Proposed Chapter

Eelco Herder* and Claus Atzenbeck

Feedback Loops and Mutual Reinforcement in Personalized Interaction

https://doi.org/..., Received ...; accepted ...

Abstract: In personalized interaction between humans and computers, not only computers and personalization algorithms learn about the users: the users also learn about the system's behavior and adapt their expectations accordingly. Particularly, as users expect systems to support their daily activities, this feedback loop may result in long-term changes in these daily activities and user decisions themselves. This can be observed in activities as different as autonomous driving and social media consumption. In this chapter, we investigate these effects by reviewing and analyzing a wide range of relevant literature.

Keywords: personalization, mutual reinforcement, bounded rationality, choice, cognitive friction

PACS: ...

Communicated by: ...
Dedicated to ...

1 Introduction

Everyday computer devices – including laptops, smartphones, smartwatches and other wearables and smart devices – are commonly used to *support* us in our daily activities. Many of these devices and corresponding apps aim to help us in decision-making and selecting between available choices. The currently dominating interface paradigm is that available choices – such as news articles, books, movies or food – are provided in a list or feed, ordered by their relevance, as deemed by the system.

***Corresponding author: Eelco Herder,** Interaction Group, Department of Information and Computing Sciences, Utrecht University, the Netherlands
Claus Atzenbeck, Institute for Information Systems (iisys), Hof University, Germany

The most common technique for creating ordered lists or feeds is by using recommender systems, such as collaborative filtering [28]. A widely recognized inherent limitation of this approach is that the recommendations are bound to be safe choices that are in line with our daily routines. This may sound attractive at first sight, but there is wide evidence that in the longer term, these safe recommendations actually *transform* our daily routines into safe routines that may end up boring, such as binge-watching [32] – or, worse, *unlearn* us to make conscious, arguably better or more satisfying, choices or even make us *unaware* of different opinions, options or perspectives, because we are comfortably stuck in our filter bubbles [44].

In this chapter, we explore how personalized systems support us in human decision making and which – potentially undesirable – effects may happen, in the short term, but particularly in the long term. In Section 2, we discuss human decision models and how this process is supported by computer systems. We then continue, in Section 3, with current insights on how computer-generated lists of choices and user responses to these (often safe) choices reinforce one another and how this eventually may lead to fundamentally different systems and types of interaction than anticipated. In Section 4, we provide and discuss examples on how this affects our current interaction with recommendations, as given by – among others – streaming providers and social media networks. Finally, in Section 5 we provide insight in current, earlier and prospective approaches and strategies for stimulating users to make conscious choices, should they want to. We end the chapter with a summary, discussion and future perspectives.

2 Computer-Supported Human Decision Making

In this section, we explore computer-supported decision making. First, we briefly introduce the concept of bounded rationality and discuss how this concept plays a role in computer-supported decisions. Then, in Section 2.2, we have a look at persuasive recommendations in the wild; we argue that all recommendations have a persuasive aspect. Finally, we look at the interaction between recommendations, user acceptance and user expectations.

2.1 Bounded Rationality and Personalized Decision Support

A commonly accepted model of human decision making is Kahneman's concept of *bounded rationality* [33]. According to this theory, spontaneous human decisions are often reflexive and intuitive – in terms of Kahneman, this is *system-1 behavior*. In many situations, this type of 'fast thinking' allows us to make quick decisions, which are surprisingly often correct. However, when one feels or suspects that there is reason for doubt, we resort to logical reasoning and conscious decision-making, so-called *system-2 behavior*.

As an example, when you drive on the highway while talking with the front-seat passenger and see the exit to your hometown that you always take, you probably would leave the highway without deeply thinking about it. Only when this routine is interrupted, you start thinking about this activity consciously; for instance, when there are construction works just before the exit, or when you actually planned to take one exit further, in order to bring your passenger home.

Arguably, (personalized) computer decision support aims to simplify or even (further) automatize our decision processes – to reduce our cognitive load or to let us make better choices by showing appropriate options [6]. In terms of bounded rationality, one could say that personalized support keeps users comfortably in their low-effort system-1 mode of behavior.

To illustrative this further, we go back to our car-driving example. Most drivers commonly use voice navigation support apps, even on routine rides, and they have gotten used to automatically following prompts to stay on the road or to take the exit. Suppose that you erroneously would have taken your usual exit and then realize your mistake: most likely, your navigation support system would have noticed it even before you have realized it yourself, so that you can comfortably follow its prompts in order to get back to the highway.

Building upon the car-driving example, it should be noted that there are several levels of support that can be provided [9]. Many cars have built-in tools that mainly *support* our actions, such as anti-lock braking, cruise control or warning systems that, for example, alert the driver when a speed limit is exceeded. On a higher level, navigation systems provide *pro-active advice* which direction to take or which route to follow, but they still leave the driver in charge. Even more automated, highway pilot and traffic jam pilot systems, *partially take over decision making* from the driver, while still leaving the driver in control. Arguably, in all above-mentioned usage contexts, users and computers work together, but with clear differences in how decisions are initiated, evaluated and executed.

2.2 Persuasive Recommendations in the Wild

A further effect of our bounded rationality is that many of our choices are largely driven by our fast system-1 thinking. This does not only yield for activities such as driving, as discussed in the previous section, but also for consumer choices, which have been observed to be choices that look attractive in the short term or that satisfy our (hedonistic) needs [42].

Persuasive recommender systems are a particular category of recommender systems that attempt to persuade users to adopt particular behavior or to make particular choices by using one or more cues [19]. Traditionally, cues considered in persuasive systems are *social*, such as a particular type of wording, or emphasizing certain desirable or undesirable social dynamics associated with particular choices. Since 2008, so-called *nudges* form a particularly popular category of cues [55]. Nudges are small interventions in a choice environment that make certain – arguably better – choices more attractive, in order to steer user behavior in a gentle way.

Increasingly, it has been recognized, though, that *any* recommender system is persuasive in some sort of way. Following Netflix, commercial recommender algorithms have been mainly designed to optimize click-through rates, and they have been evaluated based on choices that users have already made – "some of the worst possible real-world recommendations", as observed by [36]. In terms of bounded rationality, these recommendations reinforce – and do not challenge – users' system-1 behavior, the natural preference for attractive, easy choices [42].

Our natural preference for easy, intuitive choices is clearly visible in *user interaction* on the Web as well. Users have consistently been observed to click on only a few of the highest ranked search results or recommended items; items that require additional scrolling or navigating to a next page receive considerably less attention [4, 41]. Arguably, the combination of optimized recommendations that are presented in a feed-based paradigm results in a choice environment with very persuasive arguments to indulge in system-1 behavior.

It should be noted that this phenomenon is far from new. For example, it is commonly known and accepted that supermarkets design their stores in such a way that maximizes visual attention for the most profitable products [22]. On the other hand, in-store shopping strategies for reducing time and money costs in the supermarket – in other words, strategies for avoiding any 'traps' set up by the supermarket – are well-known and counterbalance this effect [58].

2.3 Recommendations, Recommendation Take-Up and User Expectations

In principle, one can see commercial recommender systems – including those used by streaming services, personalized stores and social networks – as technology that simply provides what the users expect. Given the success of these platforms, this may be a valid point of view.

Another perspective is that user expectations are largely driven by the options that are provided or promoted by these platforms. Theoretically, this would lead to a balance between supply and demand, in which platforms adjust their offerings to the users' expectations [47]. However, as discussed in the previous section, user expectations are shaped by the available choices – not just in the short-term, but also by long-term effects, as will be elaborated in the upcoming Section 3.2. A particular effect that we will discuss is that the combination of feed-based interfaces and recommender algorithms has conditioned us, the users, to conveniently consume what 'the feed' provides, and that these feeds are aligned with our (system-1) expectations. As a result, these systems are said to create what is popularly called *filter bubbles* [44].

However, as previously discussed in Section 2.1, users often feel comfortable themselves with fast, routine choices, and only engage in more active (system-2) decision-making activities when prompted that they might want or need to do so. Persuasive recommender systems and nudging techniques aim to provide such prompts for making 'prudent' or generally more informed choices [19].

Persuasive literature often concentrates on systems that aim to motivate users to make choices that society considers better – for example in terms of health or environmental impact –, but in this chapter we mainly concentrate on mechanisms that aim to support users in making active decisions that they feel more comfortable with in the long term, but not necessarily in the short term. This type of self-optimization or self-control is not limited to ambitious lifestyle changes, but also for more mundane decisions, such as actively choosing a documentary or music album, instead of just consuming what happens to be offered [51].

As we will discuss in Section 5, it is difficult to design nudges or other mechanisms that support users to be 'their better self': a related and perhaps more fundamental problem is that users themselves have difficulty to become their more effortful (system-2) better self [33] – the proverbial "I will start my diet tomorrow" trap. Being reminded of goals – even if set or planned by yourself – by an external system quickly feels patronizing [60]. Even more difficult perhaps than influencing a single choice is *changing* a users' attitude and behavior in such a way that – at some point – this changed behavior becomes natural, even without (strong) support [45].

2.4 Summary and Outlook

In this section, we have discussed the mutual interaction between users and (recommender) systems in decision making. Depending on its design, a recommender system may automate this choice process to a smaller or to a larger extent. Current commercial, often feed-based, recommender systems appear to take more control in this process than we – as users or society – consider desirable.

Persuasive systems are designed to stimulate users to adopt a particular kind of behavior. In academic contexts, this usually implies helping the user to be their better, healthier (system-2) self. Commercial recommender systems are by design persuasive as well, but largely serving system-1 goals, which may result in systems behaving as ‘friends with a bad influence’ [57]. However, poor decisions or decisions that are regretted at a later point are (almost) never (entirely) the system’s fault: it is ultimately the users’ decision, reflecting their own desires, ambitions and self-control [27]. That is to say, unless users have come in the habit of lending too much responsibility to the system. This is what will be explored in the next section.

3 Mutual Reinforcement in Classical User Modeling and HCI

A basic premise in human-computer interaction is that users aim to accomplish a goal, aided by an interactive system [17]. For this purpose, the user plans and executes several actions – such as clicking on an icon or a link, typing a search query or following a link – and then evaluates whether the system state matches the expectations, for example whether the click on the link indeed leads to the loading of a new page. This is a classic human-system feedback loop that we experience on a daily basis.

Personalized systems – including recommender systems – create even tighter feedback loops, as they “build a model of the goals, preferences and knowledge of each individual user, and use this model throughout the interaction with the user, in order to adapt to the needs of that user” [6].

3.1 Short-Term Mutual Reinforcement

Web search is a classic example in which the user and the system work together to reach a certain outcome, and respond to each others responses. Users start with an initial query, inspect the results and then, if needed, reformulate the query; for example, they may decide to add or remove some keywords or to replace a term by a synonym [49]. It has been recognized that this iterative process helps the user to build a context in which to interpret and understand the results [54].

In user modeling, particularly in the earlier days, *concept drift* was seen as potentially problematic. Users’ interests were recognized as being dynamic and likely to change. In [59], a special case of concept drift was described in which a news recommender presents news stories to the user. These stories "are assumed to directly affect the user's information needs", which raised the need for the system to anticipate and respond to this expected effect. This effect – described as early as in 2001 – is still quite short-term, but can already be seen as a forebode for current longer-term effects of misinformation and fake news in social media [52].

Short-term tensions in feedback loops between users and personalized systems have been recognized by commercial providers. Among others, it is considered problematic that systems may have different goals – such as optimizing click-through rates – than the user may have. However, longer-term effects of interventions in the algorithm or in the interface are under-investigated and hardly documented [13].

3.2 Long-Term Mutual Reinforcement

In Section 2.1, we have discussed how car automation has made decision-making a collaborative activity between human drivers and the driving support tools, such as navigation systems and cruise control. Short-term implications of human-car feedback loops in terms of driver attention and division of responsibility in case of accidents and emergency situations have been discussed in the literature [9, 24, 35]. It appears that only recently it has been recognized or acknowledged that (partial) automation reshapes the driving task and therefore the driver's behavior. As expected, it has been observed that most drivers prefer to use high levels of automation, when available [21]. For instance, car navigation systems have made us accustomed to passively following turn-by-turn instructions, which has been observed to have a negative impact on our sense of direction or our inclination to comprehend the environment in which we are driving [29].

In a similar manner, automation and personalization have changed our media consumption behavior – and consequently our expectations of what media outlets should offer and the goals these offerings should fulfill. During the first decades after the introduction of the television in households, television watching was observed to be a mix of active and passive choices: consumers had particular expectations on what programs to expect at particular times, but program choices were often an active family decision and viewers, with the possibility to select another channel – or to turn the television off –, if a program turned out not to be satisfying enough [2, 38]. In the early 2000s, streaming services such as Netflix or Amazon Prime dramatically changed watching behavior from weekly 'watching appointments' with one or more series to *binge-watching*: watching two episodes or more of the same series in a row [48]. This behavior has arguably been reinforced by optimizing the recommender algorithm based on the click-through rate – which was exactly the goal of the Netflix Prize challenge [36]. As will be discussed in Section 4, this click-rate optimization largely was achieved by creating feeds largely consisting of safe choices that reinforce routine 'system-1' behavior.

Largely, the changes as described above are designed and anticipated, but often there are also unexpected effects and (user) responses that need to be repaired or addressed. Moreover, in the long term, this can lead to systems – and associated user behavior – that are fundamentally different from their predecessors. Researchers in the field of *second-order cybernetics* fundamentally regard this as a long-term feedback loop between the system, its designers, its users and other stakeholders – and all parties involved, including stakeholders who may consider themselves as mere 'observers', play an active role in this process [37]. In other words, they treat this human-system feedback loop as a *self-organising system.*

A particular aspect of the systems discussed above is that they are used on a regular basis in order to achieve something in our daily lives. Moreover, we have observed that these systems actually have an impact on our lives, for example by simplifying or automating tasks or decisions. In terms of cybernetics, there is an ongoing process or regulation and feedback – in both directions – in order to maintain a state of relative stability [43]. This does not only imply that the algorithms and interfaces of a (personalized) system do not dramatically change, but also that the user contexts and the way users interact with the system remain fairly stable. Any disturbance of this equilibrium may lead to the need or desire for 'reparation'.

As we have observed earlier, users have a natural tendency towards easy, convenient system-1 choices. Therefore, it is tempting for (commercial) system designers to cater for these decisions. Voluntarily or involuntarily, this may lead to increased system-1 behavior – such as active television program selection that has gradually largely been replaced by a habit of picking a series that happens to be in the feed and appears acceptable enough. This is not a bad choice on itself, it may save precious time and prevent choice overload, but it may create a choice environment that users themselves – or society as a whole – may perceive as not healthy or not optimal.

Social media is a particularly problematic example, in which several iterations of design changes – responding to and anticipating user responses – have led to social media platforms being prone to polarization, misinformation and fake news [52]. As discussed earlier in Section 2.3, the feed-based design of social media may lead to filter bubbles of safe choices; however, users usually feel quite comfortable as well in a choice environment that does not challenge their current views or expectations – the so-called *echo chamber effect* [10]. Even though both the filter bubble and the echo chamber effects appear limited for individual users in a high-choice environment [18], the polarising effects on the platform as a whole are undeniable. The long-term effect of user responses to seemingly small interface design decisions can best be illustrated by Facebook's introduction of Reactions ('Love', 'Care', 'Haha', 'Wow', 'Sad', and 'Angry') in 2016, which has been shown to lead to more emotional user responses, with increased polarisation on the platform and in society in general as a result [53].

3.3 Summary and Outlook

In this section, we have explored the mutual reinforcement between users and the systems that they use. In human-computer interaction research, short-term reinforcement is a well-researched and quite well-understood process: users and systems interact with one another in order to reach a state of relative stability that allows users to achieve what they aim to achieve in a satisfactory manner. This state of relative stability does allow for minor changes in system design or usage patterns, but these changes may lead to the need of slight adjustments on either side – in the field of cybernetics, it is not uncommon to use the term *thermodynamics* for this [43].

The process of long-term mutual reinforcement is far less understood [13]. This does not come as a surprise, as these long-term effects are usually the result of a chain of planned and unplanned short-term effects. As we have seen in this chapter, users have a natural tendency towards convenience and automation. Personalized systems – and commercial recommender systems in particular – are often designed to support and encourage convenient, routine behavior and routine choices [5], either or not with click-through optimization as an explicit system goal [13].

Naturally, these longer-term effects are often easily observed in hindsight, as we have seen in the fields of recommender systems in general as well as in social media [36, 53]. Particularly when it comes to negative effects, it would be desirable to better anticipate them or to recognize them in an early stage, before 'the damage is done'. For personalized systems, it appears imperative to have a deeper understanding of these chains of feedback loops. For this purpose, in the next section we will further explore the history of academic and commercial recommender systems.

4 Everyday Recommendations from Commercial Providers

Originally, adaptive or personalized systems were thought to help users with activities such as online learning, finding relevant information, or managing personal information spaces; still, already in an early stage, opportunities for e-commerce were recognized [6]. With the advent of commercial recommender systems, the importance of metrics increased, in order to quantify sales, click-through rates or conversion rates [36]. As discussed before in Section 3.2, this is best exemplified by the data-driven Netflix Prize challenge. More fundamentally, recommender systems designed for such purposes might be considered *sales instruments* that may – or may not – act in line with the users' interests.

In this section, we explore the expectations and roles of users and recommender systems in terms of expected outcomes. We first briefly discuss how users may evaluate a situation differently in hindsight than while experiencing a situation.

4.1 Conflicting Short-Term and Long-Term User Goals

Traditional user goals for online recommender systems include exploring information spaces, finding good items – such as movies that received good reviews, or simply browsing for entertainment; other, more purposeful goals include improving oneself in one way or another, expressing oneself or helping other people [28]. It is common and natural for humans to evaluate particularly the latter activities in terms of increased satisfaction or happiness. Kahneman argues that a person might evaluate an outcome differently while *experiencing* a situation than when *remembering* a situation [34]. Arguably, the experiencing self is mainly interested in an enjoyable experience, whereas the remembering self also reflects on the purposefulness of the outcomes.

As discussed in section 3.2, users have a natural preference for automation and convenience, which is one of the reasons for the popularity of streaming services and activities such as binge-watching. However, in hindsight, the remembering self might reflect upon some lost time that could have been used for more purposeful or urgent activities – which arguably would have been the more rational choice.

4.2 Recommender Systems Acting As Salespersons

Algorithmic recommender systems that are trained with (user) data are essentially an application of machine learning. In a review of deep learning – a currently popular strand of research in machine learning – it is argued that machine learning is particularly good at *system-1* tasks that involve recognizing previously learned situations and common, expected responses; tasks that "require a deliberate sequence of steps" is observed to be still in its infancy [5].

Moreover, as many recommender systems have been deliberately designed to optimize click-through rates [36], this inherent limitation of machine learning is not countered but rather increased. In terms of Kahneman, recommender systems mainly serve the goals of our experiencing selves, not the remembering selves. Typical commercial evaluation measures – such as the cost per click or the click through rate [25] – can be considered as indications how well the 'recommender system as a salesperson' served the needs as perceived by the user at that very moment.

As discussed in Section 2.2, visual attention plays an important role in advertisements and sales. Visual cues may lead to *unplanned behavior*, largely driven by availability; for instance, common advice for preventing unintended supermarket purchases is to make use of a shopping list and to not deviate from it [1]. Building upon this perspective, it is not surprising that users who have been trained that they only need to scroll further in a feed to find new posts, stories, movies or songs, increasingly engage in such unplanned behavior – as observed by Derakshan.[1]

Of course, these differences in interest between parties do not (necessarily) imply that recommender systems are a user's adversaries: one of the premises of multi-stakeholder recommendation is that *reciprocal recommendations* are acceptable for both parties – and that ideally both parties benefit. This mutual benefit can theoretically be calculated as a multi-stakeholder evaluation measure - and attempts to do so have been made [8].

However, most – if not all – of these efforts focus on the utility from the perspective of the *experiencing users*, not of the *remembering users* – who in hindsight might wish to have spent their time in a more purposeful manner. As we have seen in Section 2.3, making 'prudent' or generally more informed choices – directed at the 'remembering user' – are the field of persuasive systems [19].

1 https://www.theguardian.com/technology/2015/dec/29/irans-blogfather-facebook-instagram-and-twitter-are-killing-the-web

5 Integrating Conscious Decision Making Into the Interface

In Section 3.2, we discussed how drivers have been accustomed – and how they accustomed themselves – to following turn-by-turn instructions from their navigation support system instead of consciously choosing directions themselves. Arguably, for this purpose, the benefits of automation outweigh drawbacks such as a drivers' reduced sense of orientation [29].

A fundamental difference with driving – an activity that is usually carried out to bring a person from a start point to a planned end point - is that online activities, such as Web search or social media use, often are meant to further shape a user's goals and expectations (as discussed in Section 3.1) or that these online activities actually are meaningful by themselves – such as keeping up with the news or politics, or interacting with family, friends and colleagues.

Already in Section 2.1, we have discussed that there are several levels of automation that systems can provide and in Section 3.2, we observed a user's natural preference for automation. Without imposing any moral views, we believe that – at least in hindsight – users often would prefer some more initiative and conscious decision making and a bit less turn-by-turn following of system recommendations [44, 27]. Sometimes we want or should want to be challenged, in order to review our current behavior or to consciously look for better – tastier, more interesting, more challenging, environment-friendlier – options.

In this section, we discuss several strands of techniques that aim to activate and support a user's conscious *system-2* behavior. Several of these approaches fall in the field of persuasive systems, but we think it is important to keep in mind that personalized systems and recommender systems are inherently persuasive by themselves – as discussed in Section 2.2.

5.1 Traditional Recommender System Approaches: Explanations and Cognitive Forcing

Explanations play an important role in recommender systems. Common goals of these explanations include helping users understand how a system works, allowing users to tell the system when it is wrong, increase users' confidence in a system and convincing users to make a certain decision [56]. Explanations and user control increase user trust and satisfaction, and make personalized systems more transparent – in the chapter "Explanations and user control in recommender systems", elsewhere in this book, Jannach and colleagues discuss several techniques for doing so.

A fundamental issue with explanations in recommender system is that users need to *engage* with them in order to experience the benefits of them. Early attempts on explainable recommender systems – then called *scrutable adaptive hypertext* – already found that it is challenging to provide an effective interface that encourage users to actively explore the rationale behind certain recommendations [14].

Moreover, explainable recommender systems usually focus on explaining why certain items are being recommended – and not on which items have *not* been recommended and which alternative directions or solutions a user may want to consider. As discussed by Baeza-Yates [4], recommender systems on the Web generally suffer from activity bias and data bias, which leads to lists of search results or recommendations that only represent a very small part of a solution space. Unfortunately, this issue is aggravated by *user interaction bias*, the already observed effect that users naturally prefer convenience and automation. According to a recent study, even if users do interact with recommendations and their explanations, this usually does not have long-term effects [20].

An alternative technique is *cognitive forcing* [7], which appears to be an extension of *cognitive friction* [12]. Cognitive forcing involves designers consciously building hurdles in the UX experiences, in order to force the user to make conscious decisions – for example by asking users to make an initial choice before showing any recommendation. The problem with such approaches is that users typically experience these interventions as unpleasant and cumbersome [7].

5.2 Mixed-Initiative or Conversational Interfaces

Rather than considering requesting or requiring user involvement as a hurdle, one could consider user involvement as a *conversation* with a personalized system. Recommender systems that involve one or more steps of interaction are called *mixed-initiative* interfaces.

Explicit user involvement may take place already before the recommendation process starts, for example by asking users to provide some initial ratings of selected items [39]. Another approach for soliciting user responses is to allow them to give explicit feedback on recommended items, for example by rating or reordering them. Similar to explainable recommenders systems – as discussed in the previous subsection – it was found that putting users into control can be challenging [31].

A problem with many mixed-initiative recommender systems is that the conversation between user and system does not flow naturally: often it still is the system that asks or requires user input, not the user who naturally reacts to a certain proposition. A promising development for conversational recommender systems is the growing popularity of voice assistants that users can talk with and that are already commercially exploited by various big tech companies [46]. However, it is still considered a challenge to initiate or stimulate natural dialogues between users and conversational (voice) assistants that go beyond simple requests and responses [30].

It should be noted that voice interaction is not mandatory for mixed-initiative interfaces. For instance, as discussed in Section 3.1, many users feel comfortable refining their queries in several iterations, in response to the list of results of their previous query [49]. In terms of second-order cybernetics (see Section 3.2), this – arguably trained – user interaction with search engines by means of queries and query results appears to be a fairly unstable equilibrium, though. Already, longer-term reinforcement effects between search engines and their users can be observed: current search engines increasingly provide direct answers and information snippets in response to user queries, which has led to a reduction in user interaction with the result sets [61]. In a sense, this development may gradually turn a search engine into an 'answer engine', with users considering the results to their initial query as an answer to their question, without further exploring the solution space [54].

5.3 Hypertext-Inspired Choice Models

In the previous sections, we discussed the role of explanations, cognitive forcing, mixed-initiative and conversational interfaces on conscious decision making. Despite the differences, all approaches have in common that the initial initiative lies on the computer side: all user considerations start with one or more items proposed by the system, with an explicit prompt to make a choice.

By contrast, the original ambitions of hypertext – the concept that lies at the core of Web interfaces and personalized systems – was to *support* human thinking by means of rich interfaces with typed and bidirectional links and visualizations, which would allow users to read, think and reason in a non-sequential manner [40]. A hypertext system was envisioned to function as a personal assistant, or a secretary that would follow a user's instructions, therewith *augmenting human intellect* [11].

Hypertext systems were explicitly designed to facilitate interaction and conscious deliberations, exploiting visual metaphors and connections, such as fixed-sized cards, landmarks, footprints or birds' eye overviews [26, 41]. A particular approach for representing and support human thinking and decision making is the concept of *spatial hypertext*, which visualizes associations between items (or nodes) by their position on a (hyper)space and the relative distances from one another – further enhanced by visual cues, such as color, shape and size. Hypertext systems encourage users to directly manipulate these positions to better reflect their current line of thinking, their preferences or their priorities.

Indeed, a premise of hypertext is that it supports associative thinking and conscious decision making, an activity that Kahneman coined as *system-2* thinking. In contrast to linear feeds, which suggest or implicate an order, spatial hypertext structures provide a choice landscape that may trigger users to explore information or solutions that they might not have considered first [16].

Particularly in the early days of the Web, it was lamented that the lack of rich links, visual metaphors and cues for direct manipulation made Web navigation only a far cry from the original intentions of hypertext researchers and designers [15]. In contrast to the associative nature of hypertext, navigation on the Web is largely hierarchical, guided by menu structures and – particularly in the domain of recommender systems – by linear feeds.

5.4 Interventions and Nudges Versus Interface Paradigms

As discussed in more detail in [3], the popularity of linear feeds and hierarchical menus can be explained by the users' natural tendency towards routines and safe choices [23], ideally supported by a system that confirms and reinforces these choices.

Despite the benefits, convenience and commercial success of recommender systems, the limitations in terms of supporting conscious decision-making and its narrowing effects on user choices has been recognized [36]. In response, several research efforts are carried out in order to build cognitive friction – as discussed in Section 5.1 into recommender interfaces: elements that stop or slow down interaction flows and that are expected to encourage users to reconsider initial impressions or decisions.

Accepting the perspective of feed-based recommender systems as decision support that largely reinforces those choices that we routinely would make, *nudges* are defined as a "*de facto* influence on choices that is 'easy and cheap to avoid'" [55]. Nudging theory builds upon psychological theory that separates the 'experiencing self' from the 'remembering self', a notion that we briefly discussed in Section 4.1. In leisure science, for instance, it has been recognized that there is a difference between what a user experiences in the moment and how the same user remembers this experience at the end of the day [62].

In the context of this chapter, we mainly focus on nudges or interventions that help users making decisions that they feel satisfied about afterwards, a goal that arguably is in the best interest of the user. However, if we consider the combination of feed-based recommenders that largely reinforce a user's system-1 behavior together with interventions that aim to reinforce system-2 choices – such as nudges, explanations or cognitive friction –, then it can be argued that every single choice is managed by the system [50]. In other words, this type of choice architecture may still seduce the user to make easy, convenient, attractive choices – such as watching the next episode of a series – and at the same time remind the user that making a different, at first sight perhaps less attractive choice – such as reading a book instead – may be more prudent and make the user more satisfied by the end of the day. Such an ambivalent choice architecture that simultaneously seduces and patronizes the user is likely to cause frustrating user experiences.

5.5 Summary and Considerations

In this section, we have discussed several approaches for integrating conscious decision making into the interface of recommender systems, as we know them now. In a sense, all these approaches aim to 'repair' a situation in which users have learned to appreciate the convenience of feed-based suggestions. Efforts to address this issue – including explanations, cognitive forcing and mixed-initiative interfaces – have an impact on the choice architecture, but still leaves the initiative for generating the options largely on the system-side.

Building upon the notions on the balance in levels of support in decision-making – which we discussed in Section 2.1 – interfaces that allow and encourage users to actively consider several options and to develop their own choice architecture would bring the users more in control. Theoretically, this would lead to more satisfactory outcomes, but – unfortunately – interfaces that attempt to do so are often considered as cumbersome.

This leads us to a somewhat paradoxical situation: as a long-term effect of mutual reinforcement, users have learned to appreciate personalized interfaces that make it easy to make safe, convenient choices that look attractive to the experiencing self – and recommender systems have been optimized for exactly this purpose; but now, we need to find mechanisms to convince users that at the end of the day they would be happier with decision support that helps them with choices that their remembering self would be happier with at the end of the day.

6 Discussion and Future Perspectives

In this chapter, we have explored and discussed feedback loops and mutual reinforcement in personalized interaction between users and systems, with a focus on systems that help users in making decisions and choices. We have observed that users appreciate the convenience of personalized decision support, including recommendations that reduce the need for making active decisions.

In the past decades, the ecosystem of personalized systems and their users has largely been shaped by commercial recommender systems, which responded to – and shaped – user expectations and user appreciation. A particular long-term effect of this co-evolution is that users have become accustomed to rather passively following suggestions rather than actively shaping and investigating their choice architecture.

This natural preference for automation and convenience, with recommender systems that support convenient system-1 decisions, works perfectly fine if these decisions do not have negative consequences at the end of the day – in other words, when the reflecting user would still be happy with, or at worst indifferent about, these choices. However, there are quite some situations where users – or society as a whole – realize too late that active decisions at an earlier stage would have led to leisure activities, social interactions or choices that in hindsight would have been more purposeful, prudent or productive.

Several approaches have been explored to combat this effect, to make recommender systems 'more fair' by means of explanations, conversations or nudges. However, it is still an open question how the short-term effects of such interventions translate to long-term effects – and how this will shape our daily activities, our social interactions and the way we approach decisions and choices.

In order to move recommender systems as a technology forward, Konstan and Terveen plead for a more holistic approach, taking user perceptions, choice overload and differences in user objectives into account. They also argue that recommender systems should learn from situations in which recommended items are *not* selected by the user [36]. While helpful, the proposed approaches still leave the initiative for decisions at the system.

As argued by Bengio et al [5], machine learning – and therewith also recommender systems – are not (yet) able to carry out tasks that "require a deliberate sequence of steps". By contrast, humans are quite good in tasks that require reasoning and active decision making, but they may only engage in such tasks when prompted to do so. As discussed in Section 5.2 and Section 5.3, designers find it hard to create interfaces that stimulate user to do so.

Throughout this chapter, we have observed that system designers as well as the users prefer recommendations and results to be presented as a list of options to choose from, already prioritized by the system. We believe it would be productive if result sets would contain elements that prompt users to treat such lists as a selection of possible choices that can be accepted, deliberated, challenged or discarded. Building in such choice moments may involve simple interventions such as breaking an 'infinite feed' into separate pages: even the single activity of clicking on a 'next' button may create a (very) brief moment for reflection. More elaborate strategies may be adopted by voice-based conversational recommender systems [30], with dialogue strategies that specifically encourage users to engage in deliberation.

To conclude, a general and arguably unforeseen long-term effect in user interaction with personalized systems is the user's inherent preference for convenience and automation, a preference that personalized systems may need to challenge more often. Particularly, it seems very useful and productive to explicitly engage users in a process of active decision making, ensuring that at the end of the day the remembering user will be satisfied with the choices made by the experiencing user, supported – and not driven – by a personalized system.

References

[1] Russell Abratt and Stephen Donald Goodey. Unplanned buying and in-store stimuli in supermarkets. *Managerial and decision economics*, 11(2):111–121, 1990.

[2] William J Adams. How people watch television as investigated using focus group techniques. *Journal of Broadcasting & Electronic Media*, 44(1):78–93, 2000.

[3] Claus Atzenbeck, Eelco Herder, and Daniel Roßner. Breaking the routine: spatial hypertext concepts for active decision making in recommender systems. *New Review of Hypermedia and Multimedia*, pages 1–35, 2023.

[4] Ricardo Baeza-Yates. Bias on the web. *Communications of the ACM*, 61(6):54–61, 2018.

[5] Yoshua Bengio, Yann Lecun, and Geoffrey Hinton. Deep learning for ai. *Communications of the ACM*, 64(7):58–65, 2021.

[6] Peter Brusilovsky. Adaptive hypermedia. *User modeling and user-adapted interaction*, 11(1):87–110, 2001.

[7] Zana Buçinca, Maja Barbara Malaya, and Krzysztof Z Gajos. To trust or to think: cognitive forcing functions can reduce overreliance on ai in ai-assisted decision-making. *Proceedings of the ACM on Human-Computer Interaction*, 5(CSCW1):1–21, 2021.

[8] Robin D Burke, Himan Abdollahpouri, Bamshad Mobasher, and Trinadh Gupta. Towards multi-stakeholder utility evaluation of recommender systems. *UMAP (Extended Proceedings)*, 750, 2016.

[9] Stephen M Casner, Edwin L Hutchins, and Don Norman. The challenges of partially automated driving. *Communications of the ACM*, 59(5):70–77, 2016.

[10] Matteo Cinelli, Gianmarco De Francisci Morales, Alessandro Galeazzi, Walter Quattrociocchi, and Michele Starnini. The echo chamber effect on social media. *Proceedings of the National Academy of Sciences*, 118(9):e2023301118, 2021.

[11] Jeff Conklin, Albert Selvin, Simon Buckingham Shum, and Maarten Sierhuis. Facilitated hypertext for collective sensemaking: 15 years on from gibis. In *Proceedings of the 12th ACM conference on Hypertext and Hypermedia*, pages 123–124, 2001.

[12] Alan Cooper and Paul Saffo. *The Inmates Are Running the Asylum*. Macmillan Publishing Co., Inc., USA, 1999.

[13] Henriette Cramer and Juho Kim. Confronting the tensions where ux meets ai. *Interactions*, 26(6):69–71, 2019.

[14] Marek Czarkowski and Judy Kay. A scrutable adaptive hypertext. In *Adaptive Hypermedia and Adaptive Web-Based Systems: Second International Conference, AH 2002 Málaga, Spain, May 29–31, 2002 Proceedings 2*, pages 384–387. Springer, 2002.

[15] David R Danielson. Web navigation and the behavioral effects of constantly visible site maps. *Interacting with computers*, 14(5):601–618, 2002.

[16] Andrew Dillon, John Richardson, and Cliff McKnight. Navigation in hypertext: a critical review of the concept. In *Proceedings of the 3rd International Conference on Human-Computer Interaction*. North Holland, 1990.

[17] Alan Dix, Janet Finlay, Gregory D Abowd, and Russell Beale. *Human-computer interaction*. Pearson Education, 2003.

[18] Elizabeth Dubois and Grant Blank. The echo chamber is overstated: the moderating effect of political interest and diverse media. *Information, communication & society*, 21(5):729–745, 2018.

[19] Brian J Fogg. Persuasive technology: using computers to change what we think and do. *Ubiquity*, 2002(December):2, 2002.

[20] Krzysztof Z Gajos and Lena Mamykina. Do people engage cognitively with ai? impact of ai assistance on incidental learning. In *27th International Conference on Intelligent User Interfaces*, pages 794–806, 2022.

[21] Pnina Gershon, Sean Seaman, Bruce Mehler, Bryan Reimer, and Joseph Coughlin. Driver behavior and the use of automation in real-world driving. *Accident Analysis & Prevention*, 158:106217, 2021.

[22] Kerstin Gidlöf, Andrey Anikin, Martin Lingonblad, and Annika Wallin. Looking is buying. how visual attention and choice are affected by consumer preferences and properties of the supermarket shelf. *Appetite*, 116:29–38, 2017.

[23] Marta C Gonzalez, Cesar A Hidalgo, and Albert-Laszlo Barabasi. Understanding individual human mobility patterns. *nature*, 453(7196):779–782, 2008.

[24] Noah J Goodall. Can you program ethics into a self-driving car? *Ieee Spectrum*, 53(6):28–58, 2016.

[25] Asela Gunawardana, Guy Shani, and Sivan Yogev. Evaluating recommender systems. In *Recommender systems handbook*, pages 547–601. Springer, 2012.

[26] Frank G Halasz. Reflections on" seven issues" hypertext in the era of the web. *ACM Journal of Computer Documentation (JCD)*, 25(3):109–114, 2001.

[27] Eelco Herder, Daniel Roßner, and Claus Atzenbeck. Reflecting on social media behavior by structuring and exploring posts and comments. *i-com*, 19(3):239–250, 2020.

[28] Jonathan L Herlocker, Joseph A Konstan, Loren G Terveen, and John T Riedl. Evaluating collaborative filtering recommender systems. *ACM Transactions on Information Systems (TOIS)*, 22(1):5–53, 2004.

[29] Toru Ishikawa. Satellite navigation and geospatial awareness: Long-term effects of using navigation tools on wayfinding and spatial orientation. *The Professional Geographer*, 71(2):197–209, 2019.

[30] Dietmar Jannach and Li Chen. Conversational recommendation: A grand ai challenge. *AI Magazine*, 43(2):151–163, 2022.

[31] Dietmar Jannach, Ahtsham Manzoor, Wanling Cai, and Li Chen. A survey on conversational recommender systems. *ACM Computing Surveys (CSUR)*, 54(5):1–36, 2021.

[32] Mareike Jenner. Is this TVIV? on Netflix, TVIII and binge-watching. *New media & society*, 18(2):257–273, 2016.

[33] Daniel Kahneman. Maps of bounded rationality: A perspective on intuitive judgment and choice. *Nobel prize lecture*, 8(1):351–401, 2002.

[34] Daniel Kahneman. *Thinking, fast and slow.* macmillan, 2011.

[35] Katja Kircher, Annika Larsson, and Jonas Andersson Hultgren. Tactical driving behavior with different levels of automation. *IEEE Transactions on Intelligent Transportation Systems*, 15(1):158–167, 2013.

[36] Joseph Konstan and Loren Terveen. Human-centered recommender systems: Origins, advances, challenges, and opportunities. *AI Magazine*, 42(3):31–42, 2021.

[37] Klaus Krippendorff. The cybernetics of design and the design of cybernetics. In *Design Cybernetics*, pages 119–136. Springer, 2019.

[38] Barbara Lee and Robert S Lee. How and why people watch tv: Implications for the future of interactive television. *Journal of advertising research*, 35(6):9–19, 1995.

[39] Sean M McNee, Shyong K Lam, Joseph A Konstan, and John Riedl. Interfaces for eliciting new user preferences in recommender systems. In *User Modeling 2003: 9th International Conference, UM 2003 Johnstown, PA, USA, June 22–26, 2003 Proceedings 9*, pages 178–187. Springer, 2003.

[40] Theodor Holm Nelson. Complex information processing: a file structure for the complex, the changing and the indeterminate. In *Proceedings of the 1965 20th national conference*, pages 84–100, 1965.

[41] Jakob Nielsen. User interface directions for the web. *Communications of the ACM*, 42(1):65–72, 1999.

[42] Harri Oinas-Kukkonen. A foundation for the study of behavior change support systems. *Personal and ubiquitous computing*, 17(6):1223–1235, 2013.

[43] Michael Paetau. Niklas luhmann and cybernetics. *Journal of sociocybernetics*, 11(1/2), 2013.

[44] Eli Pariser. *The filter bubble: How the new personalized web is changing what we read and how we think*. Penguin, 2011.

[45] Michel Peeters, Carl Megens, Elise Van Den Hoven, Caroline Hummels, and Aarnout Brombacher. Social stairs: taking the piano staircase towards long-term behavioral change. In *Persuasive Technology: 8th International Conference, PERSUASIVE 2013, Sydney, NSW, Australia, April 3-5, 2013. Proceedings 8*, pages 174–179. Springer, 2013.

[46] Marta Perez Garcia and Sarita Saffon Lopez. Exploring the uncanny valley theory in the constructs of a virtual assistant personality. In *Intelligent Systems and Applications: Proceedings of the 2019 Intelligent Systems Conference (IntelliSys) Volume 1*, pages 1017–1033. Springer, 2020.

[47] Guilherme D Pires, John Stanton, and Paulo Rita. The internet, consumer empowerment and marketing strategies. *European journal of marketing*, 40(9/10):936–949, 2006.

[48] Matthew Pittman and Kim Sheehan. Sprinting a media marathon: Uses and gratifications of binge-watching television through netflix. *First Monday*, 2015.

[49] Soo Young Rieh et al. Analysis of multiple query reformulations on the web: The interactive information retrieval context. *Information Processing & Management*, 42(3):751–768, 2006.

[50] Andreas T Schmidt and Bart Engelen. The ethics of nudging: An overview. *Philosophy compass*, 15(4):e12658, 2020.

[51] Christoph Schneider, Markus Weinmann, and Jan Vom Brocke. Digital nudging: guiding online user choices through interface design. *Communications of the ACM*, 61(7):67–73, 2018.

[52] Kai Shu, Amy Sliva, Suhang Wang, Jiliang Tang, and Huan Liu. Fake news detection on social media: A data mining perspective. *ACM SIGKDD explorations newsletter*, 19(1):22–36, 2017.

[53] Heloisa Sturm Wilkerson, Martin J Riedl, and Kelsey N Whipple. Affective affordances: Exploring facebook reactions as emotional responses to hyperpartisan political news. *Digital Journalism*, 9(8):1040–1061, 2021.

[54] Jaime Teevan, Christine Alvarado, Mark S Ackerman, and David R Karger.

The perfect search engine is not enough: a study of orienteering behavior in directed search. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 415–422, 2004.

[55] Richard H Thaler and Cass R Sunstein. *Nudge: Improving decisions about health, wealth, and happiness.* Penguin, 2009.

[56] Nava Tintarev and Judith Masthoff. Designing and evaluating explanations for recommender systems. In *Recommender systems handbook*, pages 479–510. Springer, 2010.

[57] Martin KJ Waiguny, Michelle R Nelson, and Ralf Terlutter. The relationship of persuasion knowledge, identification of commercial intent and persuasion outcomes in advergames—the role of media context and presence. *Journal of Consumer Policy*, 37(2):257–277, 2014.

[58] Rosemary Walker and Brenda Cude. In-store shopping strategies: Time and money costs in the supermarket. *Journal of Consumer Affairs*, 17(2):356–369, 1983.

[59] Geoffrey I Webb, Michael J Pazzani, and Daniel Billsus. Machine learning for user modeling. *User modeling and user-adapted interaction*, 11:19–29, 2001.

[60] Paul Weiser, Dominik Bucher, Francesca Cellina, and Vanessa De Luca. A taxonomy of motivational affordances for meaningful gamified and persuasive technologies. In *EnviroInfo and ICT for Sustainability 2015*, pages 271–280. Atlantis Press, 2015.

[61] Zhijing Wu, Mark Sanderson, B Barla Cambazoglu, W Bruce Croft, and Falk Scholer. Providing direct answers in search results: A study of user behavior. In *Proceedings of the 29th acm international conference on information & knowledge management*, pages 1635–1644, 2020.

[62] Chris AB Zajchowski, Keri A Schwab, and Daniel L Dustin. The experiencing self and the remembering self: Implications for leisure science. *Leisure Sciences*, 39(6):561–568, 2017.