Towards a Pattern Library for Algorithmic Affordances

Erik Hekman[1], Dennis Nguyen [2], Marcel Stalenhoef [1] and Koen van Turnhout[1]

[1] *Utrecht University of Applied Sciences, Heidelberglaan 15, Utrecht, The Netherlands*
[2] *Utrecht University, Padualaan 14,Utrecht, The Netherlands*

Abstract

The user experience of our daily interactions is increasingly shaped with the aid of AI, mostly as the output of recommendation engines. However, it is less common to present users with possibilities to navigate or adapt such output. In this paper we argue that adding such algorithmic controls can be a potent strategy to create explainable AI and to aid users in building adequate mental models of the system. We describe our efforts to create a pattern library for algorithmic controls: the algorithmic affordances pattern library. The library can aid in bridging research efforts to explore and evaluate algorithmic controls and emerging practices in commercial applications, therewith scaffolding a more evidence-based adoption of algorithmic controls in industry. A first version of the library suggested four distinct categories of algorithmic controls: feeding the algorithm, tuning algorithmic parameters, activating recommendation contexts, and navigating the recommendation space. In this paper we discuss these and reflect on how each of them could aid explainability. Based on this reflection, we unfold a sketch for a future research agenda. The paper also serves as an open invitation to the XAI community to strengthen our approach with things we missed so far.

Keywords

Algorithmic Affordances, Interactive Recommendation Systems, Explainable AI

1. Introduction

In the past two decades AI has become a ubiquitous and integral component of our daily interactions with computers. Users encounter the output of AI in timelines of social media, streaming media services, search engines, navigation aids, voice assistants, and e-commerce applications – often unknowingly. AI based systems are also on the rise in many professional environments such as finance and health care. Despite this proliferation of AI behind many user interfaces, the interaction design of such interfaces is not maturing at the same rate [18][33].

The dominant model for the interaction design of systems that are driven by machine learning still seems to be an 'under-the-hood-model', in which the user is only presented with the 'best' or 'optimal' outcome of the algorithm. The definition of suitability, or perfect fit, is determined by the designers of such algorithms and their assumptions about the user as well as available user data. To some extent, practitioners also consider it desirable that users are not bothered by the inner workings of a recommender system and are simply presented with valuable output after the AI has done its 'magic' [11].

In many applications, professional or otherwise, it is questionable whether this approach is desirable. Current practices raise critical questions about user autonomy, inclusion, and ethics. The metaphorical 'black box' tries to capture this multilayered problem and has triggered a lively debate about more transparency and rebalancing control in algorithmic systems [23][24].

Explainable AI (XAI) has been proposed as an alternative in which the user at least can get an

Erik Hekman, Dennis Nguyen , Marcel Stalenhoef and Koen van Turnhout. Towards a Pattern Library for Algorithmic Affordances. 2022. In Joint Proceedings of the ACM IUI Workshops 2022, March 2022, Helsinki, Finland, 10 pages.

explanation about the decisions made by the algorithm [13]. This not an easy task. It involves the technical challenge of creating machine learning models that generate output explanations but also the ‘human factors challenge’ of making such explanations fit for purpose in a certain operational context [26]. This paper is positioned towards the human factors challenge and suggests that allowing users to manipulate the parameters of an algorithm can be a viable strategy to enable them to understand and appreciate the outcomes of the algorithm. In other words: we propose that interactive recommendation systems and algorithmic controls in the user interface are well positioned to contribute to users’ understanding of the algorithm. This offers a vital avenue for XAI.

Consider the scenario of personalized health care. Many health care institutions are considering data-driven approaches for medical diagnoses. Data about the patient, collected by the physician or otherwise provided by the patient, are fed into a decision support system (DSS) that employs machine learning to aid the physician in diagnosing the patient. If the DSS has been designed with an under-the-hood-model, the physician is simply fed with a suggested diagnosis, possibly supplemented with some sort of confidence score - or a brief list of possible diagnoses. If the DSS has been designed with explainable AI in mind, the physician can also retrieve explanations about the main factors that contributed to the suggested diagnosis, thus understanding the decision support in a better way. However, if the DSS is designed with algorithm controls, the physician may be able to manipulate certain data or parameters that led to the suggested diagnosis. In this way, she can contextualize the output herself and assess its dependency on the inputs used by the algorithm. The physician could then actively explore alternative diagnoses with the aid of the system.

Interaction designers are likely to be biased towards the last solution outlined in the scenario above. They are likely to consider interactivity as the most suitable approach towards XAI in many cases for three reasons. First, it aids understanding: humans learn many things through manipulation of the world and offering action possibilities could form a basis for the formation of mental models of the AI. Second: interactivity that places emphasis on choice is desirable in diverse contexts, especially when not all users are interested in explanations and configurability (e.g., e-commerce, navigation). Third: interaction is a natural avenue for personalization, as users can explore the possibilities on their own initiative and ‘dig’ as deep as they like.

In this paper we will describe these interaction possibilities with automated, data-driven systems as algorithmic affordances. In short, algorithmic affordances cover a spectrum of design choices that allow users to interact with algorithms and steer their output more directly. To advance this agenda, we explored the state of the art and the potential of such interactive controls by compiling a pattern library: the algorithmic affordances pattern library. The idea is that we collect examples of interactive controls for algorithms from industry and academia within a single structure. We hope that such a pattern library can be an attractive asset of practical use for the industry by contributing to the solution repertoire of the field [30]. This can be expanded by adding new proposals from academia and results from user evaluations, leading into a more evidence-based practice for implementing algorithmic controls. Explainability is not the only reason for constructing the pattern library. We are interested in all potential benefits and drawbacks of algorithmic controls. However, since we think that algorithmic controls could be a potential avenue for greater transparency and autonomy of users, the proposed library may stimulate further development of applicable solutions in the field that reduce the downsides of many current algorithmic designs.

In this paper we describe our approach to develop this pattern library and the structure of its first version. We use this as a steppingstone to highlight how each category of algorithmic controls can aid explainability and outline the research agenda deriving from that. The paper is organized as follows: First, we describe the rationale for and core concepts behind the algorithmic affordances pattern library. Next, we describe our approach for developing the library, followed by the structure that emerged in the first iteration; we relate this structure to available work in academia and open questions concerning the application to explainable AI. In this section we also discuss related work in academia for each of the solution directions. Finally, we discuss open research questions for and next steps of developing the pattern library, as we invite the scientific community to aid us in developing this library further.

2. Rationale

2.1. Algorithmic affordances

Algorithmic affordances are media affordances [20] that center on controlling how an automated system uses data input to calculate an output. In this sense, algorithmic affordances describe the spectrum of explicit and implicit (hidden) interaction possibilities that enable the user to engage with and eventually control the algorithmic system directly and/or indirectly. Affordances are inherently context dependent [10]. The most crucial factors are the interface and its underlying design choices as well as an individual user's understanding of a technology and her motivations. Concerning e.g., recommender systems, this may further include algorithmic awareness and an understanding of data [7]. The perceived usefulness and/or value of an output depends on the purpose of the automated system and often also on the richness of the available data. Providing users with controls to steer the inflow of data and weighing different parameters for the underlying model is not common in current algorithmic systems but also not entirely unheard of either. Several proposals have been made, as we elaborate in the next section.

We do not consider affordances as given or accidental but follow Norman in viewing them as something that can be consciously created by designers [22]. Designers may anticipate the possible uses of a system and invite users to use it in diverse ways through the interaction design of the system. This approach is common in the interaction design of interfaces, and it is questionable whether controlling algorithms should be an exception. For example, Ellsami et al. [9] suggest that many users form mental models of algorithms that are inadequate considering the complexity of modern-day algorithms. Rather than making an argument for clearer explanations, they propose what they call 'seamful design'. Seamful design does not hide the inner workings of an algorithm behind the interface to deliver a user experience which is as smooth as possible, but instead intendedly designs the interface so that users are explicitly confronted with traces of the algorithms' operations in the background. In this way, users may become aware of the choices that are being made for them. Eventually, they not only gain algorithmic awareness but also a better conceptual understanding of how algorithms work and thus more accurate mental models. Although we work in line with this idea, in our notion of algorithmic affordances the primary objective is to increase the user's autonomy and possibilities for control, rather than consciousness of the inner workings per se.

2.2. Algorithmic affordances and XAI

Explaining the trustworthiness, causality, transferability, informativeness, confidence, fairness, accessibility, interactivity, and privacy awareness are key goals of XAI [2]. Common modes of delivering these explanations are text and graphics [34]. For example, textual explanations can uncover the inner workings of an algorithm and reveal how its results are calculated to a user who is new to a system. This communication effort may build trustworthiness and/or confidence in the algorithm. However, textual and graphical explanations remain supplementary and not necessarily central to the interaction with an algorithm. The concept of algorithmic affordances takes here a different route by focusing instead on explanation through interaction.

Zhang & Chen make a distinction between model intrinsic and model agnostic explanations [34]. Model intrinsic explanations reveal the true inner workings of the algorithm, whereas model agnostic explanations provide post-hoc rationalizations which are less tied to the actual decision process of the algorithm. In principle, algorithmic affordances follow the model intrinsic route since they would allow users to make adaptations to the system output. However, this does not mean that the full complexity of the algorithm needs to be completely exposed. The designer may be selective with the elements of the algorithm that are "freed" for user control and the respective interaction possibilities may be designed in accordance with a simplified idea of the algorithm in use.

We argue that offering controls over an algorithm invites the user to actively explore how different factors influence the outcome of an algorithm. This goes further than 'just telling' users how the algorithm works but may provide users with a deeper conceptual understanding of it that stems from their personal experience with the system. That can be considered an advantage over basic textual-graphic explanations, though both approaches could support each other. While

Arrieta et al. highlight that interactivity is a crucial part of XAI, their work mostly focus on domain experts as users and the relationship with the mental model of the user remains unexamined [2]. Note that our proposal is not to replace explanations with controls altogether, we merely suggest controls can be an asset to the repertoire of the designers of XAI.

2.3. Why a pattern library?

Pioneered by Christopher Alexander (1979) design patterns are a common way to define a design language [1]. Design patterns are reusable solutions to common problems interaction design. A pattern library comprises of a set of interrelated solutions (a pattern language) for a larger problem area. It can be seen as partly prescriptive, partly generative theory ([27]. Using pattern-libraries is a common approach in interaction design to harmonize an interaction language across different domains [4]. There were several considerations for constructing an algorithmic affordances pattern library. Our first consideration builds on the observation that the interaction language for algorithmic control was scattered across different domains, whereas in our view the interaction language could be defined in a domain independent way.

Second, while we noticed that algorithmic affordances were prevalent in diverse commercial systems, there seemed to be a disconnect between industry practice and academia. Proposals from academia found no uptake in practice, while patterns in commercial systems were insufficiently described and evaluated in academic literature. A pattern library could act as a boundary object: on the one hand, it should provide practitioners with useful practical information and concrete ideas for how to add interactivity to their algorithms. On the other hand, it should serve as a systematic overview of scientific research. More specifically, we intended to present solutions for algorithmic affordances in conjunction with the latest available evidence-based insights from academia for the effectiveness of different solutions. In this way, the pattern library aims for optimizing the knowledge transfer between academia and industry.

Finally, we also have an educational objective with the pattern library. Designers need to have a good sense of the available solutions. Offering a library may inspire young designers to expand their solution repertoire by looking at solutions they recognize from their own experience from a new angle and by being confronted with novel solutions they were previously unaware of [30].

3. Approach

Inspired by best practices for constructing pattern languages [21], the first version of this pattern library was composed by a combination and triangulation of three approaches [28][29] First, we looked for patterns in the 'wild', meaning we examined well known online services such as social media, streaming content services and dating apps for algorithmic controls. Second, we performed a scan of the literature to look for proposals for algorithmic controls and evaluation of such controls by researchers. At first glance, the literature about algorithmic controls seemed scattered across different fields, such as management science (e.g. [6]), information systems (e.g. [25]), computer supported collaborative work (e.g. [32]) and human-computer interaction (e.g. [14]). Much of this work concerns the question of whether users appreciate some form of control over algorithmic decisions. The reviewed studies come to a positive evaluation: allowing for control reduces algorithm anxiety and increases trust in the system. However, less research has been conducted on the actual design of such controls. The closest to a systematic effort that explores the solution space of algorithmic controls centers on interactive recommendation systems (i.e. [16]). Third, we initiated two student projects explicitly soliciting for algorithmic controls. These projects were executed at two different master's programs at the intersection of data science, humanities, and design. The goal of the exercise was to design a recommender system for video-on-demand services of public service media, taking public values into account. Students designed controls to empower users to make better selections within the offer, but also to invite users to explore more diverse content.

Drawing from these three sources, we composed a first version of the pattern library. We considered something to be a pattern candidate when the control occurred in two sources, for example both an academic proposal and a commercial system, and when it was sufficiently different from other controls. This led to 15 initial pattern candidates, which were subsequently clustered into four categories, signifying a fundamentally different solution direction: controls for feeding (or training) the algorithm, controls for tuning the parameters of the algorithm, controls for activating recommendation contexts and controls for navigating the recommendation space. Following this first iteration, we will publish a first version (see Figure 1) of the pattern library [17] and iterate these steps. We are planning to initiate new student projects with a different challenge and do a more systematic literature review. Also, we try to involve a wider audience of practitioners and students in the effort of identifying patterns 'in the wild'.

Figure 1: Screenshot of the algorithmic affordances pattern library [17]

4. The Algorithmic Affordances Pattern Library

In this section we describe the current structure and contents of the library.

4.1. Feeding the algorithm

The first category of algorithmic controls is intended to feed the algorithm with information of user preferences. Many social media enable this in the form of a 'like', 'favorite' or 'recommend' item. In the context of social software, such features serve the double function of informing the algorithm and informing other users of the software. For example, users using the like function in Twitter (illustrated with a little heart shape) are aware that other users are notified about this action, in particular the author of the message (see Figure 2). The latter is important for users [5] and the algorithmic output relying on 'likes data' may not be on top of the mind for users. As a result, the control may not help with building an accurate mental model of the algorithm [9].

Figure 2: The favorite and heart buttons were subtly different ways of feeding Twitter's recommender algorithm, but the change had a substantial impact on users as Bucher & Helmond [5] have shown.

Other patterns that we identified for feeding the algorithms include: cold-start solutions (where users are asked to feed the algorithm with initial information), curated lists (where users are asked to sort items in a list according to their preferences), and blacklists (in which users are allowed to ban items to prevent them from being recommended). Although these patterns may feel much more as direct controls of the algorithm, they seem to suffer from similar problems in terms of supporting the formation of a mental model of the algorithm. First, it is unclear what the scope of the actions are when the user is providing feedback about a particular item, an author, a topic, or another category. Second, the feedback of the system is delayed and indirect. A recommender may give different outputs in the future, but users are seldom aided in understanding how to relate this to their own previous actions.

Considering these problems, the patterns in this section of the library may not be the best solutions for the goals of XAI. The idea of training an algorithm by giving it regular feedback on its behavior may be a natural (i.e., anthropomorphic) model for users, but its indirect character forms a major drawback for its adoption in XAI. As controls for feeding the algorithm play a key role in many recommenders, there is an imperative for developing solutions that support the users' mental model in a better way. At least users need feedback about the impact of their actions on the algorithm [9].

4.2. Tuning algorithmic parameters

A more direct, and for XAI a more vital approach, might be to offer users direct control over parameters within the algorithm. The most straightforward solution is to enable them to open or close certain data sources as input for the algorithm. This solution was applied in our design project about recommender systems that adhere to public values conducted by several student groups. We were, however, unable to locate an example in a commercial system or a proposal in academic literature. A related idea is to allow users to change weights to elements of the decision-making algorithm such as data sources or intermediate variables included in the modelling. This is implemented in the legal search engine 'fastcase' [12], (figure 2) and it has been proposed by academics as well (e.g. [31]). Nascent studies suggest that such controls are appreciated by users. For example, Jin et al. [19] have tried to add algorithmic controls for music recommendation. They let users control the weight of six characteristics: mood, location, weather, social aspects, current activity, and time of day. This control increased perceived recommendation quality without increasing cognitive load. Users also liked to play with the system.

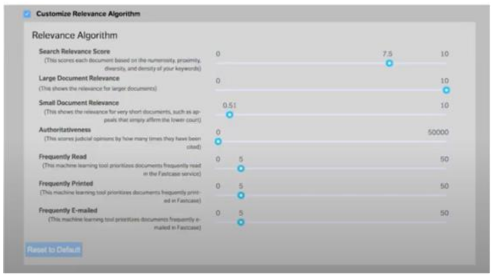

Figure 3: Screenshot of the fastcase [12] legal search engine, which allows users to give different importance to certain parameters in the search process.

There are also proposals in the literature to make the full complexity of an algorithm controllable for the user. For example, Gretarsson et al. [15] enable users to adjust decision paths (Figure 4). They built a recommender in which users can adjust the decision process in each of its layers. This solution gives users full control over the algorithm and allows them to explore the decision-making process in greater detail. However, it may not be feasible to apply this to all kinds of algorithms and in many cases the approach might be 'too direct'. It is often not necessary to completely align users' mental model with the technical implementation of the algorithm. A related idea is PeerChooser, by O'Donnavan et al. [8] which allows users to switch between recommendations crafted for them, and those crafted for other users (digital twins) at smaller or bigger distances. However, in this proposal the potential for supporting the users understanding of the algorithmic decisions is still to be explored.

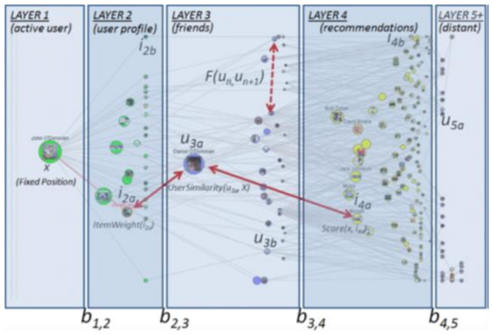

Figure 4: Gretarsson et al. (2010) propose a recommender in which users can influence each part of the decision the recommender takes, thus getting a full notion of the inner workings of the recommender.

Proposals that allow users to tune algorithmic parameters seem to have great potential to achieve explainability of the algorithms involved, because they allow for a very direct manipulation of the algorithm and users can explore the influence on the output immediately; they are the most model intrinsic approach [34]. At the same time, the proposals that we found were still very explorative and 'literal' with regards to the inner workings of the algorithm. To implement this in a way that fits the task context, and the mental model of the user will be a challenge. To us it seems insufficient to just expose the inner working of the algorithm; instead, more direct user controls should bridge between them and the decision of the algorithm in a specific task context.

4.3. Activating recommendation contexts

A third approach to allow users to give control to the algorithm is the notion of context specification. Different user contexts may ask for different settings of the algorithm and different data to be used to train the algorithm. There may be settings in which the user does not want the algorithm to learn from his actions or when the user needs different recommendations. A well-known example is Netflix's "who is watching?" function, which allows users to 'build' different recommendation profiles for e.g., their children. Similarly, the 'Incognito' function in Google Chrome allows users to avoid some of the personalization that is an integral part of Google's service. Different student projects also proposed 'reset' or 'chance' options in their recommenders, indicating a need to escape the profile that a recommender has built from time to time.

At first sight, these contextual control solutions do little to improve the explainability of algorithms and it is not the most promising avenue to explore in the context of explainable AI. Still, we should not immediately dismiss recommendation contexts as a way forward. There is a call for context sensitivity of explanations, and comparing system output for different contexts might help the user if these contexts are meaningful and designed with the right granularity.

4.4. Navigating the recommendation space

A fourth, promising, avenue for exploring XAI, may be solutions that allow users to navigate the recommendation space. Rather than treating a recommendation as a point solution - a single best outcome - the system could present the user with a 'landscape' of outcomes of the recommender and controls to navigate it. A common solution 'in the wild' is the use of ordered lists, in music and movie recommenders such as Netflix and Spotify. The user is presented with a set of tiles suggesting multiple outputs of the recommender that might be relevant and is allowed an easy choice between them. E-commerce sites also explain the social context that fed the recommendations "others who bought this item".

In the academic literature, we find more sophisticated examples of this central idea. Bakalov et al. [3] for example propose the idea of recommendation scapes (Figure 5) for controllable personalization. In their approach recommendations are not just ordered lists, but they take position in structured and interactive visualization. This helps the user to understand what alternatives the recommender may provide, and how they are related to the 'best option'.

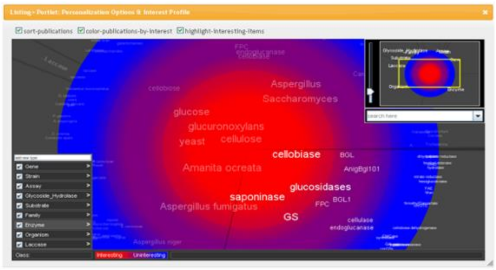

Figure 5: Bakalov et al. [3] propose a recommender in which users can influence each part of the decision the recommender takes, thus getting a full notion of the inner workings of the recommender.

It is easy to imagine how this proposal can help the physician in the fictional example above. Medical diagnoses have a structure and presenting the outcome of the decision support system with respect to alternative diagnoses, combined with putting different weights to underlying data, might be an effective way to enable the physician to make a more educated decision on how to interpret the system output. We consider alternatives for the navigation of the recommendation space as a potent avenue for XAI although a custom design for each context will be needed.

5. Conclusions and discussion

In this paper we have proposed that algorithmic controls could offer a workable solution for explainable AI. Algorithmic affordances offer a different mode for understanding the algorithm from textual explanations and graphics, possibly giving users a feeling for, - rather than only an understanding of -, the innerworkings of the algorithm. As interactive controls allow users to play with the system, they can be intrinsically tailored towards personal needs in understanding the algorithm, for a particular context of use. Algorithmic affordances, as we labeled such controls, have been explored in both industry and academia, but the current state is one of scattered exploration rather than a systematic and substantiated design research program. Moreover, little work has been done in relating the work on algorithmic controls to the substantial body of literature regarding explainable AI. We know too little about the situations in which algorithmic affordances can be a viable alternative to more conventional types of explanation and how the goals of XAI can be met through user control.

This paper modestly contributes to both problems. First, we have proposed a pattern library to draw together the currently dispersed work on algorithmic affordances, in a practical format. Second, while this work is far from complete, it is sufficiently mature to give first reflections about the potential application of algorithmic affordances to XAI. We found that certain categories of algorithmic control have potential for XAI, especially those which allow users to control algorithmic parameters directly and those which allow users to navigate the recommendation space. Other types of controls, such as those enabling users to feed the algorithm and to specify recommendation contexts seem less promising. In a next iteration, we will much more specifically examine the XAI literature to strengthen the link between the library and this field to be able to substantiate these findings. We also call for academics in this area to contribute and suggest improvements for our approach.

Schoonderwoerd et al. suggest that explainable AI should follow a human-centered design approach [26]. In their view, explanations need to be deeply rooted in the specific context of use. Indeed, with increasing complexity of algorithms, it seems a priority to make sure explanations are context specific and user-centric, rather than system centric. The user should understand why the explanation is relevant to her current interactions with the system. Our plea for interactive controls for algorithms follows the same logic. Formulating a generic interaction language such as we did in this first version of the algorithmic affordances pattern library is, however, only a necessary intermediate step. Interactive controls derive their meaning from their use-in-context. Integrating controls, such as proposed in the library, into a particular system requires profound understanding of the users and the way they will use the system and the way they give meaning to its operation in use. The pattern library can be used in the generative phases of the design process. If designers use tried and tested solutions as prototypes for specific use contexts, they have a solid basis to appropriate them and make them fit for use. This appropriation practice can in turn feed back into better pattern descriptions. We are confident that this process will yield explainable and controllable algorithms that are fit for use in real life contexts.

6. Acknowledgements

The authors would like to thank the students participating in the projects that fed into the library. This research was partly funded by a grant by the SIDN Fonds, The Netherlands.

7. References

[1] Christopher Alexander. 1979. The timeless way of building. New York: Oxford university press.

[2] Alejandro Barredo Arrieta, Natalia Díaz-Rodríguez, Javier Del Ser, Adrien Bennetot, Siham Tabik, Alberto Barbado, Salvador García, Sergio Gil-López, Daniel Molina, Richard Benjamins, and others. 2020. Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. Information Fusion 58, (2020), 82–115.

[3] Fedor Bakalov, Marie-Jean Meurs, Birgitta König-Ries, Bahar Sateli, René Witte, Greg Butler, and Adrian Tsang. 2013. An approach to controlling user models and personalization effects in recommender systems. In Proceedings of the 2013 international conference on Intelligent user interfaces, 49–56.

[4] Jan O Borchers. 2008. A pattern approach to interaction design. In Cognition, Communication and Interaction. Springer, 114–131.

[5] Taina Bucher and Anne Helmond. 2017. The affordances of social media platforms. The SAGE handbook of social media, 233-253.

[6] Berkeley J Dietvorst, Joseph P Simmons, and Cade Massey. 2018. Overcoming algorithm aversion: People will use imperfect algorithms if they can (even slightly) modify them. Management Science 64, 3 (2018), 1155–1170

[7] Leyla Dogruel, Dominique Facciorusso & Birgit Stark (2020) ‘I’m still the master of the machine.’ Internet users’ awareness of algorithmic decision-making and their perception of its effect on their autonomy, Information, Communication & Society, DOI: 10.1080/1369118X.2020.1863999

[8] John O’Donovan, Barry Smyth, Brynjar Gretarsson, Svetlin Bostandjiev, and Tobias Höllerer. 2008. PeerChooser: visual interactive recommendation. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, 1085–1088.

[9] Motahhare Eslami, Karrie Karahalios, Christian Sandvig, Kristen Vaccaro, Aimee Rickman, Kevin Hamilton, and Alex Kirlik. 2016. First I "like" it, then I hide it: Folk Theories of Social Feeds. In Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16). Association for Computing Machinery, New York, NY, USA, 2371–2382. https://doi.org/10.1145/2858036.2858494

[10] Nancy Ettlinger. 2018. Algorithmic affordances for productive resistance. Big Data & Society 5, 1 (2018), 2053951718771399.

[11] Kim Falk. 2019. Practical recommender systems. Simon and Schuster, 2019.

[12] Fastcase legal search engine. https://www.fastcase.com/,

[13] Randy Goebel, Ajay Chander, Katharina Holzinger, Freddy Lecue, Zeynep Akata, Simone Stumpf, Peter Kieseberg, and Andreas Holzinger. 2018. Explainable ai: the new 42? In International cross-domain conference for machine learning and knowledge extraction, Springer, 295–303.

[14] Ben Green and Yiling Chen. 2019. The Principles and Limits of Algorithm-in-the-Loop Decision Making. Proc. ACM Hum.-Comput. Interact. 3, CSCW, Article 50 (November 2019), 24 pages. DOI:https://doi.org/10.1145/3359152.

[15] Brynjar Gretarsson, John O’Donovan, Svetlin Bostandjiev, Christopher Hall, and Tobias Höllerer. 2010. Smallworlds: visualizing social recommendations. In Computer graphics forum, Wiley Online Library, 833–842

[16] Chen He, Denis Parra, and Katrien Verbert. 2016. Interactive recommender systems: A survey of the state of the art and future research challenges and opportunities. Expert Systems with Applications 56, (2016), 9–27.

[17] Erik Hekman, Koen van Turnhout, Marcel Stalenhoef Algorithmic Affordances Pattern Library. www.algorithmicaffordances.com

[18] Lars Erik Holmquist. 2017. Intelligence on tap: artificial intelligence as a new design material. Interactions 24.4 (2017): 28-33.

[19] Yucheng Jin, Nyi Nyi Htun, Nava Tintarev, and Katrien Verbert. "Contextplay: Evaluating user control for context-aware music recommendation." In Proceedings of the 27th ACM Conference on User Modeling, Adaptation and Personalization, pp. 294-302. 2019.

[20] Ulrike Klinger and Jakob Svensson. 2018. The end of media logics? On algorithms and agency. New Media & Society 20, 12 (2018), 4653–4670.

[21] Gerard Meszaros and Jim Doble. 1997. G. A pattern language for pattern writing. In Proceedings of International Conference on Pattern languages of program design (1997), 164.

[22] Don Norman. 2013. The design of everyday things: Revised and expanded edition. Basic books.

[23] Frank Pasquale. 2015. The black box society. Harvard University Press.

[24] Frank Pasquale. 2020. New Laws of Robotics. Harvard University Press.

[25] Raja Parasuraman, Thomas B. Sheridan, Fellow, IEEE, and Christopher D. Wickens. 2000. A model for types and levels of human interaction with automation. IEEE Transactions on systems, man, and cybernetics-Part A: Systems and Humans 30.3 (2000): 286-297.

[26] Tjeerd A.J. Schoonderwoerd, Wiard Jorritsma, Mark A. Neerincx, Karel van den Bosch. 2021. Human-centered XAI: Developing design patterns for explanations of clinical decision support systems, International Journal of Human-Computer Studies, Volume 154, 2021, 102684, ISSN 1071-5819, https://doi.org/10.1016/j.ijhcs.2021.102684.

[27] Koen van Turnhout, Sabine Craenmehr, Robert Holwerda, Mike Menijn, Jan-Pieter Zwart, and René Bakker. 2013. Tradeoffs in design research: development oriented triangulation. In Proceedings of the 27th International BCS Human Computer Interaction Conference (BCS-HCI '13). BCS

Learning & Development Ltd., Swindon, GBR, Article 56, 1–6.

[28] Koen van Turnhout, Arthur Bennis, Sabine Craenmehr, Robert Holwerda, Marjolein Jacobs, Ralph Niels, Lambert Zaad, Stijn Hoppenbrouwers, Dick Lenior, and René Bakker. 2014. Design patterns for mixed-method research in HCI. In Proceedings of the 8th Nordic Conference on Human-Computer Interaction: Fun, Fast, Foundational (NordiCHI '14). Association for Computing Machinery, New York, NY, USA, 361–370. DOI:https://doi.org/10.1145/2639189.2639220

[29] Koen van Turnhout, Marjolein Jacobs, Miriam Losse, Thea van der Geest and René Bakker. 2019. A Practical Take on Theory in HCI. Available from: http://bit.ly/TheoryHCI

[30] Koen van Turnhout, Aletta Smits (2021). Solution Repertoire. In DS 110: Proceedings of the 23rd International Conference on Engineering and Product Design Education (E&PDE 2021), VIA Design, VIA University in Herning, Denmark. 9th-10th September 2021. https:/doi.org/10.35199/EPDE.2021.41

[31] Michail Vlachos and Daniel Svonava. 2012. Graph embeddings for movie visualization and recommendation. In First International Workshop on Recommendation Technologies for Lifestyle Change (LIFESTYLE 2012).

[32] Dakuo Wang, Justin D. Weisz, Michael Muller, Parikshit Ram, Werner Geyer, Casey Dugan, Yla Tausczik, Horst Samulowitz, and Alexander Gray. 2019. Human-AI Collaboration in Data Science: Exploring Data Scientists' Perceptions of Automated AI. Proc. ACM Hum.-Comput. Interact. 3, CSCW, Article 211 (November 2019), 24 pages. DOI:https://doi.org/10.1145/3359313

[33] Qian Yang, Aaron Steinfeld, Carolyn Rosé, and John Zimmerman. 2020. Re-examining whether, why, and how human-AI interaction is uniquely difficult to design. In Proceedings of the 2020 chi conference on human factors in computing systems, 1–13.

[34] Yongfeng Zhang and Xu Chen. 2020. Explainable Recommendation: A Survey and new Perspectives. In Foundations and Trends in Information Retrieval. Vol 14. No1. Pp-1-101. P. S. Abril, R. Plant, The patent holder’s dilemma: Buy, sell, or troll?, Communications of the ACM 50 (2007) 36–44. doi:10.1145/1188913.1188915.