

## Breaking the routine: spatial hypertext concepts for active decision making in recommender systems

Claus Atzenbeck, Eelco Herder & Daniel Roßner

**To cite this article:** Claus Atzenbeck, Eelco Herder & Daniel Roßner (2023) Breaking the routine: spatial hypertext concepts for active decision making in recommender systems, New Review of Hypermedia and Multimedia, 29:1, 1-35, DOI: [10.1080/13614568.2023.2170474](https://doi.org/10.1080/13614568.2023.2170474)

**To link to this article:** <https://doi.org/10.1080/13614568.2023.2170474>



Published online: 12 Feb 2023.



Submit your article to this journal [↗](#)



Article views: 116



View related articles [↗](#)



View Crossmark data [↗](#)



# Breaking the routine: spatial hypertext concepts for active decision making in recommender systems

Claus Atzenbeck <sup>a</sup>, Eelco Herder <sup>b</sup> and Daniel Roßner <sup>a</sup>

<sup>a</sup>Institute for Information Systems (IISYS), Hof University, Hof, Germany; <sup>b</sup>Department of Information and Computing Sciences, Utrecht University, Utrecht, The Netherlands

## ABSTRACT

Recommender Systems are omnipresent in our digital life. Most notably, various media platforms guide us in selecting videos, but recommender systems are also used for more serious goals, such as news selection, political orientation and work decisions. As argued in this survey and position article, the paradigm of recommendation-based feeds has changed user behaviour from active decision making to rather passively following recommendations and accepting possibly suboptimal choices that are deemed “good enough”. We provide a historic overview of media selection, discuss assumptions and goals of recommender systems and identify their shortcomings, based on existing literature. Then, the perspective changes to hypertext as a paradigm for structuring information and active decision making. To illustrate the relevance and importance of active decision making, we present a use case in the field of TV or media selection and (as a proof of concept) carried over to another application domain: maintenance in industry. In the discussion section, we focus on categorising these actions on a spectrum of “system-1” (fast and automated) tasks and “system-2” (critical thinking) tasks. Further, we argue how users can profit from tools that combine active (spatial) structuring and categorising with automatic recommendations, for professional tasks as well as private, leisure activities.

## ARTICLE HISTORY

Received 17 June 2022  
Accepted 16 January 2023

## KEYWORDS

Recommender systems;  
hypertext; television; media;  
context; cognitive maps;  
structuring

## 1. Introduction

Recommender systems are designed to help their users to find items expected to be useful, by rating, filtering or ranking candidate items. For this purpose, many different strategies have been developed, with the following two main approaches (Ricci et al., 2015):

- *Collaborative filtering*, which assumes that users with similar preferences for items that they have seen, read, ordered or rated before, will also have similar preferences for items that they haven’t encountered yet.

**CONTACT** Claus Atzenbeck  [claus.atzenbeck@iisys.de](mailto:claus.atzenbeck@iisys.de)  Institute for Information Systems (IISYS), Hof University, Alfons-Goppel-Platz 1, Hof 95028, Germany

© 2023 Informa UK Limited, trading as Taylor & Francis Group

- *Content-based filtering* matches an item's properties or attributes or natural-language description with a user profile that represents preferences.

Further, collaborative filtering and content-based filtering can be combined into hybrid systems; recommendations can be based on the full history of the user or based on user actions in the current session; other criteria such as the user's location, time of day, day of week, inferred mood or calendar can be taken into account as well.

Arguably, all of these approaches aim, in one way or another, to reinforce users' past and current behaviour by offering them *more of the same* in a feed or stream of item recommendations. This has been recognised and partially addressed in the field of heterogeneous recommender systems (Bellogín et al., 2013), in many cases by optimising a feed of recommendations for diversity, entropy or overlap, for example, by taking items from the “long tail” or by adding in items that are highly ranked from a different perspective or angle (such as popularity, recency or regionality).

Research on *interactive recommender systems*, which uses visualisations, explanations, interactive dialogues and other techniques for soliciting user feedback (He et al., 2016) has demonstrated the importance of taking the user in the loop. However, as will be discussed and demonstrated in the remainder of this article, interactive recommender systems suffer from at least three main shortcomings:

- Users typically only can provide feedback on the items that have been recommended to them, not on other (important, surprising, interesting) items that did *not* make it to the top-ranked results.
- Because of the typical item-to-item approach of (automatic) recommender systems, user behaviour is gradually transformed to passively “consuming feeds” instead of actively exploring a wide area of possible directions and discovering things that are really new.
- Recommender systems are focussed on search activities, not persistently preserving information. They allow to “consume” media or items, but not to structure and store them in a persistent context. This has to be done with or on additional applications or materials, such as writing notes in word processing applications or on paper.

In this article, we provide an historical overview, and analyse and reflect on developments in recommender systems and the way they have influenced user behaviour and investigate how interactive hypertext techniques can bring back user experiences that largely have gone lost in the past couple of decades.

In Section 2, we introduce and illustrate these inherent limitations of recommender systems by analysing and discussing the history of television watching and programme choices. This activity has been shaped by the gradually

expanding availability of television channels as well as the advent of the remote control or the video recorder. However, the largest change seems to have happened with the introduction of recommendation-driven streaming services, such as Netflix, which has led to the phenomenon “binge-watching” and passive choices based on what happens to be in the *feed*, the list of recommendations. We argue that, even though this activity is not inherently wrong, this inevitably will lead to reduced offerings and increased boredom.

Even though television consumption may not be considered a vital activity, this change in user attitude (as a response to technological changes) takes place in many different areas, including following the news, interacting with friends (on social media), getting vital information on how to deal with ongoing pandemics or preparing for upcoming elections: increasingly, as a result of the feed paradigm and the resulting passive interaction style, we are stuck in “filter bubbles” (Pariser, 2011). These filter bubbles may not be as worrisome in terms of consequences as initially thought (Zuiderveen Borgesius et al., 2016), but still the activity of passively consuming feeds of news, updates or recommendations inherently limits our choice.

As a alternative approach, or rather design philosophy, we outline the concept of “hypertext” in Section 3. Hypertext, which allows users to create their own paths from a wide selection of possible choices, is inherently interactive and various interactive approaches have been developed and investigated. In this section, we give an historical overview of the hypertext paradigm and hypertext systems. We then discuss the role of visualisation and context for creating persistence as well as volatile, emerging structures, and explain in Section 3.3 how spatial hypertext allows users and system to work together for augmenting cognitive maps.

After providing the foundations for hypertext paradigms in the context of recommender systems, we provide two scenarios in Section 4. The first is an example for the use of such systems in the television domain. We argue that this approach is also applicable to other application domains. We illustrate this in Section 4.2 in which we describe a use case of a hypertextual recommender system used for maintenance tasks in the industry.

In Section 5, we then abstract from the given scenarios and discuss the findings. Following Kahneman (2002) and Bengio et al. (2021), we distinguish between “system-1” (i.e. fast and automated) tasks and “system-2” (i.e. critical thinking) tasks and propose a corresponding classification scheme in which we put node-link hypertext systems, cognitive maps, common recommender systems and our proposed solution in relation. Finally, Section 6 wraps up this article with some concluding remarks and future perspectives.

In summary, the contributions of this survey and position article are:

- We provide a survey of approaches and assumptions in recommender and hypertext systems and show how human intelligence and machine

automation differ from another—and influence each other—in terms of decision making.

- We discuss a set of mechanisms that combine approaches from the intersection of hypertext and recommender systems for more holistic decision making, particularly by involving users in the process and stimulating the active use of cognitive maps.
- We demonstrate the practical use and desirability of this approach in two different application domains: the traditional leisure-oriented domain of movie and series recommendation and the more professional, purposeful domain of support for maintenance in industry.
- We interpret the interaction between users and stream-based recommender systems in terms of reactive “system-1 thinking”, which ultimately forms the basis of most machine learning approaches, and how this is complemented by more rational “system-2 thinking”; we discuss how various systems can be positioned on a scale between these two extremes.

## **2. The impact of recommender system interfaces on user behaviour and behavioural choices**

In this section, we reflect on how the advance and user acceptance of recommender systems have impacted our selection and consumption behaviour and how (in reverse) this changed user behaviour reinforces the item-to-item recommendation paradigm that may lead to overly repetitive offerings and too few unexpected, useful surprises. Following Konstan and Terveen (2021), we take video recommendation as a running example, as this is the quintessential application domain representing the commercial boom of recommender systems. As will be discussed later, many observations translate to many other application domains – including recommendations for leisure and more serious and professional domains.

We will do so by briefly sketching “traditional” television consumption and how this first has been slightly influenced first by the introduction of the VCR and remote control, but then dramatically changed by the advent of recommender-based streaming services such as Netflix, which have been central to many studies and evaluations on recommender systems since the early 2000s.

In Section 2.2, we provide a brief literature overview on assumptions and implicit goals of recommender systems, as discussed in surveys and articles from the past two decades. We connect these insights to the observations from the previous section on television consumption.

We broaden our initial focus on television or media consumption in Section 2.3, where we reflect on literature and reporting on filter bubbles; it can be observed that item-to-item streams indeed seem to condition users to passively consume what is being offered, with serious implications on their world views or decision making.

Finally, we conclude the section with some historical insights and future perspectives, which will be elaborated upon in the remainder of this article.

### ***2.1. The interaction between technology changes and television watching behaviour in the past decades***

Before investigating recommender systems for items such as movies or series (Netflix-style), it is useful to first investigate how individuals or families used to determine which programmes to watch on regular television. Television broadcasts were introduced in the 1920s and television watching became rapidly commonplace between the 1950s and 1970s<sup>1</sup> and part of almost every household's favourite pastimes.

From the early beginnings, the advent of televisions in our living rooms has led to critical voices, arguing (among others) that the activity changed the population into passive “couch potatoes”, who are more interested in looks and presentation than in actual in-depth content and discussion.<sup>2</sup> Leaving that discussion aside, we focus on the activity of television watching itself, which often invokes social interaction, such as discussions, side remarks and comments (Geerts et al., 2008). This has been researched extensively in media studies.

In Lull (1982) it was investigated who determines what to watch in (very) traditional families with children. In a nutshell, the results indicated that particularly the father of the house decided this autonomously, followed by the children. The mother seemed to be happy with whatever led to consensus. Another interpretation of this study is that in the 1980s, television watching was an activity done together, and that programmes to watch were consciously decided upon, often already beforehand, based on the offerings listed in the TV guide.

In Adams (2000), focus groups reflected on their television programme choices. Often, the major channels were watched most often and there was a tendency to prefer older series to (then) current programmes. The common impression that television watching is a “passive” habit (in other words, people watch whatever comes) seems to be contradicted by the observation that participants had expectations regarding “what’s on”: if a programme was not satisfying, people would surf to other channels, if available. The latter observation appears to be confirmed by the success of the VCR, arguably the first home device for watching programmes at whatever time is convenient, which was recognised as a device to reconcile television watching with daily routines and to have a second chance to watch programmes.

The above observation that television watching was considered to be a mix of active and passive watching is confirmed by Lee and Lee (1995), who concluded that their “results suggest that interactive TV may not have a significant impact

---

<sup>1</sup>See also <https://www.cs.cornell.edu/pjs54/Teaching/AutomaticLifestyle-S02/Projects/Vlku/history.html>.

<sup>2</sup>See also <https://www.cs.cornell.edu/pjs54/Teaching/AutomaticLifestyle-S02/Projects/Vlku/social.html>.

on viewing behaviour. Findings show that viewers enjoy both low- and high-involvement viewing, and that they watch TV for relaxation and mood lift, which do not require interaction with the set”.

With the advent of online streaming services, such as Netflix or Amazon Prime, watching behaviour dramatically changed from weekly “watching appointments” with one or more series to *binge-watching*, that is watching two episodes or more of the same series in a row.<sup>3</sup> Jenner (2016) notes that Netflix “signals a significant shift in a new media landscape”. One change is that television consumption has become multi-platform, allowing people to watch series not only on a regular television, but also on their laptop or smartphone. Further, Netflix focuses solely on movies and series, abandoning “traditional” – arguably more purposeful – television genres such as news, game shows or sports. Particularly series are a genre that allows for binge-watching, an activity that is lucrative for streaming providers in terms of returning users.

Pittman and Sheehan (2015) surveyed 262 television binge-watchers to find out which factors contributed to this activity. Among the mentioned factors were relaxation, engagement and hedonism. Apart from programme quality, social aspects were also considered important.

The direct relation between recommender systems and binge-watching behaviour is investigated in Carretta (2021), who developed and compared optimal strategies for identifying series and users who are prone to binge-watching. In the following section, we will take a more theoretic view on the assumptions and implicit goals of recommender systems, such as used by streaming media providers.

In this article, we will not argue about the desirability of binge-watching behaviour and the Netflix model. However, we have already observed that binge-watching is associated with watching series, ignoring other (arguably more educational) content, such as news, documentaries or sports. In recent news articles, it is observed that users seem to turn away from binge-watching, which led stream providers to experimenting with “drip-feeding” episodes.<sup>4</sup> In a recent Dutch newspaper article, it was argued that users increasingly started to prefer scheduled programmes, in order to “slow things down and to create opportunities for discussion”<sup>5</sup>

As will be argued in more detail in Section 2.3, users of streaming media have started to recognise the limitations of the paradigm of item-to-item recommendation as used by streaming services, which typically offer “safe” (and therefore arguably less interesting, novel or exciting) options that may be considered “good enough”, but are increasingly perceived to suffer from a lack of

---

<sup>3</sup>See also <https://en.wikipedia.org/wiki/Binge-watching>.

<sup>4</sup><https://www.dailymail.co.uk/sciencetech/article-7805513/TV-binge-watching-falls-fashion-streaming-services-drip-feed-episodes-viewers.html>.

<sup>5</sup><https://www.nrc.nl/nieuws/2021/04/01/bingen-doe-je-maar-aan-het-eind-van-het-seizoen-a4038156>.

purpose. For instance, current Spotify users increasingly notice that the stream-based service “trained” them to “[use] music, rather than having it be its own experience”<sup>6</sup> We will also discuss how the above observations translate to arguably more purposeful application domains for recommender systems.

## 2.2. Assumptions and implicit goals of recommender systems

Building upon the observations from the previous section, we will now review assumptions and implicit goals of recommender systems in general, taking insights from a wide range of literature.

All broad categories of recommender systems have different assumptions, but what they have in common is that they are designed to recommend *items*, either on an individual basis or sequentially. As argued by He et al. (2016), if more than one item is to be recommended, *diversity* and *controllability* are noticed to become important.

As discussed in the introduction, collaborative filtering systems recognise commonalities between users on the basis of their ratings and recommend those items that similar users have “consumed” or “liked”. In Burke (2002), it is acknowledged that an ideal recommender would not suggest an item that a “user already owns or a movie she has already seen”. However, as most user activity is centred around a small number of disproportionately popular items, this means that most recommendations users will receive are not only *very similar to items that they are already familiar with*, but also items that are *popular in general* (Baeza-Yates, 2018).

As an alternative approach, content-based recommenders use domain-related features and, because of that, they do not suffer as much from this rich-get-richer effect. Indeed, Burke (2002) confirms that in collaborative filtering, it is very likely that only a small number of items receives sufficient attention, views and ratings in order to become recommended at all, but that in principle, using content-based recommender systems, all items are eligible to become a candidate.

These inherent limitations of traditional recommender systems are further confirmed and illustrated in Smith and Linden (2017), who reflect on two decades of recommender systems at Amazon. The article starts with the observation that “[recommender systems] simply share with you what other people have already discovered [...] and being able to explain why it recommended something”. Traditionally, Amazon recommendations took place in the domain of books and media, with collaborative filtering (which takes similarities in taste and preferences between users as a main assumption) as the most natural choice. However, even though this works well for *low cost* items

---

<sup>6</sup><https://www.theguardian.com/music/2022/sep/27/theres-endless-choice-but-youre-not-listening-fans-quitting-spotify-to-save-their-love-of-music>.



and media, for more expensive, non-media items, users behave radically differently: users actively search and compare before buying a new television and once they have bought one, they do not need a similar television anymore. Furthermore, consumables, such as toothpaste or olive oil, are bought again and again with relatively predictable intervals.

What all approaches discussed in the article have in common is that the recommender algorithm seems to be considered as a function, a one-trick pony, that is very good at learning correlations and co-occurrences of one kind or another, but fails to understand the deeper intentions that are behind (repetitive) choosing for an item. Moreover, the recommendations focus on *individual items* to be sold, not on a collection of items.

In the field of *interactive recommender systems*, it has been recognised that user input and feedback is essential for better tailoring recommendations to the actual, current user goals, rather than the estimated (general) preferences of the user. He et al. (2016) propose the use of interactive visualisation frameworks to provide the user insight into the system logic or justifications for the recommendations. Another common technique for soliciting user feedback is by explaining for what reasons an item has been recommended. Zhang and Chen (2018) recognise two main strands of explainable recommendations: those that explain the actual recommendation process and those that are generated post-hoc. The latter category arguably considers the recommender system just as much as a black box as the user, but still tries to make sense of it on behalf of the user.

In all cases, explainable recommendations aim to improve the credibility and acceptance of recommended items—not just in the context of media consumption, but also legal or healthcare recommendations. Furthermore, all explanation approaches appear to be regarded as something mechanic and logical and that all (inferred) user choices have a rational explanation. As we have seen in the previous section (and as we will discuss further), this is a very limited view on how users make decisions. As a further limitation, the explanations concern those few items that are recommended, not the very long tail of items that have *not* been recommended (Baeza-Yates, 2018).

As a final argument, most explanations seem to focus on individual items and not on the *collection as a whole*. In line with Smith and Linden (2017), it is useful to draw a comparison with regular purchases in a bricks-and-mortar supermarket. Suppose that a supermarket consumer buys a week's worth of food for a four-person family, the cart contents reflect (among others) the different likes and dislikes of each individual person, compromises made regarding shared dinners, allergies, diet preferences, current weather and seasonality and availability of items, weekend plans, obligations, decisions made in the spur of the moment and many other things. Arguably, it is not helpful to provide explanations for each individual item, but to ensure that the cart as a whole reflects the whole family's needs and plans for the upcoming week.

Similarly, sets of recommendations—such as a mix of television programmes, music tracks, or news articles—may, in many cases, be better explained and motivated as (the value of) a collection as a whole rather than on an item-by-item basis.

As a final argument on the view of typical recommender systems to be mainly item-based and focussed on learning likely correlations (not causations), we cite from a recent survey on deep learning based recommender systems (Zhang et al., 2019). In this survey, it is stated that “[d]eep learning is able to effectively capture nonlinear and nontrivial user/item relationships and enable the codification of more complex abstractions as data representations in the higher layers”. Further, the authors observed that “it is assumed that big, complex neural models are just fitting the data without any true understanding”.

Indeed, according to a recent Turing Lecture on deep learning for AI (Bengio et al., 2021), building upon theories by Kahneman (2002), deep learning systems (and therefore also deep-learning based recommender systems) are very successful at so-called *system-1 tasks* (i.e. object recognition or immediate natural language understanding, tasks that humans carry out fast and effectively, building upon routine and earlier experiences), but that *system-2 tasks* (i.e. “deliberate sequence[s] of steps which we attend to consciously”, a more effortful process of active, conscious decision making that we resort to in situations where quick and automated system-1 responses are deemed insufficient or undesirable) are still an “exciting area that is still in its infancy”.

In sum, recommender systems have their merits and purposes, but they also have inherent and conceptual limitations. First, recommendations aim to *reinforce* current behaviour and preferences. Arguably, diversification techniques aim to bring new (other) items to the user’s attention, but such approaches are typically either focussed on a small set of new or popular items or still not (radically) different from the user’s observed behaviour.

A second limitation is the focus on recommending individual items or sequences of items. This may not be a negative thing, as long as the intended use is to merely entertain or accommodate the user, for example with series to binge-watch. However, if users wish to bring *variety* into their activities or offerings, to *learn* something or *change behaviour*, they are mainly left to themselves. As illustrated in Section 2.1, we argue that the (system-2) processes needed for these purposes may actually be hindered by the interaction with recommender systems itself, as passive interaction forces users to think and decide stepwise, without much deliberation, and does not encourage more holistic thinking or planning.

### **2.3. Filter bubbles, echo chambers, bias and routine**

In the previous sections, we have discussed how media streaming services and the associated recommender systems, assumptions and paradigms have

considerably changed the media landscape and media consumption behaviour. This is arguably reinforced by the inherent behaviour and design of recommender algorithms and the widely spread item-to-item feeds of recommendations that popular streaming services such as Netflix or Amazon Prime offer.

A similar effect, in the more serious domain of freedom of expression, was observed by Hossein Derakhshan: after having spent six years in an Iranian prison, he “found the internet stripped of its power to change the world and instead serving up a stream of pointless social trivia”<sup>7</sup> His impression was that this change was not as much caused by the algorithmic nature of recommender systems as well as by changes in how we use them: (social media) feeds have trained their users that they only need to scroll further to find new content and, consequently, that items that do not reach the feed are probably not important or interesting enough.

A couple of years before Derakhshan’s release, in 2011, Eli Pariser coined the term *filter bubble* to describe the potential for online personalisation to effectively isolate people from a diversity of viewpoints or content (Pariser, 2011). Even though the detrimental effects of the filter bubble are observed to be mainly problematic in extreme cases (Zuiderveen Borgesius et al., 2016) and already polarised environments Chitra Musco (2020), the “filter bubble effect” is inherently associated with the assumptions and approaches that lie at the root of recommender systems (as discussed in the previous section) and that inherently aim to reinforce current (system-1) behaviour and preferences.

The combined observations of Derakhshan and Pariser suggest that the issues discussed in this section are the result of interaction effects between (recommender) systems and their users, both of which are guided by the responses and assumed intentions of the other party. A further complicating factor in this process is that human behaviour is largely automated and based on the choices that we are offered, combined with our existing routines—i.e. system-1 behaviour. To complicate matters even further, as discussed by Baeza-Yates in his much-cited article *Bias on the Web* (Baeza-Yates, 2018), the social media content that a typical user would see is created by a small minority of very active users (or rather user accounts); similarly, Amazon reviews and ratings, which form the basis for recommending items such as movies, are driven by a very small percentage of users. In addition, this activity is largely focussed on a very small set of highly popular items, leaving a long tail of items that remain unnoticed and, because of that, remain unpopular. These factors contribute to the issue that a potentially very large set of candidate items is reduced to a relatively small pool of candidate items to be recommended (Celma, 2010; Polatidis & Petridis, 2019).

---

<sup>7</sup><https://www.theguardian.com/technology/2015/dec/29/irans-blogfather-facebook-instagram-and-twitter-are-killing-the-web>.

Arguably, the user's natural tendency towards routines and safe choices—e.g. people spend most time on a small number of locations (Gonzalez et al., 2008; Herder et al., 2014) and tend to listen to a small set of music tracks from a small set of artists (Celma & Cano, 2008)—is *reinforced and not challenged* by this approach. The preference for known, safe choices that do not challenge our preferences or our beliefs, is commonly called the *echo chamber effect*. Similar to the filter bubble effect, it is believed that the echo chamber effect is limited for those people who are exposed to and choose to interact with a high-choice (media) environment (Dubois & Blank, 2018). As discussed in Section 2.1, in cases where users limit themselves to a reduced, limited “media diet”, this may be unsatisfying, though not very concerning in the context of series or movies. However, as argued by Pariser (2011), it is an issue of concern when it comes to news or political opinions.

It seems that the narrowing effects that are inherent to the combination of recommender systems and activity bias are not limited to individual users, but also have an effect on the choices and priorities of item *producers*, with direct, potentially narrowing effects on item production and offering. For instance, producing movies or series costs time and, above all, money that needs to be earned back by attracting a sufficiently large audience.

The advent of streaming media paradigm seems to imply that it has become more attractive and safer to produce mainstream movies and series, which appeal to a large, mainstream audience. Indeed, Amazon is reported to “pivot away from indie films toward mainstream movies”.<sup>8</sup> Similarly, Disney+ is reported to be moving towards mainstream and *consumerism*, which may offer “a good way of life”, but, as argued in the Guardian,<sup>9</sup> leads a strategy of playing safe and endlessly replicating already proven concepts.

*Mainstream* implies in most cases content that appeals to the average, white, western (American) person—most of the content on Dutch Amazon Prime is observed to be American, with European and Dutch series marked as special small sections.<sup>10</sup> Similarly, it is expected that movies and series are also targeted towards the tastes and norms of the majority group, confirming their views and not sufficiently representing minorities and minority views (Abdollahpouri, 2020; Abdollahpouri et al., 2020). These minorities would include religious and ethnic minorities, but also people who identify as being part of the LHBTQ+ community (Howard, 2021).

To summarise, in the past decade, the nature of recommender systems and the feed-based paradigm of their interfaces have been observed to have a narrowing effect on how individual users perceive the world or a particular domain, while simultaneously encouraging known, safe choices for item

<sup>8</sup><https://www.bizjournals.com/losangeles/news/2018/01/18/amazon-pivoting-away-from-indie-films.html>.

<sup>9</sup><https://www.theguardian.com/film/2021/sep/15/disney-plus-blockbuster-movie>.

<sup>10</sup>See also <https://www.engadget.com/2018-09-04-netflix-amazon-european-content-quota-eu-law.html>.

selection. Or, to phrase it differently – following the terminology of Kahneman (2002) and Bengio et al. (2021)—recommender systems facilitate and encourage routine (system-1) behaviour and choices, and therewith hinder users in engaging with more active, conscious (system-2) decision making. As a result, this is observed to lead to the risk of users being locked in a combination of algorithmic filter bubbles and self-created echo chambers; furthermore, this effect does not only impact individual users, but society as well, among others by an over-representation of mainstream choices and an underrepresentation of minority choices and viewpoints.

#### **2.4. Designers as intermediaries between user and system**

Second-order cybernetics is defined in Wikipedia as “the recursive application of cybernetics to itself and the reflexive practice of cybernetics according to such a critique”.<sup>11</sup> In other words, it involves observing and reflecting upon a self-organising system, done by stakeholders *who themselves are part of the system*. As we have seen and discussed in the preceding part of the section, the ecosystem of a recommender system does not only consist of users and the recommender algorithm and interface, but also of platform owners, platform designers, content providers, advertisers and other interested parties who all respond to developments in the platform and are able to change the dynamics in the ecosystem from within.

As argued by Krippendorff (2019), an interface that works as expected *affords* the construction (or mental model) that a user has of it. These expectations are shaped by the design of the interface, including visual elements and the selection (and order) of options offered to the user. System (or rather interface) designers play a crucial role in shaping these expectations and arguably condition users to a certain type of (designed or planned) user behaviour, such as binge-watching. In continuous development, system designers iteratively observe to what extent user behaviour corresponds to what they intended and then implement (large or small) changes to optimise any inefficiencies or repair undesirable effects. In the long-term, the effect of such small iterations may be large.

Users are not designers and, therefore, it is the duty of designers to design, foresee, observe, shape or repair such interactions (Nielsen, 1993). However, when given the right affordance, it is up to users to adopt them. Particularly, it seems that the paradigm of streams and never-ending feeds gradually has replaced our natural tendency to first obtain an overview of a domain, organise alternative choices (in other words, create a cognitive map) and only then decide upon actions to take or choices to make. This in ironic contrast to the

---

<sup>11</sup><https://en.wikipedia.org/w/index.php?title=Second-ordercybernetics&oldid=1059745059> (version of Dec 11, 2021).

original paradigm of the Web, which is rooted in the concept of hypertext, which was (and is) designed to support and stimulate associative thinking by focussing on the *connections* between items and turning them into explicit links.

In the upcoming sections, we will investigate how traditional hypertext concepts, hypertext interfaces and interaction of hypertext can guide us towards novel, more holistic and more user-centred approaches towards the paradigm of automated feeds of item-based recommendations.

### 3. Hypertext as a tool for supporting thinking and decision making

In the previous section, we have argued that the current state of recommender systems—which form the backbone of many (interactive) Web platforms, including online stores, entertainment platforms and social media—reinforces rather passive stream consumption behaviour (i.e. system-1 behaviour (Kahneman, 2002)). This is ironic, because the concept of hypertext (which forms the theoretical foundation of the World Wide Web) was actually intended to support and mimic human thinking and (system-2) decision making.

In the next subsection, we reflect on the rich history of hypertext thinking and hypertext systems and discuss the variety of assumptions and the tasks that academic and commercial systems were designed for. In Section 3.2, we discuss how different types of links, link visualisations, overviews, previews and visual connections can be used to create persistent structures. In Section 3.3, we argue how these persistent structures support and encourage users to think in a more holistic manner and create stronger cognitive maps in contrast to the volatile nature of Web browsing or the use of common recommender systems. Finally, we introduce *Mother*, a hypertext system that supports these features.

#### 3.1. A historical view on hypertext

With no doubt, the world’s largest distributed information system is the World Wide Web (Berners-Lee, 2000), which has been introduced by Berners-Lee in the late 1980s and gained popularity in the 1990s. The Web’s notion of *links* boils down to URIs (Berners-Lee et al., 2005), embedded in HTML pages. Due to the Web’s dominance in our daily lives, this notion is widely accepted by Web users worldwide. However, as argued in the previous section, the current stream and feed paradigm of large platforms (including streaming video, online shopping and social media) invites passive “consumption” of items that happen to be recommended, rather than active exploration and active choices.

Even though the Web may be called the largest hypertext system, it is neither the only nor the first one. Already in the 1940s, Bush described the *Memex*

device (Bush, 1945), which would let users create “trails” between documents. Such associations could be traversed mechanically at any time later or be even shared with others. In the vision of the Memex, the users play the most important role, as they define the trails between documents; the machine just helps in traversing those. Bush considered the Memex to serve as a “mechanized private file and library” and as “an enlarged intimate supplement to [our] memory”, to support and amplify (scientific) thinking.

With the raise of computers in the 1960s, hypertext was taken to another level: Nelson coined the term “‘hypertext’ to mean a body of written or pictorial material that is interconnected in such a complex way that it could not conveniently be presented or represented on paper” (Nelson, 1965). Nelson started working on the hypertext system *Xanadu* (Nelson, 1993) and together with van Dam on *HES* and its successor *FRESS* (Barnet, 2010).

A central concept of *Xanadu* was that the system enabled users to compare different versions of a document and to facilitate nonsequential writing and reading with rich, typed and bidirectional links and visualisations—activities that were more in line with human thinking and creative work rather than hierarchical directories and conventional files. In a 1999 article, Nelson wrote: “The World Wide Web was not what we were working toward, it was what we were trying to *prevent*” (Nelson, 1999).

Around the same time, Engelbart developed the *oN-Line System (NLS)*, which is based on his augmentation framework (Engelbart, 1962). The principle idea of this concept is that tools are primarily used for augmenting human capabilities, not for automation. This also includes “augmenting human intellect” (Barnet, 2018; Conklin, 1987; Conklin et al., 2001; van Dam, 1988).

With an increasing availability of personal computers in the 1980s, hypertext became more popular, too. It can be considered as the “high time” of hypertext systems and can be witnessed by an increased number of academic or commercial hypertext systems, including *KMS* (Akscyn et al., 1988), *Hyperties* (Shneiderman, 1987), *NoteCards* (Halasz et al., 1987), *Intermedia* (Meyrowitz, 1986), *Guide* (Brown, 1987) and *HyperCard* (Smith & Weiss, 1988). *HyperCard* was a popular application, available on the Apple Macintosh. In essence, *HyperCard* was based on the metaphor of virtual cards that users could drag-and-drop and put on stacks. Applications included narrative games (Davidson, 2008) and educational multimedia.<sup>12</sup>

Furthermore, this was the decade in which the *ACM Hypertext Conference* has been founded (Smith & Halasz, 1987). The first conference took place in 1987 and since then it has been (and still is) organised annually. Around the same time, the ACM dedicated one issue of their *Communications of the ACM* journal to hypertext (Smith & Bernhardt, 1988).

---

<sup>12</sup><https://en.wikipedia.org/wiki/HyperCard>.



In the 1990s, the Web started, a huge success that led to a de facto monoculture of hypertext systems in academia and industry, even though a small group of academics dedicated and still dedicate their research to “traditional” hypertext topics (e.g. Atzenbeck et al., 2018; Bernstein, 2010; Nürnberg et al., 2004; Tzagarakis et al., 2006). In the upcoming subsection, we will discuss how several traditional hypertext concepts and paradigms can complement the current largely recommender-driven and feed-based World Wide Web.

### 3.2. Visualisation, context, persistence and volatile structures

Visualisation has always played an important role, as hypertext is very much related to users’ interactions and, thus, requires user interfaces that enable users to actively *explore* hypertextual structures.

In hypertext, visual metaphors and connections are commonly used to complement or enhance regular links; sometimes they are even used in replacement of traditional node–link hypertext (such as linked words, as is common on the Web). Typical metaphors are the use of fixed-sized cards instead of scrollable content (Halasz, 2001), landmarks and footprints (or breadcrumbs) (Nielsen, 1990), richly described and/or coloured link anchors or link previews (Weinreich, 2012). Some hypertext applications, such as NoteCards or Xanadu, provide map views for presenting birds’ eye views on the hypertext network (Halasz et al., 1987).

As explained in the previous subsection, the core of hypertext navigation and hypertext thinking is *associative*, following a particular (individual) line of (system-2) thinking. Hierarchical structures (such as domains and menus used on the Web) are considered as a complementary—not as a primary—way to organise nodes or pages. For example, NoteCards supports so-called *file boxes*, implemented using links (Halasz, 1987), whereas the *Dexter Hypertext Reference Model* suggests *composites* for representing hierarchies (Halasz & Schwartz, 1994).

By contrast, navigation on the Web is largely hierarchical, visiting different websites (separated in different domains) and following menu structures or, as described in detail in the previous section, predefined, largely chronologically ordered feeds (Danielson, 2002). This fragmented nature of Web navigation makes the experience *volatile*, breaking a session into different perceived parts.

The introduction of tab-based browsing has made Web browsing even more volatile, as the different paths that users follow or create simultaneously or iteratively, are stored in distributed piles, each belonging to the different tabs that users happen to have opened (Weinreich et al., 2008). Once a tab is closed, the pile of navigation history associated with it is gone, or at least not readily accessible anymore.

Arguably, at least partially due to the lack of persistent structures, overviews and consistent history mechanisms offered by the Web and Web browsers, *Web*



*search* has become a prominent way of looking for information or even for returning to known places, such as frequently visited Web stores or social media sites (Broder, 2002). Actual Web (or rather site) navigation involves rather short paths in between related or unrelated queries. These paths are known to be important for placing the information, interactions or offerings found into *context* (Teevan et al., 2004), but most of them are very short and unconnected.

In sum, there are many different mechanisms for browsing, navigating and searching the Web and, as a result, many different places in which these actions are stored. For instance, the activity of creating and managing bookmarks provides structure, but this structure is not connected to search activities or recommendations.

Naturally, there are *coping mechanisms*, such as memorising URIs or keeping notes in a text document, but, ironically, this leads to even more fragmentation or volatility due to an application or media gap between browsing information (e.g. Web pages using a Web browser), search or recommendations (e.g. searching the Web) and organising (e.g. taking notes or organising URIs using a text processor), for which persistence or context creation mainly happens for the latter. This media gap can be illustrated by the search engine history, which is not aware of any (mental or actual) notes taken by the user or even user activity in the time elapsed between two subsequent queries.

The lack of easily adding or structuring new information by the user is common to a wide range of today's recommender systems or websites. Looking at previous hypertext systems, we realise that receiving information and adding new, linked nodes used to be common in many previous hypertext systems. Those systems supported readers and authors alike. For example, KMS had "no mode boundary between navigation and editing operations" (Akscyn et al., 1988). As such, readers could switch between reading nodes (called "frames" in KMS) and editing them. Furthermore, new linked frames could be instantly created by clicking on an yet unlinked object.

This support of immediately extending existing structures and modifying presented nodes directly supports the task of taking notes and, thus, differs from today's "read only" recommender systems or websites. As illustrated by the examples above, many earlier hypertext systems supported users in representing their thoughts with nodes and connecting links while reading relevant information. This shows persistence of user-generated information and contexts over the volatileness of presented data which we experience in many of today's systems.

Some more recent research aims to fix this issue by providing additional tools for commonly used systems. For example, a link plugin for some common Web browsers (Roßner et al., 2020), demonstrated by Roßner and Atzenbeck (2021), which allows Web users to add links or annotations to

any webpage. This enables Web users to author an additional hypertextual structure or comment on webpages, which otherwise would be not possible.

A particularly relevant type of structuring in the hypertext community is *spatial hypertext* (Shipman et al., 2001), which represents associations between informational units (nodes) implicitly by reflecting these associations by their absolute and relative positioning on a (hyper)space, the way they respond (or do not respond) to interaction with other nodes and/or similarities in their visual appearance (e.g. colour, shape and size).

Interaction paradigms for spatial hypertext are based on direct manipulation. As such, the recipient of such structures can also modify them. They allow for creating and changing contexts, as well as reading information and interpreting structures as part of the same iterative actions. This puts spatial hypertext systems in line with the above mentioned node-link systems that provide reading, navigation and authoring mechanisms for users alike.

Whereas linear feeds—as is commonplace on the Web and social media in particular—provide only one order for processing emerging information (such as news or ongoing conversations), spatial hypertext provides means for working with or manipulating such information. Due to its graphical appearance, spatial hypertext reduces the required cognitive overhead with respect to creation, parsing or communication of such structures (Shipman et al., 2002). In other words, spatial hypertext concepts aim to support and encourage users to actively engage with content, this in contrast to the feed-based paradigms that are paramount in recommender systems and social media, which—as argued extensively in Section 2—invite users to remain in the comfort of their routine (system-1) behaviour.

Only a few spatial hypertext applications include intelligent components, so-called *spatial parsers*, that turn the implicit structure, as created or manipulated by the user, into explicit computer knowledge. This approach opens the door for more holistic and persistent recommendations that are based on a user-created context structure, rather than fragmented item-to-item recommendations or automatic feeds (Roßner & Atzenbeck, 2018).

In the next section, we will argue that such spatial hypertext systems are well suited for supporting the creation of cognitive maps. In combination with intelligent components, they can act as a joint (and jointly created and manipulated) medium for the user's and the machine's knowledge.

### 3.3. Augmenting cognitive maps

The observation that links may serve different functions, may be represented in different ways, and may indicate different (types of) relations is acknowledged and addressed by the Semantic Web and the Linked Data movement, who focus on links as *predicates* between “subjects” and “objects” in the well-known RDF

format. Initiatives such as Schema.org<sup>13</sup> aim to create common vocabularies, centred around the naming and use of predicates, to be used by knowledge bases, including commercial knowledge bases from, among others, Google and Microsoft.

Given the importance of *typed* links in traditional hypertext as well as in current (semantic) knowledge bases, and the associative manner in which human think, it is surprising that links on the Web are largely untyped, with no (visual) indication what exact purpose they serve. Arguably, this interface decision has resulted in most links on the Web being site-internal and structural, following a menu hierarchy (Danielson, 2002). Furthermore, as discussed in the previous section, the relatively poor representation (and therefore poor use) of Web links has even diminished by a gradual paradigm change from “*actively* surfing the Web” to “*passively* following linear feeds” of recommended items.

A premise of hypertext is that it supports (human) associative thinking and decision making, an activity that can largely be considered as “system-2” thinking (see Section 2.2), whereas recommender systems and linear feeds mainly support and amplify the more reflexive, non-reflective “system-1” decision-making.

As discussed in the previous section, associations may also be expressed *spatially*, following a paper-on-desk metaphor. As with paper notes on a table, virtual information snippets get organised by the user on a (mostly 2D) space. The associations between such informational units are expressed implicitly by their arrangement or visual appearance, such as colour, size or shape.

Advantages of spatial hypertext include ease of creating *emerging structures*: by moving objects, the implicit structure adapts without further actions from the user. This is particularly beneficial for volatile structures that are in progress and change frequently. Ambiguous interpretations of structures may lead to serendipity effects, such that users get triggered toward some information or solutions which they would not have reached otherwise.

The volatility and emerging nature of this type of structure with its low requirements for users, for creation or modification tasks, makes spatial hypertext a good candidate for representing *cognitive maps*. Cognitive maps are mental representations of a physical environment, such as a city with its landmarks and street signs, or a virtual environment, such as the Web or a hypertext with its documents and the different types of relations (links) between them (Dillon et al., 1990).

More specifically, hypertext was (and is) designed to be a means for authors to collect and structure materials to reflect their own cognitive model, in anticipation of readers’ possible interests and ability to understand the relations

---

<sup>13</sup><https://schema.org/>.

created by the author (Marshall & Shipman, 1995). Spatial hypertext helps users in this process by graphically portray the structure.

Conversely, the *Semantic Web* and the Web of *Linked Data* were envisaged to be able to *reason* about the entities and relations on the Web, making use of ontologies, formal semantics, inference and reasoning (Hitzler & Van Harmelen, 2010). Such reasoning and ontologies were specifically aimed at supporting human reasoning and learning, which are typical system-2 tasks (Henze et al., 2004).

However, whereas the Semantic Web originally was designed to be built on formal logic and reasoning, arguably due to scalability issues, most reasoning on the Semantic, the Linked Web is currently performed by deep learning systems, exploiting knowledge graph embeddings and arguably surpassing human capabilities (Hitzler et al., 2020). Still, as discussed in Section 2.2, such systems are still inherently optimised for supporting and encouraging reflexive system-1 behaviour.

As will be explained in Section 3.4, spatial hypertext interfaces allow users to *augment* their (conscious) decisions and preferences with insight in the wealth of available options, and to reflect on their decisions. Rather than recommending items based on the users' past behaviour (e.g. which pages they visited or which products they bought previously), spatial representations of such recommendations complement the user's context in a direct way. Rather than automating processes (e.g. sequences of videos to be watched), these recommendations are designed to *add to the user's context* persistently.

Combining recommender functionality with the ability of context creation supports the users' holistic thinking, by allowing persistent user generated context. This context can be analysed by a recommender system, which then augments the space (i.e. the user's cognitive map) with additional information. The user may choose to pick relevant suggested items and add them to his/her own context. This forms an iterative process *in which user and machine work together* on creating a context that leads to a desired decision or a solution to a problem.

Note that this approach of generating recommendations based on a user's explicit and active actions, choices and responses differs fundamentally from recommendations based on observed (routine) behaviour. For instance, Meintanis and Shipman (2010) developed a music management environment that allows users to interactively generate playlists. The evaluation confirmed the benefits of automatic suggestions in the process, but also acknowledged that the users considered the user interface as sub-optimal and unfinished. Similarly, Park and Shipman (2014) created a personal digital library in which multiple data visualisations and annotations are placed into spatial arrangements based on the current task. Again, the results confirmed that the mixed-initiative approach was appreciated, but also required quite some effort and dedication from the users.

From these and other earlier studies, it becomes apparent that the adoption of spatial hypertext metaphors in the current Web does not only require new assumptions, recommender algorithms and interfaces, but also requires users to (re)learn to make active choices, in addition to passive choices based on mere availability. Several research lines have explored the benefits of combining recommender algorithms with spatial hypertext concepts.

In terms of media consumption, as discussed in detail in Section 2.1, this would involve stimulating users to actively search for, select, save or organise lists of programmes or series that they would like to purposely see at a later point time. This may be, for example, for pure enjoyment rather than falling back to the common, convenient and comfortable but limited pattern of binge-watching.

In the upcoming subsection, we will introduce a spatial hypertext system, *Mother*, and elaborate the concepts that we believe to contribute to this change in attitude and expectations.

### **3.4. *Mother*—combining spatial hypertext with recommender functionality**

Our system *Mother* is a so-called *component-based open hypermedia system* (CB-OHS), which allows the integration of multiple structure types, including link services, metadata or spatial hypertext. The latter supports creation and augmentation of users' cognitive maps with the goal to stimulate active reasoning and decision making, augmented by automatic recommendations.

*Mother*'s clients present the graphical interface to the users: items are represented by nodes, presented on a 2D canvas, and can be, removed or modified by the user. It follows a simple paper-on-desk metaphor. Like with paper notes on a physical desk, digital notes are arranged on the 2D space in order to form or destroy associations and build up structures. Furthermore, notes can be modified in their visual appearance, for example, size and (depending on the client) potentially also colour or shape.

A more detailed, technical, description of the overall system is provided in Atzenbeck et al. (2018). In this article we will focus primarily on the spatial structure service. With its support of emerging structures and low cognitive load it is most appropriate for creating cognitive maps. In the remainder of this section, we explain the concepts and we will demonstrate its practical use in two scenarios in the upcoming Section 4.

The structure represented as spatial hypertext gets revealed during an *interpretation process*, which aims to create *meaningful collections* of items that may be of interest to the user. Spatial distances in relation to the nodes' sizes imply strengths of associations between informational units, or items. Furthermore, nodes of similar colour, size, shape or orientation may be interpreted as related. The human perception of such a spatial structure is volatile,

emerging and very much depends on the users' background, interests, perspective and prior knowledge.

User navigation on the Web and the navigation support in the form of search results, link anchors or recommendations largely focus on the *content* of the nodes (or pages) rather than the *implied structure*. By contrast, Mother employs various specialised parsers in order to create a meaningful spatial structure (Schedel & Atzenbeck, 2016). The outcome of any parser is a weighted undirected graph, representing the strength of association between nodes. Users may switch between different parsers in order to choose between different perspectives.

The *spatial parser* analyses spatial arrangements of nodes. Primarily it identifies lists of objects and combines those as higher-order lists. Similarities of nodes' visual cues are computed by the *visual parser*. For example, same shapes or similar colours would cause strong associations. Finally, the *temporal parser*, which was introduced by Schedel (2016), analyses the sequence of user interaction with nodes. It assumes that the sequence of editing steps implies associations between nodes. Furthermore, there is an experimental implementation of a *content parser* which computes similarities of node content, as this is an aspect that Web users have learned to expect to be covered.

The various parsers produce weighted graphs independently. Those are different views of the same space. The aim is to reach an interpretation of the spatial structure that is close to what users would expect, in other words, a structure that fits their cognitive maps sufficiently. As the various parsers have different strengths, the combination of them reaches a better result compared to the individual ones (Schedel, 2016).

As a result, Mother reaches an internal representation of spatially represented user structures. The spatial hypertext becomes a medium between the user and the system for reaching a *common understanding* of the implicit associations and relations. This is the basis on which recommendations can be provided by the system and interacted with by the user.

Relevant related items are recommended to the users by presenting them as *suggestion nodes* (Roßner et al., 2019). Items are organised as nodes in a graph database, edges are weighted and undirected. These weights denote relevance and are usually pre-calculated. Suggestion nodes are identified by analysing the context, provided by parsers during the interpretation process. Sets of visually related nodes are identified and matching suggestion nodes are queried for each set. Suggestion nodes are volatile and may disappear or be rearranged if the user modifies the context. The maximum number of presented suggestion nodes is limited and depends on user provided preferences. Only the most relevant nodes are shown, which reduces the risk of information overload. Furthermore, the user may select suggestion nodes and make them part of the current, non-volatile user context. In this case, the node becomes a persistent part of the user space, that is, it will not be modified by the system, but will be considered by the various targets.

This opens an iterative process in which user and system work together on creating or augmenting an emerging context of selected items, suggested (i.e. recommended) items and their relations, which are visually reflected by the proximity of items towards other items. The system recommendations are based on knowledge coming from various sources or user expressed knowledge, interpreted by the parsers. Interaction with these suggestion nodes supports users in obtaining holistic views for problem solving or decision making over a long term and opens possibilities for information sharing and collaboration between multiple users, augmented by the machine's capability of processing huge amounts of data.

In the next section, we will show and discuss how automatic recommendations and interactive user feedback can complement one another, supported by a spatial hypertext system. For this purpose, we introduce two scenarios: one involving private, leisure activities and one involving a more purposeful, industrial setting. It will become apparent that similar principles hold in both scenarios.

## 4. Application scenarios

In Section 2, we discussed how recommender systems as well as other current forms of machine learning excel in recognising and reinforcing natural habits and spontaneous decision, so-called *system-1 behaviour*, but inherently do not support users very well in reflective decision making, that is, *system-2 behaviour*. In Section 3, we argued how the traditional concept of hypertext (more specifically spatial hypertext) could help users in this process. We also concluded that fully adopting this paradigm would not only require new systems and interfaces, but also depends on users' willingness and ability to (re)learn to make active choices.

To make the argumentation more tangible, we illustrate the interaction between users and a hypertext system in two scenarios. First, we return to the scenario of television consumption, a popular private activity that at first may not be essential, but (as discussed in Section 2.3) may have serious implications for news consumption patterns and representation of minorities in the media. In the second scenario, we explore the application of the same concepts in an industry setting.

### 4.1. Scenario 1: television

This scenario illustrates the discussion lead above in the domain video watching. This also includes users' selection of movies, series or documentaries. We compare current TV recommendation systems, such as YouTube, Netflix or Apple TV+, to hypertext supporting systems.

In an earlier publication, we described personas in different media consuming scenarios (Purucker et al., 2019). In this publication, we used four personas,



describing their attempts for finding, watching and storing relevant information. For example, one of the described personas was an elderly woman, who consults a (print) TV magazine for finding out the broadcasting times of certain TV shows about plants and gardening. She takes notes on paper during the broadcast, which she collects for further references. This shows a media gap (as described in Section 3.2) that prevents the media channel from making use of the user's notes and vice versa.

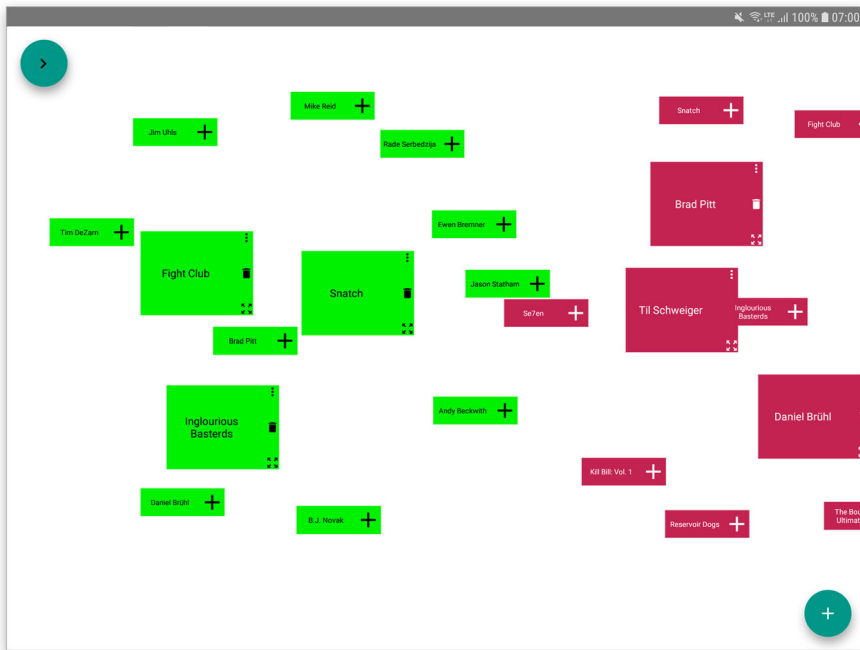
Such media gaps can be witnessed in the context of recommender systems, too. YouTube, for example, automatically recommends lists of videos to the users, either or not related to a current video, presumably to keep them on the platform as long as possible. Interactions with the video itself are possible (e.g. pause or selecting subtitles) as well as selecting videos from the list of recommended ones. Furthermore, a user can rate a video (via "thumb up/down") or share it via email, social media or messenger apps. Such interactions provide some evidence of a user's interest, appreciation or disliking, but this evidence is very indirect and presumably flows into the recommendation process without providing the user any opportunity for reflection or organisation.

As a result, *organising or categorising* videos based on a user's ideas is supported only in a very limited way. YouTube permits users to add references to videos in simple lists, but any further extended organisation can only be done with pen and paper or by using external applications. Similar functionalities can be found in other TV recommender systems, such as Netflix or Apple TV+, which mainly provide watch lists as their primary structuring paradigm for users. For any further organisation or enrichment, users would need to maintain separate lists, notes or spreadsheets, and design their own way of categorising, labelling or annotating them. As a result of these limitations, as discussed in Section 2.1, users of these platforms often make do with choices that are "good enough".

As demonstrated by Purucker et al. (2019), a spatial hypertext approach, implemented in Mother and illustrated in Figure 1, allows users to create and maintain a persistent context built from small information units, representing among others movies, actors, directors, locations, genres and keywords. By moving, grouping and labelling these units, users can express their ideas, thoughts or other considerations. The colour of the nodes should help users to understand the view of the spatial parser: it identifies related units, colours them the same way and generates individual knowledge base queries. Hence, a certain suggested unit can appear more than once, if it is considered relevant to more than one cluster. As such, it extends current recommender systems by actively involving users in the process of generating and updating suggestions.

In the suggestion process, not only the individual nodes are considered, but also their relationships. To enhance and extend the units that have been manually selected, grouped and labelled, recommendations are queried from the underlying knowledge bases, which contain relevant entities of a given domain and





**Figure 1.** Screenshot of the DemoMedia prototype, with nodes by the user and suggestions (smaller rectangles)—the colour indicates visual related notes and is controlled by the spatial parser.

their relationships. In [Figure 1](#), these recommendations are visualised as smaller boxes (i.e. suggestion nodes) surrounding the larger user-created boxes. The knowledge bases needed for this purpose may be created automatically from computational procedures, such as parsing and mining programme descriptions and metadata from the video provider. They also may be derived from parsing users' context represented as cognitive maps. This enables the system to present suggestions that are by both the computer knowledge and the users knowledge, that latter reflecting humans' experiences and opinions that are based on a combination of user experiences, choices and opinions and recommender system output, which are both integrated in one visual overview.

The discussed scenario can be applied to a single person who wants to get recommendations related to his/her interests. It can also be applied to multiple people using the system to decide upon a joint movie watching event. In such a case, the cognitive map would represent the ideas of multiple persons, which are enhanced with recommendations based on these combined ideas. This opens additional challenges for the analytical (i.e. parsing) part: specialised parsers could be introduced that compute associations between informational units based on *who* issued a specific item. They would be to the set of parsers that are specialised on nodes' spatial arrangement, visual appearance or the time of user interactions with those (Schedel & Atzenbeck, 2016).

Instead of just having to select one option from an uncategorised list of programmes to watch, interactive and exploratory interfaces allow and enforce individual users as well as groups of users to make sense of the available offerings and to actively explore or discuss which options or directions are most aligned with the user's or group's preferences, goals or intentions.

We see additional value in the activity of building contexts over time, ensuring that it is not necessary to start from scratch every time a decision has to be approached. The spatial structure allows users to group, select and categorise programmes based on whatever features they deem important, be it a particular series, genre, topic or theme, favourite actor or producer or even based on reviews or mentions encountered elsewhere that raised a user's interest. Such persistent yet emerging structures represent one's evolving interests. It relates to the user's knowledge put in a time context. The context stays visible to the users and helps them in understanding and steering the system's recommendations.

The questions about why certain successions have been presented to users is much harder to answer for current recommender systems, as those are based on algorithmic interpretations of the observed users' behaviour rather than on explicit user choices and decisions. Remembering behaviour and putting interactions into coherent relationships, however, is very hard to do for humans. As such, as discussed in this section, Mother potentially contributes to research fields related to *Explainable AI*. More specifically, the presented spatial hyper-text approach *supports and enhances* user choice behaviour, allowing and encouraging users to engage make deliberate choices—in other word system-2 decisions (Kahneman, 2002)—instead of merely amplifying reflexive routine (system-1) behaviour, which (as discussed in Section 2) may lead to effects related to filter bubbles and echo chambers.

#### **4.2. Scenario 2: maintenance in Industry**

As described in the television scenario, additional context helps the users to express their thoughts and the recommender system to derive helpful information. This fundamental assumption can also be applied to various application domains that encompass cognitive, system-2 tasks. Maintenance in industry is such a domain where people need creativity, have to research, perform complex calculations and gather a huge amount of information about processes and machines. This section illustrates how the discussion above can be carried over to other application domains and provides some pointers on how to further develop support tools for both professional and private, leisure purposes.

Predictive maintenance is a hot topic in industry (Wang, 2016). System states, given by sensors and other metadata, are monitored by software and compared to previous error states. If the monitored state suggests that an

error is likely, the operator will be informed to prevent further negative impact, like a longer downtime of the machine. This is comparable to the recommendations given by traditional video platforms: the context is derived from some input (in this case, the state of the machine) and the human operator ends up with a (ranked, prioritised, but uncategorised) list of recommended actions.

As computers are good in handling data and pattern matching, this automated approach makes sense for predictable cases. However, often maintenance happens after a process broke *unexpectedly*. Troubleshooting is then handled by humans, who are required to use their experience and knowledge to identify and solve these unexpected and possibly unfamiliar issues. A system for supporting this process would need contextual knowledge that exceeds what is given by sensors, namely the experience and thoughts of the expert users.

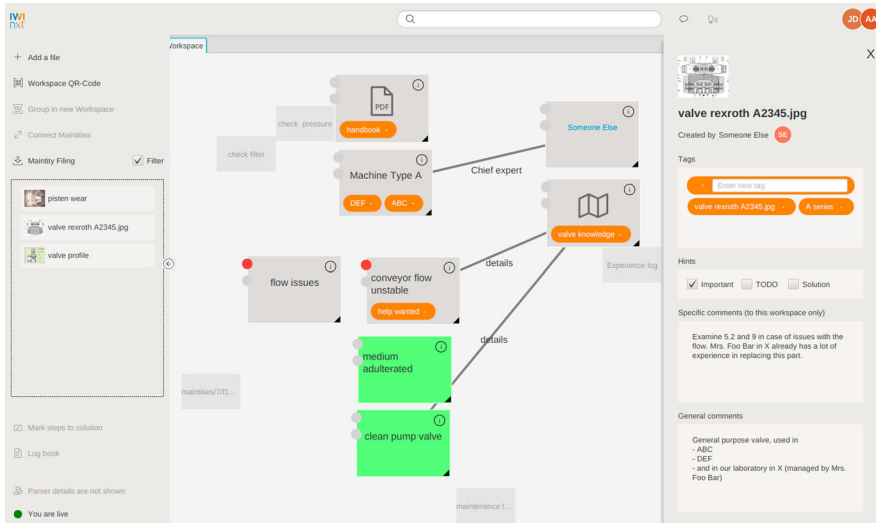
Experts are supported by state-of-the-art document management, collaboration tools and specialised applications like, for example, *Poka*.<sup>14</sup> However, those solutions are usually limited in handling user-generated contexts. The context exists by means of volatile recent queries and, thus, is not permanently available. Even worse, the whole process of searching, working on the machine, reformulating the query, asking a colleague and finally solving the problem is not persistent. This information, however, would be beneficial to derive more accurate suggestions.

As demonstrated by the *IWInxt* spatial hypertext application, which is implemented in Mother and shown in Figure 2, users can build and maintain their context on a 2D canvas. The informational units are specific to the domain and can cover documents, text snippets, pictures, co-workers and other relevant information. This enables maintenance staff to express and structure their thoughts and to save this for later usage. Such created structures get interpreted by parsers for computing suggestions from the knowledge bases, which hold domain specific knowledge. The application allows users to set a node's colour, add flags (e.g. important, todo, etc.), descriptions and other metadata. Metadata is displayed in a panel on the right side, after clicking the info button. Special symbols indicate documents and other workspaces. Furthermore, knowledge derived from the experts' structuring process gets to the system. By that, the system becomes a *learning* system, fusing human experiences and the machine's computational knowledge, which may be regarded as context-aware recommendations.

User sessions are organised in so-called *workspaces*, which can be created from scratch or loaded as already existing instances. Whenever the system recognises any visual changes, the workspace gets enqueued to a periodically running *learning process*. The computed visual parsing result is compared with the previous one (if there is any) to extract those changes. Eventually,

---

<sup>14</sup><https://www.poka.io/>.



**Figure 2.** Screenshot of a workspace within the IWI nxt spatial hypertext application.

all changed relationship weights influence the weights of the knowledge base: if a new relationship is introduced (because the user a new unit or visually related two units which were not related before), the corresponding weight in the knowledge base is raised or newly created in case it did not yet exist.

Similar to the television scenario, as discussed in Section 2.1, the interface displayed in Figure 2 reflects a combination of deliberate user choices that are enhanced with algorithmic recommendations. The core idea is to put the *human in the loop*, offering a spatial hypertext to express context and use visual parsers as a means of communicating the context towards the system.

Beside the advantages of implicit relationships created in a visual space, interviewed maintenance experts often asked for additional, more explicit options to express themselves. Imagine a typical engineering flow chart, whose traversing is defined by various activities and decisions. It would be easier to express such structures, if the application supports explicit, annotated links between informational units. By Mother's support for multiple structure domains, integrating the already implemented link service (Roßner et al., 2020) is easy. Besides adding the option of connecting lines between objects on a spatial hypertext, it also allows to link from or to any other existing resource, not limited to the focussed hypertext itself.

#### 4.3. Commonalities and differences between the two scenarios

The two scenarios have in common that they cover both expected and unexpected situations and (user) contexts. In all situations, recommendations are generated based on the current (estimated) context; these recommendations may be relevant, but they might still be suboptimal and better alternatives

may be found. Instead of automatic diversifying or critiquing, as “regular” recommender systems would do, the spatial hypertext allows users to store and organise their experiences, and reuse these experiences for finding better solutions.

As discussed in this section, this principle holds both for the maintenance industry scenario, which most readers will consider as inherently purposeful, and for more leisure-oriented scenarios such as choosing television programmes. Considering that these are two rather extreme examples on a “spectrum of usefulness”, it becomes apparent that in terms of user needs, they have much in common. The same is expected to yield for scenarios more in the middle of this spectrum, such as news consumption, information seeking or (political) decision making.

## 5. Discussion and interpretation

In Section 2, we observed that deep learning (recommender) systems are particularly good at system-1 tasks, such as predicting and supporting routine behaviour, but not in system-2 tasks, which includes conscious, active decision making (Bengio et al., 2021). In professional domains, such as medicine, critical (system-2) thinking “underlies the ability to effectively work through [...] problems to make good decisions”. Whereas system-1 is fast and intuitive and relies on experience, system-2 is more rational, deductive and analytical, which “fight[s] off the primary impulsivity of system-1 in favour of reality testing” (Croskerry, 2006). For complex or more rational problems, system-2 thinking *complements* intuitive system-1 decision making with reflection and metacognition.

Traditional node-link (i.e. navigational) hypertext is more limited compared to the expressiveness of cognitive maps, as it does not provide extensive graphical overviews or link information. By contrast, spatial hypertext is designed to support users in *strengthening their internal cognitive maps* and simultaneously recommending new items that may be of interest to the user.

Cognitive maps and other tools that support system-2 thinking, enhance metacognition and reflection, competences that are considered essential in life-long learning situations in professional contexts. Common aspects include knowledge observation, thinking strategies and judgements of self-improvement. For instance, Kuiper (2002) showed that experts are more likely to connect current situations with relevant knowledge from the past. Indeed, reflective thinking and self-reflection have been at the centre of adult educational and training activities for a very long time. The process results in “creating new pieces of cognitive structure which may also serve as ‘addresses’ to other information”, and consequently enables persons to consciously consider different alternatives (Von Wright, 1992).

However, metacognition and reflection are not only relevant in professional (work) contexts. Digital productivity tools such as dashboards, bookmarking, collaborative software and social networking are reported to both increase productivity and raise engagement. Ironically, in many cases, productivity tools that people use at home are more technologically advanced than the tools available or allowed at work (Attaran et al., 2019). This indicates that a certain level of reflection and metacognition is also incorporated in private, leisure activities.

Productivity tools, including digital and analogue PIM tools, bear the potential to simplify daily routines and choices. Such tools may include weekly planners in the kitchen, digital family calendars and lists of books-to-read or movies-to-watch. However, surprisingly, frequent use of such tools appears to be correlated with *increased* perceived busyness (Leshed & Sengers, 2011). Arguably, keeping track of everything that one might want to do, visit, read, watch or listen to, may increase the fear of missing out. Still, at the other extreme, binge-watching, which (as argued earlier) is associated with system-1 thinking, is positively correlated with feelings of depression. Particularly younger age groups and singles show the highest prevalence of marathon-viewing (Ahmed, 2017).

As argued in Section 2, the very design of recommender systems and their interfaces are largely aimed at supporting system-1 behaviour, which may be convenient, but also inherently leads to limiting effects such as filter bubbles and echo chambers. As we have discussed above, not only work-related activities but also private, leisure activities benefit from system-2 decisions in order to associate them with a particular purpose, even if this purpose is as mundane as “having a good time”. Therefore, some form of reflection, prioritising and goal-setting is very helpful and often critical in most, if not all, professional and private situations.

## 6. Conclusion

In the recommender systems community, the risks of filter bubbles and echo chambers and the need for transparent, “fair” recommendations have been widely discussed and investigated. As argued in Section 2.3, these effects appear to be the result of human responses, guided by routines and safe choices, to automatic recommendations, which, on their turn, are inherently based on observations of human (routine, system-1) behaviour. In Section 3, we have positioned hypertext as a paradigm for supporting (rational, system-2) thinking and illustratively demonstrated its benefits in professional and private contexts in Section 4.

We believe that it is important to recognise this apparent natural tendency towards “system-1 interaction” with Web interfaces, feeds and recommendations, to understand its benefits as well as its drawbacks. In addition, we observe that the feed-based paradigms consistently are preferred over

“modern” implementations of traditional hypertext concepts, such as interactive structuring, connecting, classifying and annotating, arguably because feeds and recommendations are considered as more convenient.

A running thread in this survey and position article is the mutual reinforcement of system behaviour and user responses, which appears to be driven by a natural tendency towards automation that simplifies human-computer interaction from active decision making into passively following recommendations. An important lesson to be learned from the intertwined evolution of hypertext and recommender systems is that, if we wish to counteract this natural tendency, it is not sufficient to only provide suitable user interfaces—possibly following the spatial hypertext paradigm—, but also to provide users with *incentives*, such as tangible benefits, to actually get actively involved in the process.

### Disclosure statement

No potential conflict of interest was reported by the author(s).

### Funding

This work is part of the project *Car Infotainment*, sub-project *DemoMedia* (Intelligent User Interface), funded by the Bavarian State Ministry of Science and the Arts (Bayerisches Staatsministerium für Wissenschaft und Kunst) [StWMK, grant ID 1547-RP-01] and the project *IWInxt* (Next Generation Intelligent Maintenance System for the Industry 4.0), funded by the Bavarian State Ministry of Science and the Arts [StWMK, grant ID Kap. 15 49 Tit. 547 78-2/2018].

### ORCID

Claus Atzenbeck  <http://orcid.org/0000-0002-7216-9820>

Eelco Herder  <http://orcid.org/0000-0003-0107-299X>

Daniel Roßner  <http://orcid.org/0000-0002-2539-569X>

### References

- Abdollahpouri, H. (2020). *Popularity bias in recommendation: A multi-stakeholder perspective* [PhD thesis]. University of Colorado.
- Abdollahpouri, H., Mansoury, M., Burke, R., & Mobasher, B. (2020). The connection between popularity bias, calibration, and fairness in recommendation. In *Fourteenth ACM Conference on Recommender Systems* (pp. 726–731). ACM Press.
- Adams, W. J. (2000). How people watch television as investigated using focus group techniques. *Journal of Broadcasting & Electronic Media*, 44(1), 78–93. [https://doi.org/10.1207/s15506878jobem4401\\_6](https://doi.org/10.1207/s15506878jobem4401_6)
- Ahmed, A. A. -A. M. (2017). New era of TV-watching behavior: Binge watching and its psychological effects. *Media Watch*, 8(2), 192–207. <http://doi.org/10.15655/mw/2017/v8i2/49006>.



- Akscyn, R. M., McCracken, D. L., & Yoder, E. A. (1988). KMS: A distributed hypermedia system for managing knowledge in organizations. *Communications of the ACM*, 31(7), 820–835. <https://doi.org/10.1145/48511.48513>
- Attaran, M., Attaran, S., & Kirkland, D. (2019). The need for digital workplace: Increasing workforce productivity in the information age. *International Journal of Enterprise Information Systems (IJEIS)*, 15(1), 1–23. <https://doi.org/10.4018/IJEIS>
- Atzenbeck, C., Roßner, D., & Tzagarakis, M. (2018). Mother – an integrated approach to hypertext domains. In *Proceedings of the 29th ACM Conference on Hypertext and Social Media* (pp. 145–149). ACM Press.
- Baeza-Yates, R. (2018). Bias on the web. *Communications of the ACM*, 61(6), 54–61. <https://doi.org/10.1145/3209581>
- Barnet, B. (2010). Crafting the user-centered document interface: The hypertext editing system (HES) and the file retrieval and editing system (FRESS). *Digital Humanities*, 4(1). <http://www.digitalhumanities.org/dhq/vol/4/1/000081/000081.html>.
- Barnet, B. (2018). Boosting human capability. In *Proceedings of the 1st Workshop on Human Factors in Hypertext, HUMAN'18* (p. 1). ACM Press.
- Belloçin, A., Cantador, I., & Castells, P. (2013). A comparative study of heterogeneous item recommendations in social systems. *Information Sciences*, 221, 142–169. <https://doi.org/10.1016/j.ins.2012.09.039>
- Bengio, Y., Lecun, Y., & Hinton, G. (2021). Deep learning for ai. *Communications of the ACM*, 64(7), 58–65. <https://doi.org/10.1145/3448250>
- Berners-Lee, T. (2000). *Weaving the web: The original design and ultimate destiny of the world wide web*. Harper Business.
- Berners-Lee, T., Fielding, R., & Masinter, L. (2005). *Uniform resource identifier (URI): Generic syntax* [Technical report]. Internet Engineering Task Force.
- Bernstein, M. (2010). Criticism. In *Proceedings of the 21st ACM Conference on Hypertext and Hypermedia* (pp. 235–244). ACM Press.
- Broder, A. (2002). A taxonomy of web search. In *ACM Sigir forum* (Vol. 36, pp. 3–10). ACM Press.
- Brown, P. J. (1987). Turning ideas into products: The Guide system. In *Proceedings of the ACM Conference on Hypertext* (pp. 33–40). ACM Press.
- Burke, R. (2002). Hybrid recommender systems: Survey and experiments. *User modeling and User-adapted Interaction*, 12(4), 331–370. <https://doi.org/10.1023/A:1021240730564>
- Bush, V. (1945). As we may think. *The Atlantic Monthly*, 176(1), 101–108. <https://cdn.theatlantic.com/media/archives/1945/07/176-1/132407932.pdf>
- Carretta, M. (2021). *Identifying binge-watching behavior on serial content and exploring its relevance in recommender systems*.
- Celma, Ò. (2010). The long tail in recommender systems. In *Music Recommendation and Discovery* (pp. 87–107). Springer.
- Celma, Ò., & Cano, P. (2008). From hits to niches? Or how popular artists can bias music recommendation and discovery. In *Proceedings of the 2nd KDD Workshop on Large-Scale Recommender Systems and the Netflix Prize Competition* (pp. 1–8). ACM Press.
- Chitra, U., & Musco, C. (2020). Analyzing the impact of filter bubbles on social network polarization. In *Proceedings of the 13th International Conference on Web Search and Data Mining* (pp. 115–123). ACM Press.
- Conklin, J. (1987). Hypertext: an introduction and survey. *Computer*, 20(9), 17–41. <https://doi.org/10.1109/MC.1987.1663693>
- Conklin, J., Selvin, A., Buckingham Shum, S., & Sierhuis, M. (2001). Facilitated hypertext for collective sensemaking: 15 years on from gIBIS. In *Proceedings of the 12th ACM Conference on Hypertext and Hypermedia* (pp. 123–124). ACM Press.



- Croskerry, P. (2006). Critical thinking and decisionmaking: Avoiding the perils of thin-slicing. *Annals of Emergency Medicine*, 48(6), 720–722. <https://doi.org/10.1016/j.annemergmed.2006.03.030>
- Danielson, D. R. (2002). Web navigation and the behavioral effects of constantly visible site maps. *Interacting with Computers*, 14(5), 601–618. [https://doi.org/10.1016/S0953-5438\(02\)00024-3](https://doi.org/10.1016/S0953-5438(02)00024-3)
- Davidson, D. (2008). *Stories in between: Narratives and mediums @ play*. ETC Press.
- Dillon, A., Richardson, J., & McKnight, C. (1990). Navigation in hypertext: A critical review of the concept. In *Proceedings of the 3rd International Conference on Human-Computer Interaction*, North Holland.
- Dubois, E., & Blank, G. (2018). The echo chamber is overstated: The moderating effect of political interest and diverse media. *Information, Communication & Society*, 21(5), 729–745. <https://doi.org/10.1080/1369118X.2018.1428656>
- Engelbart, D. C. (1962). *Augmenting human intellect: A conceptual framework summary report AFOSR-3233*. Stanford Research Institute.
- Geerts, D., Cesar, P., & Bulterman, D. (2008). The implications of program genres for the design of social television systems. In *Proceedings of the 1st International Conference on Designing Interactive User Experiences for TV and Video* (pp. 71–80). ACM Press.
- Gonzalez, M. C., Hidalgo, C. A., & Barabasi, A. -L. (2008). Understanding individual human mobility patterns. *Nature*, 453(7196), 779–782. <https://doi.org/10.1038/nature06958>
- Halasz, F. G. (1987). Reflections on NoteCards: Seven issues for the next generation of hypermedia systems. In *Proceedings of the ACM Conference on Hypertext*, (pp. 345–365). ACM Press.
- Halasz, F. G. (2001). Reflections on “Seven issues”: Hypertext in the era of the web. *ACM Journal of Computer Documentation (JCD)*, 25(3), 109–114. <https://doi.org/10.1145/507317.507328>
- Halasz, F. G., Moran, T. P., & Trigg, R. H. (1987). NoteCards in a nutshell. In *Proceedings of the SIGCHI/GI Conference on Human Factors in Computing Systems and Graphics Interface (CHI '87)* (pp. 45–52). ACM Press.
- Halasz, F. G., & Schwartz, M. (1994). The dexter hypertext reference model. *Communications of the ACM*, 37(2), 30–39. <https://doi.org/10.1145/175235.175237>
- He, C., Parra, D., & Verbert, K. (2016). Interactive recommender systems: A survey of the state of the art and future research challenges and opportunities. *Expert Systems with Applications*, 56, 9–27. <https://doi.org/10.1016/j.eswa.2016.02.013>
- Henze, N., Dolog, P., & Nejdl, W. (2004). Reasoning and ontologies for personalized e-learning in the semantic web. *Journal of Educational Technology & Society*, 7(4), 82–97. <https://www.jstor.org/stable/jeductechsoci.7.4.82>
- Herder, E., Siehdnrl, P., & Kawase, R. (2014). Predicting user locations and trajectories. In *International Conference on User Modeling, Adaptation, and Personalization*, (pp. 86–97). Springer.
- Hitzler, P., Bianchi, F., Ebrahimi, M., & Sarker, M. K. (2020). Neural-symbolic integration and the semantic web. *Semantic Web*, 11(1), 3–11. <https://doi.org/10.3233/SW-190368>
- Hitzler, P., & Van Harmelen, F. (2010). A reasonable semantic web. *Semantic Web*, 1(1–2), 39–44. <https://doi.org/10.3233/SW-2010-0010>
- Howard, K. T. (2021). Queerness and modification in mainstream and indie games: Examining problems with queer representation in video games and exploring design solutions. In *Proceedings of the 32nd ACM Conference on Hypertext and Social Media* (pp. 101–109). ACM Press.
- Jenner, M. (2016). Is this TVIV? on Netflix, TVIII and binge-watching. *New Media & Society*, 18(2), 257–273. <https://doi.org/10.1177/1461444814541523>

- Kahneman, D. (2002). Maps of bounded rationality: A perspective on intuitive judgment and choice. *Nobel Prize Lecture*, 8(1), 351–401. <https://www.nobelprize.org/uploads/2018/06/kahnemann-lecture.pdf>
- Konstan, J., & Terveen, L. (2021). Human-centered recommender systems: Origins, advances, challenges, and opportunities. *AI Magazine*, 42(3), 31–42. <https://doi.org/10.1609/aimag.v42i3.18142>
- Krippendorff, K. (2019). The cybernetics of design and the design of cybernetics. In *Design Cybernetics* (pp. 119–136). Springer.
- Kuiper, R. (2002). Enhancing metacognition through the reflective use of self-regulated learning strategies. *The Journal of Continuing Education in Nursing*, 33(2), 78–87. <https://doi.org/10.3928/0022-0124-20020301-11>
- Lee, B., & Lee, R. S. (1995). How and why people watch tv: Implications for the future of interactive television. *Journal of advertising Research*, 35(6), 9–19.
- Leshed, G., & Sengers, P. (2011). “I lie to myself that I have freedom in my own schedule”. Productivity tools and experiences of busyness. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 905–914). ACM Press.
- Lull, J. (1982). How families select television programs: A mass-observational study. *Journal of Broadcasting & Electronic Media*, 26(4), 801–811. <https://doi.org/10.1080/08838158209364049>
- Marshall, C. C., & Shipman, F. M. (1995). Spatial hypertext: Designing for change. *Communications of the ACM*, 38(8), 88–97. <https://doi.org/10.1145/208344.208350>
- Meintanis, K., & Shipman, F. M. (2010). Visual expression for organizing and accessing music collections in musicwiz. In *International Conference on Theory and Practice of Digital Libraries* (pp. 80–91). Springer.
- Meyrowitz, N. (1986). Intermedia: The architecture and construction of an object-oriented hypemedia system and applications framework. *ACM SIGPLAN Notices*, 21(11), 186–201. <https://doi.org/10.1145/960112.28716>
- Nelson, T. H. (1965). Complex information processing: A file structure for the complex, the changing and the indeterminate. In *Proceedings of the ACM 20th National Conference*, pp. 84–100. ACM Press.
- Nelson, T. H. (1993). *Literary Machines: The report on, and of, Project Xanadu concerning word processing, electronic publishing, hypertext, thinkertoys, tomorrow's intellectual revolution, and certain other topics including knowledge, education and freedom*, (1st ed., vol. 93). Mindful Press.
- Nelson, T. H. (1999). Xanalogical structure, needed now more than ever: Parallel documents, deep links to content, deep versioning, and deep re-use. *ACM Computing Surveys (CSUR)*, 31(4es), 33–es. <https://doi.org/10.1145/345966.346033>
- Nielsen, J. (1990). The art of navigating through hypertext. *Communications of the ACM*, 33(3), 296–310. <https://doi.org/10.1145/77481.77483>
- Nielsen, J. (1993). Iterative user-interface design. *Computer*, 26(11), 32–41. <https://doi.org/10.1109/2.241424>
- Nürnberg, P. J., Wiil, U. K., & Hicks, D. L. (2004). A grand unified theory for structural computing. In D. L. Hicks (Ed.), *Proceedings of the International Metainformatics Symposium 2003 Volume 3002 of Lecture Notes in Computer Science* (pp. 1–16). Springer.
- Pariser, E. (2011). *The filter bubble: How the new personalized web is changing what we read and how we think*. Penguin.
- Park, S. I., & Shipman, F. (2014). Percon: A personal digital library for heterogeneous data. In *IEEE/ACM Joint Conference on Digital Libraries* (pp. 97–106). IEEE.

- Pittman, M., & Sheehan, K. (2015). Sprinting a media marathon: Uses and gratifications of binge-watching television through netflix. *First Monday*.
- Polatidis, N., & Petridis, M. (2019). A switching approach that improves prediction accuracy for long tail recommendations. In *Intelligent Systems and Applications: Proceedings of the 2019 Intelligent Systems Conference (IntelliSys) Volume 1* (Vol. 1037, pp. 18). Springer Nature.
- Purucker, S., Atzenbeck, C., & Roßner, D. (2019). Intelligent hypertext for video selection: A design approach. In *Proceedings of the 2nd International Workshop on Human Factors in Hypertext (HUMAN'19)* (pp. 19–26). ACM Press.
- Ricci, F., Rokach, L., & Shapira, B. (2015). Recommender systems: Introduction and challenges. In *Recommender systems handbook* (pp. 1–34). Springer.
- Roßner, D., & Atzenbeck, C. (2018). Spatial hypertext for end-user development tools. In *Proceedings of the 1st Workshop on Human Factors in Hypertext* (pp. 9–15). ACM Press.
- Roßner, D., & Atzenbeck, C. (2021). Demonstration of Weblinks: A rich linking layer over the web. In *Proceedings of the 32nd ACM Conference on Hypertext and Social Media (HT '21)* (pp. 283–286). ACM Press.
- Roßner, D., Atzenbeck, C., & Gross, T. (2019). Visualization of the relevance: Using physics simulations for encoding context. In *Proceedings of the 30th ACM Conference on Hypertext and Social Media (HT '19)* (pp. 67–76). ACM Press.
- Roßner, D., Atzenbeck, C., & Urban, D. (2020). Weblinks: Augmenting web browsers with enhanced link services. In *Proceedings of the 3rd Workshop on Human Factors in Hypertext* (pp. 283–286). ACM Press.
- Schedel, T. (2016). *Spatio-temporal parsing in spatial hypermedia* [PhD thesis]. Aalborg University].
- Schedel, T., & Atzenbeck, C. (2016). Spatio-temporal parsing in spatial hypermedia. In *Proceedings of the 27th ACM Conference on Hypertext and Social Media* (pp. 149–157). ACM Press.
- Shipman, F. M., Airhart, R., Hsieh, H., Maloor, P., Moore, J. M., & Shah, D. (2001). Visual and spatial communication and task organization using the visual knowledge builder. In *Proceedings of the 2001 International ACM SIGGROUP Conference on Supporting Group Work* (pp. 260–269). ACM Press.
- Shipman, F. M., Moore, J. M., Maloor, P., Hsieh, H., & Akkapeddi, R. (2002). Semantics happen: Knowledge building in spatial hypertext. In *Proceedings of the 13th Conference on Hypertext and Hypermedia* (pp. 25–34). ACM Press.
- Shneiderman, B. (1987). User interface design for the hyperties electronic encyclopedia (panel session). In *Proceedings of the ACM Conference on Hypertext* (pp. 189–194). ACM Press.
- Smith, B., & Linden, G. (2017). Two decades of recommender systems at Amazon.com. *IEEE Internet Computing*, 21(3), 12–18. <https://doi.org/10.1109/MIC.2017.72>
- Smith, J. B., & Halasz, F. (Eds.). (1987). *Hypertext '87: Proceedings of the ACM conference on hypertext*. ACM Press.
- Smith, J. B., & Weiss, S. F. (Eds.). (1988). *Communication of the ACM* (Vol. 31). ACM Press.
- Smith, T., & Bernhardt, S. (1988). Expectations and experiences with HyperCard: A pilot study. In *Proceedings of the 6th Annual International Conference on Systems Documentation (SIGDOC '88)* (pp. 47–56). ACM Press.
- Teevan, J., Alvarado, C., Ackerman, M. S., & Karger, D. R. (2004). 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* (pp. 415–422). ACM Press.

- Tzagarakis, M., Vaitis, M., & Karousos, N. (2006). Designing domain-specific behaviours in structural computing. *New Review of Hypermedia and Multimedia (Special issue on Structural Computing)*, 12(2), 113–142. <https://doi.org/10.1080/13614560601053527>
- van Dam, A. (1988). Hypertext '87: Keynote address. *Communication of the ACM*, 31(7), 887–895. <https://doi.org/10.1145/48511.48519>
- Von Wright, J. (1992). Reflections on reflection. *Learning and Instruction*, 2(1), 59–68. [https://doi.org/10.1016/0959-4752\(92\)90005-7](https://doi.org/10.1016/0959-4752(92)90005-7)
- Wang, K. (2016). Intelligent predictive maintenance (IPdM) system – Industry 4.0 scenario. *WIT Transactions on Engineering Sciences*, 113, 259–268. <https://doi.org/10.2495/IWAMA150301>.
- Weinreich, H. (2012). *Zur Benutzbarkeit assoziativer Verknüpfungen in verteilten Informationssystemen: Entwicklung und Evaluation von Konzepten zur Optimierung der Navigation in offenen, verteilten Hypertext-Informationssystemen gezeigt am World Wide Web* [PhD dissertation]. Universität Hamburg.
- Weinreich, H., Obendorf, H., Herder, E., & Mayer, M. (2008). Not quite the average: An empirical study of web use. *ACM Transactions on the Web (TWEB)*, 2(1), 1–31. <https://doi.org/10.1145/1326561.1326566>
- Zhang, S., Yao, L., Sun, A., & Tay, Y. (2019). Facial expression analysis under partial occlusion. *ACM Computing Surveys*, 51(2), 1–49. <https://doi.org/10.1145/3158369>
- Zhang, Y., & Chen, X. (2018). Explainable recommendation: A survey and new perspectives. *arXiv preprint arXiv:1804.11192*.
- Zuiderveen Borgesius, F., Trilling, D., Möller, J., Bodó, B., De Vreese, C. H., & Helberger, N. (2016). Should we worry about filter bubbles? *Internet Policy Review. Journal on Internet Regulation*, 5(1), 1–16. <https://doi.org/10.14763/2016.1.401>.