Recommender Systems for Self-Actualization

Bart P. Knijnenburg, Saadhika Sivakumar, Daricia Wilkinson
Clemson University School of Computing
215 McAdams Hall
Clemson, SC 29634
bartk@clemson.edu, ssivaku@g.clemson.edu, dariciw@g.clemson.edu

ABSTRACT
Every day, we are confronted with an abundance of decisions that require us to choose from a seemingly endless number of choice options. Recommender systems are supposed to help us deal with this formidable task, but some scholars claim that these systems instead put us inside a "Filter Bubble" that severely limits our perspectives. This paper presents a new direction for recommender systems research with the main goal of supporting users in developing, exploring, and understanding their unique personal preferences.

CCS Concepts
• **Information systems→Recommender systems**
• **Human-centered computing→Interaction paradigms.**

Keywords
Recommender Systems; Filter Bubble; Choice Overload; Self-Actualization.

1. INTRODUCTION
Recommender systems were invented in the 1990s to help users find useful and attractive items among the large assortments that came available with the growth of the Internet [31]. Such systems are now embedded in a wide range of online applications that help us find desirable products, and increasingly permeate our online interactions. As Eric Schmidt, CEO of Google, has pointed out, recommendation techniques are now employed in virtually every online service, including search engines and social networks. As people experience most of the Web through these services, it becomes very hard for them to watch or consume something that has not in some sense been tailored to their needs [15].

While the move to a personalized Web has been welcomed by most, some scholars have voiced an interesting critique against recommender systems: they argue that recommender systems put users inside a *filter bubble* that severely limits their perspectives and that may make them complacent consumers of easy-to-consume items [28].

What causes this pushback against recommender systems? Is the filter bubble simply a consequence of our psychology, or is there something wrong with the way recommenders operate? And if so, what are the consequences of this shortcoming? And how can we solve it? As part of the discussion about the past, present, and future of recommender systems, this paper attempts to start a dialogue surrounding these questions. Particularly, it acknowledges some of the shortcomings—in recommender systems as well as their users—that have led to the filter bubble, and suggests a new direction for recommender systems research to address these shortcomings. This leads us to propose the development of *Recommender Systems of Self-Actualization*: personalized systems that have the explicit goal to not just present users with the best possible items, but to *support* users in developing, exploring, and understanding their own unique tastes and preferences.

Such deep understanding of one's own tastes is a particularly important goal in decisions that have a resounding impact on one's life—e.g. choosing an education, a job, a health insurance plan, or a retirement fund. For these types of decisions, rather than have people choose the easiest option, we wish to have them develop a strong sense of determination of having selected the right path. A deep understanding of one's own tastes is also important for cultural diversity—we want people to make lifestyle choices (e.g., music, movies and fashion) based on carefully developed personal tastes, rather than blindly followed recommendations.

RecSys '16, September 15-19, 2016, Boston, MA, USA

DOI: http://dx.doi.org/10.1145/2959100.2959189

2. BEYOND THE ALGORITHM
"The algorithm accounts for only 5% of the commercial success of our recommender systems […] The interactive components of a recommender account for about 50%" — Francesco Martin [22].

Traditionally, the field of recommender systems focused on developing more accurate algorithms [20, 31]. This goal appears reasonable: the more accurate the algorithm, the better the system can predict the best recommendations for the user, which in turn should lead to a better user experience. Researchers have come to realize, though, that recommenders should go well beyond making accurate predictions. McNee et al. [25], for example, argued that "being accurate is not enough", and that recommender systems should be studied "from a user-centric perspective to make them not only accurate and helpful, but also a pleasure to use". They also [26] suggested that researchers should investigate the interactive components of the recommender system, i.e., the mechanism through which users indicate their preferences ("preference elicitation"), and the interface that displays the recommendations.

Subsequent work has indeed demonstrated that the algorithms that test best offline are not always the most successful in real life [8, 24], especially when focusing on users' subjective evaluation of the system [35]. Inspired by these findings and the need to thoroughly evaluate recommender systems from a user-centric standpoint, researchers have developed conceptual frameworks for the user-centric evaluation of recommender systems (cf. [19, 29]), and are increasingly evaluating the effects of *all* aspects of a recommender system (not just the algorithm, but also the preference elicitation method and the presentation of the recommendation list) on *all* aspects of the user's interaction experience (not just the accuracy of the algorithm, but also subjective aspects such as system satisfaction and choice satisfaction) [18].

This more inclusive perspective has uncovered several interesting problems that escape the attention of traditional recommender systems research. One of these problems is the inadequacy of existing preference elicitation methods [17]. Current recommenders rely on either implicit or explicit feedback for preference elicitation. Implicit feedback is easy to gather, but can result in "overspecialization", because the system only recommends items that it thinks the user likes: Even if the users' actual preferences are wider than the provided set of recommendations, the system will end up targeting a very specific preference. Diversifying the recommendation can prevent this [30, 36] but the main downside of diversifications is that it depends on the system's interpretation of diversity rather than the users'. Explicit feedback fares slightly better, since users can rate items negatively, thereby preventing overspecialization. However, research shows that users' ratings are often inaccurate [2, 14], arguably because consumers' preferences are often constructed on the spot [3]. If users are often unable to accurately express their preferences, then how much can be gained by accurately predicting said preferences? This conundrum demonstrates that the traditional recommender goal of accurately predicting these preferences may very well be a chimera.

Another problem is *choice overload* [4, 13]: Given that consumers construct their preference on the spot, it is no surprise that they encounter difficulties in selecting items from the Top-N recommendations. Overcoming choice overload is one of the challenges of research on the presentation of recommendations, and a good solution to this problem has yet to be devised.

3. THE FILTER BUBBLE

"[Computers] are useless. They can only give you answers."
— Pablo Picasso [10].

As the pervasiveness of recommender systems increases, arguments have emerged that attack the very nature of recommender systems. Spearheaded by Eli Pariser, these voices claim that by filtering all but the top predicted items, recommender systems provide a very myopic view of the world. Pariser argues that users get stuck in a *filter bubble*: recommenders isolate us from a diversity of viewpoints, content, and experiences, and thus make us less likely to discover and learn new things [28]

A careful examination of the Filter Bubble phenomenon in recommender systems has validated the occurrence of this effect, albeit to a lesser extent than suggested by Pariser's claims [27]. Regardless of the actuality of the effect, the *idea* of the Filter Bubble has gained a lot of traction in popular opinion and it is interesting to analyze why this may be the case.

Psychologically, the Filter Bubble plays into our tendency for *loss aversion* and our *fear of missing out*. For recommender system users, it means that in certain situations the sum of the (presumed) missed opportunities presented by all the items that are ignored by the recommender, may loom larger than the benefits of receiving a short-list of items tailored to a specific subset of their preferences. In other words, the joy of getting recommendations may be spoiled by our worry of missing out on other enjoyable items that were not recommended. This looming loss may decrease their satisfactions with the system—and even their satisfaction with their choices—because decision-making research shows that the mere thought of missed opportunities may reduce one's decision confidence [7], and cause one to regret one's decision [13].

The existence of the Filter Bubble may have a long-term consequence that is arguably worse than the fear of missing things: The possibility that users will eventually *embrace* it. This is not an unlikely scenario, because recommender systems have been shown to have *persuasive* qualities: users are prone to agree with a recommender's predicted ratings [8] and to follow a recommender's advice [12]. This creates what Lanier calls a "positive feedback loop" [28]: users will unknowingly make themselves better "fit in" with a system, i.e., make themselves more easily targetable by the algorithm. Rather than going through the trouble of developing our own unique taste, we take the default setting—something we are prone to do [33]—and simply consume whatever the recommender serves us.

The positive feedback loop leads to the very worrying concern that recommender algorithms may gradually replace human creativity and understanding [21]: if we embrace the Filter Bubble, we run the risk of getting locked in by the algorithm, which subsequently becomes a self-fulfilling prophecy. When this happens, recommenders do not just inhibit discovery and learning, they actively work against it. Pariser is afraid that personalization will create "self-fulfilling identities": Your identity shapes your recommendations, and your recommendations then shape what you believe, and what you care about.

If recommender algorithms indeed turn into self-fulfilling prophecies, then what will they recommend? Pariser argues that the items that tend to make it past the filter bubble are usually the kinds of things that are "easy" to like or consume [28]. Psychologically, this phenomenon is based on a human tendency called *temporal discounting* [9]: We tend to discount future gains, and are thus likely to choose guilty pleasures that provide instant gratification (a funny Internet meme or a spectacular action movie) over substantive educational experiences of long-term value (a complex essay or an acclaimed period piece). This leads to what Boyd calls "the psychological equivalent of obesity", where all recommended content is the cerebral equivalent of junk food [5].

4. MOVING FORWARD

"In order to find his own self, [a person] needs to live in a milieu where the possibility of many different value systems is explicitly recognized and honored" — Christopher Alexander et al. [1].

The Filter Bubble persists, despite the fact that recommender systems researchers have taken several steps in a more user-centric direction [20, 30]. One reason for this is that virtually all recommender systems are built with the goal of recommending good items to the user. If we are to solve the Filter Bubble problem, we will have to build recommender systems with a different goal in mind: a "Recommender System for Self-Actualization" (RSSA), which supports users in developing, exploring, and understanding their unique personal tastes. Below we outline how the operating principles of RSSAs differ from traditional recommenders:

RSSAs support rather than replace decision-making. Traditional recommenders turn preferences into choice options, but research shows that user preferences are fleeting, constructed on the fly and vulnerable to distorting influences, rather than well-defined, fixed, and invariant [2, 3, 14]. RSSAs take the additional step to help users develop and express their preferences.

RSSAs focus on exploration rather than consumption. RSSAs do not focus on optimizing the probability that the user will like recommendations, but instead focus on exploring underdeveloped tastes. Consequently, their recommendations can be likened to "samples" or "pathways", rather than "alternatives".

RSSAs attempt to cover users' tastes, plural. Research has shown that users' preferences are not singular, but rather multi-faceted and only loosely connected [14]. Whereas traditional recommenders are targeted to fit any part of a user's preferences, RSSAs endeavor to help the user discover all of these preferences.

Implementing these operating principles will likely require a combination of new innovations in recommender system features, interfaces, and algorithms. In this paper we highlight the most straightforward innovation: Presenting recommendations that are *not* part of the Top-N. Existing research on critiquing [6] and diversification [36] already expand the notion of the Top-N to offer a better alternatives, but these techniques still focus on providing "good" recommendations. In contrast, we suggest four completely different recommendation lists, displayed alongside but separately from the Top-N. Each new list is next discussed in detail.

"Things we think you will hate" A recommender may mistakenly predict a very low rating for some of the items that the user actually likes. Those mistakes will be hard to correct, since the system never recommends them. We propose to present a list of things the system predicts the user will hate. This allows users to either confirm or correct these predictions, thereby mitigating loss aversion. To resolve mistakes more quickly, corrections can be given a higher weight, which counters the unwanted persuasive effect of the recommender.

"Things we have no clue about" Another cause for "gaps" in the recommender's knowledge of users' tastes is the fact that certain preferences simply remain unexpressed when the system hones in too quickly on a presumptive Top-N. We propose to show a list of hard-to-predict items that may be used identify unexpressed preferences. This involves modifying existing *active learning* approaches [16] to detect not just *some* but *all* of the user's preferences.

"Things you'll be among the first to try" Solutions for the *item cold-start problem* are abound [32], but they ignore the fact that certain users may (at times) actually be excited to try out new items. We propose to present a list of yet-to-be-rated items to users that are identified (using a "hipster measure") as having a high willingness to try out new items.

"Things that are polarizing" Nearest-neighbor recommender algorithms often give recommendations that the neighbors unanimously like. It is possible that certain polarizing items divide these neighbors into rivaling camps; some of them may absolutely love a controversial item, while others absolutely hate it. Experiencing controversial items could have an important value to the user though, because such items would allow the user to develop unique tastes. We therefore propose to detect such items (e.g. by measuring the rating variability of items among the neighbors, or by sub-clustering the neighbors, and then selecting items that best discriminate between clusters) and to present them to the user.

There are several reasons why displaying items that are unrelated to the Top-N can help overcome the Filter Bubble problem. Showing items outside the Top-N is arguably the only way to combat the fear of missing things, and mitigating this fear may increase the users' satisfaction with the system and overall choice satisfaction [4]. Furthermore, by getting more feedback on items outside the Top-N, recommenders can get a better idea of the users' tastes. It can also help users to better understand their own tastes, because developing one's tastes means trying new things, even if this includes things that one may not like [34].

That said, there are other, more interaction-related features that could also contribute to the support of self-actualization. One of these features is to **connect people**. An unfortunate side effect of recommender systems is that advice-giving has become passive and indirect: Users have no idea how exactly their tastes are being used to help other users, and they have no active say in the process. We therefore propose a feature for users to actively recommend items to other users. This can contribute to a sense of fulfilment (helping others) and pride (being called upon for expertise). An algorithm that uses advances in the field of *people recommendation* can be used to drive this process. This feature is expandable beyond simple one-to-one connections between users, to recommend groups of users to come together and develop "taste-based communities" that are based on shared preferences, e.g. regarding certain controversial items.

Another suggestion is to construct a human-readable **taste profile** to help users explore and understand their own tastes. The developers of some commercial recommender systems (e.g. OkCupid, The EchoNest) have recently started to share fascinating insights into consumer tastes, using compelling infographics to highlight surprising preference dynamics, sometimes broken down by state, gender, age or other demographic dimensions. Could such analyses be personalized? For example, a simple analysis could be conducted to figure out which of your tastes are predictable (e.g. the fact that you like both Mozart and Bach), and which are unique (e.g. the fact that besides these two, you also like Nicki Minaj). This feature allows users to explore the common and unique sides of their identity, and—if comparable across users—provide a starting point for establishing sub-cultures of uniquely like-minded individuals.

5. CONCLUSION

"We need help overcoming our rationality sometimes, and allow our thoughts to wander." — David Gelernter [11].

This paper has unpacked the *filter bubble* critique of recommender systems, and proposed a new path for research: to support rather than replace human decision-making. By making us better understand our own preferences, *Recommender Systems for Self-Actualization* will improve our potential to have confidence in (and take ownership over) our life decisions. They allow us to each develop a unique personal style, thereby supporting Maslow's need for Esteem and Self-Actualization [23], and preventing the erosion of our autonomy as consumers. This would usher the field of recommender systems into a new era of computing, where systems move from serving our basic needs (e.g. "find item X") to supporting us to reach our full potential (e.g. helping us understand and reflect upon our own desires). We are actively pursuing the RSSA features presented in this paper, and we encourage others to join us in this exciting quest to pop the filter bubble.

6. ACKNOWLEDGMENTS

This research was supported in part by the NSF award IIS 1565809.

7. REFERENCES

[1] Alexander, C., Ishikawa, S. and Silverstein, M. 1977. *A Pattern Language: Towns, Buildings, Construction*. Oxford University Press.

[2] Amatriain, X., Pujol, J.M., Tintarev, N. and Oliver, N. 2009. Rate It Again: Increasing Recommendation Accuracy by User Re-rating. *Proceedings of the Third ACM Conference on Recommender Systems* (New York, NY, 2009), 173–180.

[3] Bettman, J.R., Luce, M.F. and Payne, J.W. 1998. Constructive consumer choice processes. *Journal of consumer research*. 25, 3 (1998), 187–217.

[4] Bollen, D., Knijnenburg, B.P., Willemsen, M.C. and Graus, M. 2010. Understanding choice overload in recommender systems. *Proceedings of the fourth ACM conference on Recommender systems* (Barcelona, Spain, 2010), 63–70.

[5] Boyd, D. 2010. Streams of Content, Limited Attention: The Flow of Information through Social Media. *Educase Review*. (Oct. 2010), 26–36.

[6] Chen, L. and Pu, P. 2011. Critiquing-based recommenders: survey and emerging trends. *User Modeling and User-Adapted Interaction*. 22, 1-2 (Oct. 2011), 125–150.

[7] Chernev, A. 2003. When More Is Less and Less Is More: The Role of Ideal Point Availability and Assortment in Consumer Choice. *Journal of Consumer Research*. 30, 2 (Sep. 2003), 170–183.

[8] Cosley, D., Lam, S.K., Albert, I., Konstan, J.A. and Riedl, J. 2003. Is Seeing Believing?: How Recommender System Interfaces Affect Users' Opinions. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Ft. Lauderdale, FL, 2003), 585–592.

[9] Doyle, J.R. 2013. Survey of time preference, delay discounting models. *Judgment and Decision Making*. 8, 2 (2013), 116–135.

[10] Fifield, W. 1964. Pablo Picasso: A Composite Interview. *The Paris Review*. 32, (1964), 37.

[11] Gelernter, D. 2010. Time to start taking the Internet seriously. *Edge*. (Mar. 2010), https://www.edge.org/conversation/david_gelernter–time–to–start–taking–the–internet–seriously.

[12] Gretzel, U. and Fesenmaier, D.R. 2006. Persuasion in Recommender Systems. *International Journal of Electronic Commerce*. 11, 2 (Dec. 2006), 81–100.

[13] Iyengar, S.S. and Lepper, M.R. 2000. When choice is demotivating: Can one desire too much of a good thing? *Journal of Personality and Social Psychology*. 79, 6 (2000), 995–1006.

[14] Jameson, A., Willemsen, M.C., Felfernig, A., de Gemmis, M., Lops, P., Semeraro, G. and Chen, L. 2015. Human Decision Making and Recommender Systems. *Recommender Systems Handbook*. 611–648.

[15] Jenkins, H.W. 2010. Google and the Search for the Future. *Wall Street Journal*.

[16] Karimi, R., Freudenthaler, C., Nanopoulos, A. and Schmidt-Thieme, L. 2012. Exploiting the characteristics of matrix factorization for active learning in recommender systems. *Proceedings of the sixth ACM conference on Recommender systems* (2012), 317–320.

[17] Knijnenburg, B.P., Reijmer, N.J.M. and Willemsen, M.C. 2011. Each to his own: how different users call for different interaction methods in recommender systems. *Proceedings of the fifth ACM conference on Recommender systems* (Chicago, IL, 2011), 141–148.

[18] Knijnenburg, B.P. and Willemsen, M.C. 2015. Evaluating Recommender Systems with User Experiments. *Recommender Systems Handbook*. F. Ricci, L. Rokach, and B. Shapira, eds. Springer US. 309–352.

[19] Knijnenburg, B.P., Willemsen, M.C., Gantner, Z., Soncu, H. and Newell, C. 2012. Explaining the user experience of recommender systems. *User Modeling and User-Adapted Interaction*. 22, 4-5 (2012), 441–504.

[20] Konstan, J.A. and Riedl, J. 2012. Recommender systems: from algorithms to user experience. *User Modeling and User-Adapted Interaction*. 22, 1-2 (Apr. 2012), 101–123.

[21] Lanier, J. 2010. *You Are Not a Gadget: A Manifesto*. Thorndike Press.

[22] Martin, F.J. 2009. Top 10 Lessons Learned Developing Deploying and Operating Real-world Recommender Systems. *Proceedings of the Third ACM Conference on Recommender Systems* (New York, NY, USA, 2009), 1–2.

[23] Maslow, H.A. 1943. A theory of human motivation. *Psychological Review*. 50, 4 (1943), 370–396.

[24] McNee, S.M., Albert, I., Cosley, D., Gopalkrishnan, P., Lam, S.K., Rashid, A.M., Konstan, J.A. and Riedl, J. 2002. On the recommending of citations for research papers. *Proceedings of the 2002 ACM conference on Computer supported cooperative work* (New Orleans, LA, 2002), 116–125.

[25] McNee, S.M., Riedl, J. and Konstan, J.A. 2006. Being accurate is not enough: how accuracy metrics have hurt recommender systems. *Extended abstracts on Human factors in computing systems* (Montreal, Canada, 2006), 1097–1101.

[26] McNee, S.M., Riedl, J. and Konstan, J.A. 2006. Making Recommendations Better: An Analytic Model for Human-recommender Interaction. *Extended Abstracts on Human Factors in Computing Systems* (Montreal, Canada, 2006), 1103–1108.

[27] Nguyen, T.T., Hui, P.-M., Harper, F.M., Terveen, L. and Konstan, J.A. 2014. Exploring the filter bubble: the effect of using recommender systems on content diversity. *Proceedings of the 23rd international conference on World wide web* (2014), 677–686.

[28] Pariser, E. 2012. *The filter bubble: how the new personalized Web is changing what we read and how we think*. Penguin Books.

[29] Pu, P., Chen, L. and Hu, R. 2011. A User-centric Evaluation Framework for Recommender Systems. *Proceedings of the Fifth ACM Conference on Recommender Systems* (Dublin, Ireland, 2011), 157–164.

[30] Resnick, P., Garrett, R.K., Kriplean, T., Munson, S.A. and Stroud, N.J. 2013. Bursting your (filter) bubble: strategies for promoting diverse exposure. *Proceedings of the 2013 conference on Computer supported cooperative work companion* (2013), 95–100.

[31] Resnick, P., Iacovou, N., Sushak, M., Bergstrom, P. and Riedl, J. 1994. GroupLens: An Open Architecture for Collaborative Filtering of Netnews. *ACM Conference on Computer Supported Cooperative Work* (Chapel Hill, NC, 1994), 175–186.

[32] Schein, A.I., Popescul, A., Ungar, L.H. and Pennock, D.M. 2002. Methods and Metrics for Cold-start Recommendations. *Proceedings of the 25th International Conference on Research and Development in Information Retrieval* (Tampere, Finland, 2002), 253–260.

[33] Thaler, R.H. and Sunstein, C. 2008. *Nudge : improving decisions about health, wealth, and happiness*. Yale University Press.

[34] Tian, K.T., Bearden, W.O. and Hunter, G.L. 2001. Consumer's Need for Uniqueness: Scale Development and Validation. *Journal of Consumer Research*. 28, 1 (2001), 50–66.

[35] Torres, R., McNee, S.M., Abel, M., Konstan, J.A. and Riedl, J. 2004. Enhancing digital libraries with TechLens+. *Proceedings of the 2004 joint ACM/IEEE conference on Digital libraries* (Tuscon, AZ, USA, 2004), 228.

[36] Ziegler, C.-N., McNee, S.M., Konstan, J.A. and Lausen, G. 2005. Improving recommendation lists through topic diversification. *Proceedings of the 14th international conference on World Wide Web* (Chiba, Japan, 2005), 22.