

Biases in Decision Making

Alexander Felfernig

Institute for Software Technology,
Inffeldgasse 16b, 8010 Graz, Austria
`alexander.felfernig@ist.tugraz.at`

Abstract. Decisions are typically taken on the basis of heuristics that are a door opener for different types of decision biases. Such biases can be interpreted as a tendency to decide in certain simplified ways which can often lead to suboptimal decision outcomes. Recommender systems support users in different types of decision making tasks and thus should be aware of such biases. In this paper we provide a short overview of different types of decision biases and their impacts on recommender systems. We also discuss some issues for future work.

Keywords: Decision Making, Biases, Recommender Systems

1 Introduction

Recommender systems [5] support users in identifying relevant candidates from an item assortment. Collaborative filtering [14] is based on word-of-mouth promotion where ratings of users with similar preferences are exploited for recommending items. Content-based filtering [24] recommends items that are similar to those the user has experienced in the past. Knowledge-based approaches [4, 7] rely on semantic item knowledge that is exploited for determining recommendations. For example, constraint-based recommenders [6, 7] rely on an explicit set of constraints that support the determination of a recommendation. Finally, group recommenders determine recommendations for groups of users on the basis of group decision heuristics [20]. In this paper we focus on knowledge-based and also group and collaborative recommendation approaches. An example of a knowledge-based recommendation environment is WEEVIS [12] which is a MediaWiki extension for the definition and execution of recommender applications. Furthermore, CHOICLA [25, 26] is an environment that supports group decision tasks on the basis of group recommendation technologies. In the remainder of this paper we will discuss different types of decision biases and their role in recommendation scenarios.

2 Decision Biases

In most of the cases when users are interacting with recommender systems, they do not know their preferences beforehand but rather construct and frequently

adapt them [17, 23]. In this context, users do not optimize their decisions but apply decision heuristics which can act as a door opener for different cognitive (decision) biases [23]. In the following we provide an overview of example decision biases (100's of these exist) and their role in recommender systems.

Decoy Effects. A decision is taken depending on the context in which alternatives are presented. Thus, completely inferior alternatives can trigger changes in choice behaviors. An overview of decoy effects (context effects) is provided in Figure 1.¹ Item T is denoted as target item for which we want to increase the selection share. Item C is the competitor of T and D is assumed to be the decoy item which can be used to increase the selection share of item T. A target T is a compromise to decoy item D if it is less expensive and has a slightly lower quality. Furthermore, the attraction effect denotes a situation where T is slightly more expensive but has a significantly higher quality. Finally, asymmetric dominance denotes a situation where T is cheaper than D and has a higher quality.

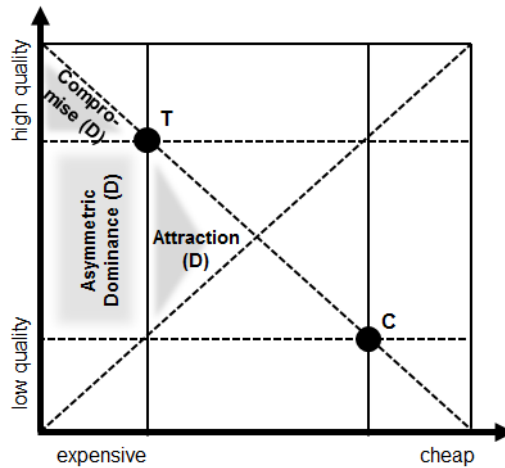


Fig. 1. An overview of different decoy effects.

An example of an asymmetric dominance effect in the evaluation of Internet connection alternatives is depicted in Figure 2. In this example, item A (the target item) dominates the decoy item in two dimensions whereas item B (the competitor) dominates the decoy item in only one dimension. The decision heuristic often applied in this context is a pairwise comparison of attributes [23]. In our example, T is the clear winner since it dominates D in two dimensions.

Impacts of decoy effects on recommender applications can be summarized as follows. First, decoy items could be exploited for increasing the selection share of specific target items [30] (see also the above example) – for sure, this application comes along with ethical issues. An empirical study related to decoy

¹ Note that the dimensions *cost* and *quality* are examples – other dimensions could be used as well [28].

Product	A (T)	B	D
price per month	30	15	50
download limit	10GB	5GB	9GB

Fig. 2. An example of asymmetric dominance. The target item (A) dominates the decoy item (D) in both dimensions whereas the competitor item (B) dominates D in only one dimension.

effects in the financial services domain is presented in [29]. Further decoy-related studies in the context of recommender systems are reported in [27] (hotel rooms) and [28] (game characters). Knowledge about decoy items can also be exploited for de-biasing purposes. Such a scenario is discussed in [10] where a dominance model is introduced to figure out dominance relationships between different items in a candidate set. On the basis of identified dominance relationships decoy items can be eliminated from the result set. Finally, decoys can also trigger the construction of explanations that are related to decision heuristics (e.g., *attribute-wise comparison*): item A (see Figure 2) is the clear winner since it dominates D in both dimensions.

Primacy and Recency. Primacy/recency effects describe situations in which items presented at the beginning and the end of a list are evaluated significantly more often than others. Since users are not interested in evaluating large lists to identify relevant items, they often focus their evaluations to the beginning and end of a list (interpretation of primacy/recency is a decision phenomenon). Murphy et al. [21] show this effect in the context of an analysis of the clicking behavior of users. Primacy/recency has also a cognitive aspect: information units at the beginning (primacy) and at the end (recency) of a list are recalled more often than information units in the middle of a list. Felfernig et al. [8, 11] show the existence of primacy/recency effects in the context of recommendation dialogs. The outcome of their analysis is that product properties at the beginning and the end of a recommendation dialog are recalled more often and are then also preferred as selection criteria when selecting items from a consideration set. This also holds for unfamiliar product properties (see Figure 3).

Impacts of primacy/recency effects on recommender applications can be summarized as follows. Similar to decoy effects, primacy/recency effects can be exploited for controlling item selection behavior when interacting with a recommender system (on the basis of attribute orderings). Further related studies are a major issue for future work, for example, the impact of different attribute orders on product comparison pages, different orderings of argumentations in item reviews, and different orderings of repair proposals [9] in knowledge-based recommendation.

Framing. The way a decision alternative is presented influences the decision behavior of the user. For example, users will prefer meat that is 80% lean compared to meat that is 20% fat. Another example is *price framing* [3]: if a user has to choose between two companies selling wood pellets (X,Y) where X sells




Fig. 3. Primacy/recency effects in the recall of product properties [8]. Properties at the beginning and the end of a list are recalled more often – also in cases where unfamiliar items were positioned at the beginning and the end (continuous line).

pellets for €24.50 per 100kg and gives a discount of €2.50 if the customer pays with cash and Y sells pellets for €22.00 per 100kg and charges a €2.50 surcharge if the customer uses a credit card, users will prefer the first alternative. This selection behavior can, for example, be explained by *prospect theory* [15] which suggests that alternatives are evaluated with regard to gains and losses where losses have a higher negative value compared to equal gains. In the price framing example, the loss would be the surcharge, in the first example the loss is associated with the 20% fat meat.

Impacts of framing effects on recommender applications can be summarized as follows. Positive framing can increase the selection probability of items. Price framing can trigger a potential focus shift from quality attributes of items to so-called secondary attributes (e.g., payment services) associated with items and – as a consequence – can change the item selection behavior of a user. In this context it must be pointed out that not every item property is equally salient at decision time and this can lead to significant shifts in selection behavior [3].

Further Effects. Priming [16, 22] represents the idea of making some properties of a decision alternative more accessible in memory such that this setting directly influences user evaluations. An example is *background priming* that exploits the fact that different page backgrounds can directly influence the decision-making process [16]. People often tend to favor the status quo compared to other decision alternatives. If *defaults* are used, users are reluctant to change predefined settings due to the fact that they are loss-averse [18, 19]. Loss in this context can mean, for example, additional costs resulting from inconsistent item settings

triggered by de-selecting a default [7]. Finally, *anchoring* denotes the effect that users often heavily rely on the first information (anchor) when evaluating decision alternatives. For example, item ratings (e.g., in collaborative filtering) of other users manipulated to be higher result in higher ratings of the current user [1]. A similar phenomenon has been observed in the context of release planning scenarios where initial evaluations manipulated the follow-up evaluations of requirements [13]. An approach to de-bias ratings in collaborative filtering recommendation scenarios is presented in [2].

3 Conclusions and Future Work

Human decisions are often not based on optimization functions but on decision heuristics that are door-openers for different decision biases. We discussed a small set of example biases and their (potential) impact on recommender applications. There are a couple of issues for future research which include an in-depth investigation of possibilities for de-biasing recommendations, the development of consensus-fostering recommendations in group decision making, and the general investigation of the properties of decision biases in group decision making.

References

1. G. Adomavicius, J. Bockstedt, S. Curley, and J. Zhang. Recommender systems, consumer preferences, and anchoring effects. In *Decisions@RecSys11*, pages 35–42, Chicago, IL, USA, 2011.
2. G. Adomavicius, J. Bockstedt, S. Curley, and J. Zhang. De-biasing user preference ratings in recommender systems. In *IntRS 2014 Workshop*, pages 2–9, 2014.
3. M. Bertini and L. Wathieu. The framing effect of price format. In *Working Paper, Harvard Business School*, pages 1–24, 2006.
4. R. Burke. Knowledge-based recommender systems. *Encyclopedia of Library and Information Systems*, 69(32):180–200, 2000.
5. R. Burke, A. Felfernig, and M. Goeker. Recommender systems: An overview. *AI Magazine*, 32(3):13–18, 2011.
6. A. Felfernig. Koba4MS: Selling Complex Products and Services Using Knowledge-Based Recommender Technologies. In *CEC 2005*, pages 92–100, 2005.
7. A. Felfernig and R. Burke. Constraint-based Recommender Systems: Technologies and Research Issues. In *10th ACM Intl. Conference on Electronic Commerce (ICEC'08)*, pages 17–26, Innsbruck, Austria, 2008.
8. A. Felfernig, G. Friedrich, B. Gula, M. Hitz, T. Kruggel, G. Leitner, R. Melcher, D. Riepan, S. Strauss, E. Teppan, and O. Vitouch. Persuasive recommendation: Serial position effects in knowledge-based recommender systems. In Y. Kort, W. IJsselstein, C. Midden, B. Eggen, and B. Fogg, editors, *Persuasive Technology*, volume 4744 of *LNCS*, pages 283–294. Springer Berlin Heidelberg, 2007.
9. A. Felfernig, G. Friedrich, M. Schubert, M. Mandl, M. Mairitsch, and E. Teppan. Plausible repairs for inconsistent requirements. In *IJCAI'09*, pages 791–796, Pasadena, California, USA, 2009.
10. A. Felfernig, B. Gula, G. Leitner, M. Maier, R. Melcher, S. Schippel, and E. Teppan. A dominance model for the calculation of decoy products in recommendation environments. In *AISB Symposium on Persuasive Technologies*, pages 43–50, 2008.

11. A. Felfernig, B. Gula, G. Leitner, M. Maier, R. Melcher, and E. Teppan. Persuasion in Knowledge-Based Recommendation. In *PERSUASIVE 2008*, pages 71–82, 2008.
12. A. Felfernig, S. Reiterer, M. Stettinger, and M. Jeran. An Overview of Direct Diagnosis and Repair Techniques in the WeeVis Recommendation Environment. In *25th Intl. Workshop on Principles of Diagnosis*, pages 1–6, Graz, Austria, 2014.
13. A. Felfernig, C. Zehentner, G. Ninaus, H. Grabner, W. Maaleij, D. Pagano, L. Weninger, and F. Reinfrank. Group decision support for requirements negotiation. In *Decisions@RecSys11*, volume 7138 of *LNCIS*, pages 105–116, 2012.
14. J. L. Herlocker, J. A. Konstan, L. G. Terveen, and J. T. Riedl. Evaluating collaborative filtering recommender systems. *ACM Trans. Inf. Syst.*, 22(1):5–53, 2004.
15. D. Kahneman and A. Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–291, 1979.
16. N. Mandel and E. Johnson. Constructing preferences online: Can web pages change what you want? In *Association for Consumer Research Conference*, pages 1–37, Montreal, Canada, 1998.
17. M. Mandl, A. Felfernig, E. Teppan, and M. Schubert. Consumer decision making in knowledge-based recommendation. *Journal of Intelligent Information Systems (JIIS)*, 37(1):1–22, 2010.
18. M. Mandl, A. Felfernig, and J. Tiihonen. Evaluating Design Alternatives for Feature Recommendations in Configuration Systems. In *CEC 2011*, pages 34–41, 2011.
19. M. Mandl, A. Felfernig, J. Tiihonen, and K. Isak. Status quo bias in configuration systems. In *24th International Conference on Industrial, Engineering and Other Applications of Applied Intelligent Systems*, pages 105–114, Syracuse, New York, 2011.
20. J. Masthoff. Group recommender systems: Combining individual models. *Recommender Systems Handbook*, pages 677–702, 2011.
21. J. Murphy, C. Hofacker, and R. Mizerski. Primacy and recency effects on clicking behavior. *Computer-Mediated Communication*, 11:522–535, 2012.
22. A. North, D. Hargreaves, and J. McKendrick. In-store music affects product choice. *Nature*, 390:132, 1997.
23. J. W. Payne, J. R. Bettman, and E. J. Johnson. *The Adaptive Decision Maker*. Cambridge University Press, Cambridge, UK, 1993.
24. M. Pazzani and D. Billsus. Learning and revising user profiles: The identification of interesting web sites. *Machine Learning*, 27:313–331, 1997.
25. M. Stettinger and A. Felfernig. Configuring Decision Tasks. In *16th International Workshop on Configuration*, pages 17–22, Novi Sad, Serbia, 2014.
26. M. Stettinger, G. Ninaus, M. Jeran, F. Reinfrank, and S. Reiterer. WE-DECIDE: A Decision Support Environment for Groups of Users. In *IEA/AIE'13*, pages 382–391, 2013.
27. E. Teppan and A. Felfernig. The asymmetric dominance effect and its role in e-tourism recommender applications. In *Wirtschaftsinformatik (WI'2009)*, pages 791–800, Vienna, Austria, 2009.
28. E. Teppan and A. Felfernig. Minimization of decoy effects in recommender result sets. *Web Intelligence and Agent Systems*, 10(4):385–395, 2012.
29. E. Teppan, A. Felfernig, and K. Isak. Decoy effects in financial service e-sales systems. In *RecSys'11 Workshop on Human Decision Making in Recommender Systems (Decisions@RecSys'11)*, pages 1–8, Chicago, IL, 2011.
30. A. Tversky and I. Simonson. Context-dependent preferences. *Management Science*, 39(10):1179–1189, 1993.