., decision-making and problem-solving). An attractive planed work for future could be to implement a time-aware deep learning model to construct and analyse learners' profiles to better understand students' performance and skills. The learning models would be enable Recommender Systems to identify similar performing students which may facilitate personalizing the learning process, subject selections, and recruitment.

3.2.3. Recruitment

Talent acquisition and recruitment processes are examples of adhoc processes that are controlled by knowledge workers aiming to achieve a business objective/goal. Attracting and recruiting the right talent is a key differentiator in modern organizations and Recommender Systems can play an important role in assisting recruiters in the recruitment process. For example, consider a recommendation engine that has access to LinkedIn[3] profiles, is able to extract data and knowledge from business artifacts, e.g., candidates' CV and position description, have access to curation algorithms to contextualize the data and knowledge, and is able to link them to the facts in the recruitment domain Knowledge Base.

Artificial Intelligence (AI) has enabled organizations to create business leverage by applying cutting edge automation techniques can help Shahbaz, Beheshti, Nobari, Qu, Paik, and Mahdavi (2018): (i) improving the overall quality and effectiveness of the recruitment process; (ii) extracting relevant information from a candidate's CV automatically; (iii) aggregation of different candidate evaluations and relevant information; (iv) understanding the best practices used by recruiters; (v) extracting personality traits and applicant's attitude from social media sites which was traditionally only possible through interviews. All these techniques can be leveraged by Recommender systems in building effective ranking algorithms that can optimize the recommendations and help maintain a priority list of talent pool. AI-enabled Recommender Systems would be able to help in matching the behaviours of best talented people in their organizations help businesses recruit the right candidates for open jobs by aggregating information from different sources and then ranking them based on their overall score. Another future line of work would be to use computer vision algorithms to assess interviews of the potential candidates and compares them to the best organization's talent and makes recommendations Hir (2019).

4. Challenges

Recommender systems typically exploit datasets that contains user feedbacks (e.g., likes/dislikes, or ratings) that represent preferences produced by a big crowd of interconnected users to a large list of items Desrosiers and Karypis (2011). Exploiting such a data empowers the recommender systems to learn the patterns and connections among users and use them to estimate the missing assessments (likes/dislikes or ratings) of users for the unexplored items and suggest items that shall be attractive to a target user Koren and Bell (2011).

The above-mentioned procedure is oversimplified and there are many grand concerns that have not been fully addressed so far. Hereafter, we briefly explain some of these concerns.

4.1. Cold Start

Recommender systems may still encounter a wide range of challenges due to different reasons such as the lack of rating data for some of the users or items Adomavicius and Tuzhilin (2005); Schein, Popescul, Ungar, and Pennock (2002). One of the problematic issues in building personalized recommendation, namely, *Cold Start* which is

[3]LinkedIn (https://www.linkedin.com/) is an American business and employment-oriented online service that operates via websites and mobile apps.

strongly related to the low quality or quantity of the input data. A sub-problem of cold start is called as the *New User* situation which refers to once a new user begins to use system by demanding suggestions prior to giving any preferences to any existing item. Similarly, *New Item* situation, a new item is introduced to the item catalog waiting to obtain assessments from existing users (in terms of rating, reviews, or tags). In addition to that, in spite of the fact that these are typically situations in a cold starting recommender systems, there exists another problem called *Sparsity*. Originally, sparsity is a measure of the data density and proportional to the number of available feedbacks (e.g., ratings) over the overall possible feedbacks:

$$1 - \frac{\#\ of\ existing\ feedbacks}{\#\ of\ all\ possible\ feedbacks} \tag{5}$$

In some of the acute situations of sparsity, the effectiveness of the recommender systems can be strongly deteriorated consequently resulting in a significant decrease in the performance of the system. In such a condition, the quantity of available user feedbacks is largely smaller than the number of missing feedbacks and the operating system has to build predictions with satisfactory level of quality. Adomavicius and Tuzhilin (2005); Braunhofer, Elahi, and Ricci (2014).

Different cold start conditions can take place in the actual applications, namely, *Extreme* cold start and *Moderate* cold start conditions.

- *Extreme Cold Start* take place when a user starts using the system and asks for recommendation before producing any feedback, describing her affinities. The problem can also happen when a brand-new product is inserted into the catalog and has not associated data that can represents that item. This can in turn lead to the failure in suggesting that new product to an existing users. Both of these situations are critical issues and has to be handled by the system, actively.
- *Mild Cold Start* happens once a small number of feedbacks are produced by a user to existing products and the system can use this limited data to generate recommendation. This problem may also take place for a new product when little amount of content data are not fully produced. Mild cold start may take place as a combined condition of extreme and warm start. This can still lead to a failure if not be promptly addressed by operating system.

4.2. Context-awareness

Although context may have different meanings in various domains It typically emphasises on the situation describing an event which will be helpful in order to be understood Beheshti et al. (2020b). In general, context is represented by factors which influence the computation and still be different from the the input and output data Lieberman and Selker (2000). Time, location and social relations are examples of a context which motivated researchers to focus on contest-aware techniques such as location-aware and time-aware Recommender systems.

For instance, time-aware Recommender systems can be beneficial for the system to understand the development of users' affinities over a certain period of time and provide contextual recommendations for different times (e.g., seasonal, monthly, weakly, or daily with different weather conditions). As another example, Social-aware RSs can benefit from the big data generated on social media in different aspects, from the social characterisation of users (including social relationships, followers, shared contents) to identify the personality, behaviour and attitude of social users Ghafari, Yakhchi, Beheshti, and Orgun (2018). For example, features such as intimacy, emotional intensity, along with location and time-aware context can be calculated and used to provide accurate and context-aware recommendations.

4.3. Style-awareness

Modern recommender systems in various application domains are getting more and more aware of *style* of the items and products. Integration of product style within the recommendation process is in turn becoming increasingly important. In multimedia domain, examples of style elements are lighting, colorfulness, and movement and sound.

There are diverse reasons for having lighting in films and movies. The most important one is to enable the perceived understanding of the space and to build observable objects the audience are seeing. But lighting can perhaps alter the way an occurrence is to be seen, acting in a way that it goes beyond logical perception of human.

Colors express a similar capability by setting up the emotions derived by an encounter condition. The specific quality of colors disallows them to be perceived separately from space lighting. Likewise, colors tend to contribute in making a unique 'perception' of a space in comparison to the other aesthetic characteristics functioning in the similar manner. Experts in media has a common belief that the impact of colors becomes larger they are predisposed in making a particular emotional goal.

A number of research works on recommender systems have reported that users' preferences can be impacted greater by the low-level descriptors in comparison to the high-level descriptors (expressing the semantic or syntactic forms in films) (Elahi, Deldjoo, Bakhshandegan Moghaddam, Cella, Cereda, and Cremonesi, 2017; He, Fang, Wang, and McAuley, 2016; Messina, Dominquez, Parra, Trattner, and Soto, 2018; Rimaz, Elahi, Bakhshandegan Moghadam, Trattner, Hosseini, and Tkalcic, 2019; Roy and Guntuku, 2016). Examples of such low-level descriptors can be color energy, shot duration, and lighting key (Wang and Cheong, 2006) have a proved to influence on user mood and emotion (Roberts, Hager, and Heron, 1994). In addition to that, vaious forms of motion (such as camera movement) can play a significant role and hence are commonly adopted by film makers when aiming to affect the perception of movie watchers (Heiderich, 2018). A range of methods and techniques have been adopted to address the task of learning visual descriptors from films (Ewerth, Schwalb, Tessmann, and Freisleben, 2004; Savian, Elahi, and Tillo, 2020; Tan, Saur, Kulkami, and Ramadge, 2000).

Despite of the importance of low-level descriptors, the usage of them has not drawn much consideration in recommendation systems (e.g., an example is Messina et al. (2018)). However, these audiovisual descriptors thoroughly investigated in the related areas, namely computer vision community (Naphide and Huang, 2001; Snoek and Worring, 2005).

5. Advanced Topics

In the times of AI, an intelligent recommender system should be highly data-driven, knowledge-driven, and cognitive-driven. Cognitive Recommender Systems Beheshti et al. (2020b) proposed as a new type of intelligent Recommender Systems that aim to analyze and understand users' preferences and explore mechanisms to intelligently understand the compound and changing environments. In this context, we categorize the advanced topics in recommender systems into the following categories: AI-Powered Recommendations, Cognition-aware, Intelligent Personalization, Intelligent Ranking, and Intelligent customer engagement.

5.0.1. AI-Enabled Recommendations

As discussed in Section 3, AI-Enabled Recommendations should benefit from data-driven and knowledge-driven approaches such as Data Lakes Beheshti et al. (2017a) and Knowledge Lakes Beheshti et al. (2018, 2019).

Data-driven Recommendations will enable leveraging Machine Learning technologies to contextualize the Big Data aiming to enhance the precision of automatic suggestions, facilitating the use of content/context, and moving from statistical modeling to advanced models based on multi layer neural networks. These will improve mining patterns between item and user descriptors to build better suggestions for users;

Knowledge-driven Recommendations will empower simulating the expertise of the domain experts (e.g., using crowdsourcing methods Beheshti et al. (2019)) and to adopt methods such as reinforcement learning to enhance making relevant and accurate recommendations.

5.0.2. Cognition-aware

Cognition refers to the mental procedure of eliciting knowledge and learning by thinking, experience, and the diverse sources of senses Beheshti et al. (2020a). In this context, a Cognition-aware recommendation will enable recognizing the users' personality and emotion as well as studying their particular characteristics and affinities over time.

Social Cognition-aware Recommendations can refer to the processes that enable a Recommender System to interpret social information and provide context-aware recommendations. This task could be quite challenging as social cognition means different things to different people, and can be quite subjective. Moreover, it can be involved in social interactions at a group level or on a one-to-one basis. This in turn makes Group-aware recommendations a challenging task as Group-aware Recommender Systems support scenarios in which a *group* of users might be target for items to be consumed.

Attention Mechanism would be also an interesting line of future work in Recommender Systems as it can be used to simulate neural activities of a person. This task could quite be challenging as the recommendation methods and algorithms should take into consideration the variation in the user's preferences. Applications could include in case-based and session-based recommendations which may require taking the user's previous cases/sessions into account. A future line of work may focus in a hybrid model, leveraging the attention mechanism and context-aware approaches to learn the user's stable preferences and linking that with dynamic intents into account to make the best use of awareness of the variation in the user's contextual preference to cope with the cold-start problem within sessions.

5.0.3. Intelligent Personalization

Personalization, i.e., the action of designing or producing something to meet someone's individual requirements, can support analytics around users' cognitive activities. This, in turn, can support intelligent and time-aware recommendations. In particular, product and content recommendations tailored to a user's profile and habits can facilitate understanding and analyzing users' behavior, preferences, and history. This, in turn, will boost users' engagement and satisfaction in real-time, in contrast to a uniform experience. This process requires automatically process and curate the data, identify meaningful features, select the right algorithms, and use them to train a proper personalization model.

In this context, Feature Engineering Anderson, Antenucci, Bittorf, Burgess, Cafarella, Kumar, Niu, Park, Ré, and Zhang (2013), i.e., the procedure of using domain knowledge to obtain diverse forms of describing features from the raw data via data mining methods, would be a key future work for Intelligent Personalization and to improve and optimize a personalization model. To achieve this goal, features (i.e., the observations or characteristics on which a model is built) can be engineered by decomposing or splitting features (e.g., from external data sources open, private, social, and IoT data islands) or aggregating or combining features to create new features (e.g., to understand users' behavior). However, what makes the task of effective feature engineering to meet someone's individual requirements hard would be the countless number of options of transformations that could be performed for training and optimizing a personalization model.

5.0.4. Intelligent Ranking

The business's priorities, i.e., its value proposition and profit formula, are quite important and they will require the organizations' recommendation engine to promote specific content or products, such as trending news or a promotional offer, over time. Learning the personalized recommendation list can be cast as a ranking problem. Ranking could be quite subjective as it requires determining the causal connections among various features and more importantly, calculating and assigning weights and importance scores to the connections between features.

An intelligent ranking algorithm should be capable of being trained on top of the domain experts' knowledge and experience to understand the context, extract related features, and to determine the causal connections among various features over time. The goal here would be to concentrate on the *change* from statistical modeling to novel forms of modeling such as those based on deep learning to improve potential similarities among descriptors to build a more accurate ranking. A future work direction may focus on empowering the exploitation of personality, behaviour and attitudes when generating recommendation that can be used to improve correlations discovery among features Yakhchi, Beheshti, Ghafari, and Orgun (2020).

5.0.5. Intelligent Customer Engagement

Customer engagement, i.e., the emotional connection between a customer and a brand, is a vital process for organizations as highly engaged customers buy more, promote more, and demonstrate more loyalty. Powering real-time personalized product and content recommendations can improve customer engagement. Most of the existing recommender systems have been designed to recommend the right products to users, however, very few recommender systems focused on the emotional connection between a customer and a brand, to better understand their needs, feelings, and environment.

Provenance Beheshti, Motahari-Nezhad, and Benatallah (2012), i.e., the logged records of an object (e.g., service, or product) or the documentation of procedure in an object's life stage that logs the sequences that the object was derived and evolved, can help in tracing the customer activities and their interactions with specific products over time. A provenance-aware Recommender System would be ab interesting future direction as it will facilitate understanding and analyzing customer behaviour and activities. Another future plan can be to adopt Gamification methods, i.e., the application of game-design units and game principles in non-game contexts, to learn from the users' activities Beheshti et al. (2020b).

6. Conclusion

Choosing the right product, in the age of big data and AI, becoming a growing challenge due to the enormous volume and variety of products and services. Recommender Systems can deal with this problem providing personalized suggestions of items that will likely match the user's requirements. In particular, recommender systems can improve the user experience in online applications by assisting the users to explore different types of items. In this book chapter, we provided an summary of different types of real-world RSs, their methods, and applications along with challenges and opportunities in the age of big data and AI. We discussed how the Machine Learning algorithms are adopted by recommender systems to obtain knowledge from diverse sources of input information, e.g., user data (ratings, reviews, and comments) and item data (categories, descriptions, and images) when generating a personalized recommendation for users. We presented applications and motivating scenarios and discussed how increasing growth of the cognitive technology mixed with development in areas such as AI, Data Science, and User Personalization and Engagement, shows a clear opportunity to take the Recommender Systems the next level.

References

Hirevue; job search engine. https://www.hirevue.com/products/video-interviewing, 2019. Accessed: 2019-09-14.

Gediminas Adomavicius and Alexander Tuzhilin. Toward the next generation of recommender systems: a survey of the state-of-the-art and possible extensions. *Knowledge and Data Engineering, IEEE Transactions on*, 17(6):734–749, June 2005. ISSN 1041-4347. .

Gediminas Adomavicius and Alexander Tuzhilin. Context-aware recommender systems. In *Recommender Systems Handbook*, pages 217–253. Springer, 2011.

Jae-wook Ahn, Peter Brusilovsky, Jonathan Grady, Daqing He, and Sue Yeon Syn. Open user profiles for adaptive news systems: help or harm? In *Proceedings of the 16th international conference on World Wide Web*, pages 11–20. ACM, 2007.

Chris Anderson. *The Long Tail*. Random House Business, 2006.

Michael R Anderson, Dolan Antenucci, Victor Bittorf, Matthew Burgess, Michael J Cafarella, Arun Kumar, Feng Niu, Yongjoo Park, Christopher Ré, and Ce Zhang. Brainwash: A data system for feature engineering. In *Cidr*, 2013.

Marko Balabanović and Yoav Shoham. Fab: Content-based, collaborative recommendation. *Commun. ACM*, 40(3):66–72, March 1997. ISSN 0001-0782. . URL `http://doi.acm.org/10.1145/245108.245124`.

Linas Baltrunas, Bernd Ludwig, Stefan Peer, and Francesco Ricci. Context relevance assessment and exploitation in mobile recommender systems. *Personal Ubiquitous Comput.*, 16(5):507–526, June 2012. ISSN 1617-4909. . URL `http://dx.doi.org/10.1007/s00779-011-0417-x`.

Amin Beheshti, Boualem Benatallah, Reza Nouri, Van Munin Chhieng, HuangTao Xiong, and Xu Zhao. Coredb: a data lake service. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management, CIKM 2017, Singapore, November 06 - 10, 2017*, pages 2451–2454, 2017a.

Amin Beheshti, Boualem Benatallah, Reza Nouri, and Alireza Tabebordbar. Corekg: a knowledge lake service. *Proc. VLDB Endow.*, 11(12):1942–1945, 2018.

Amin Beheshti, Boualem Benatallah, Quan Z. Sheng, and Francesco Schiliro. Intelligent knowledge lakes: The age of artificial intelligence and big data. In Leong Hou U, Jian Yang, Yi Cai, Kamalakar Karlapalem, An Liu, and Xin Huang, editors, *Web Information Systems Engineering - WISE 2019 Workshop, Demo, and Tutorial, Hong Kong and Macau, China, January 19-22, 2020, Revised Selected Papers*, volume 1155 of *Communications in Computer and Information Science*, pages 24–34. Springer, 2019.

Amin Beheshti, Vahid Moraveji Hashemi, Shahpar Yakhchi, Hamid Reza Motahari-Nezhad, Seyed Mohssen Ghafari, and Jian Yang. personality2vec: Enabling the analysis of behavioral disorders in social networks. In James Caverlee, Xia (Ben) Hu, Mounia Lalmas, and Wei Wang, editors, *WSDM '20: The Thirteenth ACM International Conference on Web Search and Data Mining, Houston, TX, USA, February 3-7, 2020*, pages 825–828. ACM, 2020a.

Amin Beheshti, Shahpar Yakhchi, Salman Mousaeirad, Seyed Mohssen Ghafari, Srinivasa Reddy Goluguri, and Mohammad Amin Edrisi. Towards cognitive recommender systems. *Algorithms*, 13(8):176, 2020b.

Seyed-Mehdi-Reza Beheshti, Hamid Reza Motahari-Nezhad, and Boualem Benatallah. Temporal provenance model (tpm): model and query language. *arXiv preprint arXiv:1211.5009*, 2012.

Seyed-Mehdi-Reza Beheshti, Boualem Benatallah, and Hamid Reza Motahari-Nezhad. Scalable graph-based OLAP analytics over process execution data. *Distributed Parallel Databases*, 34(3):379–423, 2016.

Seyed-Mehdi-Reza Beheshti, Alireza Tabebordbar, Boualem Benatallah, and Reza Nouri. On automating basic data curation tasks. In Rick Barrett, Rick Cummings, Eugene Agichtein, and Evgeniy Gabrilovich, editors, *Proceedings of the 26th International Conference on World Wide Web Companion, Perth, Australia, April 3-7, 2017*, pages 165–169. ACM, 2017b.

Dmitry Bogdanov and Perfecto Herrera. How much metadata do we need in music recommendation? a subjective evaluation using preference sets. In *ISMIR*, pages 97–102, 2011.

Dmitry Bogdanov, Joan Serrà, Nicolas Wack, Perfecto Herrera, and Xavier Serra. Unifying low-level and high-level music similarity measures. *Multimedia, IEEE Transactions on*, 13 (4):687–701, 2011.

Dirk Bollen, Bart P Knijnenburg, Martijn C Willemsen, and Mark Graus. Understanding choice overload in recommender systems. In *Proceedings of the fourth ACM conference on Recommender systems*, pages 63–70. ACM, 2010.

Matthias Braunhofer, Mehdi Elahi, and Francesco Ricci. Techniques for cold-starting context-aware mobile recommender systems for tourism. *Intelligenza Artificiale*, 8(2):129–143, 2014.

Darin Brezeale and Diane J Cook. Automatic video classification: A survey of the literature. *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on*, 38(3):416–430, 2008.

Robin Burke. Knowledge-based recommender systems, 2000.

Robin Burke. Hybrid recommender systems: Survey and experiments. *User Modeling and User-Adapted Interaction*, 12(4):331–370, 2002. URL `./papers/burke-umuai-ip-2002.pdf`.

Luca Canini, Sergio Benini, and Riccardo Leonardi. Affective recommendation of movies based on selected connotative features. *Circuits and Systems for Video Technology, IEEE Transactions on*, 23(4):636–647, 2013.

Iván Cantador, Martin Szomszor, Harith Alani, Miriam Fernández, and Pablo Castells. Enriching ontological user profiles with tagging history for multi-domain recommendations. 2008.

Xiaofei Chao, Mark J Huiskes, Tommaso Gritti, and Calina Ciuhu. A framework for robust feature selection for real-time fashion style recommendation. In *Proceedings of the 1st international workshop on Interactive multimedia for consumer electronics*, pages 35–42, 2009.

Guanliang Chen and Li Chen. Recommendation based on contextual opinions. In *User Modeling, Adaptation, and Personalization*, pages 61–73. Springer, 2014.

Mark Claypool, Anuja Gokhale, Tim Miranda, Pavel Murnikov, Dmitry Netes, and Matthew

Sartin. Combining content-based and collaborative filters in an online newspaper. In *Proceedings of the ACM SIGIR '99 Workshop on Recommender Systems: Algorithms and Evaluation*, Berkeley, California, 1999. ACM.

Yashar Deldjoo and Reza Ebrahimi Atani. A low-cost infrared-optical head tracking solution for virtual 3d audio environment using the nintendo wii-remote. *Entertainment Computing*, 12:9–27, 2016.

Yashar Deldjoo, Mehdi Elahi, Paolo Cremonesi, Franca Garzotto, Pietro Piazzolla, and Massimo Quadrana. Content-based video recommendation system based on stylistic visual features. *Journal on Data Semantics*, pages 1–15, 2016.

Christian Desrosiers and George Karypis. A comprehensive survey of neighborhood-based recommendation methods. In Francesco Ricci, Lior Rokach, Bracha Shapira, and Paul B. Kantor, editors, *Recommender Systems Handbook*, pages 107–144. Springer, 2011.

Mehdi Elahi. Adaptive active learning in recommender systems. In *Proceedings of the 19th International Conference on User Modeling, Adaption and Personalization*, pages 414–417, 2011.

Mehdi Elahi. *Empirical Evaluation of Active Learning Strategies in Collaborative Filtering.* PhD thesis, Ph. D. Dissertation. Free University of Bozen-Bolzano, 2014.

Mehdi Elahi and Lianyong Qi. Fashion recommender systems in cold start. 2020.

Mehdi Elahi, Valdemaras Repsys, and Francesco Ricci. Rating elicitation strategies for collaborative filtering. In *Proceedings 12th International Conference on E-Commerce and Web Technologies*, pages 160–171, 2011.

Mehdi Elahi, Francesco Ricci, and Neil Rubens. Adapting to natural rating acquisition with combined active learning strategies. In *ISMIS'12: Proceedings of the 20th international conference on Foundations of Intelligent Systems*, pages 254–263, Berlin, Heidelberg, 2012. Springer-Verlag. ISBN 978-3-642-34623-1. .

Mehdi Elahi, Mouzhi Ge, Francesco Ricci, David Massimo, and Shlomo Berkovsky. Interactive food recommendation for groups. In *Poster Proceedings of the 8th ACM Conference on Recommender Systems, RecSys 2014, Foster City, Silicon Valley, CA, USA, October 6-10, 2014*. 2014.

Mehdi Elahi, Yashar Deldjoo, Farshad Bakhshandegan Moghaddam, Leonardo Cella, Stefano Cereda, and Paolo Cremonesi. Exploring the semantic gap for movie recommendations. In *Proceedings of the Eleventh ACM Conference on Recommender Systems*, pages 326–330. ACM, 2017.

Ralph Ewerth, Martin Schwalb, Paul Tessmann, and Bernd Freisleben. Estimation of arbitrary camera motion in MPEG videos. In *Proceedings of the 17th International Conference on Pattern Recognition, 2004. ICPR 2004.* IEEE, 2004. . URL `https://doi.org/10.1109/icpr.2004.1334181`.

A. Felfernig and R. Burke. Constraint-based recommender systems: Technologies and research issues. In *Proceedings of the 10th International Conference on Electronic Commerce*, ICEC '08, pages 3:1–3:10, New York, NY, USA, 2008. ACM. ISBN 978-1-60558-075-3. . URL `http://doi.acm.org/10.1145/1409540.1409544`.

Ignacio Fernández-Tobías and Iván Cantador. Exploiting social tags in matrix factorization models for cross-domain collaborative filtering. In *CBRecSys@ RecSys*, pages 34–41, 2014.

Jill Freyne and Shlomo Berkovsky. Intelligent food planning: personalized recipe recommendation. In *IUI*, pages 321–324. ACM, 2010.

Jill Freyne and Shlomo Berkovsky. Evaluating recommender systems for supportive technologies. In *User Modeling and Adaptation for Daily Routines*, pages 195–217. Springer, 2013.

Simon Funk. Netflix update: Try this at home, 2006. `http://sifter.org/~simon/journal/20061211.html`.

Daniel Gallego, Wolfgang Woerndl, and Gabriel Huecas. Evaluating the impact of proactivity in the user experience of a context-aware restaurant recommender for android smartphones. *Journal of Systems Architecture*, 59(9):748–758, 2013.

Seyed Mohssen Ghafari, Shahpar Yakhchi, Amin Beheshti, and Mehmet A. Orgun. Social context-aware trust prediction: Methods for identifying fake news. In *Web Information*

Systems Engineering - WISE 2018 - 19th International Conference, Dubai, United Arab Emirates, November 12-15, 2018, Proceedings, Part I, volume 11233 of *Lecture Notes in Computer Science*, pages 161–177. Springer, 2018.

Mohammad Hammoud, Dania Abed Rabbou, Reza Nouri, Seyed-Mehdi-Reza Beheshti, and Sherif Sakr. DREAM: distributed RDF engine with adaptive query planner and minimal communication. *Proc. VLDB Endow.*, 8(6):654–665, 2015. . URL `http://www.vldb.org/pvldb/vol8/p654-Hammoud.pdf`.

Negar Hariri, Bamshad Mobasher, and Robin Burke. Context-aware music recommendation based on latenttopic sequential patterns. In *Proceedings of the sixth ACM conference on Recommender systems*, pages 131–138. ACM, 2012.

Naieme Hazrati and Mehdi Elahi. Addressing the new item problem in video recommender systems by incorporation of visual features with restricted boltzmann machines. *Expert Systems*, page e12645.

Ruining He and Julian McAuley. Ups and downs: Modeling the visual evolution of fashion trends with one-class collaborative filtering. In *proceedings of the 25th international conference on world wide web*, pages 507–517, 2016.

Ruining He, Chen Fang, Zhaowen Wang, and Julian McAuley. Vista: A visually, socially, and temporally-aware model for artistic recommendation. In *Proceedings of the 10th ACM Conference on Recommender Systems*, RecSys '16, pages 309–316, New York, NY, USA, 2016. ACM. ISBN 978-1-4503-4035-9. . URL `http://doi.acm.org/10.1145/2959100.2959152`.

Timothy Heiderich. Cinematography techniques: The different types of shots in film. *Ontario Mining Assosiation. Recuperado de https://www. oma. on. ca/en/contestpages/resources/free-report-cinematography. pdf*, 2018.

Weiming Hu, Nianhua Xie, Li Li, Xianglin Zeng, and Stephen Maybank. A survey on visual content-based video indexing and retrieval. *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on*, 41(6):797–819, 2011.

Dietmar Jannach, Markus Zanker, Alexander Felfernig, and Gerhard Friedrich. *Recommender Systems: An Introduction*. Cambridge University Press, 2010.

Marius Kaminskas, Francesco Ricci, and Markus Schedl. Location-aware music recommendation using auto-tagging and hybrid matching. In *Proceedings of the 7th ACM conference on Recommender systems*, pages 17–24. ACM, 2013.

Peter Knees, Tim Pohle, Markus Schedl, and Gerhard Widmer. A music search engine built upon audio-based and web-based similarity measures. In *Proceedings of the 30th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 447–454. ACM, 2007.

Yehuda Koren. Factorization meets the neighborhood: A multifaceted collaborative filtering model. In *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '08, pages 426–434, New York, NY, USA, 2008a. ACM. ISBN 978-1-60558-193-4. . URL `http://doi.acm.org/10.1145/1401890.1401944`.

Yehuda Koren. Factorization meets the neighborhood: a multifaceted collaborative filtering model. In *KDD '08: Proceeding of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 426–434, New York, NY, USA, 2008b. ACM. ISBN 978-1-60558-193-4. .

Yehuda Koren and Robert Bell. Advances in collaborative filtering. In Francesco Ricci, Lior Rokach, Bracha Shapira, and Paul Kantor, editors, *Recommender Systems Handbook*, pages 145–186. Springer Verlag, 2011.

Taras Lehinevych, Nikolaos Kokkinis-Ntrenis, Giorgos Siantikos, A Seza Dogruöz, Theodoros Giannakopoulos, and Stasinos Konstantopoulos. Discovering similarities for content-based recommendation and browsing in multimedia collections. In *Signal-Image Technology and Internet-Based Systems (SITIS), 2014 Tenth International Conference on*, pages 237–243. IEEE, 2014.

Qing Li and Byeong Man Kim. An approach for combining content-based and collaborative filters. In *Proceedings of the Sixth International Workshop on Information Retrieval with*

Asian Languages - Volume 11, AsianIR '03, pages 17–24, Stroudsburg, PA, USA, 2003. Association for Computational Linguistics. . URL http://dx.doi.org/10.3115/1118935.1118938.

Henry Lieberman and Ted Selker. Out of context: Computer systems that adapt to, and learn from, context. *IBM Syst. J.*, 39(3&4):617–632, 2000.

Pasquale Lops, Marco De Gemmis, and Giovanni Semeraro. Content-based recommender systems: State of the art and trends. In *Recommender systems handbook*, pages 73–105. Springer, 2011.

Yucheng Low, Danny Bickson, Joseph Gonzalez, Carlos Guestrin, Aapo Kyrola, and Joseph M Hellerstein. Distributed graphlab: a framework for machine learning and data mining in the cloud. *Proceedings of the VLDB Endowment*, 5(8):716–727, 2012.

Marcelo Garcia Manzato. gsvd++: Supporting implicit feedback on recommender systems with metadata awareness. In *Proceedings of the 28th Annual ACM Symposium on Applied Computing*, SAC '13, pages 908–913, New York, NY, USA, 2013. ACM. ISBN 978-1-4503-1656-9. . URL http://doi.acm.org/10.1145/2480362.2480536.

Pablo Messina, Vicente Dominquez, Denis Parra, Christoph Trattner, and Alvaro Soto. Exploring content-based artwork recommendation with metadata and visual features. *UMUAI*, 2018.

Raymond J Mooney and Loriene Roy. Content-based book recommending using learning for text categorization. In *Proceedings of the fifth ACM conference on Digital libraries*, pages 195–204. ACM, 2000.

H. R. Naphide and Thomas Huang. A probabilistic framework for semantic video indexing, filtering, and retrieval. *IEEE Transactions on Multimedia*, 3(1):141–151, March 2001. ISSN 1520-9210. .

Nagarajan Natarajan, Donghyuk Shin, and Inderjit S Dhillon. Which app will you use next?: collaborative filtering with interactional context. In *Proceedings of the 7th ACM conference on Recommender systems*, pages 201–208. ACM, 2013.

Hai Thanh Nguyen, Thomas Almenningen, Martin Havig, Herman Schistad, Anders Kofod-Petersen, Helge Langseth, and Heri Ramampiaro. Learning to rank for personalised fashion recommender systems via implicit feedback. In *Mining Intelligence and Knowledge Exploration*, pages 51–61. Springer, 2014.

Michael J. Pazzani. A framework for collaborative, content-based and demographic filtering. *Artif. Intell. Rev.*, 13(5-6):393–408, December 1999. ISSN 0269-2821. . URL http://dx.doi.org/10.1023/A:1006544522159.

Michael J. Pazzani and Daniel Billsus. The adaptive web. chapter Content-based Recommendation Systems, pages 325–341. Springer-Verlag, Berlin, Heidelberg, 2007. ISBN 978-3-540-72078-2. URL http://dl.acm.org/citation.cfm?id=1768197.1768209.

Person. Definition of fashion, 2019. URL https://fashion-history.lovetoknow.com/alphabetical-index-fashion-clothing-history/definitionn-fashion.

Hua Quanping. Analysis of collaborative filtering algorithm fused with fashion attributes. *International Journal of u-and e-Service, Science and Technology*, 8(10):159–168, 2015.

Zeeshan Rasheed, Yaser Sheikh, and Mubarak Shah. On the use of computable features for film classification. *Circuits and Systems for Video Technology, IEEE Transactions on*, 15 (1):52–64, 2005.

Al Mamunur Rashid, Istvan Albert, Dan Cosley, Shyong K. Lam, Sean M. Mcnee, Joseph A. Konstan, and John Riedl. Getting to know you: Learning new user preferences in recommender systems. In *Proceedings of the 2002 International Conference on Intelligent User Interfaces, IUI 2002*, pages 127–134. ACM Press, 2002.

Steffen Rendle. Factorization machines with libfm. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 3(3):57, 2012.

Paul Resnick and Hal R Varian. Recommender systems. *Communications of the ACM*, 40(3): 56–58, 1997.

Francesco Ricci and Quang Nhat Nguyen. Acquiring and revising preferences in a critique-based mobile recommender system. *Intelligent Systems, IEEE*, 22(3):22–29, 2007.

Francesco Ricci, Lior Rokach, and Bracha Shapira. Recommender systems: Introduction and challenges. In *Recommender Systems Handbook*, pages 1–34. Springer US, 2015.

Mohammad Hossein Rimaz, Mehdi Elahi, Farshad Bakhshandegan Moghadam, Christoph Trattner, Reza Hosseini, and Marko Tkalcic. Exploring the power of visual features for the recommendation of movies. In *Proceedings of the 27th ACM Conference on User Modeling, Adaptation and Personalization*, pages 303–308, 2019.

Ralph J. Roberts, Lisa D. Hager, and Christine Heron. Prefrontal cognitive processes: Working memory and inhibition in the antisaccade task. *Journal of Experimental Psychology: General*, 123(4):374–393, 1994. . URL https://doi.org/10.1037/0096-3445.123.4.374.

Sujoy Roy and Sharath Chandra Guntuku. Latent factor representations for cold-start video recommendation. In *Proceedings of the 10th ACM Conference on Recommender Systems*, RecSys '16, pages 99–106, New York, NY, USA, 2016. ACM. ISBN 978-1-4503-4035-9. . URL http://doi.acm.org/10.1145/2959100.2959172.

Neil Rubens, Mehdi Elahi, Masashi Sugiyama, and Dain Kaplan. Active learning in recommender systems. In *Recommender Systems Handbook - chapter 24: Recommending Active Learning*, pages 809–846. Springer US, 2015. .

Stefano Savian, Mehdi Elahi, and Tammam Tillo. Optical flow estimation with deep learning, a survey on recent advances. In *Deep Biometrics*, pages 257–287. Springer, 2020.

Andrew I. Schein, Alexandrin Popescul, Lyle H. Ungar, and David M. Pennock. Methods and metrics for cold-start recommendations. In *SIGIR '02: Proceedings of the 25th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 253–260, New York, NY, USA, 2002. ACM. ISBN 1-58113-561-0. .

Klaus Seyerlehner, Markus Schedl, Tim Pohle, and Peter Knees. Using block-level features for genre classification, tag classification and music similarity estimation. *Submission to Audio Music Similarity and Retrieval Task of MIREX 2010*, 2010.

Usman Shahbaz, Amin Beheshti, Sadegh Nobari, Qiang Qu, Hye-Young Paik, and Mehregan Mahdavi. irecruit: Towards automating the recruitment process. In *Service Research and Innovation - 7th Australian Symposium, ASSRI 2018, Sydney, NSW, Australia, September 6, 2018, and Wollongong, NSW, Australia, December 14, 2018, Revised Selected Papers*, volume 367 of *Lecture Notes in Business Information Processing*, pages 139–152. Springer, 2018.

Cees G.M. Snoek and Marcel Worring. Multimodal video indexing: A review of the state-of-the-art. *Multimedia Tools and Applications*, 25(1):5–35, Jan 2005. ISSN 1573-7721. . URL https://doi.org/10.1023/B:MTAP.0000046380.27575.a5.

Xiaoyuan Su and Taghi M. Khoshgoftaar. A survey of collaborative filtering techniques. *Adv. in Artif. Intell.*, 2009:4:2–4:2, January 2009. ISSN 1687-7470. . URL http://dx.doi.org/10.1155/2009/421425.

Yap-Peng Tan, D. D. Saur, S. R. Kulkami, and P. J. Ramadge. Rapid estimation of camera motion from compressed video with application to video annotation. *IEEE Trans. Cir. and Sys. for Video Technol.*, 10(1):133–146, February 2000. ISSN 1051-8215. . URL https://doi.org/10.1109/76.825867.

Chun-Yuen Teng, Yu-Ru Lin, and Lada A Adamic. Recipe recommendation using ingredient networks. In *Proceedings of the 4th Annual ACM Web Science Conference*, pages 298–307. ACM, 2012.

LLC Timely Development. Netflix prize, 2008. http://www.timelydevelopment.com/demos/NetflixPrize.aspx.

Michele Trevisiol, Luca Chiarandini, and Ricardo Baeza-Yates. Buon appetito: recommending personalized menus. In *Proceedings of the 25th ACM conference on Hypertext and social media*, pages 327–329. ACM, 2014.

Qingqing Tu and Le Dong. An intelligent personalized fashion recommendation system. In *2010 International Conference on Communications, Circuits and Systems (ICCCAS)*, pages 479–485. IEEE, 2010.

TURI. Graphlab recommender factorization recommender, 2018. URL https://turi.com/products/create/docs/generated/graphlab.recommender.factorization_ recommender.FactorizationRecommender.html.

Hee Lin Wang and Loong-Fah Cheong. Affective understanding in film. *IEEE Transactions on Circuits and Systems for Video Technology*, 16(6):689–704, jun 2006. . URL https://doi.org/10.1109/tcsvt.2006.873781.

Shoujin Wang, Liang Hu, Yan Wang, Longbing Cao, Quan Z Sheng, and Mehmet Orgun. Sequential recommender systems: challenges, progress and prospects. *arXiv preprint arXiv:2001.04830*, 2019.

Yuanyuan Wang, Stephen Chi-Fai Chan, and Grace Ngai. Applicability of demographic recommender system to tourist attractions: A case study on trip advisor. In *Proceedings of the The 2012 IEEE/WIC/ACM International Joint Conferences on Web Intelligence and Intelligent Agent Technology - Volume 03*, WI-IAT '12, pages 97–101, Washington, DC, USA, 2012. IEEE Computer Society. ISBN 978-0-7695-4880-7. . URL http://dx.doi.org/10.1109/WI-IAT.2012.133.

Robert West, Ryen W White, and Eric Horvitz. From cookies to cooks: Insights on dietary patterns via analysis of web usage logs. In *Proceedings of the 22nd international conference on World Wide Web*, pages 1399–1410. International World Wide Web Conferences Steering Committee, 2013.

Shahpar Yakhchi, Amin Beheshti, Seyed Mohssen Ghafari, and Mehmet Orgun. Enabling the analysis of personality aspects in recommender systems. *arXiv preprint arXiv:2001.04825*, 2020.

Bo Yang, Tao Mei, Xian-Sheng Hua, Linjun Yang, Shi-Qiang Yang, and Mingjing Li. Online video recommendation based on multimodal fusion and relevance feedback. In *Proceedings of the 6th ACM international conference on Image and video retrieval*, pages 73–80. ACM, 2007.

Xiaojian Zhao, Guangda Li, Meng Wang, Jin Yuan, Zheng-Jun Zha, Zhoujun Li, and Tat-Seng Chua. Integrating rich information for video recommendation with multi-task rank aggregation. In *Proceedings of the 19th ACM international conference on Multimedia*, pages 1521–1524. ACM, 2011.