

# Explanation in Multi-Stakeholder Recommendation for Enterprise Decision Support Systems\*

Giandomenico Cornacchia<sup>1</sup>[0000-0001-5448-9970],  
Francesco M. Donini<sup>2</sup>[0000-0003-0284-9625],  
Fedelucio Narducci<sup>1</sup>[0000-0002-9255-3256],  
Claudio Pomo<sup>1</sup>[0000-0001-5206-3909], and  
Azzurra Ragone<sup>3</sup>

<sup>1</sup> Politecnico di Bari – Via E. Orabona 4, Bari (I-70125), Italy  
{name.surname}@poliba.it

<sup>2</sup> Università della Tuscia – Via S.M. in Gradi 4, Viterbo (I-01100), Italy  
donini@unitus.it

<sup>3</sup> EY – Via Oberdan 40, Bari (I-70125), Italy  
azzurra.ragone@it.ey.com

**Abstract.** Business agility requires support from recommendation systems, but explaining recommendations may yield information disclosure. We analyze how to provide explanations in the scenario of Multi-Stakeholder Recommendation where the sensible information of one stakeholder should not be disclosed in the explanation to another stakeholder. Among the several types of explanations analyzed, counterfactual explanations come off best as they allow the system to preserve each stakeholder’s privacy and sensitive information in terms of preferences.

**Keywords:** Decision Support Systems. Multi-Stakeholder Recommender Systems. Explanations. Counterfactual explanations.

## 1 Introduction

Business Agility requires tools to support personalized information access, most notably Recommendation Systems. Business Decision Support System (DSS) like e-commerce or streaming platforms, social media, multimedia applications, booking systems, among others, exploit Recommendation Systems to help users find their way through the multitude of information available to them. We can consider a Recommendation System as a specific DSS, evaluated by its ability to propose appealing items to a user. However, recently, this kind of systems evolved to match a better user experience. This aspect is addressed by the latest research beyond accuracy metrics. In this direction, a significant role is played by explanation.

---

\* Accepted manuscript. The final authenticated version is available online at [https://doi.org/10.1007/978-3-030-79022-6\\_4](https://doi.org/10.1007/978-3-030-79022-6_4)

With the social emergence of Recommendation Systems as a well-established tool for orienting oneself in scenarios involving a choice among a substantial number of options, users became increasingly choosy. To make a user feel comfortable about the provided suggestion, it is necessary to explain why he received that recommendation from the system. Hence, explaining a recommendation provides a clear tentative to engage the user, gain her trust, and give her the best user experience. Moreover, in the context of a real Recommendation System, the agents playing a role are not just the users. In the last few years, a new research field has emerged that considers different stakeholders involved in the recommendation process, leading to Multi-Stakeholder Recommendation Systems (MS-RS). This perspective is also acknowledged in the literature as Reciprocal Recommendation, and researchers have deeply investigated how to make a recommendation acceptable to both parties involved in the recommendation process. Thanks to emerging studies about the effect of recommendations on different user clusters and what impact specific groups of items have on the overall process, MS-RS attract an increasing interest. Furthermore, MS-RS are involved in many real scenarios based on transactions between a consumer who is looking for a product/service and a provider who wants to match the preferences of as many consumers as possible to sell them his items.

In this paper, we investigate the explanations that can be given to two kinds of users of an MS-RS: *consumers* and *providers*. In the MS-RS context, each agent has some personal information that should not be revealed to other agents. We review different kinds of explanations and determine that a counterfactual explanation comes off best, since it explains to a generic consumer how her preferences drive the recommendation process, and to a generic provider how a different strategy could change consumers' recommendation lists.

The rest of the paper is organized as follows: in the next section we summarize the state of the art for MS-RS, pointing to the (very few) works about explanation in this field. Then in Section 3 we set up some notation that formalizes profiles, utilities, and recommendations in MS-RS. In Section 4 we review the available methods that can be used to compute balanced recommendations in MS-RS, and then in Section 5, after considering several types of explanations, we focus our attention on counterfactual explanations. The final section summarizes our conclusions and draws some future directions of research.

## 2 Related Work

This work cuts across two main topics in the research area of Recommender Systems: Explanations and Multiple stakeholders. From a general point of view, recommender systems are linked to the idea of learning-to-rank. Ideally, a generic user would like to receive a list of recommendations of the most appealing items for her. Generally, this list ranks items in descending order, starting from the most important item to the least important one. Hence, a Recommendation System estimates the user's *utility* of a set of items by optimizing an accuracy metric. However, this kind of metrics do not take into account some other aspects of the

utility score computation [19,24]. A new evaluation perspective tries to put them in the loop beyond accuracy metrics [25]. This family of metrics are helpful to estimate, for example, the novelty [7], the diversity [15], or the serendipity [19] of a recommendation with the aim of improving the user experience. In this context, a crucial role is played by explanation [22]. Explainable recommendations are a research field that emerged years ago when early models whose aim was to suggest items appeared in the scientific literature [12,21,28]. Tintarev *et al.* [23] provide seven different dimensions to consider when an explanation is provided: user’s trust, satisfaction, persuasiveness, efficiency, effectiveness, scrutability, and transparency. Accordingly, when a Recommendation System suggests an item to a generic user, she could ask why she receives that suggestion. A good explanation could impact at least one of the above-mentioned dimensions. Related to this aspect, Gedikli *et al.* [10] study how different explanation types and strategies affect the final result of the process, and provide guidelines to evaluate each of these aspects.

However, providing an explanation is particularly challenging and difficult when more than one kind of user is involved in the recommendation process, as in the case of Multi-Stakeholder Recommender Systems (MS-RS). Such systems are useful in a real recommendation scenario like e-commerce, where also the provider of products is involved in the recommendation process. Another classical scenario is dating, in which the recommendation has to be acceptable to both kinds of users of the transaction [27]. Following this idea, group Recommendation Systems were proposed, with the aim of maximising the utility of each stakeholder in the group [18]. In this direction, it is clear that the MS-RS approach is to devise a strategy that includes the utility of different stakeholders (like in a multi-side approach) and this approach was generalised to every recommendation task [2,3]. Abdollahpouri *et al.* [6] propose a general model for MS-RSs, which considers three kinds of users in the loop: the consumer who receives the recommendation, the system that supports the recommendation process, and the provider who feeds the system catalogue. Naturally, in the MS-RS scenario, all involved kinds of users must be taken into account in the explanation process.

To the best of our knowledge, research in MS-RS did not deeply address yet such explanation aspects. Verdeaux *et al.* [26] consider counterfactual explanations in MS-RS scenario, but only from the consumer’s viewpoint. In that work, the authors adopt a causality-based approach for the counterfactual explanation, but they do not consider explanations from the provider’s perspective. Conversely, in our work we propose a counterfactual explanation both to the consumer and to the provider, based on each stakeholder’s own utility function. In this way, the private/confidential information of each stakeholder is never revealed to the counterpart during the explanation process.

### 3 Notation

In this section, we formally define the viewpoint of each stakeholder, in terms of both her profile and the recommendation the MS-RS gives her. This will set

our notation for a formalization of counterfactual explanations of such recommendations in Section 5.

In the present study, we envisage two types of stakeholders: consumers and providers (we leave the inclusion of the MS-RS utility for future work). We denote the set of all consumers as  $\mathcal{C} = \{c_1, c_2, \dots\}$ , and the set of providers as  $\mathcal{P} = \{p_1, p_2, \dots\}$ .

Items are enumerated into a set  $\mathcal{I} = \{I_1, I_2, \dots\}$ . To simplify the formulas in the paper, we represent an item just by its index in  $\mathcal{I}$ , so that a list of items  $\langle I_3, I_7, I_2 \rangle$  (e.g., a recommendation) will be just a list of natural numbers  $\langle 3, 7, 2 \rangle$ . We do not delve in this paper into the characteristics of items—i.e., their features.

In general, the MS-RS keeps a *profile* for each consumer  $c$  that collects her preferences or requirements. In this paper, we consider a consumer profile as a list of items in decreasing consumer preference order:  $P_c = \langle i_1, i_2, \dots \rangle$ —i.e., an ordered list of items the consumer has chosen (or preferred in the past) the most. The recommendation process consists of a utility function  $u_c : \mathcal{C} \times \mathcal{I} \rightarrow \mathbb{R}^+$ . Such a utility can be represented by an accuracy, diversity, serendipity metric, or any other consumer utility, with the constraint that  $u_c$  is such that  $u_c(c, i_1) \geq u_c(c, i_2) \geq u_c(c, i_3) \geq \dots$ , i.e., the utility is coherent with the consumer’s profile. A *recommendation for a consumer  $c$*  is an ordered list of items, denoted by  $R_c = \langle i_1, i_2, \dots \rangle$ , meaning that the MS-RS suggests the consumer the new item  $i_1$  as most suitable, then item  $i_2$  as a second choice, etc. The recommendation must be coherent with the consumer’s utility, i.e.,  $u_c(c, i_1) \geq u_c(c, i_2) \geq u_c(c, i_3) \geq \dots$ .

Similarly, the *profile* of a provider  $p$  is a collection of her requirements. In this case, the provider’s requirements represent some strategy that could maximize e.g., profits, stock clearance, budget allocation, some other objectives, or a combination of some of them. Each strategy yields an ordered list of items,  $P_p = \langle i_1, i_2, i_3 \dots \rangle$ , with the meaning that the provider would prefer to sell item  $i_1$  the most, then item  $i_2$ , etc. Observe that such a set of strategies could be as large as needed, taking into account all possible choices the provider could make. Similarly to consumers, the recommender implements a utility function  $u_p : \mathcal{P} \times \mathcal{I} \rightarrow \mathbb{R}^+$ , giving a value to items from the provider’s point of view, with the constraint that such utility is coherent with the provider’s strategy, that is,  $u_p(p, i_1) \geq u_p(p, i_2) \geq u_p(p, i_3) \geq \dots$ .

In the next section, we summarize methods combining the above utilities into recommendations that balance between different stakeholders’ objectives.

## 4 Computing recommendations in MS-RS

Traditional Recommendation Systems are built to recommend to end users a ranked list of items based on the user’s tastes and preferences. Accordingly, the development of conventional collaborative filtering [16] algorithms has been centred on minimising an error to maximise unilateral utility metrics (i.e., the consumer’s point of view). However, it is now recognized that the recommendation task is not unilateral.

Considering only the user utility in the recommendation task raises a problem called ‘‘Popularity Bias’’ [4] in which the Recommendation System suggests the most popular items with higher probability than less frequent ones. In this case, the problem was addressed by spreading diversity in the recommendation task [17]. Yet, approaches that promote diversity still lack the provider’s perspective.

In a more recent work [1], Abdollahpouri *et al.* proposed another way to implement the recommendation task in a MS-RS setting by using learning-to-reranking methodologies. The core problem is to compose the (sometimes) diverging interests of the two principal stakeholders: consumers and providers. Consumers want a personalized recommendation list that maximizes their utility, whereas providers want their products to have a higher probability of being sold. To find a new recommendation list which reflects a possible equilibrium point between consumer and provider utility functions a possible approach is to introduce a maximization problem of log-likelihood estimation. Following the same direction adopted by Abdollahpouri *et al.* [1] the problem becomes

$$\max_{\beta} \mathcal{L}(u_p|R_c, \mathcal{I}) = \sum_{j=1}^m \log(u_c(c, i_j)) + \beta \times \log(u_p(p, i_j))$$

In this formulation  $\mathcal{L}$  denotes the loss of the log-likelihood estimation,  $m$  is the number of items presented in the recommendation list  $R_c$ . This maximization problem aims to fine-tune the parameter  $\beta$  to generate the new list of recommendation  $R_c^*$  optimized for both consumer and provider utility functions. Furthermore, the idea is to provide a new recommendation list that is not disruptive from the consumer’s viewpoint. Hence,  $R_c^*$  is expected to be as similar as possible to  $R_c$  and this similarity could be expressed by a distance measure like the Kendall tau. This metric operates on the relative pairwise order of the items between the two lists to measure their difference.

Considering these two aspects, it is possible to introduce a new formulation for the generation of  $R_c^*$  in the form of

$$\min_{\beta, \gamma} \mathcal{L}(u_p|R_c, \mathcal{I}) = \mathcal{L}(u_p|\mathcal{I}) + \gamma(1 - \hat{K}(R_c, R_c^*))$$

The first term is referred to the optimization problem for generation  $R_c^*$  considering both consumer- and provider utilities. The term  $\hat{K}(R_c, R_c^*)$  is the kernel-ized version [14] of the Kendall tau distance that regularized the loss as a similarity-based distance of  $R_c^*$  from the original  $R_c$ , while  $\beta$  are the weights of the optimized functions and  $\gamma$  is a hyper-parameter responsible for balancing the effect between the two terms of optimization.

## 5 Explanations for MS-RS

This analysis does not consider *white-box* explanations, since—by exploiting the inner mechanisms implemented by the recommendation algorithm for generating

the explanation—they might reveal preferences and private information that the part (*i.e.*, each stakeholder) does not desire to disclose. Hence in this section, we focus on *black-box* approaches, analyzing two of the most prominent ones, namely, counterfactual and contrastive explanations.

### 5.1 Counterfactual explanations.

We now discuss counterfactual explanations in the context of MS-RS. Counterfactual explanations follow the causality theory by Halpern&Pearl [11] for generating an explanation. Explanations depending on causality have not yet stood out in the Recommendation System research area, but recently they are starting to attract interest.

In their work, Halpern&Pearl identify two kinds of events, *exogenous* and *endogenous*. The former are determined by external factors and define the context. The latter are the factors an agent can change to influence a result and are in this way the expected causes of that result. In our MS-RS scenario, we consider that events are exogenous or endogenous based on a stakeholder’s perspective: namely, each stakeholder sees her actions as endogenous events, while all events corresponding to other stakeholder choices are exogenous.

Clearly, in a MS-RS scenario, the only choices stakeholders can make are about their profile: a consumer might change her list of preferred items, while a provider might change his strategy. Consequently, we consider as events of our causal theory of counterfactuals the stakeholders’ profiles.

We consider this kind of approach as the most suitable for MS-RS, since we can distinguish between consumer-side and provider-side explanations, where each explanation does not reveal to a stakeholder the other stakeholders’ preferences—as they are seen as exogenous causes.

Depending on the granularity of events, the computation of an explanation could change considerably. In the approach by Verdeaux *et al.* [26], the events that cause a change in a consumer’s recommendation list are purchases of single items; eliminating a suitable subset of such events would cause a rearrangement in the recommendation list, pushing lower items upwards. In that case, choosing a minimal set of purchases that change the consumer’s preferences can be a computationally intractable problem [9]. However, for simplicity, in this preliminary paper we treat the entire profile as an event, simplifying the search for a counterfactual cause of the recommendation to a simple rearrangement of the profile—in the simplest case, just a change in the first item. In this way, we decouple our analysis from computational problems, which we will deal with in future works. Explanations from the provider’s perspective follow a similar approach: the endogenous cause of a particular recommendation to the consumer is the provider’s strategy, that is, his profile as an ordered list of items. A counterfactual explanation looks for another strategy the provider could have chosen, which would have changed the recommendation.

More formally, a *counterfactual explanation* of a recommendation  $R_c^*$  for a consumer  $c$ , with profile  $P_c$ , is a pair  $(P'_c, R_c^{*'})$ , where both  $P_c \neq P'_c$  and

$R_c^* \neq R_c^{*'}$ , to be interpreted as follows: “Had Consumer  $c$  the profile  $P_c'$ , the MS-RS would recommend  $R_c^{*'}$  instead of  $R_c^*$ ”.

In the simplest case, the recommender could focus on the first item of each list, providing an explanation of the following form: “I recommended you *Apple Phone XS* because based on your profile, you preferred *Samsung Galaxy S21* the most; if your most preferred item were *Samsung Galaxy S10*, I would suggest you *Google Pixel 5* instead”.

On the provider’s side, supposing the provider chose the strategy  $P_p$ , a *counterfactual explanation* of a recommendation  $R_c^*$  given to a consumer  $c$ , is a pair  $(P_p', R_c^{*'})$ , where both  $P_p \neq P_p'$  and  $R_c^* \neq R_c^{*'}$ , to be interpreted as follows: “Had provider  $p$  a different strategy  $P_p'$ , the MS-RS would have recommended to  $c$  the new list  $R_c^{*'}$ ”.

Again, the simplest of such explanations would be to focus on one element only; for example: “I recommended to *Early adopter #1* the item *Google Pixel 5* because *Google Pixel 5* was the first one in your priority list; had you chosen Strategy  $P_p'$ , whose most prominent item is *Samsung Galaxy S10*, I would put this item in *Early adopter #1*’s recommendation”.

Summarizing, counterfactual explanations never reveal to a stakeholder the other stakeholder’s preferences, since they refer always to each stakeholder’s own choices.

## 5.2 Contrastive explanations.

Exploiting the formal models of causation by Halpern& Pearl and extending the causal chain definition provided by Hilton [13], Miller [20] proposed contrastive explanations in the context of classical eXplainable Artificial Intelligence (XAI) for classification tasks. With the contrastive explanation, one wants to answer the question “Why P and not Q?”. For example, a XAI system classifying pictures of animals should be able to justify its outcome by answering questions like, “Why did you classify that photo as a *spider* and not as a *crab*?” Of course, a contrastive explanation presumes that the user of the system already knows in some way the items to be contrasted.—in the previous example, the classes of spiders and crabs.

While this type of approach is claimed by Miller to be very effective in the context of XAI, when moving to the context of MS-RS, however, it seems unsuitable because it may reveal indirectly other stakeholder’s preferences. To make an example, suppose that a consumer already knows items *Apple Phone XS* and *Samsung Galaxy S21*, and suppose such items are completely equivalent from the consumer’s perspective; yet the MS-RS recommended *Apple Phone XS* in a privileged position over *Samsung Galaxy S21*, just because this ordering meets the preferences of the provider. A contrastive explanation to the consumer question “Why did you put *Samsung Galaxy S21* so lower than *Apple Phone XS* if I like them both?” would have in this case no reason to put forward, but the provider’s preferences. No possible answer to the consumer seems both adequate and trustworthy here. The provider’s side contrastive explanations suffer from the same

drawback: answering about the reasons of a big discrepancy in the recommendation of very similar—from the provider’s preferences—items may reveal some consumer’s preferences that she might have declared as private knowledge—information that MS-RS is not authorized to reveal, adhering to EU GDPR, or other non-EU legislation.

## 6 Conclusion

Recommendation Systems can play an important role as a tool of Business Agility. Since a business context always presumes the existence of more than one stakeholder, we studied which kind of explanation is suitable for recommendations in an MS-RS scenario, where each stakeholder provides to the system private information that may not be disclosed to other stakeholders. It turns out that an explanation based on counterfactuals comes off best, since it can be based on the choices of each stakeholder without revealing the other’s reserved information.

Our analysis leads to two future extensions: *(i)* consider items as described by a set of features, leading to Content-Based recommendation in the MS-RS scenario, where stakeholder preferences could be expressed as preferred feature values, and counterfactual explanations should be expressed in terms of such feature preferences; *(ii)* apply different optimization functions to address the recommendation list re-rank problem, following either Game Theory optimization [8] or Pareto frontier derivation [5].

## References

1. Abdollahpouri, H., Adomavicius, G., Burke, R., Guy, I., Jannach, D., Kamishima, T., Krasnodebski, J., Pizzato, L.A.: Beyond personalization: Research directions in multistakeholder recommendation (2019)
2. Abdollahpouri, H., Adomavicius, G., Burke, R., Guy, I., Jannach, D., Kamishima, T., Krasnodebski, J., Pizzato, L.A.: Multistakeholder recommendation: Survey and research directions. *User Model. User Adapt. Interact.* **30**(1), 127–158 (2020)
3. Abdollahpouri, H., Burke, R.: Multi-stakeholder recommendation and its connection to multi-sided fairness. In: *RMSE@RecSys*. *CEUR Workshop Proceedings*, vol. 2440. CEUR-WS.org (2019)
4. Abdollahpouri, H., Mansoury, M., Burke, R., Mobasher, B.: The unfairness of popularity bias in recommendation. In: *RMSE@RecSys*. *CEUR Workshop Proceedings*, vol. 2440. CEUR-WS.org (2019)
5. Abdou, W., Bloch, C., Charlet, D., Spies, F.: Multi-pareto-ranking evolutionary algorithm. In: *EvoCOP*. *Lecture Notes in Computer Science*, vol. 7245, pp. 194–205. Springer (2012)
6. Burke, R.D., Abdollahpouri, H., Mobasher, B., Gupta, T.: Towards multi-stakeholder utility evaluation of recommender systems. In: *UMAP (Extended Proceedings)*. *CEUR Workshop Proceedings*, vol. 1618. CEUR-WS.org (2016)
7. Castells, P., Hurley, N.J., Vargas, S.: Novelty and diversity in recommender systems. In: *Recommender Systems Handbook*, pp. 881–918. Springer (2015)

8. Celli, A., Marchesi, A., Farina, G., Gatti, N.: No-regret learning dynamics for extensive-form correlated equilibrium. In: *NeurIPS* (2020)
9. Eiter, T., Lukasiewicz, T.: Complexity results for structure-based causality. *Artificial Intelligence* **142**(1), 53–89 (2002). [https://doi.org/https://doi.org/10.1016/S0004-3702\(02\)00271-0](https://doi.org/https://doi.org/10.1016/S0004-3702(02)00271-0)
10. Gedikli, F., Jannach, D., Ge, M.: How should I explain? A comparison of different explanation types for recommender systems. *Int. J. Hum. Comput. Stud.* **72**(4), 367–382 (2014)
11. Halpern, J.Y., Pearl, J.: Causes and explanations: A structural-model approach - part II: explanations. In: *IJCAI*. pp. 27–34. Morgan Kaufmann (2001)
12. Herlocker, J.L., Konstan, J.A., Riedl, J.: Explaining collaborative filtering recommendations. In: *CSCW*. pp. 241–250. ACM (2000)
13. Hilton, D., McClure, J., Slugoski, B.: The course of events: counterfactuals, causal sequences, and explanation (2005)
14. Hofmann, T., Schölkopf, B., Smola, A.J.: Kernel methods in machine learning. *Annals of Statistics* **36**(3), 1171–1220 (2008)
15. Hurley, N., Zhang, M.: Novelty and diversity in top-n recommendation - analysis and evaluation. *ACM Trans. Internet Techn.* **10**(4), 14:1–14:30 (2011)
16. Koren, Y., Bell, R.M.: Advances in collaborative filtering. In: Ricci, F., Rokach, L., Shapira, B., Kantor, P.B. (eds.) *Recommender Systems Handbook*, pp. 145–186. Springer (2011)
17. Kunaver, M., Porl, T.: Diversity in recommender systems a survey. *Know.-Based Syst.* **123**(C), 154–162 (May 2017)
18. Masthoff, J.: Group recommender systems: Combining individual models. In: *Recommender Systems Handbook*, pp. 677–702. Springer (2011)
19. McNee, S.M., Riedl, J., Konstan, J.A.: Being accurate is not enough: how accuracy metrics have hurt recommender systems. In: *CHI Extended Abstracts*. pp. 1097–1101. ACM (2006)
20. Miller, T.: Explanation in artificial intelligence: Insights from the social sciences. *Artif. Intell.* **267**, 1–38 (2019)
21. Pu, P., Chen, L.: Trust-inspiring explanation interfaces for recommender systems. *Knowl. Based Syst.* **20**(6), 542–556 (2007)
22. Tintarev, N., Masthoff, J.: A survey of explanations in recommender systems. In: *ICDE Workshops*. pp. 801–810. IEEE Computer Society (2007)
23. Tintarev, N., Masthoff, J.: Explaining recommendations: Design and evaluation. In: *Recommender Systems Handbook*, pp. 353–382. Springer (2015)
24. Vargas, S.: Novelty and diversity enhancement and evaluation in recommender systems and information retrieval. In: *SIGIR*. p. 1281. ACM (2014)
25. Vargas, S., Castells, P.: Rank and relevance in novelty and diversity metrics for recommender systems. In: *RecSys*. pp. 109–116. ACM (2011)
26. Verdeaux, W., Moreau, C., Labroche, N., Marcel, P.: Causality based explanations in multi-stakeholder recommendations. In: *EDBT/ICDT Workshops. CEUR Workshop Proceedings*, vol. 2578. CEUR-WS.org (2020)
27. Xia, P., Liu, B., Sun, Y., Chen, C.X.: Reciprocal recommendation system for online dating. In: *ASONAM*. pp. 234–241. ACM (2015)
28. Zhang, Y., Chen, X.: Explainable recommendation: A survey and new perspectives. *Found. Trends Inf. Retr.* **14**(1), 1–101 (2020)