Making Business Partner Recommendation More Effective: Impacts of Combining Recommenders and Explanations through User Feedback

Oznur Alkan, Massimiliano Mattetti, Sergio Cabrero Barros and Elizabeth M. Daly

IBM Research, Europe

Abstract

Business partnerships can help businesses deliver on opportunities they might otherwise be unable to facilitate. Finding the right business partner (BP) involves understanding the needs of the businesses along with what they can deliver in a collaboration. BP recommendation meets this need by facilitating the process of finding the right collaborators to initiate a partnership. In this paper, we present a real world BP recommender application which uses a similarity based technique to generate and explain BP suggestions, and we discuss how this application is enhanced by integrating a solution that *1. dynamically combines different recommender algorithms, and 2. enhances the explanations to the recommendations*, in order to improve the user's experience with the tool. We conducted a preliminary focus group study with domain experts which supports the validity of the enhancements achieved by integrating our solution and motivates further research directions.

Keywords

explanation, heterogeneous data sources, orchestration, interaction

1. Introduction and Background

Strategic partnerships are important for businesses to grow and explore more complex opportunities [1, 2], since these partnerships can open up possibilities to new products, services, markets and resources [2]. However, finding the right business partner (BP) with whom to form a partnership is challenging, since one has to face a large space of possible partners and process many different data sources to find the BPs that best suit ones requirements. BP recommendation systems can be a solution as they help to analyze the available information around BPs. In this paper, our focus is on *BP C*onnector, a real-world application that provides company to company recommendations, where the companies themselves become the *subject* items to recommend to each other, and the recommendations must suit the preferences of both parties involved. This setting is studied under the *reciprocal recommender systems* research [3]; these systems have arisen as an extension to classical item-based recommendation processes to deal with scenarios where users become the item being recommended to other users. In this context,

Joint Proceedings of the ACM IUI 2021 Workshops, April 13-17, 2021, College Station, USA
oalkan2@ie.ibm.com (O. Alkan); massimiliano.mattetti@ibm.com (M. Mattetti); sergiocabrerobarros@gmail.com (S.C. Barros); elizabeth.daly@ie.ibm.com (E.M. Daly)

both the end user and the user being recommended should accept the *matching* recommendation to yield a successful recommender performance [4]. Hence, for BP recommendations, both the users who ask for recommendations and the recommendation items themselves are BPs, and the goal is to satisfy the interests of the two sides of the partnership.

BP Connector has already been deployed by an organization with a large ecosystem of BPs to foster collaborations among them in order to create a virtuous cycle, where a successful engagement between BPs promotes the business interest of the instigating organization itself. The system defines two roles for the partnership: the *beneficiary* and the *helper*. Beneficiary refers to the company who is seeking assistance in a specific territory, technology, etc., whereas the helper refers to the company who states that it can provide assistance. The system allows companies to first specify whether they are seeking help or asking for help, and then asks them to fill in a form to specify the details around their interests and expertise. Both the *beneficiary* and the *helper* complete the same forms, therefore providing information around the same features. These features constitute the BP profiles and are used as both the user and the item profiles by the underlying recommender to generate BP recommendations [5]. More specifically, a beneficiary requesting a BP connection is the user who is seeking for recommendations of helper BPs, where the helper BPs constitute the items of this recommendation setting. The *initial solution* used a content-based recommender [6] which is based on the similarity between the profiles of the beneficiaries and the helpers to generate both the recommendations and the explanations, where the explanations reveal the degree of the similarity between the two profiles. Therefore, the quality of the recommendations depends on the quality of the information that is completed through the web forms. However, the information entered may not always be *complete* (users might have missed out some fields or sections), *accurate* (users might have mistakenly provided incorrect information) or *recent* (users might have provided information some time ago which may be outdated). This results in user and item profiles not reflecting the current interests and actual expertise of the BPs, which may degrade not only the quality of the recommendations but also the explanations. However, the organization deploying BP Connector has access to data around BPs such as the historical sales records and product certifications, which, if integrated into the recommender logic, would improve the quality of the recommendations and the explanations, and this can help users to make better decisions [7].

Although using more data has benefits, one important challenge is that data around BPs exists in different heterogeneous sources and these data sources have different coverage. Moreover, there is a possibility that additional data sources may become available over time. To handle this, hybrid recommendation approaches can be used, which can essentially fuse the benefits of multiple data sources and leverage the complementary knowledge in order to provide better recommendations [8, 9]. Hybrid recommenders support combining different recommenders built on different data sources. For example, one model might be a collaborative filtering recommender that uses a ratings matrix including the feedback provided by the companies regarding their previous partnerships, whereas another model could be a content-based recommender. In such cases, it would be important to combine explanations generated from different recommenders as well, which will assist users' in the decision making process.

Motivated by these discussions, in this paper, we present our solution called *M*ulti *S*ource *E*vidence *R*ecommender, henceforth referred to as MSER, which is built to enhance the recommendation and the explanation facilities of BP Connector. MSER can ensemble different recommendation algorithms that are built on top of different data sources. Moreover, it can receive explanations from these different recommenders, which are presented to the user to support their decision making process. MSER can re-rank and post-process the recommendations based on pre-configured business rules. When we developed MSER, we were aware that different companies may have different goals when seeking a partnership, where these goals strongly influence which features and which data sources may be the most relevant to support the recommendation process. For example, *company A* may need a local presence for a sales opportunity, therefore the location information may be the most important factor, whereas *company B* may be looking for an expert in a specific technology, therefore, accurate information on product certifications and sales performance could be the most important factor. To support this, we designed MSER to *enable users to provide feedback around the data sources they are interested in*, in order to better align the recommendations with the users' dynamic interests.

Integrating MSER to the BP Connector application leads to substantial changes over the initial version. These changes lead us to initially formulate two research questions: *1. What is the difference in subjective recommendation quality between the recommendations generated by a single recommender and recommendations generated by MSER? 2. How do the users perceive the explanations generated by MSER?* In order to investigate these research questions, a preliminary focus group study with domain experts is conducted which motivates us for further research.

In the rest of the paper, we first present a brief review of the related art, and then describe our solution we designed for BP Connector application in order to enhance its recommendation and explanation capabilities. Then, we present the initial focus group study and discuss our findings. We conclude with proposals for future research.

2. Related Work

BP recommendations have been studied considering different sources of data and different types of methods [10]. [1] presents a solution for recommending BPs to individual business users through combining item-based fuzzy semantic similarity and collaborative filtering techniques. In [11], authors discuss the reciprocity aspect of the BP recommendations, where they propose a machine learning approach to predict customer-supplier relationships. As discussed before, BP recommendations fall into the category of reciprocal recommender systems, which have been applied to many online social services such as online dating [12, 13], social media [14], recruitment [15] and online mentoring systems [16]. All these domains including business partnership increasingly rely on the concept of matching users with the right users. They differ from the traditional items-to-users recommendation as their goal is not just to predict a user's preference towards a passive item, but to find a match where preferences of both sides are satisfied [17].

Our solution, MSER, orchestrates different recommender algorithms that run on disparate data sources, which relates our work to hybrid recommenders [18, 9]. MSER can be considered as a recommender ensemble [19], which is a particular type of hybrid recommenders in which the recommender algorithms to combine are treated as black boxes. In [20], authors present several approaches for generating an ensemble of collaborative models based on a single collaborative filtering algorithm. In [21], authors presented a hybrid recommender with an interactive interface which allows users to adjust the weights assigned to each recommender through sliders. This proposed system is designed to provide recommendations on media content leveraging multiple social sources. With the enhancements designed for BP Connector, we aim to enable users to interact with the recommenders. In this regard, our initial choice for a new interactive UI fell on a chatbot system. Among the possible interaction models, chatbot systems have seen a steep increase in popularity in the recent years driven by the wide adoption of mobile messaging applications [22]. They also represent a natural interface for conversational recommenders which provide recommendations to the users through dialogue [23].

Considering the explanations, in [24], authors reviewed the literature on explanations in decision-support systems, where they distinguished between variables such as the length of the explanations, their vocabulary and their presentations, and they concluded that additional studies are necessary to assess the impact of these variables. In [25], authors introduced the concept of reciprocal explanation where the user who is looking for a connection is also presented with an explanation on what would be the interest of the other party in establishing a mutual connection. Kouki et al. [26] studied how to provide useful hybrid explanations that capture informative signals from a multitude of data sources, and conducted a crowd sourced user study to evaluate several different design approaches for hybrid explanations. In another work [27], authors proposed a taxonomy that categorizes different explainable recommenders and the authors argue that future research should create new kinds of information, interaction, and presentation styles. To this end, MSER is designed to support combining explanations generated by different recommenders through dynamic user feedback, and it can support different explanation styles.

The primary contribution of this paper is to describe how the recommendation and explanation generation facilities of an existing recommender application, BP Connector, are enhanced through designing a solution called MSER, which combines recommendations and explanations through user feedback.

3. Proposed Solution: Multi-Source Evidence Recommender (MSER)

The enhancements designed for BP Connector are encapsulated within MSER, which is designed around four main components, *Controller, Connector Layer, Rank-Combiner and Post-Processor*, as depicted in Figure 1. The figure shows the high-level view of the components in which the components' interactions are labelled in sequence to show the execution flow. Below, we summarize the details of these components.

Controller connects the client application with the underlying recommender logic, thus making it responsible for orchestrating the execution flow of MSER. It exposes a *get_recommendations* method, which takes two parameters: 1. *query parameters*, which specifies the properties of the recommendation request, 2. *recommender weights*, which determines the weights that should be assigned to different recommender algorithms, where a weight of 0 indicates that the corresponding recommender should be excluded from the recommendation process.

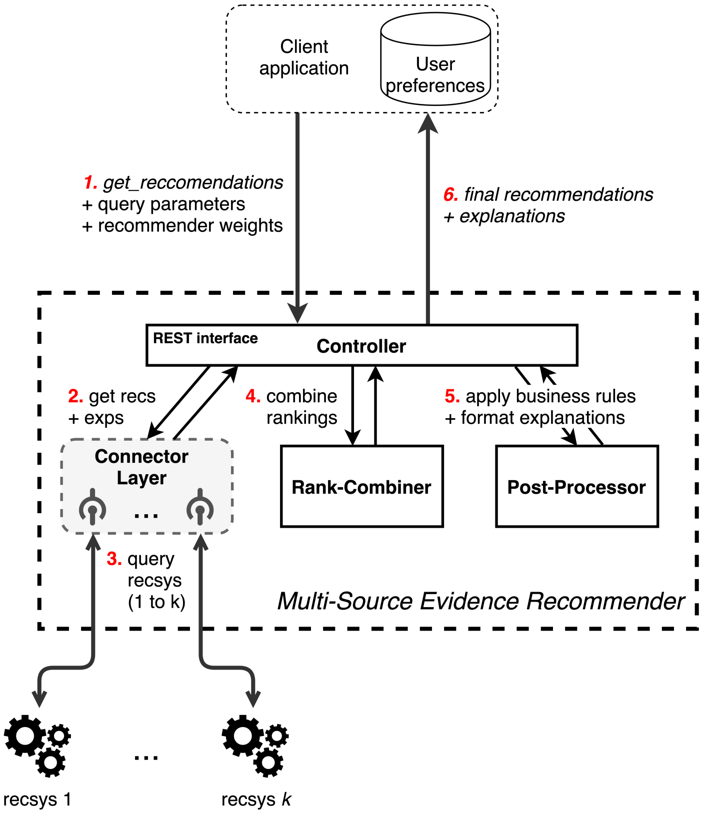

Figure 1: MSER - System architecture.

Once a recommendation request is received from the client application through calling the *get_recommendations* method of the *Controller* (1), *Controller* first forwards this request to the *Connector Layer* (2) which in turn calls the configured recommender systems to receive the recommendations and the explanations (3). The responses received from the recommenders are then handed over to the *Rank-Combiner* together with the recommender weights. *Rank-Combiner* computes the ranking of the final recommendation list using a linear combination of the recommendation scores [9, 28], where it adjusts the weighting based on the *recommender weights* it receives from the *Controller* (4). The ranked list is then processed by the *Post-Processor* which applies the business rules (5). *Avoid recommending a firm to another firm if their business needs do not coincide or if they operate in different geographies* is an example of a business rule that BP Connector enforces. Each recommender can send an explanation associated with the recommended BP, which is also combined by *Post-Processor* to present the final explanation in a way that is pre-configured within the solution. Lastly, the final list is returned to the client application (6).

Integrating MSER to BP Connector. The

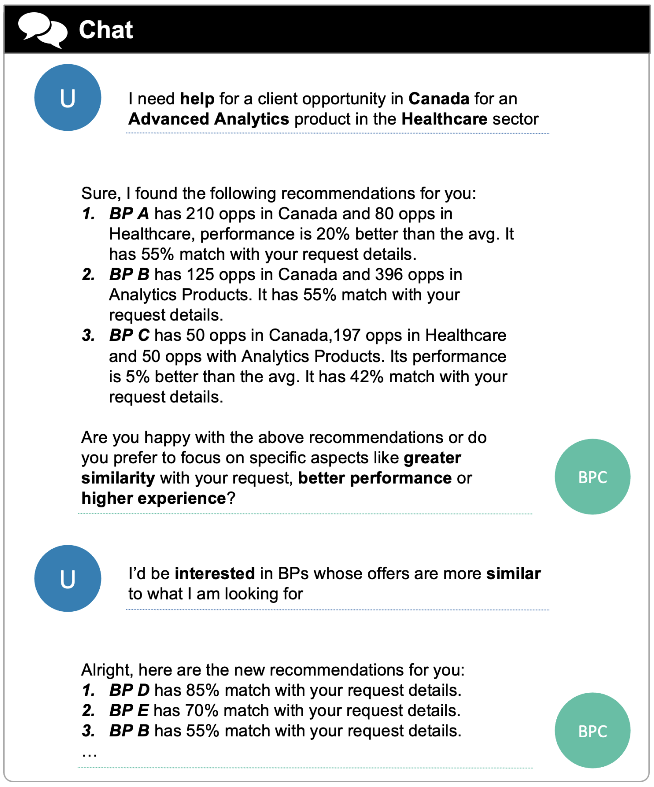

Figure 2: BP Connector - Sample screenshot for the dialogue-based interface.

initial version of BP Connector used a single *Similarity-Based Recommender (SBR)*, and through the adoption of MSER, the solution has been enhanced with two additional recommenders: *Expert Recommender (ER)* and *Performance Recommender (PR)*. ER has been serving a production application in the sales domain for more than two years, therefore, the existing recommender service was plugged into the BP Connector solution, whereas PR is specifically designed for BP Connector.

SBR computes the similarity between the features that the *beneficiary* and the *helper* specified in the initial web forms. To achieve this, SBR first represents the form parameters as a vector of weights, and then it computes the similarity between these vectors using the *Cosine Similarity* metric. The web form data represents a kind of *explicit* user profile [13], and SBR tries to connect a *proactive* user (beneficiary) with a *reactive* one (helper), so that the reciprocal recommendation satisfies the preferences of both sides.

ER formulates the recommendation problem as an Information Retrieval process [29], where the sales history of a BP corresponds to a document, an attribute of a sales opportunity is a field of the document (e.g. country, sector, product), and an attribute value corresponds to a term (e.g. United States for country; banking for sector). The beneficiary request form plays the role of the query, and a *TF-IDF Similarity score*[1] is computed for each document, which represents the *proficiency score* of the helper BP corresponding to the document.

PR uses a machine learning model to predict the probability of an opportunity being won or lost by considering the expertise of a BP. It computes a probability score for a helper to win an opportunity whose characteristics match the requirements defined in the beneficiary request form. This recommender builds a *Gradient Boosting Classifier* [30] for each helper BP in the dataset using historical sales data.

Explanations. In addition to the recommendations, each of the three recommenders provide its own set of *explanations* which is combined by MSER. As for the explanations, *SBR* provides the similarity score between the helper request and the beneficiary request as an explanation. Moreover, it provides four other scores, which represent the overlap between the beneficiary request and the helper request in terms of *technology* (e.g. Analytics, Cloud, Security, etc.), *business need* (e.g. Consulting, Marketing, Sales, etc.), *industry* (e.g. Banking, Education, Healthcare, etc.) and *assistance type* (e.g. developing new sales relationship, creating new services, supporting new solutions, etc.). *ER*, on the other hand, provides the number of deals that a helper had in the past in the sector, industry, country, etc. listed in the beneficiary request form. *PR* establishes a baseline win rate given the parameters specified in the beneficiary request form. As explanation, the performance of a helper is provided as a relative increment of the win rate over the baseline's. As it is relative to a baseline value, performance can assume negative values as well.

User Interaction. The original form-based user interface of BP Connector limited the users to following a predefined set of steps. We aimed to increase the interactivity between the user and the application by designing a dialogue-based interface that sits next to original interface. From this dialogue, beneficiaries can perform the following interactions: 1. fill in request details, 2. receive recommendations, 3. guide MSER to use the required recommenders, and 4. receive explanations. A sample screenshot for the third interaction listed is given in Figure 2. The dialogue is designed to be able to elicit user preferences towards the recommendation algorithms. It assigns a weight of 1 to a recommender if the user expresses *interest* in it, or a weight of 0 if the user shows *no interest* towards it. At the beginning of the conversation a weight of 1 is assigned to each recommender. The dialogue is built using Watson Assistant[2], an existing service which is one of the natural language understanding services for conversational question answering [31].

4. Evaluation

Setup and Participants. We evaluated MSER as the new recommender behind BP Connector with two different groups. The first group involved *7 domain experts*, and the second group included *5 active users* of the application. Domain experts were employed by the organization deploying BP Connector and they worked directly with BPs. They operated at a global scale (2 in North America, 1 in Europe, 1 in Middle East and 3 in Asia). Active users included the users of the initial BP Connector before MSER deployment. Domain experts participated in a remote briefing meeting to get information about the user study. Afterwards, they filled in a survey, which was the same for all of them, and then

[1]https://lucene.apache.org/core/8_7_0/core/org/apache/lucene/search/similarities/TFIDFSimilarity.html

[2]https://www.ibm.com/cloud/watson-assistant/

Engagement details: You asked for help in business development in the areas of Sales and Development in the United Kingdom. Specifically, you asked for help in Internet of Things, Blockchain and Cloud.

Recommended Partners:

BP1

- Match score: **75%**

BP2

- Match score: **70%**

BP3

- Match score: **75%**

(a) Match score explanation

Recommended Partners:

BP1

- This partner is a **75%** match to your request.
- This partner has worked on **8** deals related to requested products. Their primary expertise is in **Cloud & Data Platform**.

BP2

- This partner is a **70%** match to your request.
- This partner has worked on **115** deals related to requested products. Their primary expertise is in **Cloud & Data Platform**.
- Based on our internal sales data and with regards to the expertise you requested, we estimate **5%** better performance compared to all registered business partners in the system.

BP3

- This partner is a **75%** match to your request.
- This partner has worked on **1** deal related to requested products. Their primary expertise is in Cognitive Applications. Rate the following statements.

(b) Short explanation

BP2

[Similarity-based explanation]

- This partner is a **70%** match to your request.
  - There is a **100%** overlap between your requested expertise and the expertise this business partner can help with.
  - There is a **0%** overlap between your requested industry and the industry this business partner wants to help with.
  - There is an **80%** overlap between your requested assistance type and the assistance this business partner provides.
  - There is a **100%** overlap between the solutions specified in your request and this business partner is providing help with.

[Experience-based explanation]

- Based on our internal sales data this business partner worked on **82** deals on **Cloud & Data Platform**, **29** deals on **Cognitive Applications** and **4** deals on **HPC Hardware**. Moreover, it has **124** opportunities in the **United Kingdom**. Finally, it worked on **12** deals of size **250k - 1M USD** and **1** deal of size over **1M USD**.

[Predicted success-based explanation]

- Based on our internal sales data and with regards to the expertise you requested, we estimate **5% higher** performance compared to all registered business partners in the system.

(c) Detailed explanation

Figure 3: Screenshots from BP Connector User Study - Examples of match score (a), short (b) and detailed (c) explanations for the recommendations generated for a sample connection request. For the detailed explanation, explanations for only BP2 is displayed.

participated in a remote focus group to discuss the results and provide further feedback. Active users, on the other hand, answered a survey personalized to their company. This was performed through selecting one of their former requests made to the initial BP Connector and generating a new set of recommendations and explanations using MSER.

The surveys were similar for both groups. During the surveys, a partnership request was explained, and three companies were recommended as potential partners, where each recommended company had one explanation accompanying it. We experimented on three types of explanations with different levels of details: 1. *match score*, 2. *short explanation*, and 3. *detailed explanation. Match score* explanation includes only the percentage value representing how much the offer of help from a company fits the help request, which is generated by SBR, whereas *short explanation* and *detailed explanation* are formed using the explanations from all three recommenders, SBR, ER and PR. For the explanations generated by SBR, *short explanation* includes only the percentage of match, (same with the the *match score* explanation), whereas the *detailed explanation* presents the details of the overlap between the offer and the request of help considering the four dimensions; *technology, business need, industry and assistance type*, as discussed in Section 3. For the explanations generated by ER, *short explanation* includes the total number of opportunities the helper BP had in the past with the products listed in the beneficiary request form, together with the product family that represents the main area of expertise of the helper BP. The *detailed explanation*, on the other hand, includes the details of this expertise, specifically, the number of opportunities for the different products, countries, sectors and the deal sizes requested by the beneficiary. Finally, the explanation generated by PR is the same for both types. Examples of the three types of explanations for the same request are given in Figure 3. If a recommender did not recommend a specific BP that appeared in the final recommendation list, its explanation was omitted from both the short and the detailed explanations.

A page of the survey showed all three companies with the same type of explanation. Subsequent survey pages showed different types. However, the order was always kept the same as follows: 1. match score, 2. short explanation, and 3. detailed explanation, since each of the next explanations adds more information to the previous one. This allows us to explore the *completeness principle* as defined in [32], where each explanation includes more information than the previous one in order to detect where information overload starts generating a problem.

Table 1
Recommendation quality perceived by the experts for each type of explanation

Exp. Type	Very good	Good	Neutral	Bad
Match score	0	5	1	1
Short	2	5	0	0
Detailed	1	4	0	1

Results and Discussion. To evaluate how participants perceive the recommendations from MSER, we examined their evaluation of the recommendations with each of the explanations provided with them. Table 1 summarizes the results for the group of experts. As can be seen from the table, the majority of the experts ranked the recommendations as *Good* independent of which explanation type was provided. However, when they were presented with more than just a match score, their ratings improved. One of the experts said *"I like that I can understand the size of their experience."*. Users, on the other hand, responded as *Neutral* when a match score was provided to them; however, receiving either a short or a detailed explanation helped them to build more confidence in the recommendations. We observed that evaluating recommendations without explanations is difficult in this context, as one cannot quantify if a partnership worked or not after it really happens. In our evaluation, however, we could only evaluate the judgment that the users made of a potential partnership; therefore, providing users with valuable explanations was key to support their decisions.

Regarding the amount of information provided *(explanation completeness)*, the preference of short versus detailed explanations was not homogeneous among participants. One participant mentioned: *"Of little value just showing a name and a percentage match"* for the *match score* type, and another one said *"I can get an idea of the experience and type of work of each partner."* for the detailed type. Some declared that the detailed explanation shows too much information and is difficult to process, whereas others mentioned that they would like to have as much information as possible to decide on future partnerships. This aligns with findings in [25] about how the cost of the decision influences the explanation effectiveness. Apart from the personal preferences, the presentation mode was also important for our participants. When they were asked about interaction and visualization, personal preferences played an important role. Participants mentioned that interactivity with the system and graphical representations of the data presented for each company are desirable. The design could therefore include an interactive interface in which users initially receive a match score, ask for a short explanation, and are able to explore the detailed explanation of each dimension individually. This would allow users to find their own balance in the explanation completeness and information overload scale.

5. Conclusion

We presented MSER which is built to enhance the recommendation and the explanation facilities of a real-world application, BP Connector that provides company to company recommendations. An initial user study revealed that the extensions enabled by MSER can improve both the recommendation and the explanation capabilities of BP Connector, and the results motivates further research. As a future work, we aim to evaluate the *scalability* of the solution by enlarging the recommender engine behind BP Connector with additional recommender systems based on additional data sources such as data around product certifications, ratings given by the beneficiaries to the helpers they connected with, and implicit preferences based on users' behaviour [33] such as requests of connections and responses to matches.

6. Acknowledgements

We would like to acknowledge the support and collaboration provided by IBM CAO team: Sanjmeet Abrol, Cindy Wu and Alice Chang.

References

[1] J. Lu, Q. Shambour, Y. Xu, Q. Lin, G. Zhang, a web-based personalized business partner recommendation system using fuzzy semantic techniques, Computational Intelligence 29 (2013) 37–69.

[2] W. Bergquist, J. Betwee, D. Meuel, Building strategic relationships: How to extend your organization's reach through partnerships, alliances, and joint ventures, in: Building strategic relationships: how to extend your organization's reach through partnerships, alliances, and joint ventures, 1995, pp. 246–246.

[3] J. Neve, I. Palomares, Hybrid reciprocal recommender systems: Integrating item-to-user principles in reciprocal recommendation, in: Companion Proceedings of the Web Conference 2020, WWW '20, Association for Computing Machinery, New York, NY, USA, 2020, p. 848–854. URL: https:

//doi.org/10.1145/3366424.3383295. doi:10.1145/3366424.3383295.

[4] I. Palomares, C. Porcel, L. Pizzato, I. Guy, E. Herrera-Viedma, Reciprocal recommender systems: Analysis of state-of-art literature, challenges and opportunities towards social recommendation, Information Fusion 69 (2021) 103–127.

[5] J. Leskovec, A. Rajaraman, J. D. Ullman, Recommendation Systems, 2 ed., Cambridge University Press, 2014, p. 292–324. doi:10.1017/CBO9781139924801.010.

[6] M. J. Pazzani, D. Billsus, Content-based recommendation systems, in: P. Brusilovsky, A. Kobsa, W. Nejdl (Eds.), The Adaptive Web, volume 4321 of *Lecture Notes in Computer Science*, Springer, Berlin/Heidelberg, 2007, pp. 325–341. URL: http://dx.doi.org/10.1007/978-3-540-72079-9_10. doi:10.1007/978-3-540-72079-9_10.

[7] D. Jannach, M. Jugovac, I. Nunes, Explanations and user control in recommender systems, in: Proceedings of the 23rd International Workshop on Personalization and Recommendation on the Web and Beyond, ABIS '19, Association for Computing Machinery, New York, NY, USA, 2019, p. 31. URL: https://doi.org/10.1145/3345002.3349293. doi:10.1145/3345002.3349293.

[8] C. C. Aggarwal, Ensemble-Based and Hybrid Recommender Systems, Springer International Publishing, Cham, 2016, pp. 199–224. URL: https://doi.org/10.1007/978-3-319-29659-3_6. doi:10.1007/978-3-319-29659-3_6.

[9] R. Burke, Hybrid recommender systems: Survey and experiments, User Modeling and User-Adapted Interaction 12 (2002). doi:10.1023/A:1021240730564.

[10] J. Bivainis, Development of business partner selection, Ekonomika 73 (2006) 7–18.

[11] J. Mori, Y. Kajikawa, H. Kashima, I. Sakata, Machine learning approach for finding business partners and building reciprocal relationships, Expert Systems with Applications 39 (2012) 10402–10407.

[12] P. Xia, B. Liu, Y. Sun, C. Chen, Reciprocal recommendation system for online dating, in: 2015 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM), 2015, pp. 234–241.

[13] L. Pizzato, T. Rej, T. Chung, I. Koprinska, J. Kay, Recon: A reciprocal recommender for online dating, in: Proceedings of the Fourth ACM Conference on Recommender Systems, RecSys '10, Association for Computing Machinery, New York, NY, USA, 2010, p. 207–214. URL: https://doi.org/10.1145/1864708.1864747. doi:10.1145/1864708.1864747.

[14] X. Cai, M. Bain, A. Krzywicki, W. Wobcke, Y. S. Kim, P. Compton, A. Mahidadia, Learning to make social recommendations: a model-based approach, in: International Conference on Advanced Data Mining and Applications, Springer, 2011, pp. 124–137.

[15] R. Liu, W. Rong, Y. Ouyang, Z. Xiong, A hierarchical similarity based job recommendation service framework for university students, Frontiers of Computer Science 11 (2016) 912–922.

[16] C.-T. Li, Mentor-spotting: recommending expert mentors to mentees for live trouble-shooting in codementor, Knowledge and Information Systems 61 (2019) 799–820.

[17] F. Vitale, N. Parotsidis, C. Gentile, Online reciprocal recommendation with

theoretical performance guarantees, in: Advances in Neural Information Processing Systems, 2018, pp. 8257–8267.
[18] C. Aggarwal, Recommender Systems, 2016. doi:10.1007/978-3-319-29659-3.
[19] R. Cañamares, M. Redondo, P. Castells, Multi-armed recommender system bandit ensembles, in: Proceedings of the 13th ACM Conference on Recommender Systems, RecSys '19, Association for Computing Machinery, New York, NY, USA, 2019, p. 432–436. URL: https://doi.org/10.1145/3298689.3346984. doi:10.1145/3298689.3346984.
[20] A. Bar, L. Rokach, G. Shani, B. Shapira, A. Schclar, Improving simple collaborative filtering models using ensemble methods, in: International Workshop on Multiple Classifier Systems, Springer, 2013, pp. 1–12.
[21] S. Bostandjiev, J. O'Donovan, T. Höllerer, Tasteweights: a visual interactive hybrid recommender system, in: Proceedings of the sixth ACM conference on Recommender systems, 2012, pp. 35–42.
[22] P. B. Brandtzaeg, A. Følstad, Why people use chatbots, in: International Conference on Internet Science, Springer, 2017, pp. 377–392.
[23] D. Jannach, A. Manzoor, W. Cai, L. Chen, A survey on conversational recommender systems., CoRR abs/2004.00646 (2020). URL: http://dblp.uni-trier.de/db/journals/corr/corr2004.html#abs-2004-00646.
[24] I. Nunes, D. Jannach, A systematic review and taxonomy of explanations in decision support and recommender systems, User Modeling and User-Adapted Interaction 27 (2017) 393–444. URL: https://doi.org/10.1007/s11257-017-9195-0. doi:10.1007/s11257-017-9195-0.
[25] A. Kleinerman, A. Rosenfeld, S. Kraus, Providing explanations for recommendations in reciprocal environments, in: Proceedings of the 12th ACM conference on recommender systems, 2018, pp. 22–30.
[26] P. Kouki, J. Schaffer, J. Pujara, J. O'Donovan, L. Getoor, User preferences for hybrid explanations, in: Proceedings of the Eleventh ACM Conference on Recommender Systems, 2017, pp. 84–88.
[27] G. Friedrich, M. Zanker, A taxonomy for generating explanations in recommender systems, AI Magazine 32 (2011) 90–98.
[28] M. Claypool, A. Gokhale, T. Miranda, P. Murnikov, D. Netes, M. Sartin, Combining content-based and collaborative filters in an online newspaper, 1999.
[29] A. Costa, F. Roda, Recommender systems by means of information retrieval, in: Proceedings of the International Conference on Web Intelligence, Mining and Semantics, 2011, pp. 1–5.
[30] A. Natekin, A. Knoll, Gradient boosting machines, a tutorial, Frontiers in neurorobotics 7 (2013) 21. doi:10.3389/fnbot.2013.00021.
[31] D. Braun, A. Hernandez-Mendez, F. Matthes, M. Langen, Evaluating natural language understanding services for conversational question answering systems, in: Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue, Association for Computational Linguistics, 2017, pp. 174–185.
[32] T. Kulesza, S. Stumpf, M. Burnett, S. Yang, I. Kwan, W.-K. Wong, Too much, too little, or just right? ways explanations impact end users' mental models, in: 2013 IEEE Symposium on Visual Languages and Human Centric

Computing, IEEE, 2013, pp. 3–10.
[33] L. Pizzato, T. Chung, T. Rej, I. Koprinska, K. Yacef, J. Kay, Learning user preferences in online dating, in: Proceedings of the Preference Learning (PL-10) Tutorial and Workshop, European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML PKDD), Citeseer, 2010.