

**Measuring the impact of online personalisation
Past, present and future**

Zanker, Markus; Rook, Laurens; Jannach, Dietmar

DOI

[10.1016/j.ijhcs.2019.06.006](https://doi.org/10.1016/j.ijhcs.2019.06.006)

Publication date

2019

Document Version

Final published version

Published in

International Journal of Human Computer Studies

Citation (APA)

Zanker, M., Rook, L., & Jannach, D. (2019). Measuring the impact of online personalisation: Past, present and future. *International Journal of Human Computer Studies*, 131, 160-168.
<https://doi.org/10.1016/j.ijhcs.2019.06.006>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.



Measuring the impact of online personalisation: Past, present and future

Markus Zanker^{*,a}, Laurens Rook^b, Dietmar Jannach^c

^a Free University of Bozen-Bolzano, Italy

^b Delft University of Technology, the Netherlands

^c University of Klagenfurt, Austria



ARTICLE INFO

Keywords:

Web personalisation
Adaptive systems
Recommender systems

ABSTRACT

Research on understanding, developing and assessing personalisation systems is spread over multiple disciplines and builds on methodologies and findings from several different research fields and traditions, such as Artificial Intelligence (AI), Machine Learning (ML), Human-Computer Interaction (HCI), and User Modelling based on (applied) social and cognitive psychology. The fields of AI and ML primarily focus on the optimisation of personalisation applications, and concentrate on creating ever more accurate algorithmic decision makers and prediction models. In the fields of HCI and Information Systems, scholars are primarily interested in the phenomena around the use and interaction with personalisation systems, while Cognitive Science (partly) delivers the theoretical underpinnings for the observed effects. The aim and contribution of this work is to put together the pieces about the impact of personalisation and recommendation systems from these different backgrounds in order to formulate a research agenda and provide a perspective on future developments.

1. Introduction

Nowadays, online personalisation encompasses all aspects of individualising the interaction and information content a system exchanges with its users. Different approaches and definitions towards personalisation exist, as summarised by Fan and Poole (2006). The main focus of the present paper is online personalisation, where a system: (1) makes assumptions on an individual's goals, interests and preferences, (2) in order to tailor interaction and content, (3) so as to provide the most relevant user experience. In this perspective, the personalisation process consists of a: (1) learning, (2) matching, and (3) evaluation stage – as has been proposed, for instance, by Murthi and Sarkar (2003). In this paper, we will primarily concentrate on the evaluation stage, and discuss measurement and evaluation approaches to determine the impact of personalisation mechanisms. Specifically, we will describe how personalisation research has evolved over time – from activities aimed at making systems adaptable for better usability to development of ever more accurate prediction mechanisms. Research into personalisation essentially is multidisciplinary in nature – i.e., formed and influenced mainly by academic disciplines such as Artificial Intelligence (AI) and Machine Learning (ML), Human-Computer Interaction (HCI) and Information Systems (IS), and User Modelling based on (applied) social and cognitive psychology. We acknowledge that each of these domains and traditions has made its own respective

contribution to personalisation research, and we will elaborate on specific theory contributions, evaluation methodologies and high-impact results in separate subsections.

Also, we will provide an outlook on future opportunities and challenges for research into online personalisation. We will argue that future research endeavours should aim at resolving the disaccord between chasing ever more opportunities for measurement and learning on the one hand, and, on the other hand, raising awareness about the importance of information privacy and regulatory requirements, to improve transparency in the context of a post-GDPR Europe (see Section 4).

2. Dimensions of personalisation research

Representative areas, such as adaptive hypertext and hypermedia (Brusilovsky, 1998), recommender systems (Jannach et al., 2010; Ricci et al., 2015), web personalisation (Mobasher et al., 2000), information filtering (Foltz and Dumais, 1992), and personalised information retrieval (Ghorab et al., 2013), all shaped the concepts of personalising system behaviour and/or its output towards users. As a result, the design space for adaptation and personalisation mechanisms nowadays consists primarily of the following three dimensions:

1. *User interface*: Different layouts and presentation styles (i.e.,

* Corresponding author.

E-mail address: mzanker@unibz.it (M. Zanker).

adaptation of available control functionality and input elements) are ways to tailor the interaction space between a system and its users. Dating back to the pre-Web era and traditional (Web 1.0) desktop user interfaces (UIs), this research adapts control structures and menu navigation based on system monitoring and assumptions about users' imminent needs (Greenberg and Witten, 1985). The purpose of this type of adaptivity primarily lies in making users more efficient in using information systems, as has been measured, for instance, through visual search time and required motor movements (Findlater and Gajos, 2009). Consensus exists on achievable benefits in terms of user satisfaction and performance of adaptive UI elements, but also negative impacts – even of highly accurate adaptation mechanisms – have been identified. Findlater and McGrenere (2010), for instance, showed that a user's awareness to new features and the likelihood of using those features later on dramatically reduced performance on new tasks – i.e., personalised adaptation was detrimental to incidental learning of system features. Interestingly, this finding is called the serendipity problem in content personalisation, and refers to the fact that highly accurate content recommendations may reduce the likelihood to experience unexpected and fortuitous items (McNee et al., 2006a).

2. **Content:** Content traditionally encompasses items or objects such as news articles, products, or media content in the broadest sense of the word; it may however also refer to price tags, service offerings, or highly specific, fine-grained differences in textual wordings. The first research endeavours on selective filtering of information objects for different users date from the late 1950s (Hensley, 1963), but the application-oriented research domain on recommender systems (RS) emerged in the early stages of Web 1.0 (Jannach et al., 2016). Research on RS truly gained momentum in the subsequent Web 2.0 era, when large amounts of (socially) networked data became available (Chen et al., 2012). RS produce personalised rankings of large sets of items based on their presumed relevance to recipients. A multitude of evaluation approaches have been proposed for measuring the impact of personalised content, ranging from purely accuracy-driven AI and ML applications, as well as marketing research on customer value and customer churn minimisation, to cognitive and social psychological study of user satisfaction and user engagement.
3. **Interaction process:** The ubiquity of information access opportunities, and the pervasiveness of data collection – also outside of the traditional browser window – in the Web 3.0 era of today marks the potential for novel interaction processes and modalities across devices and environments (Chen et al., 2012). Apart from the determination of *which* user interface functionality and *what* content to present, algorithms now also are capable of deciding *when* and *how* to approach users. This additional dimension of personalisation obviously leads to an increase of information privacy concerns, and a discourse on the ethical and societal implications of algorithmic decision making.

As discussed above, personalisation can affect the way software systems are used – as in the case of GUI personalisation – but may also influence the content and information provided and displayed. The recent ubiquity paradigm not only allows personalisation approaches to evolve beyond the content itself, but also to *proactively* determine the point in time plus the situational context, in which personalised information can reach (or more accurately: may be targeted towards) the (unsuspecting) user. It goes without saying that the novel personalisation mechanisms currently being developed fuel the already heated debate on issues and challenges related to invasiveness and user-control.

3. Perspectives from different areas

In this section, we will focus on the particular methodological

approaches as well as the theoretical orientations and peculiarities of established research fields such as Machine Learning (ML) / Artificial Intelligence (AI), Human-Computer Interaction (HCI), Information Retrieval (IR) and Information Systems (IS) as well as Cognitive & Social Psychology that made contributions to online personalisation research. Under the Machine Learning perspective we summarise primarily data-driven and algorithmic research exploiting an offline experimentation methodology, as it is typically applied to develop novel recommendation and retrieval techniques. In contrast, the section on human-centred research methods primarily discusses personalisation research in the fields of HCI and IS as well as on personalised IR that directly investigates the perception of personalisation on users. Next, Section 3.3 takes a complementary viewpoint from industry, while finally Section 3.4 discusses the theoretical underpinnings of the impact of personalisation from the cognitive science perspective.

3.1. Machine learning perspective on personalisation

Machine Learning (ML) research in the tradition of pragmatic Artificial Intelligence (AI) has developed into a thriving research field. From a highly simplified perspective, the primary focus regarding personalisation applications within ML lies on optimisation – i.e., to create ever more accurate algorithmic decision makers and prediction models (Jannach et al., 2012). In particular, the recent advancements in learning techniques in neural networks have also triggered research on deep learning-based recommendation systems (Zhang et al., 2019). Besides such learning techniques, however, the central elements in ML research are data and methodologies to train and learn models or functions from these data to fit as well as possible the portions of unseen data. Recommender systems are generally regarded as the most prominent personalisation applications in the ML field, and all the major online platforms, including Amazon, YouTube, Facebook and Netflix, rely on ML-based RS technology to adapt and personalise presented contents.

The key contributions in ML-based personalisation research include the development of new algorithms, the extension of existing algorithms towards domain-specific aspects, and validation/replication studies based on novel datasets. In order to demonstrate the effectiveness of a newly proposed technique, researchers typically use historical datasets. They either contain (a) explicit ratings provided by a user community to the recommendable items (*explicit feedback*), or (b) a log of recorded user interactions – e.g., purchases, item views or listening events (*implicit feedback*). The subsequent comparison of different algorithms is based on “offline” evaluation procedures and measures common in Machine Learning and Information Retrieval. First, a fraction of the data representing *ground truth* is withheld for later validation (*test set*). Then, a *model* is learned on the remaining data (*training set*), which is used to predict the initially held-out data in the test set. Assessment of the performance of an algorithm, finally, is primarily done by comparing the prediction error and the accuracy of the model with those of the test data. There are basically two different families of measures, depending on whether an algorithm actually predicts a quantitative relevance score (i.e. prediction task) or ranks a set of items based on their presumed relevance (i.e. ranking task). For rating prediction tasks error measures like the Root Mean Squared Error (RMSE) quantify the deviation between predicted and ground truth values. In case of a ranking task, accuracy measures such as Normalised Discounted Cumulative Gain (NDCG), Mean Average Precision (MAP), or the F-measure, indicate how well a ranked list of items corresponds to an idealised ranking, with all the ground truth content items (i.e., those items that have been actually clicked on, bought or liked) ranked at the highest possible positions. For more details on the methodological underpinnings and the computation of the concrete measures we refer to the following chapters on evaluating recommendation applications (Gunawardana and Shani, 2015; Jannach et al., 2010).

The intuitive assumption underlying the evaluation approach

described above is that a more *accurate* prediction of the relevance of an item for the user directly translates into a better user experience. This should be the case, because highly relevant items will be placed at screen positions that will receive more attention from users, whereas less relevant items may be even filtered out. As a consequence, offline experiments in ML tend to approximate the expected *impact* of personalisation by tapping into accuracy measures. However, research on user interface adaptations (Findlater and McGrenere, 2010) indicates that the sole reliance on accuracy measures may be insufficient for content personalisation (McNee et al., 2006b). Accurate prediction (or rather “post-diction” (Jannach et al., 2016)) of historical ratings and/or user actions may not be very helpful for users – as they probably are already aware of the existence of these items, or perhaps regard the resulting list as too similar to items selected previously.

In recent years, researchers have, therefore, come up with a variety of computational measures other than accuracy, so as to capture additional quality aspects, such as list diversity, novelty, and serendipity (surprise) (Vargas and Castells, 2011; Ziegler et al., 2005) that rely on additional (ground-truth) data such as the content, popularity or freshness of an item. One typical challenge in consideration of such additional measures is that they may merely represent trade-offs. As a point in case, higher levels of serendipity or surprise can typically only be obtained by making a compromise on accuracy. Furthermore, considerable evidence exists that improvements in terms of accuracy either do not translate into a better user experience, or are not perceived as such (Beel and Langer, 2015; Cremonesi et al., 2012; Ekstrand et al., 2014; Garcin et al., 2014; Maksai et al., 2015; McNee et al., 2002; Rossetti et al., 2016). These concerns on limited predictive power have also been raised by Netflix researchers (Gomez-Uribe and Hunt, 2015).

3.2. Human-centred perspectives on personalisation

Particularly scholars in the fields of Human–Computer Interaction (HCI), Information Systems (IS) as well as research on personalised Information Retrieval (IR) share an empirical research methodology and research interest in phenomena concerning computing devices and information systems such as usability and effective system use. Personalisation is seen as a means for minimising a user’s costs in accessing information, improving the user experience, and rendering people more efficient in their use of computer devices. However, many of these developments were in fact caused by commercial market forces around end-1990s and beginning of the 2000s. That is, the technical abilities for personalisation developed largely in parallel with the rise of electronic commerce, when web users turned into clients and further transformed into customers. In this context, the concept of customer relationship management (CRM) gained momentum – leading to a new type of relationship-oriented marketing, in which highly personalised marketing offerings generate high levels of loyalty and engagement among clients (Reinartz et al., 2004; Verhoef et al., 2010). Personalised systems thus provide the technical infrastructure to realise highly targeted one-to-one marketing and individualised CRM strategies at affordable automation costs, capable of turning (online) customers into loyal brand ambassadors (Godin, 1999). The ultimate goal of personalisation from a commercial standpoint lies in excelling in business metrics so as to build customer value and realise low churn rates.

Tam and Ho (2006) did widely recognised work on the impact of (web) personalisation on users’ cognitive processes. They measured the attention items or information received as well as users’ ability to recall the content later on. The authors manipulated self-referent content (i.e., directly addressing users with their name) and relevant content that matched the task users were up to, and found that the content was perceived as personalised under both conditions – even though only the relevant content was also remembered later on. This clearly showed that personalised communication as such impacts perception, and further indicated that also accuracy in the proposed content is crucial for achieving cognitive processing beyond the initial attention stimulus.

This fundamental issue of a ‘placebo effect’, in which a message is perceived to be personalised when, in fact, it is not was also researched by Li (2016). They found that users’ perception of personalised messages did not necessarily depend on a prior personalisation process of the messages, but on the extent to which the received content matched the receivers’ expectations.

Xiao and Benbasat made an inventory of empirical research on the impact of recommendation agents on commercial platforms (Xiao and Benbasat, 2007). They looked into the influence recommendation agents exert on the users’ decision making processes, how they influence the outcome of decisions, and the way in which such personalisation mechanisms are perceived and evaluated. The authors discussed how a variety of system characteristics, such as explicit preference elicitation, helps to increase decision quality (or transparency) of a system’s reasoning logic. This leads to more trust in the system. Subsequent work on how the impact of different source characteristics can be purposefully exploited so as to create more effective and influential advisory agents (so-called persuasive recommender systems) has, for instance, been summarised here (Yoo et al., 2012; 2015); its theoretical underpinnings will be discussed subsequently in Section 3.4.

Interestingly, Knijnenburg and Willemsen (2015) propose a framework with constructs and questionnaires to guide future user-centred evaluations. Their contribution lies in the identification of an extensive set of measurement constructs grouped into “objective” system aspects, perceptions of the user about the system and the interaction, the situational context as well as personal user characteristics. In addition, the authors offer pointers for relating these aspects to various outcome variables of system usage.

However, controlled laboratory experiments also have their limitations, and need a very careful design and evaluation. Most importantly, the recommendation and decision scenarios are typically artificial in contrast to the field tests performed in industrial environments. Not surprisingly, clear mismatches have been identified between user-centred studies and offline experimentation. For instance, Ekstrand and colleagues observed that highly accurate algorithms in offline experimentation partly obscured (niche) recommendations in a user study (Ekstrand et al., 2014). Likewise, Rossetti and colleagues noticed that users differently ranked the accuracy of algorithms in offline vs. online comparison (Rossetti et al., 2016). Analogously, Ghorab et al. (2013) note in their survey on personalised information retrieval techniques the challenges of comparing results derived from different studies. User experiments are therefore crucial for advancing the state-of-the-art, and in order to discover to which extent the computational measures from artificial offline experiments hold in more ecologically valid settings – such as those provided by industry.

3.3. Industry perspective on personalisation

The strong economic interest in the subject of online personalisation beyond the pure *cognitive interest* is the border between academic and industrial evaluation practice. Research labs in industry tend to consent to a three-tiered evaluation approach that includes offline experimentation, exploratory study with beta testers, and so-called A/B testing designs with a representative share of the user base. Corporate research laboratories typically probe the (series of) preliminary results of these approaches before settling on wide-scaled fielding of newly developed algorithm variants. Jannach and Adomavicius (2016) proposed a conceptual framework intended to help practitioners (managers and engineers) to align strategic business goals with the most appropriate *metrics* for assessment of the actual impact of RS technology. Furthermore, it should be emphasised that corporate RS service providers hardly ever reveal the true business value of their RS personalisation technologies. Still, some tentative evidence exists of the business impact and common success measures employed in the field:

- *Click-Through Rates*: The Click-Through Rate (CTR) measures the

proportion of presented items that actually received a click. The measure is commonly used in online advertising, but also applied in other domains such as news platforms (cf., [Das et al., 2007](#); [Kirshenbaum et al., 2012](#)) with the underlying assumption that more clicks equal more relevant recommendations for the users. Increases of around 35% in terms of the CTR are not uncommon, when personalised recommendations are field-tested against simple popularity-based techniques ([Garcin et al., 2014](#)).

- **Adoption and conversion rates:** The CTR is not always the right proxy for assessing the true relevance of recommended and clicked items. Sudden spikes of CTR can also refer to catchy or insufficient link texts that actually misguided and annoyed users. In the online (streaming) media domain, a recommended link is, therefore, only considered a successful hit, if a certain portion of the video or music track was actually played ([Davidson et al., 2010](#)). Auction platforms such as eBay measure how often users make bids on recommended items ([Chen and Canny, 2011](#)), whereas dating sites measure how many conversations actually followed with a recommended partner ([Wobcke et al., 2015](#)). Obviously, adoption and conversion rates are platform-specific, as a result of which the reported impact of (personalised) recommendations varies across domains. Adoption goes beyond basic CTR, as it considers the use of a service or functionality – like inspecting item details, or watching a video in case of media content. Conversion is typically associated with a commercial or transactional meaning, such as adding to a basket or actually checking out an item on a shopping platform.
- **Sales and revenues:** Precise revenues and profit margins can be determined for transparent and accessible corporate sales data, such as traditional e-commerce shops. Jannach and Hegelich, for instance, reported a sales increase of around 3–4% for a mobile commerce platform with game apps ([Jannach and Hegelich, 2009](#)). A much more pronounced increase in sales was observed, when a recommender was introduced on an online retail Website and the alternative experimental condition did not involve a recommender at all ([Lee and Hosanagar, 0000](#)). Some evidence exists that RS may yield indirect effects on sales, such as when recommendations apparently inspired more purchases in other domains of an online grocery store ([Dias et al., 2008](#)). The increases in hard business figures are often modest, but represent a considerable return-on-investment compared with the technological investment of the service provider.
- **User engagement and behaviour:** User engagement, or more general, the impact on user behaviour, is another commonly used indirect measure. Various studies exist showing that RS leads to more user activity and longer sessions, cf. ([Domingues et al., 2013](#); [Garcin et al., 2014](#)). In many cases, user engagement measures are application-specific and therefore relate, for instance, to the number of answers provided on a query-answering Website, or to the number of email messages exchanged on a recruiting platform ([Szpektor et al., 2013](#); [Xu et al., 2014](#)).
- **Effects on sales distributions:** Some evidence exists that personalisation mechanisms positively impact overall corporate sales distributions. Recommendation agents have been found to improve the quality of matching between items and customers in electronic commerce ([Xiao and Benbasat, 2007](#)). Also, several studies observed a smoothing effect on long-tail distributions, which meant that RS actually enabled businesses to create additional business value from niche items ([Lawrence et al., 2001](#); [Zanker et al., 2006](#)). Consistent with this long-tail distribution argument, Netflix measures its “Effective Catalogue Size” ([Gomez-Urbe and Hunt, 2015](#)) to assess to what extent their recommendations help users explore larger portions of their catalogue. The catalogue size is the share of items that actually gets ordered or is viewed by a minimum number of users within a specified time frame. Especially online companies with business models based on flat rate subscription fees exploit the long-tail effect to reduce customer attrition and increase loyalty.

Given the observed effects of personalisation in the previous subsections, we will discuss the impact of personalised content and interaction in the next section from its fundamental behavioural principles.

3.4. Cognitive science perspective on personalisation

Research into (applied) social and cognitive psychology provides the scientific underpinning of why users may (or may not) appreciate the delivery of personalised services and content in one form or another. Theoretical frameworks such as Similarity Attraction Theory ([Byrne, 1997](#)) and the classic Belief-Congruence Theory ([Rokeach, 1968](#)) from pre-Web times state that humans are attracted by similar others, and this principle has also been explored in the Web 1.0 era for settings, in which human beings interacted with televisions and personalised computers ([Nass and Lee, 2001](#)). Together with the related finding from Cognitive Dissonance Theory ([Festinger, 1962](#)), that being confronted with facts contradicting personal beliefs and values leads to attitude change, this observation may still hold in the Web 2.0 era that is now rapidly evolving into an ubiquitous Web 3.0.

Over the past decades, *Media Equation Theory* ([Nass et al., 1996](#)) became an influential theoretical perspective among researchers in human–computer interaction to study how people use and interact with computer devices. Derived from a series of explorations of human–computer interaction as a social psychological phenomenon, Media Equation Theory posits that *computers are social actors*, and that social rules from traditional human-to-human interaction also apply to people’s interaction with computer devices. Studies reveal that people occasionally praise their desktop computer ([Nass et al., 1994](#)), affiliate with and conform to their computer (when it provides advice) ([Nass et al., 1996](#)), and are sensitive to the desktop computer’s praise and flattery ([Nass and Lee, 2001](#)). These findings have also been found for other technologies, such as televisions ([Nass and Moon, 2000](#); [Reeves and Nass, 1996](#)).

A significant portion of research in the Media Equation Theory paradigm concerns test of the similarity-attraction effect. Classic research in social psychology ([Festinger, 1954](#); [Newcomb, 1961](#)), such as Belief-Congruence Theory ([Insko et al., 1983](#); [Roccas and Schwartz, 1993](#); [Rokeach, 1968](#)), but especially the well-known *attraction paradigm*, provides compelling evidence that people are attracted to others who are similar to them in terms of attitudes, age, personality traits, and many other factors ([Byrne, 1961](#); [1962](#)). This so-called similarity-attraction link is “[o]ne of the most robust phenomena in social psychology” ([Montoya and Horton, 2004](#), p. 696), and has been replicated across a range of (cross-cultural) situations as well as for a wide range of populations; for review, see [Byrne \(1997\)](#); for meta-analysis, see [Montoya et al. \(2008\)](#) and [Montoya and Horton \(2013\)](#). The attraction paradigm also features prominently in network studies via the principle of *homophily* – i.e., the amount of homogeneity (vs. heterogeneity) of a person’s social network in terms of demographic, behavioural and other factors. In a social network, homophily likely yields a larger number of connections with like-minded others, such that “similarity breeds connection” ([McPherson et al., 2001](#), p. 415).

In Media Equation Theory, the similarity-attraction effect is explored and reported primarily in response to computer-synthesised speech using similarity matching between a user’s personality and an extrovert or introvert computer voice, both in spoken and written form ([Nass and Lee, 2001](#)). This similarity-attraction effect appears to hold under conditions in which people do not have free choice to select voice-interface modes ([Lee et al., 2011](#)), and in studies on anthropomorphism with regard to (psychophys(iolog)ical) similarity ([de Visser et al., 2016](#)).

However, the role of machines has changed dramatically in recent years. Web 2.0 and the currently emerging mobile Web 3.0 era enable the provision of networked information, accessed through social media ([Chen et al., 2012](#)). Media Equation Theory seems to apply to human–chatbot interactions on Twitter ([Edwards et al., 2014](#)), but

scholars increasingly question its viability for present and future times. Advances in ML/AI have made machinery, decision support systems, and mobile devices appearing so (artificially) intelligent that they increasingly outperform humans in problem-solving capabilities – rendering it necessary to reconsider traditional HCI approaches. To deal with the newly arisen *media inequality*, cf. (Mou and Xu, 2017), scholars are now starting to turn to state-of-the-art process models in (applied) cognitive and social psychology that are better capable of explaining conditions under which people are implicitly and/or explicitly prone to, feel like, or are motivated to, process information and interact with Web 2.0 and Web 3.0 intelligent machinery, cf. (Culley and Madhavan, 2013). The implications for future research into personalisation from a cognitive process point of view will be outlined in the next session.

4. What's next? Research roadmap

In this final section, we will suggest an agenda for future research, and offer an informed discussion on novel measurement approaches – among others involving (neuro)physiological tools and sensors. We will address the necessity and potential of incorporating such measurement instruments in future research endeavours, together with elaborated prospects on other future developments.

4.1. Future perspectives on ML personalisation research

Whether or not small improvements in terms of accuracy – and in terms of other abstract quality measures such as diversity or novelty – actually have a measurable impact in reality is largely under-investigated. Said and Bellogin (2014) recently found that implementations of the same algorithms in different recommender system libraries, which were tested against the same measures and according to the same methodology, did not yield the same results. This was due to different data management procedures and slight deviations in the interpretation of the evaluation methodology and algorithmic steps. Similar observations of lacking repeatability, reproducibility and generalisability of experimental results and using weak baselines are also reported from the field of IR (Arguello et al., 2015; Ferro et al., 2018). The issue of replication of ML results is therefore a key issue, and should receive much more attention than it currently does in the various RS-related research communities (Beel et al., 2016).

Generally, we postulate a more holistic view when it comes to assessing the characteristics of personalisation mechanisms. Interestingly, regulatory requirements, such as the General Data Protection Regulation (GDPR) in the European Union, have already begun to put the focus on transparency of algorithmic decision making, on explainability of computed outcomes as well as on ethical considerations in the context of Machine Learning. Therefore, regulatory aspects and questions about the fairness of algorithms will definitely arise more often in the future, and, for instance, more firmly set the research agenda for improving our understanding on how the recommended information diet of users can become and remain sufficiently diverse (Helberger et al., 2018).

4.2. Future perspectives on human-centred personalisation research

In recent years, a politically-inspired controversy has materialised on the potentially harmful societal and Web-specific effects of online RS-produced personalisation that facilitates the spreading of misinformation to their target groups (Fernandez and Alani, 2018). Some people fear that ever more targeted recommendations inevitably put online users in highly specific, profiled grids (or 'filter bubbles' Pariser (2011)). The supposed effect of being in a filter bubble would be a dramatic limitation of the diversity of information received and potentially scrutinised, as well as the eventual loss of connection to other social groups in society, cf. (Bozdag and van den Hoven, 2015;

Sunstein, 2018). This brings to mind classic pre-Web debates on the societal implications of selective exposure to (news) media, cf. (Sears and Freedman, 1967), and on group polarisation in (applied) social psychology, cf. (Myers and Lamm, 1976) – albeit with a twist to the Web 2.0 and Web 3.0 era (Messing and Westwood, 2014). Also, it resonates with the observation that higher education systems tend to turn into 'academic tribes' – i.e., highly inward-oriented disciplines with distinct research practices, conventions, and behaviours, in which actors become resistant to alternative approaches and perspectives (Becher and Trowler, 2001).

This coincides with criticism from within the Intelligent User Interface community that intelligent systems "dumb down" the user when a portion of the user's cognitive load is offloaded to the system (Lanier, 1995). Intelligent systems may inadvertently harm the development of individual users by offering them an overly selective number of learning opportunities in a specific domain – which may lead to a problematic reduction in the users breadth of experience. Importantly, first evidence has been generated in the setting of HCI that ongoing online RS personalisation – regardless of the occurrence of a narrowing down in the variety of recommendations a person receives over a longer time period – *does not produce* filter bubbles (Nguyen et al., 2014). Nonetheless, it remains important to conduct further research on the potential 'dark side of online personalisation', and to debunk erroneous opinions and beliefs, whenever necessary. A further line of future research will likely develop at the intersection of understanding the persuasive traits of a system generating recommendations and exploiting individual user differences in order to tailor persuasion strategies to such user traits (Kaptein et al., 2015).

4.3. Future industry perspective on personalisation

When it comes to the future industry perspective, the challenges already mentioned in Section 4.1 will in general also apply to the large-scale industrial application of ML techniques. Thus, practitioners will increasingly struggle to comply with privacy legislation, especially when they seek to provide an enhanced user experience as well. Put differently, what is technically possible to offer to the customer may not be legally allowed, or otherwise be regulatory constrained.

In addition, given the progressing and ageing of currently deployed applications and algorithms, the life-time perspective on personalisation applications will become increasingly relevant. This may include novel quantitative measurements to understand how metrics on adoption, conversion, user engagement, or the sales distribution itself, actually evolve over time. In particular, industrial data scientists do have perfect opportunities to identify common patterns for different life-cycle stages of personalisation applications over time, should they exist. Nevertheless, it will be hard to perform (let alone, replicate) such research endeavours outside of industrial labs.

4.4. Future perspectives on cognitive & psychological science personalisation research

In pre-Web days, similarity attraction was understood in terms of an affect-based reinforcement model. Similar others were considered to make us feel good (they offered consensual validation for our own viewpoints), whereas dissimilar others made us feel bad (they challenged our attitudes) (Byrne, 1997; Byrne and Clore, 1967). Nowadays, similarity attraction is regarded a *cognitive process* determined by information-processing. People judge a person's similarity (or dissimilarity) in attitudes as positive (or negative) information concerning someone else's assumed qualities, cf., (Montoya and Horton, 2004; 2013; Montoya et al., 2008; Singh et al., 2008). The role of such *cognitive evaluation* (of a person's qualities) on similarity attraction has not yet been studied in human-computer studies, but it will surely have an impact on human interaction with highly personalised and potentially invasive Web 2.0 and Web 3.0 technologies. It stands to reason that

scholars in HCI will turn to Cognitive Dissonance Theory (Festinger, 1962) to account for unexplained findings on similarity-attraction, cf. (Lee et al., 2011).

Festinger (1962) famous theory originally stated that people experience cognitive dissonance when displaying behaviours in violation of how they ought to behave, cf. (Aronson and Carlsmith, 1962; Aronson et al., 1999), e.g., when contravening normative standards, societal norms and conventions (Cooper and Fazio, 1984; Stone and Cooper, 2001). Attitudes and behaviour would usually be altered in favour of the cognition a person found the hardest to change (Festinger, 1962; Festinger and Carlsmith, 1959; Festinger et al., 1956); see also (Beauvois et al., 1996; Harmon-Jones and Harmon-Jones, 2007). The Action-Based Model of Cognitive Dissonance Processes (Harmon-Jones, 1999) is a recent neuropsychological modification, capable of explaining *why* cognitive dissonance occurs, and *why* it makes people alter initially held beliefs; for review, see (Harmon-Jones et al., 2009; Harmon-Jones and Harmon-Jones, 2007; Harmon-Jones et al., 2015) A key insight of the Action-Based Model is that inconsistent action tendencies produce dissonance (Harmon-Jones and Harmon-Jones, 2007) caused by goal-directed eagerness towards gains, rewards and non-punishment, cf. (Carver and White, 1994; Gollwitzer, 1999; Gollwitzer and Sheeran, 2006). Dissonance (or: discrepancy) reduction thus takes place in settings, where prospects of gains, wins, successes, and achievements reign (Harmon-Jones et al., 2015). This insight holds great promise for future research into personalised Web 2.0 and Web 3.0 technologies, as the (partial) removal of competitive persuasion triggers seems to be sufficient in overcoming inconsistent action tendencies due to intrusiveness, infringements of information privacy, and filter bubbles.

4.5. Towards an integrated psychoinformatics / neuroIS approach

In recent years, scholars in social and cognitive psychology have launched a call for more interdisciplinary research into *psychoinformatics*, aimed at exploration of problems and challenges on the verge of psychological and computer science (Yarkoni, 2012). Researchers increasingly turn to mobile technologies, such as the smartphone and the tablet computer, to study human behaviour and mental processes outside of the laboratory (Dufau et al., 2011; Miller, 2012). Also an electronically activated recorder (EAR) has been developed to unobtrusively record snippets of natural language, cf. (Mehl et al., 2001; 2010; 2007). Likewise, several linguistic inventory and word count (LIWC) software packages have been assembled – and are constantly modified and updated – to accompany speech-sampling methods, and to enable the behavioural analysis of computerised texts derived from Web 2.0 and Web 3.0 social media platforms, cf. (Pennebaker et al., 2003; Tausczik and Pennebaker, 2010). Such endeavours have put the study of ‘*actual behaviour*’ more firmly on the research agenda (i.e., they observe real behaviours of real people under real circumstances), and as such move beyond the mere understanding of social and cognitive psychology as “the science of self-reports and finger movements” (Baumeister et al., 2007). Applied to Human–Computer Studies, these and other psychoinformatics tools surely hold promise for ecologically valid future research into the study of actual human behaviour in interaction with Web 2.0 and Web 3.0 technologies, and for the assessment of actual behavioural responses to such technologies in realistic – i.e., real-world – settings.

Interestingly, and in parallel, scholars in the information systems research community have in recent years made comparable endeavours to incorporate behavioural and neuropsychological theorising and measurement techniques to better capture IS phenomena (Dimoka et al., 2012) such as system usage or cognitive overload. Like their counterparts in psychological science, IS scholars also are fascinated by the rise of Web 2.0 and Web 3.0, and are fully aware of the necessity to combine the theories, tools and techniques of the behavioural research domain with those of computer science. Some speak of

a “*computing science social media research*” paradigm (Shneiderman et al., 2011, p. 26). Others have coined the label NeuroIS (Dimoka et al., 2011) instead. In the field of NeuroIS, opportunities especially lie in the development of user interfaces based on fundamental behavioural principles (Reinecke and Bernstein, 0000). A cornerstone study using behavioural as well as functional neuroimaging (fMRI) equipment, for instance, found behavioural evidence that high (vs. low) trust (vs. distrust) in the seller on an auction platform impacts price-related buying intentions. In a second stage, also the neural correlates of trust and distrust were mapped, and distinct brain areas associated with trust or distrust were identified (Dimoka, 2010). The neural correlates of trust in e-commerce platforms were also explored for gender differences, cf. (Riedl et al., 2010), which seems to suggest that the role of trust in relation to information systems is now well-documented. Apart from fMRI, also less costly neurophysiological tools such as eye-trackers and skin conductance response tools are widely used; cf. (Dimoka et al., 2012). Some have even turned to self-report measures developed within cognitive neuroscience – i.e., the well-known BIS/BAS scales for Behavioural Inhibition and Behavioural Activation (Carver and White, 1994), which have also been used to validate the Action-Based Model of Cognitive Dissonance Processes discussed above (Harmon-Jones et al., 2015); for an application to online personalisation in the context of recommender systems, see Rook et al. (2018). In spite of its theoretical and methodological diversity (Riedl et al., 2014), however, it is nevertheless intriguing that the field of NeuroIS – in sharp contrast to psychoinformatics in psychological science – seems to move away from real-world applications, and into the laboratory – perhaps to become the next-generation “science of self-reports and finger movements”.

4.6. Putting it all together

Overall, we observe that the vast bulk of personalisation research these days still takes place in the distinct silos of the respective research fields and with separate methodological traditions. Interestingly, this is exploited by the work of Salatino et al. (2017), who identify the emergence of novel topics based on intensified collaborations and interactions between such topical silos.

Academic ML research relies largely on offline evaluation measures based on historical data, impact assessments in organisational practice are often based on business value and domain-specific measures. A particularly worrying aspect in the latter case is that offline experiments then may not be very indicative of the success of a recommender in practice. As a future research agenda we, therefore, postulate the need for a more explicit interdisciplinary research approach that would draw a bow from the underpinnings of cognitive and psychological science, to the principled observation and experimentation from HCI, and towards an ecologically valid implementation of results via novel models and efficient machine learning techniques.

In the future, more work is required following common HCI research principles – in particular in the form of user studies – so as to obtain a clearer picture of the impact of content personalisation, and to gain insights for developing algorithms optimising on multiple objectives. The work of Coba et al. (2019) could be considered as a first minor step towards this direction. Even though recommendations largely depend on the aggregated opinions, preferences and behaviours of users, this HCI-inspired study showed that the different influential properties of rating summary statistics (such as the total number of ratings, the mean rating value, or the skewness of the distribution) should also be analysed in order to more thoroughly understand how they influence users in their decision making (under a *ceteris paribus* condition). Drawing from Cognitive Sciences, eye-tracked observations of users’ gazes in this interdisciplinary study disclosed that different decision making styles produced significant differences in users’ choices. Specifically, users with the tendency to follow maximising behavioural tendencies (Schwartz, 2016) were more likely to follow

compensatory decision making strategies (where different attributes are carefully weighted against each other (Payne, 1976)), and base their decisions on the highest mean rating (Coba et al., 2019). The outcome of this interdisciplinary research is currently being exploited as input for developing algorithms, where recommended items are better justified by their rating summaries in the eyes of users. In the end, our understanding of the relation between rating summary statistics and user decision making will thus be deepened thanks to the multiple and integrated perspectives of HCI, User Modelling and AI & ML.

In this context, it is of vital importance to validate the correspondence of abstract computational measures with the users' perceptions, and to find new measures that are more predictive of the adoption of the recommendations. It is surprising to see that the well-established instrument of customer satisfaction surveys from the context of fielded applications does not seem to play a major role in research endeavours in the field of computer science and information systems. Such surveys about the user experience of fielded personalised services should receive more attention in the future to complement indirect measures of the success of recommenders, and we consider them as a means to obtain impact in practice.

5. Conclusions

In the present paper, we provided an overview of the genesis of research into online personalisation, and described how this led to a large, rich, vibrant, and potentially multidisciplinary body of knowledge spanning Machine Learning (ML) & Artificial Intelligence (AI), Human–Computer Interaction (HCI), Information Retrieval (IR), Information Systems (IS), and User Modelling based on (applied) social and cognitive psychology. Scholars can nowadays choose to study online personalisation from strong accuracy-driven viewpoints and applications, investigate the establishment and maintenance of customer value and customer churn minimisation in electronic commerce settings, or opt for the cognitive and social psychological study of user engagement and satisfaction moderated by explicit and implicit cognitive processes. Future research endeavours in Human–Computer Studies should aim at nurturing the intellectual diversity embodied in these different theoretical and methodological traditions and positions regarding personalisation, while, at the same time, keeping a keen eye for the scientific and societal pitfalls and challenges represented by the ongoing ubiquity of personalisation mechanisms, now and in the future.

Funding

This work was supported by the Open Access Publishing Fund provided by the Free University of Bozen-Bolzano.

Conflict of interest

None.

References

- Arguello, J., Diaz, F., Lin, J., Trotman, A., 2015. Sigir 2015 workshop on reproducibility, inexplicability, and generalizability of results (rigor). Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '15. ACM, New York, NY, USA, pp. 1147–1148.
- Aronson, E., Carlsmith, J.M., 1962. Performance expectancy as a determinant of actual performance. *J. Abnorm. Soc. Psychol.* 65 (3), 178.
- Aronson, J., Cohen, G., Nail, P.R., 1999. Self-affirmation theory: an update and appraisal. In: Harmon-Jones, E., Mills, J. (Eds.), *Cognitive Dissonance Theory: Progress on a Pivotal Theory in Social Psychology*. American Psychological Association, Washington, DC, pp. 127–147.
- Baumeister, R.F., Vohs, K.D., Funder, D.C., 2007. Psychology as the science of self-reports and finger movements: whatever happened to actual behavior? *Perspect. Psychol. Sci.* 2 (4), 396–403.
- Beauvois, J.L., Joule, R.V., Joule, R., 1996. *A Radical Dissonance Theory*. Taylor & Francis.
- Becher, T., Trowler, P.R., 2001. *Academic Tribes and Territories*. SHRE, UK.

- Beel, J., Breitinger, C., Langer, S., Lommatzsch, A., Gipp, B., 2016. Towards reproducibility in recommender-systems research. *User Model. User-Adapt. Interact.* 26 (1), 69–101.
- Beel, J., Langer, S., 2015. A comparison of offline evaluations, online evaluations, and user studies in the context of research-paper recommender systems. *Proc. of the 22nd International Conference on Theory and Practice of Digital Libraries, TPD L'15*. pp. 153–168.
- Bozdog, E., van den Hoven, J., 2015. Breaking the filter bubble: democracy and design. *Ethics Inf. Technol.* 17 (4), 249–265.
- Brusilovsky, P., 1998. *Methods and techniques of adaptive hypermedia*. Adaptive Hypertext and Hypermedia. Springer, pp. 1–43.
- Byrne, D., 1961. Interpersonal attraction and attitude similarity. *J. Abnorm. Soc. Psychol.* 62 (3), 713.
- Byrne, D., 1962. Response to attitude similarity-dissimilarity as a function of affiliation need 1. *J. Personal.* 30 (2), 164–177.
- Byrne, D., 1997. An overview (and underview) of research and theory within the attraction paradigm. *J. Soc. Personal. Relatsh.* 14 (3), 417–431.
- Byrne, D., Clore Jr, G.L., 1967. Effectance arousal and attraction. *J. Personal. Soc. Psychol.* 6 (4p2), 1.
- Carver, C.S., White, T.L., 1994. Behavioral inhibition, behavioral activation, and affective responses to impending reward and punishment: the bis/bas scales. *J. Personal. Soc. Psychol.* 67 (2), 319.
- Chen, H., Chiang, R.H., Storey, V.C., 2012. Business intelligence and analytics: from big data to big impact. *MIS Q.* 1165–1188.
- Chen, Y., Canny, J.F., 2011. Recommending ephemeral items at web scale. Proceedings of the 34th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '11. pp. 1013–1022.
- Coba, L., Rook, L., Zanker, M., Symeonidis, P., 2019. Decision making strategies differ in the presence of collaborative explanations: two conjoint studies. Proceedings of the 24th International Conference on Intelligent User Interfaces, IUI '19. ACM, New York, NY, USA, pp. 291–302.
- Cooper, J., Fazio, R.H., 1984. A new look at dissonance. *Adv. Exp. Soc. Psychol.* 17, 229–268.
- Cremonesi, P., Garzotto, F., Turrin, R., 2012. Investigating the persuasion potential of recommender systems from a quality perspective: an empirical study. *ACM Trans. Interact. Intell. Syst.* 2 (2), 11:1–11:41.
- Culley, K.E., Madhavan, P., 2013. A note of caution regarding anthropomorphism in HCI agents. *Comput. Hum. Behav.* 29 (3), 577–579.
- Das, A.S., Datar, M., Garg, A., Rajaram, S., 2007. Google news personalization: scalable online collaborative filtering. Proceedings of the 16th International Conference on World Wide Web, WWW '07. pp. 271–280.
- Davidson, J., Liebold, B., Liu, J., Nandy, P., Vleet, T.V., Gargi, U., Gupta, S., He, Y., Lambert, M., Livingston, B., Sampath, D., 2010. The youtube video recommendation system. Proceedings of the Fourth ACM Conference on Recommender Systems, RecSys '10. pp. 293–296.
- Dias, M.B., Locher, D., Li, M., El-Deredy, W., Lisboa, P.J., 2008. The value of personalised recommender systems to e-business: a case study. Proceedings of the 2008 ACM Conference on Recommender Systems, RecSys '08. pp. 291–294.
- Dimoka, A., 2010. What does the brain tell us about trust and distrust? evidence from a functional neuroimaging study. *MIS Q.* 373–396.
- Dimoka, A., Davis, F.D., Gupta, A., Pavlou, P.A., Banker, R.D., Dennis, A.R., Ischebeck, A., Müller-Putz, G., Benbasat, I., Gefen, D., et al., 2012. On the use of neurophysiological tools in research: developing a research agenda for neurois. *MIS Q.* 679–702.
- Dimoka, A., Pavlou, P.A., Davis, F.D., 2011. Research commentary neurois: the potential of cognitive neuroscience for information systems research. *Inf. Syst. Res.* 22 (4), 687–702.
- Domingues, M.A., Gouyon, F., Jorge, A.M., Leal, J.P., Vinagre, J., Lemos, L., Sordo, M., 2013. Combining usage and content in an online recommendation system for music in the long tail. *Int. J. Multimed. Inf. Retr.* 2 (1), 3–13.
- Dufau, S., Duñabeitia, J.A., Moret-Tatay, C., McGonigal, A., Peeters, D., Alario, F.X., Balota, D.A., Brysbaert, M., Carreiras, M., Ferrand, L., et al., 2011. Smart phone, smart science: how the use of smartphones can revolutionize research in cognitive science. *PLoS One* 6 (9), e24974.
- Edwards, C., Edwards, A., Spence, P.R., Shelton, A.K., 2014. Is that a bot running the social media feed? Testing the differences in perceptions of communication quality for a human agent and a bot agent on twitter. *Comput. Hum. Behav.* 33, 372–376.
- Ekstrand, M.D., Harper, F.M., Willemsen, M.C., Konstan, J.A., 2014. User perception of differences in recommender algorithms. Proceedings of the 8th ACM Conference on Recommender Systems, RecSys '14. pp. 161–168.
- Fan, H., Poole, M.S., 2006. What is personalization? Perspectives on the design and implementation of personalization in information systems. *J. Org. Comput. Electron. Commer.* 16 (3–4), 179–202.
- Fernandez, M., Alani, H., 2018. Online misinformation: challenges and future directions. Companion of the The Web Conference 2018 on The Web Conference 2018, International World Wide Web Conferences Steering Committee. pp. 595–602.
- Ferro, N., Fuhr, N., Rauber, A., 2018. Introduction to the special issue on reproducibility in information retrieval: evaluation campaigns, collections, and analyses. *J. Data Inf. Qual.* 10 (3), 9:1–9:4.
- Festinger, L., 1954. A theory of social comparison processes. *Hum. Relat.* 7 (2), 117–140.
- Festinger, L., 1962. *A Theory of Cognitive Dissonance*, Vol. 2. Stanford University Press.
- Festinger, L., Carlsmith, J.M., 1959. Cognitive consequences of forced compliance. *J. Abnorm. Soc. Psychol.* 58 (2), 203.
- Festinger, L., Riecken, H., Schachter, S., 1956. *When prophecy fails*. Minneapolis, mn, us.
- Findlater, L., Gajos, K.Z., 2009. Design space and evaluation challenges of adaptive graphical user interfaces. *AI Mag.* 30 (4), 68.
- Findlater, L., McGrenere, J., 2010. Beyond performance: feature awareness in

- personalized interfaces. *Int. J. Hum. Comput. Stud.* 68 (3), 121–137.
- Foltz, P.W., Dumais, S.T., 1992. Personalized information delivery: an analysis of information filtering methods. *Commun. ACM* 35 (12), 51–60.
- Garcin, F., Faltings, B., Donatsch, O., Alazzawi, A., Bruttin, C., Huber, A., 2014. Offline and online evaluation of news recommender systems at swissinfo.ch. *Proceedings of the 8th ACM Conference on Recommender Systems, RecSys '14*. pp. 169–176.
- Ghorab, M.R., Zhou, D., O'Connor, A., Wade, V., 2013. Personalised information retrieval: survey and classification. *User Model. User-Adapt. Interact.* 23 (4), 381–443.
- Godin, S., 1999. Permission Marketing: Turning Strangers into Friends and Friends Into Customers. Simon and Schuster.
- Gollwitzer, P.M., 1999. Implementation intentions: strong effects of simple plans. *Am. Psychol.* 54 (7), 493.
- Gollwitzer, P.M., Sheeran, P., 2006. Implementation intentions and goal achievement: a meta-analysis of effects and processes. *Adv. Exp. Soc. Psychol.* 38, 69–119.
- Gomez-Uribe, C.A., Hunt, N., 2015. The netflix recommender system: algorithms, business value, and innovation. *Trans. Manag. Inf. Syst.* 6 (4), 13:1–13:19.
- Greenberg, S., Witten, I.H., 1985. Adaptive personalized interfaces a question of viability. *Behav. Inf. Technol.* 4 (1), 31–45.
- Gunawardana, A., Shani, G., 2015. Evaluating recommender systems. *Recommender Systems Handbook*. Springer, pp. 265–308.
- Harmon-Jones, E., 1999. Toward an understanding of the motivation underlying dissonance effects: is the production of aversive consequences necessary? In: Harmon-Jones, E., Mills, J. (Eds.), *Cognitive Dissonance Theory: Progress on a Pivotal Theory in Social Psychology*. American Psychological Association, Washington, DC, pp. 71–99.
- Harmon-Jones, E., Amodio, D.M., Harmon-Jones, C., 2009. Action-based model of dissonance: a review, integration, and expansion of conceptions of cognitive conflict. *Adv. Exp. Soc. Psychol.* 41, 119–166.
- Harmon-Jones, E., Harmon-Jones, C., 2007. Cognitive dissonance theory after 50 years of development. *Z. Sozialpsychol.* 38 (1), 7–16.
- Harmon-Jones, E., Harmon-Jones, C., Levy, N., 2015. An action-based model of cognitive-dissonance processes. *Curr. Dir. Psychol. Sci.* 24 (3), 184–189.
- Helberger, N., Karppinen, K., DACunto, L., 2018. Exposure diversity as a design principle for recommender systems. *Inf. Commun. Soc.* 21 (2), 191–207.
- Hensley, C., 1963. Selective dissemination of information (sdi): state of the art in may, 1963. *Proceedings of the May 21–23, 1963, Spring Joint Computer Conference*. ACM, pp. 257–262.
- Insko, C.A., Nacoste, R.W., Moe, J.L., 1983. Belief congruence and racial discrimination: review of the evidence and critical evaluation. *Eur. J. Soc. Psychol.* 13 (2), 153–174.
- Jannach, D., Adomavicius, G., 2016. Recommendations with a purpose. *Proceedings of the 10th ACM Conference on Recommender Systems (RecSys 2016, Boston, Massachusetts, USA)*. pp. 7–10. <https://doi.org/10.1145/2959100.2959186>.
- Jannach, D., Hegelich, K., 2009. A case study on the effectiveness of recommendations in the mobile internet. *Proceedings of the 10th ACM Conference on Recommender Systems, RecSys '09*. pp. 205–208.
- Jannach, D., Resnick, P., Tuzhilin, A., Zanker, M., 2016. Recommender systems beyond matrix completion. *Commun. ACM* 59 (11), 94–102.
- Jannach, D., Zanker, M., Felfernig, A., Friedrich, G., 2010. *Recommender Systems: An Introduction*. Cambridge University Press.
- Jannach, D., Zanker, M., Ge, M., Gröning, M., 2012. Recommender systems in computer science and information systems—a landscape of research. *International Conference on Electronic Commerce and Web Technologies*. Springer, pp. 76–87.
- Kaptein, M., Markopoulos, P., De Ruyter, B., Aarts, E., 2015. Personalizing persuasive technologies: explicit and implicit personalization using persuasion profiles. *Int. J. Hum. Comput. Stud.* 77, 38–51.
- Kirshenbaum, E., Forman, G., Dugan, M., 2012. A live comparison of methods for personalized article recommendation at forbes.com. *Machine Learning and Knowledge Discovery in Databases*. pp. 51–66.
- Knijnenburg, B.P., Willemsen, M.C., 2015. Evaluating recommender systems with user experiments. In: Ricci, F., Rokach, L., Shapira, B. (Eds.), *Recommender Systems Handbook*. Springer, US, pp. 309–352.
- Lanier, J., 1995. Agents of alienation. *J. Conscious. Stud.* 2 (1), 76–81.
- Lawrence, R., Almasi, G., Kotlyar, V., Viveros, M., Duri, S., 2001. Personalization of supermarket product recommendations. *Data Min. Knowl. Discov.* 5 (1), 11–32.
- Lee, D., Hosanagar, K., 2015. How do recommender systems affect sales diversity? a cross-category investigation via randomized field experiment. In: *Information Systems Research*.
- Lee, K.M., Jung, Y., Nass, C., 2011. Can user choice alter experimental findings in human-computer interaction? Similarity attraction versus cognitive dissonance in social responses to synthetic speech. *Int. J. Hum.-Comput. Interact.* 27 (4), 307–322.
- Lí, C., 2016. When does web-based personalization really work? The distinction between actual personalization and perceived personalization. *Comput. Hum. Behav.* 54, 25–33.
- Maksai, A., Garcin, F., Faltings, B., 2015. Predicting online performance of news recommender systems through richer evaluation metrics. *Proceedings of the 9th ACM Conference on Recommender Systems*. ACM, pp. 179–186.
- McNee, S.M., Albert, I., Cosley, D., Gopalkrishnan, P., Lam, S.K., Rashid, A.M., Konstan, J.A., Riedl, J., 2002. On the recommending of citations for research papers. *Proceedings of the 2002 ACM Conference on Computer Supported Cooperative Work, CSCW '02*. pp. 116–125.
- McNee, S.M., Riedl, J., Konstan, J.A., 2006. Being accurate is not enough: how accuracy metrics have hurt recommender systems. *CHI'06 Extended Abstracts on Human Factors in Computing Systems*. ACM, pp. 1097–1101.
- McNee, S.M., Riedl, J., Konstan, J.A., 2006. Being accurate is not enough: how accuracy metrics have hurt recommender systems. *CHI '06 Extended Abstracts on Human Factors in Computing Systems*. CHI EA '06. pp. 1097–1101.
- McPherson, M., Smith-Lovin, L., Cook, J.M., 2001. Birds of a feather: homophily in social networks. *Annu. Rev. Sociol.* 27 (1), 415–444.
- Mehl, M.R., Pennebaker, J.W., Crow, D.M., Dabbs, J., Price, J.H., 2001. The electronically activated recorder (ear): a device for sampling naturalistic daily activities and conversations. *Behav. Res. Methods Instrum. Comput.* 33 (4), 517–523.
- Mehl, M.R., Vazire, S., Holleran, S.E., Clark, C.S., 2010. Eavesdropping on happiness: well-being is related to having less small talk and more substantive conversations. *Psychol. Sci.* 21 (4), 539–541.
- Mehl, M.R., Vazire, S., Ramírez-Esparza, N., Slatcher, R.B., Pennebaker, J.W., 2007. Are women really more talkative than men? *Science* 317 (5834), 82–82.
- Messing, S., Westwood, S.J., 2014. Selective exposure in the age of social media: endorsements trump partisan source affiliation when selecting news online. *Commun. Res.* 41 (8), 1042–1063.
- Miller, G., 2012. The smartphone psychology manifesto. *Perspect. Psychol. Sci.* 7 (3), 221–237.
- Mobasher, B., Cooley, R., Srivastava, J., 2000. Automatic personalization based on web usage mining. *Commun. ACM* 43 (8), 142–151.
- Montoya, R.M., Horton, R.S., 2004. On the importance of cognitive evaluation as a determinant of interpersonal attraction. *J. Personal. Soc. Psychol.* 86 (5), 696.
- Montoya, R.M., Horton, R.S., 2013. A meta-analytic investigation of the processes underlying the similarity-attraction effect. *Journal of Social and Personal Relationships*. 30, pp. 64–94.
- Montoya, R.M., Horton, R.S., Kirchner, J., 2008. Is actual similarity necessary for attraction? a meta-analysis of actual and perceived similarity. *J. Soc. Personal. Relatsh.* 25 (6), 889–922.
- Mou, Y., Xu, K., 2017. The media inequality: comparing the initial human-human and human-ai social interactions. *Comput. Hum. Behav.* 72, 432–440.
- Murthi, B., Sarkar, S., 2003. The role of the management sciences in research on personalization. *Manag. Sci.* 49 (10), 1344–1362.
- Myers, D.G., Lamm, H., 1976. The group polarization phenomenon. *Psychol. Bull.* 83 (4), 602.
- Nass, C., Fogg, B., Moon, Y., 1996. Can computers be teammates? *Int. J. Hum. Comput. Stud.* 45 (6), 669–678.
- Nass, C., Lee, K.M., 2001. Does computer-synthesized speech manifest personality? Experimental tests of recognition, similarity-attraction, and consistency-attraction. *J. Exp. Psychol.* 7 (3), 171.
- Nass, C., Moon, Y., 2000. Machines and mindlessness: social responses to computers. *J. Soc. Issues* 56 (1), 81–103.
- Nass, C., Steuer, J., Henriksen, L., Dryer, D.C., 1994. Machines, social attributions, and ethopoeia: performance assessments of computers subsequent to "self-" or "other-" evaluations. *Int. J. Hum. Comput. Stud.* 40 (3), 543–559.
- Newcomb, T.M., 1961. *The Acquaintance Process*. Holt, Rinehart & Winston.
- Nguyen, T.T., Hui, P.M., Harper, F.M., Terveen, L., 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*. ACM, pp. 677–686.
- Pariser, E., 2011. *The Filter Bubble: What the Internet is Hiding from You*. Penguin UK.
- Payne, J.W., 1976. Task complexity and contingent processing in decision making: an information search and protocol analysis. *Organ. Behav. Hum. Perform.* 16 (2), 366–387.
- Pennebaker, J.W., Mehl, M.R., Niederhoffer, K.G., 2003. Psychological aspects of natural language use: our words, our selves. *Annu. Rev. Psychol.* 54 (1), 547–577.
- Reeves, B., Nass, C.I., 1996. *The Media Equation: How People Treat Computers, Television, and New Media Like Real People and Places*. Cambridge University Press.
- Reinartz, W., Krafft, M., Hoyer, W.D., 2004. The customer relationship management process: its measurement and impact on performance. *J. Market. Res.* 41 (3), 293–305.
- Reinecke, K., Bernstein, A., 2015. Knowing what a user likes: a design science approach to interfaces that automatically adapt to culture. *MIS Q.* 37(2).
- Ricci, F., Rokach, L., Shapira, B., 2015. *Recommender Systems: Introduction and Challenges*. Recommender Systems Handbook. Springer, pp. 1–34.
- Riedl, R., Davis, F.D., Hevner, A.R., 2014. Towards a neurois research methodology: intensifying the discussion on methods, tools, and measurement. *J. Assoc. Inf. Syst.* 15 (10), 1.
- Riedl, R., Hubert, M., Kenning, P., 2010. Are there neural gender differences in online trust? An fMRI study on the perceived trustworthiness of ebay offers. *MIS Q.* 34 (2), 397–428.
- Roccas, S., Schwartz, S.H., 1993. Effects of intergroup similarity on intergroup relations. *Eur. J. Soc. Psychol.* 23 (6), 581–595.
- Rokeach, M., 1968. A theory of organization and change within value-attitude systems I. *J. Soc. Issues* 24 (1), 13–33.
- Rook, L., Sabic, A., Zanker, M., 2018. Reinforcement sensitivity and engagement in proactive recommendations: experimental evidence. *Information Systems and Neuroscience*. Springer, pp. 9–15.
- Rossetti, M., Stella, F., Zanker, M., 2016. Contrasting offline and online results when evaluating recommendation algorithms. *Proceedings of the 10th ACM Conference on Recommender Systems*. ACM, pp. 31–34.
- Said, A., Bellogin, A., 2014. Comparative recommender system evaluation: benchmarking recommendation frameworks. *Proceedings of the 8th ACM Conference on Recommender Systems*. ACM, pp. 129–136.
- Salatino, A.A., Osborne, F., Motta, E., 2017. How are topics born? Understanding the research dynamics preceding the emergence of new areas. *PeerJ Comput. Sci.* 3, e119.
- Schwartz, B., et al., 2016. On the meaning and measurement of maximization. *Judgm. Decis. Mak.* 11 (2), 126.
- Sears, D.O., Freedman, J.L., 1967. Selective exposure to information: a critical review.

- Public Opin. Q. 31 (2), 194–213.
- Shneiderman, B., Preece, J., Pirolli, P., 2011. Realizing the value of social media requires innovative computing research. *Commun. ACM* 54 (9), 34–37.
- Singh, R., Ng, R., Ong, E.L., Lin, P.K., 2008. Different mediators for the age, sex, and attitude similarity effects in interpersonal attraction. *Basic Appl. Soc. Psychol.* 30 (1), 1–17.
- Stone, J., Cooper, J., 2001. A self-standards model of cognitive dissonance. *J. Exp. Soc. Psychol.* 37 (3), 228–243.
- Sunstein, C.R., 2018. *Republic: Divided Democracy in the Age of Social Media*. Princeton University Press.
- Szpektor, I., Maarek, Y., Pelleg, D., 2013. When relevance is not enough: promoting diversity and freshness in personalized question recommendation. *Proceedings of the 22nd International Conference on World Wide Web, WWW '13*. pp. 1249–1260.
- Tam, K.Y., Ho, S.Y., 2006. Understanding the impact of web personalization on user information processing and decision outcomes. *MIS Q.* 30 (4), 865–890.
- Tausczik, Y.R., Pennebaker, J.W., 2010. The psychological meaning of words: liwc and computerized text analysis methods. *J. Lang. Soc. Psychol.* 29 (1), 24–54.
- Vargas, S., Castells, P., 2011. Rank and relevance in novelty and diversity metrics for recommender systems. *Proceedings of the fifth ACM Conference on Recommender Systems*. ACM, pp. 109–116.
- Verhoef, P.C., Reinartz, W.J., Krafft, M., 2010. Customer engagement as a new perspective in customer management. *J. Serv. Res.* 13 (3), 247–252.
- de Visser, E.J., Monfort, S.S., McKendrick, R., Smith, M.A., McKnight, P.E., Krueger, F., Parasuraman, R., 2016. Almost human: anthropomorphism increases trust resilience in cognitive agents. *J. Exp. Psychol.* 22 (3), 331.
- Wobcke, W., Krzywicki, A., Sok, Y., Cai, X., Bain, M., Compton, P., Mahidadia, A., 2015. A deployed people-to-people recommender system in online dating. *AI Mag.* 36 (3), 5–18.
- Xiao, B., Benbasat, I., 2007. E-Commerce product recommendation agents: use, characteristics, and impact. *MIS Q.* 31 (1), 137–209.
- Xu, Y., Li, Z., Gupta, A., Bugdayci, A., Bhasin, A., 2014. Modeling professional similarity by mining professional career trajectories. *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '14*. pp. 1945–1954.
- Yarkoni, T., 2012. Psychoinformatics: new horizons at the interface of the psychological and computing sciences. *Curr. Dir. Psychol. Sci.* 21 (6), 391–397.
- Yoo, K.H., Gretzel, U., Zanker, M., 2012. *Persuasive Recommender Systems: Conceptual Background and Implications*. Springer Science & Business Media.
- Yoo, K.H., Gretzel, U., Zanker, M., 2015. Source factors in recommender system credibility evaluation. *Recommender Systems Handbook*. Springer, pp. 689–714.
- Zanker, M., Bricman, M., Gordea, S., Jannach, D., Jessenitschnig, M., 2006. Persuasive online-selling in quality and taste domains. *E-Commerce and Web Technologies, EC-Web '06*. pp. 51–60.
- Zhang, S., Yao, L., Sun, A., Tay, Y., 2019. Deep learning based recommender system: a survey and new perspectives. *ACM Comput. Surv.* 52 (1), 5.
- Ziegler, C.N., McNee, S.M., Konstan, J.A., Lausen, G., 2005. Improving recommendation lists through topic diversification. *Proceedings of the 14th International Conference on World Wide Web*. ACM, pp. 22–32.