Article 25fa pilot End User Agreement

This publication is distributed under the terms of Article 25fa of the Dutch Copyright Act (Auteurswet) with explicit consent by the author. Dutch law entitles the maker of a short scientific work funded either wholly or partially by Dutch public funds to make that work publicly available for no consideration following a reasonable period of time after the work was first published, provided that clear reference is made to the source of the first publication of the work.

This publication is distributed under The Association of Universities in the Netherlands (VSNU) 'Article 25fa implementation' pilot project. In this pilot research outputs of researchers employed by Dutch Universities that comply with the legal requirements of Article 25fa of the Dutch Copyright Act are distributed online and free of cost or other barriers in institutional repositories. Research outputs are distributed six months after their first online publication in the original published version and with proper attribution to the source of the original publication.

You are permitted to download and use the publication for personal purposes. All rights remain with the author(s) and/or copyrights owner(s) of this work. Any use of the publication other than authorised under this licence or copyright law is prohibited.

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please contact the Library through email: copyright@ubn.ru.nl, or send a letter to:

University Library
Radboud University
Copyright Information Point
PO Box 9100
6500 HA Nijmegen

You will be contacted as soon as possible.

The Need for Identifying Ways to Monetize Personalization and Recommendation

Eelco Herder

eelcoherder@acm.org

Radboud Universiteit

Institute for Computing and Information Sciences

Nijmegen, The Netherlands

ABSTRACT

Research on user modeling and personalization typically only serves the needs of end-users. However, when applied in real-world, commercial contexts, recommendations should also serve the (often monetary) interests of other parties, such as platform providers, sellers and advertisers. This paper provides a brief historical perspective on the research field, contrasts this with the commercial context, and investigates the topics currently addressed at the UMAP and RecSys conferences. The paper concludes with a discussion on the need for the research community to take multi-stakeholder interests into account in the design and evaluation of adaptive systems. This would allow us to foresee unwanted effects, such as online filter bubbles, and to pro-actively find strategies to prevent them.

CCS CONCEPTS

• **Information systems** → **Personalization; Recommender systems.**

KEYWORDS

multi-sided recommender systems, fair recommender systems, commercial applications, academic research

ACM Reference Format:

Eelco Herder. 2019. The Need for Identifying Ways to Monetize Personalization and Recommendation. In *27th Conference on User Modeling, Adaptation and Personalization Adjunct (UMAP'19 Adjunct)*, June 9–12, 2019, Larnaca, Cyprus. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3314183.3323844>

1 INTRODUCTION

In the early days, User Modeling and Personalization has long been an academic pursuit without much take-up in non-academic contexts, except for the narrow field of product recommendations in web stores. However, gradually the Web arguably has become far more personalized than we might have envisaged. Social media newsfeeds are adapted to our interests, search results take our past queries, preferred language and location into account, and product recommendations have become a standard way of interacting with

customers. Unfortunately, many users seem not to appreciate the whole concept of personalization at all, but instead feel that their information freedom is limited by echo chambers, filter bubbles and ‘the algorithm’.

Much has been written about the abuse of personal data for dynamic pricing and targeted advertisements, which obviously are types of personalization that do not directly serve the interests of the end-user. This will not be the focus of this paper. Instead, we concentrate on personalization strategies and recommendation of items that are (claimed to be) primarily meant to serve the interests of the user – either because they serve typical user goals (such as discussed by Brusilovsky [3]) or recommend items that a user actually requested for. A main problem appears to be that users experience that ‘the algorithm’ would hide things for us, keeping us locked in our own ‘filter bubbles’ [12] or that recommended items seem not to match our tastes, but instead serve other goals than the user’s benefit.

In this paper, we reflect on the original purposes of user modeling and personalization and how these have been implemented in practice. First, we provide a very brief historical perspective on the original vision and goals of adaptive hypermedia, which were mainly user centered. We continue with an overview of recommendations in a commercial context, where also other stakeholders’ (monetary) interests need to be served. In the third section, we investigate current research issues in the research community by analyzing the contents of the proceedings of UMAP 2018 and RecSys 2018. Finally, we discuss the role of the research community in further shaping the application of personalization and recommendation in commercial contexts.

2 A BRIEF HISTORICAL PERSPECTIVE

In 1996, Peter Brusilovsky published a well-cited survey on adaptive hypermedia [3]. In this survey, he discussed various goals and application areas of such systems. In the field of educational hypermedia, goals would include the adaptation of course content to a learner’s knowledge level and the provision of navigational help, in order to prevent learners getting lost in the non-linear hypermedia structure. Personalized online help systems should provide context-sensitive help, in order to improve local orientation. Brusilovsky also envisaged that information retrieval systems should support users in finding the required documents, among others by suggesting relevant links to follow. In companies and other institutions, adaptive systems should be designed to support work.

A particular issue that many believed needed to be solved was ‘lostness in hyperspace’, a situation where users wouldn’t know where to go, would fail to return from interesting sidepaths, or

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

UMAP'19 Adjunct, June 9–12, 2019, Larnaca, Cyprus

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6711-0/19/06...\$15.00

<https://doi.org/10.1145/3314183.3323844>

would be so preoccupied with navigation that they were likely to forget their original goals [11].

In a 2008 survey [5], Anthony Jameson identified, among others, the following application areas and goals for adaptive systems: personalization could be used for taking over routine tasks, such as organizing email and scheduling meetings, for helping users to interact with a system by providing advice and support, and for mediating interaction with the real world – for example by filtering incoming emails or by delaying incoming messaging to avoid interrupting the user's workflow. More general, adaptive systems would help users to find information, tailor the presentation of the information, support collaboration and enhance learning.

Already in the early stages, it was (tacitly) acknowledged that personalization and user modeling have significant privacy implications, due to the large amount of user data needed for this purpose. In 2007, Alfred Kobsa published an inventory of privacy issues and user concerns [7], including user tracking and the use of cookies. Still, it was assumed that the value of personalization would motivate users to give up some privacy in exchange for improved services. He concluded that 'further advances are likely to only take place if privacy plays a much more important role in the future'.

A long-term topic in user modeling is *scrutability*, introduced by Judy Kay [6]. Scrutable user models provide insight on which user data is collected, how it is interpreted and how it contributes to adaptation decisions; scrutable user models should also provide users with means to control each step in this process. An alternative strategy for increasing user acceptance and trust are explanations, which address aspects such as transparency, scrutability, trust, persuasiveness, effectiveness, efficiency and satisfaction [15].

3 RECOMMENDATIONS IN A COMMERCIAL CONTEXT

As is commonly known, in commercial contexts, recommendations are often not only provided for the user's benefits, but also to serve a platform's goal, which is typically to make money. For example, web stores recommend products in the hope that users will buy these products and video platforms recommend videos that aim to keep the user entertained – and exposed to personalized advertisements. Recommender systems that balance the interests of several parties have been coined *multi-sided fair recommender systems* [4].

Item recommendation also takes place within social media, for example in Facebook. These recommendations are often a mix of 'genuine' recommendations intertwined with sponsored or promoted items; within the regular newsfeed, Facebook also shows sponsored posts that promote pages, groups or causes. The exact mix of these items – and the objectives behind them – may have a huge impact on the users: Facebook is an immersive platform that has been shown to have the power to influence the users' moods [8]. Any of such positive or negative effects may be subtle. For instance, it is still unclear whether and, if so, to what extent, targeted advertising on Facebook by Cambridge Analytica had any impact on the results of the Brexit referendum¹.

As much as platforms aim to (also) satisfy the needs of the end user, it is far from clear what these user needs actually are. As argued in a very entertaining CHI'18 Extended Abstract[9], simple measures of 'engagement' (such as a click-through rate) rather measure to what extent users engage in activities such as watching cat videos, instead of doing something more useful (which leads to the question whether it is a recommender system's responsibility to determine this for the user).

The practice of blending sponsored recommendations or advertisements within the actual content – be it a news feed, search results or a list of 'genuine' recommendations provides additional challenges. As long as the sponsored recommendations are in line with the user context and actual user goals, they might sufficiently blend in, but as soon as the users sense a *mismatch* [10], and the sponsored content is not sufficiently labelled as such, users start to feel uneasy and might start to distrust the user context as a whole, particularly if the user context is considered private, such as a chat window. Note that for users it is hard to recognize the reason why an item has been recommended: 'odd' items might just as well have been introduced by diversification techniques [16].

An important difference between 'academic' recommender systems and recommender systems in practice is that the latter category often lacks transparency. For instance, it is well known that Google personalized search takes more than 50 factors into account in the ranking process², but how exactly is unknown and topic of various reverse-engineering studies³.

4 LEGAL AND ETHICAL REQUIREMENTS

Arguably, the most influencing legal framework for privacy concerns is the European General Data Protection Regulation (GDPR)⁴. The GDPR is expected to influence the collection of user data, as it now can only be collected with a specific given purpose. However, it is unlikely that it will have much impact on actual recommendations or recommender algorithms. It is argued that the GDPR states that decisions 'based solely on automated processing' will only be allowed under specific conditions (including the provision of transparency and explanations [14]), but article 22.2c states that this does not yield if personalization is based on the data subject's (i.e. the user's) explicit consent⁵.

Even though it may be unlikely that Google or Facebook will be legally required to provide explanations on the algorithmic decisions that led to personalized search results or news feeds, the academic community has been pressing for transparent, bias-free personalization from various directions [13–15]. Tintarev et al [15] defined a good number of principles for the design of such explanations, but realistically most users will not extensively interact with them. In a recent keynote⁶, Mireille Hildebrandt argued that the term 'explanations' might be suboptimal: users would not be helped by insight in the technicalities behind a recommendation

¹Information Commissioner's Office (ICO). (2018). Investigation into the use of data analytics in political campaigns - Investigation update 11 July 2018. <https://ico.org.uk/media/action-weve-taken/2259371/investigation-into-data-analytics-for-political-purposes-update.pdf>

²https://en.wikipedia.org/wiki/Google_Personalized_Search

³A 2017 study on queries related to German politics reported that the effects of personalization are relatively small, <http://www.spiegel.de/netzwelt/netzpolitik/google-projekt-von-algorithmwatch-filterblase-welche-filterblase-a-1219981.html>

⁴<https://gdpr-info.eu/>
⁵<https://gdpr-info.eu/art-22-gdpr/>

⁶<https://www.universiteitleiden.nl/en/events/2019/03/fairness-and-transparency-symposium>

or a profile, but rather in a *justification*, ‘an acceptable reason for doing something’⁷.

A more philosophical but also pressing question: can we actually separate between genuine, sponsored and promoted recommendations, or could seemingly promoted results be genuine recommendations that just happen to be advertorials or otherwise sponsored content? In most countries, there are legal requirements for this content to be labeled as such⁸. In practice, advertisers often seem to prefer the label to be as invisible as possible⁹. If users cannot spot the difference between ‘real’ and ‘sponsored’ content, how could recommender systems separate between the two categories?

5 A COMPARISON OF UMAP 2018 AND RECSYS 2018

In order to get an overview on research topics currently investigated related to personalization and recommendation, we analyzed the full and short papers of ACM UMAP 2018 [2] and ACM RecSys 2018 [1]. The goals discussed in this section were explicitly stated in the abstract, introduction and/or conclusions of these papers.

The proceedings of UMAP 2018 contained a total of 36 full and short papers, in which 46 recommender goals were identified – 8 papers addressed two or more goals, 1 paper addressed even four goals. Of the 81 RecSys papers, 56 addressed a recommender system and corresponding goal, in only one paper more than one goal was explicitly investigated.

Of the 46 identified personalization goals addressed in UMAP 2018, 15 concerned *item recommendations* (32%), directly followed by *learning-related personalization* (12 times) and, to a somewhat smaller extent *health* (six items). Other personalization areas included *contacts* (focusing on kindness and trust), *information* (personalized news, reducing overload, work efficiency) and *travel* (route suggestions and venue recommendation).

For item recommendations, the goals or evaluation criteria are varied and include “good choices”, diversification, eudaimonic (meaningful in terms of self-realization) recommendations, explainability and trust (4 times) and group recommendation (2x). Goals for personalized learning include self-reflection (2x), pedagogical strategies (2x), self-regulated learning (2x), dropout prediction and scaffolding. Health goals involved mental wellbeing, personalized training programs, self-regulation and healthy shopping.

By contrast, the overwhelming majority of recommender goals mentioned in RecSys 2018 papers concerned item recommendation (44 out of 57, 77%), out of which 14 focused on improving performance. Two papers directly addressed *improving click-through rates* (CTR), which can be interpreted as a commercial goal, and another paper concerned ‘assessment after consumption’. Other goals included budget-aware recommendations, calibration and coherence, explanations (3x), gender-aware recommenders, package recommendation and multi-stakeholder recommendation (2x). Recommender systems that did not directly address items or products, aimed at finding information (2x), contacts (2x), answers to questions, learning, locations and routes.

In addition, the RecSys 2018 Proceedings included 10 industry presentations. The conference had twenty sponsors, including major players in online shopping, audio streaming and other online media. UMAP 2018 did not have an industry track and of the sponsors, only one commercial company could be identified.

From this bird’s eye overview it seems that most academic papers, both in UMAP 2018 and in RecSys 2018, largely address recommender systems that mainly or exclusively serve the end-user goals; the interests of other stakeholders (such as platform owners or sellers) are not directly addressed, at least not presented as a main research objective. However, papers in RecSys seem to use more ‘commercial’ evaluation measures.

6 DISCUSSION AND CONCLUSIONS

Historically, the UMAP community has focused on the many opportunities for user modeling and personalization for the benefit of benefit to the end user, in many different user contexts. Personalization encompasses far more than item recommendation and serves goals such as improving learning outcomes that are not always easy to be (directly) monetized.

By contrast, the RecSys community largely focuses on item recommendation, where items can be products, locations, contacts, movies or music. For this type of personalization, commercial exploitation models are usually easily recognized: companies can directly earn money by selling items piece by piece or via a subscription, or receive a fee from the producer for each item that a user has taken up (clicked upon, watched, listened to, read, bought). Free services, such as social media and video channels, earn money by bundling the actual services or items with (targeted) advertisements. From last year’s RecSys conference it can be observed that several research papers employ ‘commercial’ evaluation measures, such as the click-through rate.

When looking at recommender systems in practice, it appears that it is unavoidable – and arguably desirable – that personalization and recommendation methods are (also) applied for commercial purposes. After all, a functioning marketplace should satisfy both buyers and sellers. To make the concept of multi-sided fair recommender systems [4] more concrete, we should realize that monetary benefits or other (commercial) incentives are often essential for fostering take-up.

An important lesson that can be learned is that if the research community does not think about ways to monetize personalization methods, we might fail to (timely) recognize when and how personalization serves other interests than the benefit of the end users. One could argue that exactly this has happened in the past decade: the current discussions on filter bubbles and algorithmic bias are arguably mainly caused by too strong an (algorithmic) focus on optimizing user engagement by addressing hedonistic needs and immediate satisfaction.

Therefore, in order to turn the tide and to make personalization ‘good’ again, we think it is important that the research community will more actively envisage, develop and experiment with revenue models – taking into account the interests of all stakeholders and balancing them.

⁷<https://www.merriam-webster.com/dictionary/justification>

⁸<https://en.wikipedia.org/wiki/Advertorial>

⁹The problem of not properly labeled sponsored content is beyond the scope of this paper, I wrote more about it in <https://www.eelcoherder.com/organization-and-outreach/research-blog/19-is-this-a-recommendation-or-an-advertisement>

REFERENCES

- [1] 2018. *RecSys '18: Proceedings of the 12th ACM Conference on Recommender Systems*. ACM, New York, NY, USA.
- [2] 2018. *UMAP '18: Proceedings of the 26th Conference on User Modeling, Adaptation and Personalization*. ACM, New York, NY, USA.
- [3] Peter Brusilovsky. 1996. Methods and techniques of adaptive hypermedia. *User modeling and user-adapted interaction* 6, 2-3 (1996), 87–129.
- [4] Robin Burke. 2017. Multisided fairness for recommendation. In *Proc. FAT/ML - 4th Workshop on Fairness, Accountability, and Transparency in Machine Learning, co-located with KDD 2017*.
- [5] Anthony Jameson. 2008. Adaptive interfaces and agents. *The human-computer interaction handbook: Fundamentals, evolving technologies and emerging applications* 305330 (2008).
- [6] Judy Kay. 2006. Scrutable adaptation: Because we can and must. In *International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems*. Springer, 11–19.
- [7] Alfred Kobsa. 2007. Privacy-enhanced web personalization. In *The adaptive web*. Springer-Verlag, 628–670.
- [8] Adam DI Kramer, Jamie E Guillory, and Jeffrey T Hancock. 2014. Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences* (2014), 201320040.
- [9] Ulrik Lyngs, Reuben Binns, Max Van Kleek, and Nigel Shadbolt. 2018. So, Tell Me What Users Want, What They Really, Really Want!. In *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, alt04.
- [10] Francis T McAndrew and Sara S Koehnke. 2016. On the nature of creepiness. *New Ideas in Psychology* 43 (2016), 10–15.
- [11] Malcolm Otter and Hilary Johnson. 2000. Lost in hyperspace: metrics and mental models. *Interacting with computers* 13, 1 (2000), 1–40.
- [12] Eli Pariser. 2011. *The filter bubble: What the Internet is hiding from you*. Penguin UK.
- [13] P Resnick. 2011. Personalized filters yes; bubbles no. In *Keynote. UMAP 2011*.
- [14] Maartje ter Hoeve, Mathieu Heruer, Daan Odijk, Anne Schuth, and Maarten de Rijke. 2017. Do News Consumers Want Explanations for Personalized News Rankings. In *FATREC Workshop on Responsible Recommendation Proceedings*.
- [15] Nava Tintarev and Judith Masthoff. 2011. Designing and evaluating explanations for recommender systems. In *Recommender systems handbook*. Springer, 479–510.
- [16] Cai-Nicolas Ziegler, Sean M McNee, Joseph A Konstan, and Georg Lausen. 2005. Improving recommendation lists through topic diversification. In *Proceedings of the 14th international conference on World Wide Web*. ACM, 22–32.