Examining the Impact of Algorithm Awareness on Wikidata's Recommender System Recoin

Jesse Josua Benjamin
Human-Centered Computing
Freie Universität Berlin
jesse.benjamin@fu-berlin.de

Claudia Müller-Birn
Human-Centered Computing
Freie Universität Berlin
clmb@inf.fu-berlin.de

Simon Razniewski
Max Planck Institute for Informatics
srazniew@mpi-inf.mpg.de

ABSTRACT

The global infrastructure of the Web, designed as an open and transparent system, has a significant impact on our society. However, algorithmic systems of corporate entities that neglect those principles increasingly populated the Web. Typical representatives of these algorithmic systems are recommender systems that influence our society both on a scale of global politics and during mundane shopping decisions. Recently, such recommender systems have come under critique for how they may strengthen existing or even generate new kinds of biases. To this end, designers and engineers are increasingly urged to make the functioning and purpose of recommender systems more transparent. Our research relates to the discourse of algorithm awareness, that reconsiders the role of algorithm visibility in interface design. We conducted online experiments with 105 participants using MTurk for the recommender system Recoin, a gadget for Wikidata. In these experiments, we presented users with one of a set of three different designs of Recoin's user interface, each of them exhibiting a varying degree of explainability and interactivity. Our findings include a positive correlation between comprehension of and trust in an algorithmic system in our interactive redesign. However, our results are not conclusive yet, and suggest that the measures of comprehension, fairness, accuracy and trust are not yet exhaustive for the empirical study of algorithm awareness. Our qualitative insights provide a first indication for further measures. Our study participants, for example, were less concerned with the details of understanding an algorithmic calculation than with who or what is judging the result of the algorithm.

KEYWORDS

Algorithm awareness, recommender system, transparency, peer production, Wikidata.

1 MOTIVATION AND BACKGROUND

After three decades of continuous growth [42], the Web has become an integral part of our society. It is designed as an open and transparent system, but more recently algorithmic systems that neglect these principles populate the Web. Exemplary of these systems are recommender systems, which have different impacts ranging from societal discourses (e.g., the British EU referendum or the U.S. presidential election of 2016) to profane details of everyday life, such as when choosing a product or service, a place to spend the holidays, or consuming personalized entertainment products. On the whole range of their usage, recommender systems are often interpreted as algorithmic decision support systems that frequently include discussions on bias (e.g., [28]). One of these frequently raised issues

is that of *algorithmic bias*, where specific groups of people based on gender, ethnicity, class or ideology are systematically discriminated by algorithmic decisions [2]. These discussions indicate a growing discomfort with the algorithmic advances both used in and facilitated by the Web.

To this end, ongoing design discourses urge engineers to consider explaining the presence and function of algorithms to end users [24]. Lawmakers, too, are increasingly called upon to respond to issues such as algorithmic bias. Exemplary of the latter is the General Data Protection Regulation (GDPR) by the European Union, which seeks to curtail violations of data privacy and unregulated profitisation from user data. Selbst and Powles go so far as to say that the GDPR effectively guarantees a "right to explanation" of algorithmic processes to end users [35].

This right to explanation is understood as an explicit challenge to the discipline of Human-Computer Interaction (HCI); to be met with concrete means of representing and explaining algorithmic processes. In HCI, this challenge aligns with the discourse of *algorithm awareness*. Hamilton and colleagues define algorithm awareness as the extent to which a user is aware of the existence and functioning of algorithms in a specific context of use [18].

However, the scope of the term algorithm awareness is not yet clearly defined, partially a result of the lack of experimental results associated with the discourse. As a consequence, it is unresolved whether algorithm awareness is the result of unearthing new methods of interaction, novel forms of representation, finding means of explaining algorithmic processes, or all of these aspects taken together. Similarly, its methodological perspective is vaguely defined. Are algorithm-aware designs, for example, a result of a critical technical practice [36], or are they a new form of human-centered design? If algorithm awareness as a principle is to contribute to an understanding of web-based algorithmic systems today (and tomorrow), these methodological shortcomings need to be addressed.

In our research, we focus on one specific aspect from the discourse of algorithm awareness, the aspect of algorithm representation. We first discuss related work from the areas of HCI, Computer-Supported Collaborative Work (CSCW), and Science and Technology Studies (STS). We argue that algorithm awareness should be understood within the context of human-technology relations since algorithmic systems increasingly impact how we see the world. We then introduce a use case which allows for studying different representations for algorithm awareness because of its open design. The use case is situated in the peer production system Wikidata, in which the completeness recommender Recoin is being used by editors to receive recommendations for their next edits. As opposed to commercial web-based systems, the design principles of Wikidata give as access to all necessary information available regarding

Recoin's engineering and usage. We can, thus, reflect on the various decisions made during Recoin's development and can suggest different modes of representing the algorithmic system by considering the dimensions of explainability and interactivity.

Our research makes the following contributions: We provide experimental results to be used in the continuing development of the discourse on algorithm awareness. This concerns insights on design measures, namely textual explanation of algorithmic logic and interactive affordances, respectively. Our results suggest that the provision of more interactive means for exploring the mechanism of the algorithm for users has significant potential, but that more research is needed in this area. As for conducting experiments in this context, we provide first methodological insights, which suggest that measures of comprehension, fairness, accuracy and trustworthiness employed in the field are not yet exhaustive for the concerns of algorithm awareness. Our qualitative insights provide a first indication for further measures. The study participants, for example, were less concerned with the details of understanding an algorithmic calculation than with who or what is judging the result of the algorithm.

In the following section, we discuss existing research from three perspective. First, we discuss the role of automation in peer production communities which leads to an increased usage of algorithmic systems in this context. Second, we review existing approaches that attempt to make these algorithmic systems transparent. Third, we combine these insights to argue for the urgency of researching algorithm awareness. This theoretical section is followed by a detailed introduction to Recoin, our use case, including its technical design as well as how it connects specifically to the topic at hand. Subsequently, we showcase our experiment, in which we detail the setup, design, results and analyses involved in our experimental study. Then, we proceed to discuss our insights. Finally, we conclude by outlining future work, in which we seek to undertake qualitative studies into how the evaluated modes of representation affect the relations between humans and algorithmic systems.

2 RELATED WORK

Automation in Peer Production Communities. In contrast to the predominant commercial platforms on the Web, peer production communities, such as Wikipedia, OpenStreetMap, or Linux, provide a valuable alternative for people to share their ideas, experiences, and their collaboratively created knowledge openly [4, 5]. In these communities automation is an integral component in order to handle especially "mindless, boring and often reoccurring tasks" [27]. In Wikipedia, for example, various forms of algorithmic support exist; recommender systems, such as the SuggestBot help people to find suitable tasks [8], neural networks, such as employed by ClueBots help to revert obvious vandalism [7], or semi-automated user interfaces, such as Snuggle help editors to socialize more efficiently with newcomers [17]. Wikidata as Wikipedia's sister project profited from existing experiences in these automation efforts, thus, tools for vandalism detection were highly sophisticated from the beginning [34]. However, depending on how this automation is being used, the outcome goes in both directions. The unreflected use of automation can suppress participation of good-faith newcomers [16], and on the other hand, recommender systems on Wikipedia can

significantly improve editor engagement and content creation [41]. Existing research shows, how the openness of peer production systems, such as the various Wikimedia projects (Wikipedia, Wikidata, etc.) enable researchers to investigate the manifold facets of automation in a real-world setting, and simultaneously support these projects in their goals of providing free high quality content.

Approaches to Algorithm Awareness. With regards to related discourses such as on Fairness, Accountability and Transparency (FAT) [23] or Explainable Artificial Intelligence (XAI) [15]; algorithm awareness is more aligned to the study of lay's persons experiences of algorithmic systems. As with FAT and XAI, the concerns of the discourse are illustrative of pressing socio-cultural, economic and political needs. However, and similarly to FAT and XAI, algorithm awareness so far suffers from the lack of a methodological definition. Both in terms of design and engineering, implementation of algorithm-aware designs is challenged by two fundamental issues which can be derived from Hamilton and colleagues definition [18]: (1) the perceivability of an algorithm (e.g., results, logic, data) and (2) an actionable mode of representation that leads to informed usage.

So far, the context of conducted algorithm awareness studies differs greatly. Studies have included, for example, both attempts at reverse-engineering web-based systems such as the Facebook's newsfeed [11] as well as manipulating online peer grading systems [21]. In the former, Eslami specifies that an algorithm-aware design should provide an actionable degree of transparency to algorithmic processes in order to promote a more informed and adaptive use of a specific system by its users [10]. In her study, Eslami operationalizes the approach of *seamful design*[1] to display results of the Facebook newsfeed algorithm that usually do not get displayed.

In the latter, Kizilcec proposes another dimension to algorithm awareness [21]: the question of how much transparency of an algorithmic system is actually desirable to ensure understandability and usage. For his study, Kizilcec exposes participants in a peer-graded massive open online course to three kinds of transparency when confronted with their course grades. For each kind of transparency, he asked participants to rate their comprehension of the user interface and to what extent they evaluate the provided information as fair, accurate and trustworthy. These measures provide a first set of measures to empirically study how humans understand algorithmic systems. His results suggest a medium degree of transparency (in this case, textually disclosing the result and logic) as most effective. A high transparency (of the result, the underlying logic and raw peer grading data), he finds, is in fact detrimental to trust in the algorithmic system – whether or not the received grade was lower than expected.

A particular focus in algorithm awareness, as well as in XAI and FAT, are the concrete means by which humans may become more informed about algorithmic systems. A frequently deployed solution is the use of textual explanation of algorithmic processes or outputs across all discourses, featuring in contexts such as social media [30], aviation assistants [25], online advertising [12], classification algorithms in machine learning [33] and in online peer grading as discussed above [21]. The prevalence of this solution

---

[1]The term seamful design is created as opposite concept to seamless design, where the algorithmic system fades into the background of human's perception.

may be interpreted as a clear indication for textual explanation being most suitable for establishing algorithm awareness. Within the aforementioned studies, various versions of textual explanations were studied comparatively. For example, even though Kizilcec questioned how much information a user may require, his various conditions of transparency all feature textual, explanatory solutions only [21]. This may be considered a gap in the discourse. Returning to Hamilton and colleagues, the complexities of contemporary algorithmic systems do not only pose the question of how much humans may need to understand, but also in what way [18]. This suggests, for example, that experimenters should also explore differences between textual explanation of algorithmic logic and interactive, non-declarative solutions in the same context.

Urgency for Algorithm Awareness. Due to the increase of automation on the Web, finding means for a better understanding of algorithms both by experts and lay users is particularly urgent. With algorithms, existing biases may become amplified substantially. In the discourse on recommender systems, bias has been observed as a challenge early on, and a major line of recommender systems research investigates how to avoid popularity bias, i.e., providing recommendations that are already known to satisfy a large number of users [13, 14]. More recently, several works investigate the explainability of recommender systems [19, 43]. Even open peer production systems such as Wikidata need to be seen in this context. That is, if there is a pre-existing bias in a knowledge base such as Wikidata, a recommender system may cause this bias to become self-perpetuating. Additionally, encoded bias may spread into the outputs of Wikidata APIs–thereby opaquely influencing the standard in domains that rely on Wikidata services. In his overview of bias on the Web, Baeza-Yates concludes that an awareness of bias (whether algorithmic or cultural) is the primary precondition for designers and engineers to mitigate potentially negative effects on users [2]. The developer perspective as advanced by Baeza-Yates suggests that an engineering solution may be found with the potential to eliminate bias, whether by way of analyzing biased tendencies in the data used by a Web platform or running of extensive A/B-testing of subgroups [22].

However, as repeatedly noted by Wiltse and Redström, the complexity of algorithmic systems in the modern Web troubles this suggestion. In their words, the Web is populated not by clear developer-client relations, but by *fluid assemblages*, i.e. socio-technical configurations that change in various context of use [32, 40]. Bias, therefore, is not necessarily a definitive phenomenon for either human or machine. Accordingly, counting on purely technical solutions to eliminate bias needs to be up for debate. Instead, and as called for by various researchers from algorithm awareness, FAT and XAI, empirical studies that provide insights into how algorithmic systems (and the biases encoded therein) may be made more transparent.

In the next section, we introduce the context of the open peer production system Wikidata, in which our use case, Recoin – a property recommender-system – is used.

3 RECOIN: PROPERTY RECOMMENDER

Wikidata is an open peer production system [39]. Its structured data is organized in entities, where two types of entities exist: items

and properties. Items represent real-world objects (individuals) or abstract concepts (classes).

Each item is described by statements; for example, the entity Q1076962 represents the human *Chris Hadfield*. Each item is described by statements that follow a *subject-predicate-object* structure (e.g., *Chris Hadfield* (Q1076962) "is instance of" (P31) "human" (Q5)). Thus, a property, i.e. predicate, describes the data value, i.e. object, of a statement. In October 2018, the community has more than 200k registered contributors with 19K active on a monthly base. They have created more than 570m statements on more than 50m entities.

Even though Wikidata was founded to serve as a structured data hub across all Wikimedia projects, today, it is utilized for many other purposes; for example, researchers apply Wikidata as authoritative for interlinking external datasets, such as for gene data [6] or digital preservation data [37], or companies use Wikidata's knowledge graph for improving their search results, such as Google or Apple. A significant issue for Wikidata's community is consequently the quality of the data. Data quality is a classical problem in data management, however, in peer production settings such as in Wikidata, data quality assessment is complicated because of the continuous, incremental data insertions by its users, the distributed expertise and interest of the community, and the absence of a defined boundary in terms of its scope. Over the past years, the community has introduced many tools that address this challenge, that range from visualizing constraint violations to de-duplication tools and translation tools. One of these tools is Recoin that we present in more detail in the next section.

3.1 Technical Design

Recoin is a recommender system for understanding and improving the completeness of entities in Wikidata [1, 3]. A main motivation for implementing Recoin is Wikidata's openess, since it allows anyone to add nearly any kind of entities - items and properties. The latter led to a huge space of possible properties (4,859 properties as of July 2nd, 2018), with many applying only to a very specific context (e.g., "possessed by spirit" (P4292) or "doctoral advisor" (P184)). Consequently, even experienced editors in Wikidata may lose track of which properties are relevant and important to a given item which might hinder them to improve data quality in Wikidata [31].

Recoin is a gadget - an interface element - on Wikidata[2]. A visual indicator informs a person about the relative completeness of an item and, moreover, it provides an editor with concrete recommendations about potentially missing properties on this item. Figure 1) shows the gadget on an item page on Wikidata. A visual indicator (icon on the top right) shows a color-coded progress bar that has five levels ranging from empty to complete. On the top of an item page, the recommendations are provided in an expandable list that shows up to ten of the most relevant missing properties.

The idea of relative completeness is motivated by the situation that in absolute terms, measuring the completeness of an item is impossible in an open setting. The relative completeness, thus, considers completeness in relation to other, similar items. The relatedness function of Recoin considers two items as similar if they
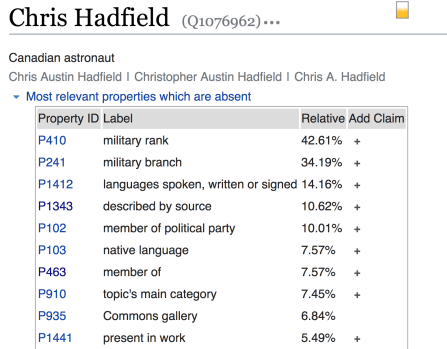
Figure 1: Recoin for the astronaut Chris Hadfield.

share the same class[3]. The visual indicator of Recoin should not be understood as an absolute statement, i.e. level 5 (=complete) means, that all possible statements a given on the item page, but should rather be interpreted as a comparative measure, i.e. the statements on this item are more complete than on similar items.

The completeness levels in Recoin are based on thresholds that are manually determined. It considers the average frequency of the 5 most frequent properties among the related ones to consider an item as *most complete* (0%-5% average frequency), *quite complete* (5%-10% average frequency), and so on. Furthermore, each user is shown 10 recommendations in order to avoid an overwhelming user experience.

3.2 Need for Algorithm Awareness

As of September 25, 2018, Recoin is enabled by 220 editors on Wikidata[4], who created, based on Recoin's recommendations, 7,062 statements on 4,750 items. Even though Recoin is a straightforward approach for improving data quality on Wikidata, editors are hesitating to apply Recoin. Moreover, after persons used Recoin, they have raised a number of concerns. Based on existing discussions on Recoin's community page[5] and on the mailing list, we identified three typical issues.

Editors, for example, posed questions regarding the *scope of the recommender*: *"Not sure if Recoin looks at qualifiers & their absence; if not, might this be something to think about?"*. The information provided by Recoin hindered editors to understand which data was being used to compute the recommendations. In another case, an editor was wondering about Recoin's underlying algorithm: *"Something weird going on on Jan Jonker (Q47500903), stays on least complete."*. In this case, the unchanging visual indicator of Recoin caused the user to *question the functionality* of Recoin. Another user was concerned about the provided recommendation and its suitability for specific items: *"How is Property:P1853 "Blood type" on this list, is that relevant (or even desirable) information for most*

people?". The user was not able to include its personal preferences - world view - in Recoin's recommendation.

However, the third typical issue exemplified by Wikidata's editors raises a more genuine concern over the *impact of Recoin on an already biased knowledge base* (e.g., the predominance of the English language [20]). One editor stated: *"This tool has its attractions but it does entrench cultural dominance even further as it stamps "quality" on items. The items with the most statements are the ones that are most likely to get good grades. Items from Indonesia let alone countries like Cameroon or the Gambia are easily substandard."*.[6]

On a surface read, this quote further substantiates the misunderstood nature of Recoin's function, as it is not intended as a unilateral absolute grading of the completeness of a particular item, but rather as a comparative tool that recommendations depend on the activities of the editors on similar items. However, and much more significant, the concern raised about cultural dominance is a very contemporary problem in algorithmic system design. Recoin fails to address this concerns by its current design and mediated function. In other words, the cultural bias in the recommended properties, even if not intended, seem to affect the usage of Recoin.

Based on these insights, we wanted to better understand how a re-design of Recoin that considers algorithm awareness by focusing on explainability and interactivity can address the aforementioned issues. As opposed to existing research in this context, for example carried out by Eslami et al. [11], we do not require methods such as reverse engineering to understand the algorithmic system we are dealing with on a technical level. This knowledge is key to understanding the intricacies of the Web platforms today; as the ways in which an algorithm operates within a larger socio-technical context arguably also shapes the extent to which humans can or should be aware of it. Therefore, with an openly available recommender system in an open peer production system, we can conduct experiments that are closely tied to the actual practice of Wikidata editing activities, i.e., we can reflect on the technical and the social system similarly.

In the following section, we introduce our experimental setup that help us to examine the impact of varying degrees of explainability and interactivity of the UI of the recommender system on humans. Following the concept of Recoin, our experiment featured a task of data completion. By measuring the interactions of participants with various designs during a data completion task and by eliciting self-reports, we sought to understand which design measures increased task efficiency while at the same time were most effective in increasing understanding of the algorithmic system.

4 EXPERIMENTAL SETUP

Informed by previous research, we designed two alternative UIs where each represents another degree of explainability and interactivity. Each user interface extends or replaces the previous version by specific design elements. In the following, we differentiate the original Recoin design (R1), the textual explanation design (RX), and the interactive redesign (RIX). The RX design is mainly inspired by previous work, from which we adapted the explanation design [21]. The RIX design follows an interactive approach, where the user

---

[3]An exception are items that are instance a of the class *human*. In this case, the class "occupation" is used.

[4]Further information are provided on the following page https://www.wikidata.org/wiki/Special:GadgetUsage.

[5]https://www.wikidata.org/wiki/Wikidata_talk:Recoin

[6]The corresponding thread can be found on Wikidata's mailing list archive https://lists.wikimedia.org/pipermail/wikidata/2017-December/011576.html.

Chris Hadfield (Q1076962)

Canadian astronaut
Chris Austin Hadfield | Christopher Austin Hadfield | Chris A. Hadfield

This entry is *very basic*, because it misses information about *educated at (74.11% relevant for astronauts)*, *astronaut mission (68.25%)*, *award received (53.85%)*, *part of (46.64%)*, and *employer (46.52%)*.

▾ Most relevant properties which are absent

Property ID	Label	Relative	Add Claim
P69	educated at	74.11%	✚
P450	astronaut mission	68.25%	✚
P166	award received	53.85%	✚
P361	part of	46.64%	✚
P108	employer	46.52%	✚
P734	family name	44.57%	✚
P410	military rank	42.61%	✚
P241	military branch	34.19%	✚
P1412	languages spoken, written or signed	14.16%	✚
P1559	name in native language	13.31%	✚

Figure 2: Recoin with Explanation *(RX).*

Chris Hadfield (Q1076962)

Canadian astronaut
Chris Austin Hadfield | Christopher Austin Hadfield | Chris A. Hadfield

▾ ▬▬▬▬▬ This Astronaut is *very basic* by comparison.

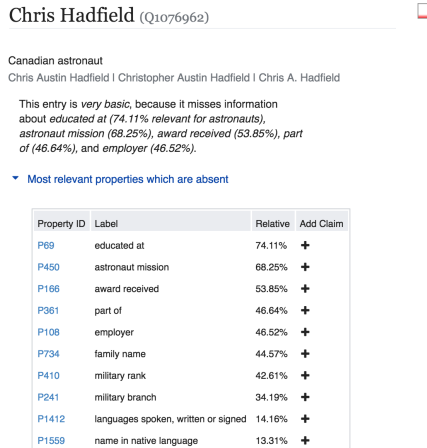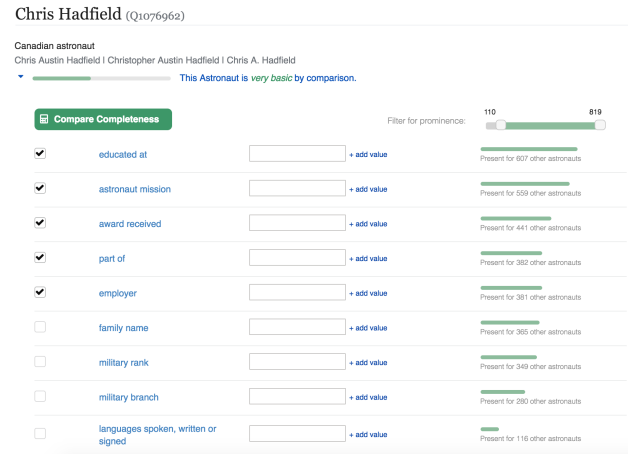| 🖫 Compare Completeness		Filter for prominence: 110 — 819
☑ educated at	[ ] + add value	Present for 607 other astronauts
☑ astronaut mission	[ ] + add value	Present for 559 other astronauts
☑ award received	[ ] + add value	Present for 441 other astronauts
☑ part of	[ ] + add value	Present for 382 other astronauts
☑ employer	[ ] + add value	Present for 361 other astronauts
☐ family name	[ ] + add value	Present for 365 other astronauts
☐ military rank	[ ] + add value	Present for 349 other astronauts
☐ military branch	[ ] + add value	Present for 280 other astronauts
☐ languages spoken, written or signed	[ ] + add value	Present for 116 other astronauts

Figure 3: Recoin Interactive Redesign *(RIX).*

can interact with the outcome of the algorithm, thus, can explore how the algorithm outcome changes based on specific settings. The three UI designs are used in five experimental conditions (C1 to C4), supplemented by a baseline where participants only used the regular Wikidata interface. We explain these conditions in more detail in one of the following paragraphs.

Based on these designs, participants should solve the same task across all conditions: adding data to a Wikidata item. We recruited 105 participants via Amazon Mechanical Turk (MTurk)[7], whereby each participant had a minimum task approval rate of 95%, and a minimum amount of HITs of 1,000. Each participant received USD 3.50 (equivalent to USD 14.00 hourly wage) for full participation. We recruited only U.S. participants to reduce cultural confounds. We randomly and evenly distributed participants over our five conditions (i.e. 21 participants), also ensuring that no participant could re-do the task or join another condition by associating participants with qualifications. Each participants was given 10 minutes for task completion.

In each condition, participants went through the same general procedure during task completion. At first, we provided a brief on-boarding, then we provided a task briefing. After the study participant carried out the task, she had to fill out an explicative self-report which contained the dimensions comprehension, fairness, accuracy and trust. Additionally, all participants obtained a task completion score, which we correlated with server activity to ensure that our final study corpus featured no invalid contributions. All data and results of our study will be available under an open license[8].

In the following, we outline the design decisions that have led to our three designs for Recoin in more detail. Then, we describe the task and the experimental design.

4.1 Design Rationales

In the following, we describe each UI approach of the recommender Recoin in more detail. For each user interface, we provide a corresponding visual representation.

4.1.1 Recoin User Interface R1. The original design of Recoin (cp. Figure 1) was primarily informed by existing UI design practices in Wikipedia. The status indicator icon was chosen to mirror the article quality levels on Wikipedia, such as "Good article" or "Featured article" [9]. The used progress bar was motivated by existing visualizations in Wikipedia projects.[10]

Some parameters to represent the results of the Recoin recommender were determined without further considerations, such as the thresholds that represent the five levels of completeness.

4.1.2 Recoin User Interface RX. Textual explanation of algorithmic logic is a wide-spread measure in the related work, and has been deployed in contexts such as social media [30], aviation assistants [25], on-line advertising [12] and classification algorithms in machine learning [33]. For our design of RX, we drew inspiration from Kizilcec, who tested three states of transparency to understand algorithm awareness in on-line peer grading [21]. Since the algorithm's function can be compared to Recoin (i.e., a rating algorithm), we adapted the format of Kizilcec's best solution and added a textual explanation to Recoin's user interface that describes the logic behind Recoin's calculation (cp. Figure 2).

4.1.3 Recoin User Interface RIX. Our interactive user interface (RIX, cp. Figure 3) is based on insights gained from user feedback from Recoin's current users (cp. Section 3.2) and from the philosophy of technology as discussed by Verbeek [38]. Concerning the latter, we posit that Recoin *actively* transforms the relationship an editor has with Wikidata and the entities therein. Through Recoin, Wikidata items that formerly were objects containing knowledge


The Canadian astronaut Chris Hadfield rose to prominence by playing "Major Tom" on his guitar on the International Space Station (ISS). Ever since returning to Earth, he's been an active public figure as well. However, Wikidata doesn't know very much about him at all!

Your job is to add important data to Chris' Wikidata entry. Search this text **for the data. Wikidata has autocomplete— please only use suggested options. You may use the attributes from** Buzz Aldrin **as a guideline.**

If you need more examples of astronauts, have a look at this list.

By clicking "Begin" at the bottom of this page, you will start your task. You will have 10 minutes *(the timer will continue to run even when you are in another tab)* to add data to Chris Hadfield's entry.

Here is a brief video tutorial (2:30min) on how to add statements and values to a Wikidata entry:
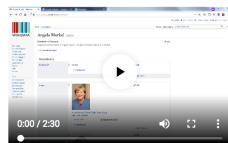
Figure 4: Briefing page, with material added and resources for manually carrying out Recoin's functions and tutorial.

are now also objects that are rated. Technically, this rating is not an indication of absolute qualities, but one of community-driven standards, i.e. how the Wikidata community currently views a specific class of items.

However, as illustrated in the various responses to Recoin, this mediation is not adequately communicated by Recoin's current design. Furthermore, in reflecting on Recoin with the original developers, we found that the comparative parameter of dividing the relevancy of the top five properties was arbitrarily chosen. In line with Mager [26], we consider transparency for this result of developer decision-making as essential.

We operationalized these insights for RIX by considering how the community-driven aspect of Recoin could not only be displayed, but made interactively explorable. To this end we (1) included a reference to the class of the displayed entity (e.g., "astronaut" in our running example) in the drop-down title. This was designed to convey that this particular item is rated based on its class. Next, we augmented the drop-down itself extensively. We (2) substituted the relevance percentage with a numerical explanation for each suggested property (e.g., a relevance for the property 'time in space' of 67.03% means that 549 out of 819 astronauts have this property). In contrast to a percentage, it was our intuition that relating to the class would highlight the community-driven aspect of Recoin. To strengthen this aspect further, we (3) included a range slider which allows filtering properties based on their prominence in the class (i.e., compare this entity based on their occurrence in a minimum/maximum of n astronauts). Finally, we offered a way for directly interacting with Recoin's calculation: we (4) allowed our participants to reconfigure the relevancy comparison by (de-)selecting individual properties. Thereby, we wished to show that relevancy can be a dynamic, community-driven attribute in this algorithmic system.

4.2 Task

For the study participants, we defined a typical editing task on Wikidata. We presented each study participant with a copy of Wikidata's user interface, to provide a most realistic task setting. First, the participants received a brief on-boarding for Wikidata, and, depending on the condition, for Recoin as well. Participants then proceeded to the task briefing page (cp. Figure 4). The participants were asked to add further properties and data to a Wikidata item. Additionally, we supplied participants with a short video tutorial that explained how properties can be added to an item on Wikidata.

In each condition, the Wikidata item to be edited was *Chris Hadfield*, a Canadian astronaut[11]. This item was chosen because it has a number of missing statements that are easily retrievable, and on the other hand, the item describes an astronaut who is probably well-known by our study participants which are U.S. based. Additionally, the occupation of astronauts was thought to be relatively neutral, as opposed to, for example, politicians or soccer players.

We provided study participants with source material for the task composed of comparatively relevant and irrelevant pieces of information about *Hadfield*. We also supplied a link to a very detailed Wikidata item with the same occupation, the US-American astronaut *Buzz Aldrin*[12], and a link to a Wikidata query of the occupation "astronaut"[13], both with the intention to allow study participants to compare the given item with other items, i.e. we encouraged our participants to perform the functionality of Recoin manually.

In addition, we provided a short video tutorial on how to add statements to Wikidata items. Following the task briefing, participants could choose to commence the task, which lead to the reconstructed Wikidata page for *Hadfield*. Within a 10 minute limit, participants could then add statements to the item.

We randomly assigned each participant to one of the conditions. Once the 10 minutes passed, participants were alerted that time was up, and that they should proceed to the self-report. Here, participants were confronted with a grade (from A-F) of their task. This grade was calculated through the difference in completeness before and after participants added information to the Wikidata item (e.g., when a participant additions increased the relative completeness of *Hadfield* by more than 20% but less than 30%, they received a "B").

In correspondence to this grade, participant's were asked to rate their comprehension (5-point Likert scale), feelings of accuracy, fairness and trust (7-point Likert scale) of the recommender system. Again, due to substantial methodological and contextual similarities to Kizilcec's online study [21], we adopted the aforementioned measures to our study. Participants were also asked to expand on their ratings using free text fields. Upon submitting their ratings, participants were returned to the MTurk platform.

4.3 Study Design

We conducted a between-subject study with five conditions. In the following, we define each condition and explain each measure we collected during the study.

Relevance	Difference of the completeness value of the item before and after task completion.
Usage	Number of times the recommender Recoin was used during task completion.
Compre-hension	To what extent do you understand how your task has been graded? *(1) No understanding at all to (5) Excellent understanding.*
Fairness	How fair or unfair is your grade? *(1) Definitely unfair to (7) Definitely fair.*
Accuracy	How inaccurate or accurate is the grade? *(1) Very inaccurate to (7) Very accurate.*
Trust	How much do you trust or distrust Wikidata to fairly grade your task? *(1) Definitely distrust to (7) Definitely trust.*

Table 1: Overview of measures employed in our online experiment.

4.3.1 Conditions.

4.3.1 Conditions. The first three conditions (baseline, C1, C2) were designed to test usage and understanding of the current version of Recoin, i.e. R1. We then proceeded to test the collected baseline against textual explanation (C3 with RX) as found in related work [15, 21, 29] and a redesign motivated by the shortcomings found therein (C4 with RIX). By comparing the results of the conditions, we aimed to gather insights on how design impacted human understanding of Recoin's function.

All conditions are described in more detail next, followed by a description of the collected measures.

- *Baseline*: Participants can add data on a Wikidata item *without* Recoin being present in the user interface.
- *Condition 1*: Participants can add data on a Wikidata item *with* Recoin (R1) being present in the user interface.
- *Condition 2*: C1 but Recoin is mentioned during the on-boarding process.
- *Condition 3*: C2 but Recoin (RX) with explanation interface.
- *Condition 4*: C2 but Recoin (RIX) with interactive interface.

4.3.2 Task Measures. **Relevance**: As the improvement of data quality is the primary goal of Recoin, we wanted to ensure that we understood how each condition affected the change in completeness; independent of the quantity of contributions. Thus, we defined the metric *relevance* as our dependent variable. Relevance is defined as the difference of the completeness values of Recoin before and after a participant added properties to the item.

Recoin Usage: As a recommender system, it is particularly important to understand how each condition (aside from the baseline) affected the number of times Recoin was used directly to add information to an item. This is expressed by the measure *usage* which serves as a dependent variable.

Time: We fixed the time participants can add properties to an item to ensure that our conditions are comparable. The measure *time* serves as our control variable.

Demographics: All study participants were recruited via MTurk. While we assume that the majority of participants are US-Americans, we did not further specify our demographics. Thus, as is typical, demographics were our covariates.

Thank you for adding data to Chris Hadfield's Wikidata entry! We've had a look at the relevance of your additions, and give your performance the grade (A-F):

B

Your additions made Chris Hadfield 25 % more relevant.

To help us improve Wikidata, please answer the following questions. You can use the scales to answer and also provide more specific feedback in the textfields.
After completion, clicking "Submit Questionnaire" at the bottom of the page will enable you to finish the HIT.

1 To what extent were you familiar with Wikidata prior to this HIT?

Not familiar at all ○ ○ ○ ○ ○ Very Familiar

2 To what extent do you understand how your task has been graded? (Please expand)

Figure 5: Questionnaire after adding properties that lead to an increase in relevance of 25%.

Condition	# All Edits	# Recoin Usage	SD Recoin Usage
Baseline	249	-	-
C1	319	61	3.10
C2	382	91	8.56
C3	301	55	3.50
C4	281	71	4.25

Table 2: Number of edits, i.e. contributions in each condition, with the number of Recoin usage and standard deviation.

4.3.3 Self-Report Measures. Upon completion of the task, participants were directed to a self-report page (cp. Figure 5). The page prominently featured a grading of the participant task performance, which was calculated by normalizing the average comparative relevance, i.e. Recoin's assessment, of each contribution per participant. We graded the participants performance in their task (A-F) in order to elicit a reaction to the task even if participants did not notice, use or understood Recoin. We purposefully designed this grading to encourage study participants to reflect on the task, for example a participant may receive the grade F despite many additions to the item.

Furthermore, we included ratings with the four factors from previous research on algorithm awareness [21]: comprehension (5-point Likert scale), accuracy, fairness and trustworthiness (7-point Likert scale) of the algorithmic system. These measures should be strongly correlated according to the procedural justice theory in the related work. Low ratings of all measures, for example, would stem from violated expectations in an outcome [21]. We asked all participants in addition, to expand on their ratings via text-fields for collecting qualitative data as well.

4.4 Hypotheses

For our hypotheses, we were interested in testing the impact of Recoin on data completeness. Having provided participants with equal opportunities to add relevant data to the item *Hadfield*, we

examined whether Recoin improves the completeness of an item or not.

Based on our analysis of the status quo (cp. Section 3.2), we did not expect study participants to actively use Recoin which lead to the following hypothesis: H_1: Using Recoin does not lead to significantly higher relevance in terms of data completeness.

Based on the discussed literature on algorithm awareness, we assume that an user interface that conveys explainability and interactivity of the underlying recommender system leads to higher usage rates: H_2: The interface design of Recoin impacts the number of time participants used Recoin.

Furthermore, we assumed that the effectiveness of algorithm aware designs would be captured most succinctly by the comprehension measure, which would accordingly allow us to distinguish the impact of the RX and RIX designs. Given the results of textual explanation employed in related work (cp. Section 2), we therefore hypothesized that: H_3: A textual explanation of the algorithmic logic leads to higher comprehension than the interactive redesign.

Finally, to gain insights on methodological procedure, we sought to test the experimental self-report measures employed by Kizilcec [21]. According to this research, the self-report measures should exhibit a high degree of correlation (Cronbach's α = 0.83). We therefore hypothesized: H_4: The correlation of self-report measures for textual explanation solutions will equally hold for testing the interactive solution.

4.5 Results

We recruited 21 participants for each condition (n = 105). Overall, we received 1,532 edits (cp. Table 2), with participants in the *C2*-condition providing the most. In the *C4*-condition, our interactive redesign, participants used Recoin most frequently, with more than half (61.09%) of participants adding data via the Recoin interface at least once. This condition also included the most relevant contributions, with a median increase in completeness for the Hadfield item of 21%. The median values of task performance, i.e., received grade and average increase in completeness, as well as the ordinal Likert-scales from the participant self-report, i.e., comprehension, fairness, accuracy and trust, can be seen in Table 3.

We expected only a small amount of qualitative data. However, we found that displaying a grade in the self-report provided a highly effective trigger. Overall, 82 of our 105 participants chose to expand on their self-reported ratings via the provided text fields. This allowed us to probe participant statements for insights on specific subjective perspectives.
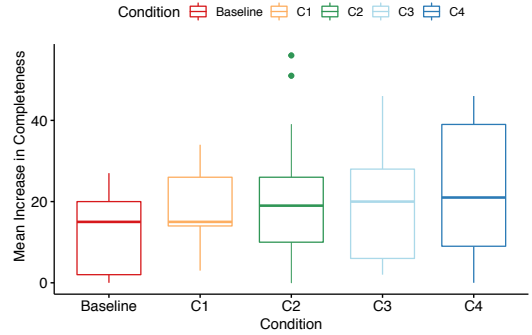
Figure 6: Boxplot of the mean increase of completeness per condition.

In the following, we show the results of our analysis for each hypothesis by using the Kruskal-Wallis test for ordinal data and ANOVA for numerical data. We report the results of the algorithm awareness measures with Spearman correlation tests. Finally, we provide findings of our qualitative analysis of participant statements.

H_1: *Using Recoin does not lead to significantly higher relevance in terms of data completeness.* We reject this hypothesis. An increase of comparative relevance for the Hadfield item is highly dependent on using Recoin at least once ($p_{rel,rec}$ < 0.001). Additionally, when looking at the increase of comparative relevance per participant as a function of the numbers of additions made via Recoin, we can see that the most significant difference occurs around 7 additions made ($p_{rel,numUse}$ = 0.02). This shows that Recoin is highly efficient, as adding a majority of the ten recommended properties should lead to the highest increase in relevance.

H_2: *The interface design of Recoin impacts Recoin usage.* Even though the redesign (C4) slightly outperformed the other conditions in terms of the goal of the set task, we could not find any significant difference for the number of additions made via Recoin between C1, C2, C3, or C4 (p = 0.74). Therefore, we cannot confirm this hypothesis with statistical significance.

H_3: *Textual explanation of the algorithmic logic leads to higher comprehension than the interactive redesign.* We could not find statistically significant differences between ratings of comprehension between condition C4 (RX) and C5 (RIX) ($p_{comp,con}$ = 0.98). Therefore, this hypothesis is not confirmed.

H_4: *The correlation of self-report measures for textual explanation solutions will equally hold for testing the interactive solution.* We had to reject this hypothesis as well. Reacting to the large variance

Condition	Grade	Rel.	Comp.	Fair.	Acc.	Trust
Baseline	C	11	2	4	4	4
C1	C	15	3	4	5	4
C2	C	19	3	5	5	5
C3	B	20	3	4	4	4
C4	B	21	3	6	6	5

Table 3: Median values for (1) task performance: *Grades* dependent on the increase of *Relevance*; (2) self-report: *Comprehension* (1 - 5), *Fairness, Accuracy* and *Trust* (1-7).

Condition	Base-line	C1	C2	C3	C4
Cronbach's α	0.79	0.75	0.67	0.65	0.51

Table 4: Cronbach's α for questionnaire measures across all conditions.

Factor	Comp.	Fair.	Acc.	Trust
Comp.	-	0.19	0.15	**0.33***
Fair.	0.19	-	**0.40***	**0.60***
Acc.	0.15	**0.40***	-	0.16
Trust	**0.33***	**0.60***	0.16	-

Table 5: Spearman correlation coefficients and p-values for C2-C4 for self-report measures *Comp.=Comprehension, Fair.=Fairness, Acc.=Accuracy* and *Trust* with * for $p < 0.05$.
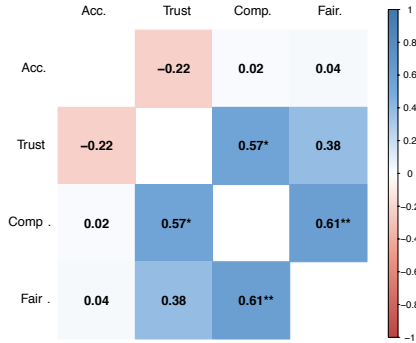
Figure 7: Spearman correlation coefficient matrix for C4 with * for $p < 0.05$; ** for $p < 0.005$.

in C4 (cp. Figure 6), we tested the validity of the questionnaire measures. As opposed to previous research [21], the self-reported measures are not correlated, instead they differ significantly (cp. Table 4). This especially concerns C4, our redesign (RIX), where variance was very high (Cronbach's $\alpha = 0.51$).

As a reaction to the high variance, we conducted Spearman's correlation tests for the ordinal Likert scales used in the participant self-report measures.

The most statistically significant correlative finding from our data is that trust and fairness share a medium to strong positive relationship in our experiment. This is shown across all conditions ($r_{t,f} = 0.65, p < 0.01$), as well as between those wherein Recoin was introduced during onboarding (C2, C3, C4) ($r_{t,f} = 0.60, p < 0.01$) (cf. table 5).

The predominant relationship of trust and fairness was reaffirmed for C2 and C3, the original design and the additional textual explanation respectively, as the strongest and most significant relationship (C2: $r_{t,f} = 0.73, p < 0.01$; C3: $r_{t,f} = 0.63, p < 0.01$).

A further relationship of fairness and accuracy was found for C2 and C3, as a medium positive relationship (C2: $r_{f,a} = 0.52, p = 0.01$; C3: $r_{f,a} = 0.53, p = 0.01$).

When participants used the interactive design of Recoin (RIX in C4), two different relationships emerged (cp. Figure 7). We found that the relationship between comprehension and fairness was strongest ($r_{c,f} = 0.61, p < 0.01$), closely followed by the relationship comprehension and trust ($r_{c,t} = 0.57, p = 0.01$). Surprisingly, the strong relationships found across all conditions were not present for the interactive design (cp. Figure 7).

4.5.1 Qualitative Analysis. Due to the high variance encountered in C4, and the unexpected lack of correlation between our self-report measures, we also expanded our analysis to participant statements. Accordingly, we sampled participant statements in order to probe specific subjective viewpoints. In this section, we showcase some preliminary insights.

Base Trust in Open Knowledge Platforms. A recurring theme, when participants chose to expand on their rating of trust, was a certain *base trust* in open knowledge platform. This occurred even when no explanation element was provided, and also when participants received a poor grade in the condition that did not feature Recoin (Baseline):

> *"Considering there was not a good definition of how we would be judged, it is tough to know if the judging was actually fair or unfair. However, I tend to trust Wikipedia so Wikidata is probably trustworthy."* Baseline-P18; graded *D*

This *base trust* was also extended to the algorithm specifically, as long as it abides by platform standards:

> *"I assume that an algorithm is used to grade the task, in which case I assume that it's free of bias, which is why I do trust Wikidata a good deal when it comes to fairness. Provided, the algorithm itself works as it's supposed to."* C2-P15; rated *Trust* at *6 (High)*

High task efficiency may not indicate algorithm awareness. The qualitative data also suggests that task efficiency in terms of the algorithm does not necessarily indicate algorithm awareness. On the contrary, the only participant that offered a fundamentally accurate account of Recoin received the second lowest grade possible:

> *"My only theory is that it's graded based on the relevance of entries made in regards to his occupation (astronaut) while most of my entries concerned his family, his awards and etc, rather than his activity as an astronaut."* C2-P15; graded *D*

The commentary of a well-performing participant (graded *B*) furthermore suggests that there may be a difference in understanding algorithmic logic and understanding the integration into the algorithmic system:

> *"It seems odd that I would be the one putting in the data and it is grading me considering why couldn't it just put the data in itself if it is accurate enough to grade."* C1-P17; rated *Accuracy* at *6 (Very accurate)*

Finally, and in a similar fashion, a participant formulated the key question they had to the algorithmic system as follows:

> *"I understand that the relevance is graded, I'm not sure exactly how relevance is judged."* C2-P2; rated *Comprehension* at *2 (Low Understanding)*

In summary, the unexpectedly high variance in the C4 condition, combined with the difference in correlative relationships across conditions, as well as our qualitative data, allow us to gather relevant insights for further research. In the next section, we will

discuss limitations to our experiment, before concluding with the contributions as well as the implications for future work.

5 DISCUSSION

First, we found no significant differences between the conditions in terms of average increases in completeness. However, this also suggests that the solution of textual explanation found in related work is not an inherently clear choice for algorithm awareness. This indicates that the design decisions for algorithm awareness are still methodologically unrefined.

Additionally, we sought to understand if our alternative to textual explanation, one taking an interactive and non-declarative approach, could be measured according to the existing self-report measures as suggested by previous research [21]. We found that the measures "Comprehension", "Fairness", "Accuracy" and "Trust" were not equally distributed across our experimental conditions. On the contrary, divergent correlative relationships emerged. The status quo design (R1) as well as the addition of textual explanation design (RX) featured the same strong relationships of trust and fairness as well as fairness and accuracy. In contrast, our redesign (RIX) did not exhibit these relationships, but rather suggested that comprehension was most influential. This was shown in the medium to strong correlation between comprehension and fairness as well as comprehension and trust. We therefore posit that expanding on these self-report measures for algorithm awareness is another, distinct area requiring further research.

Moreover, the qualitative data we gathered also included insightful statements made by our participants. The phenomenon of *base trust* that we encountered in participant statements is relevant for future algorithm awareness studies. If verified, it needs to be taken into account in cases where researchers may wish to abstract from platforms to look at specific problems.

In a broader context, experiments on transparency in algorithmic systems, especially in recommender systems, are frequently undertaken in order to minimize or even eliminate bias. However, as also found by Ekstrand and Tian in experimenting with various recommendation algorithms [9], a complete solution to the problem of bias is improbable. That is to say: bias is inevitable, and is a result of humans and technology interacting. This position is echoed in the work of the philosopher of technology Verbeek, who argues that technology fundamentally *mediates* human relations to a particular "world", i.e., groups of other humans, values, practices etc. [38]. Biases, especially those commonly not aware of, play an instrumental role here. The solution, then, may not be finding the best measure for an elimination of bias, but rather finding the most actionable measure for making bias transparent. Our experimental results align with this assertion insofar that participants had issues with understanding the algorithmic system not on the basis of whether or not something is correctly calculated, but rather who or what has the *agency* for judging the result (e.g., the platform itself, the algorithm as a contained unit, peer review etc.). This, along with a lack of significant differences between conditions, indicates that our intuition to design for an interactive mediation of the community-driven basis for Recoin was useful. Therefore, we posit that promoting algorithm awareness by interactivity is a promising research area.

Our study has a number of limitations that should be considered. As opposed to other work (e.g. [15]), our research focuses on non-technical experts. Furthermore, by recruiting our study participants over MTurk, it can be easily asserted that the demographics of the platform predispose the experiment to cultural bias. Additionally, online experiments in general are limited in two ways. On the one hand, observation of the subtleties of human-technology relations is not possible, such as the non-linguistic ways in which interaction expresses itself and decision-making occurs. On the other, by using MTurk we did not study Wikidata editors, but novices which might have never before come into contact with Wikidata. This means that, while we certainly could infer insights on algorithm awareness and human-technology relations, studying the lived practice of Wikidata editors may reveal other or even contradictory results.

6 CONCLUSION

Our research was motivated by a wish to deepen our understanding of existing design parameters for algorithm awareness. We used the recommender system Recoin, employed in the online peer production system Wikidata, as a use case for our online experiments. In five different conditions, we provided the study participants with a varying degree of explanations and interactivity while using the recommender system. We were able to gather experimental data on the effect of various algorithm aware design measures, and to reflect on the validity of measures used in related work. However, our experiments alone are not yet exhaustive enough for us to reason more substantially about what human awareness means when algorithms are involved. Partly, this is due to the lack of longitudinal, qualitative data gathered from extensive and sustained use of Recoin. The participants of our experiments were predominantly unaware of Wikidata, and the task itself was both brief and controlled in terms of the knowledge that was provided. Wikidata lives and breathes from enthusiasts and domain experts that contribute extensively in their areas of interest. Thus, in future work, we seek to conduct studies that complement these results by probing individual and subjective use over time. This will allow us to understand more deeply, for example, how algorithm aware designs impact the relation between the Wikidata editors and the platform. From such studies, we plan to expand our framework to other use cases. Thereby, we hope to contribute to the urgent need for understanding how the increasingly ubiquitous algorithmic systems shape everyday life for and from the Web.

REFERENCES

[1] Albin Ahmeti, Simon Razniewski, and Axel Polleres. 2017. Assessing the completeness of entities in knowledge bases. In *ESWC*. 7–11.

[2] Ricardo Baeza-Yates. 2018. Bias on the Web. *Commun. ACM* 6 (2018), 54–61.

[3] Vevake Balaraman, Simon Razniewski, and Werner Nutt. 2018. Recoin: relative completeness in Wikidata. In *Wiki workshop @ The Web conference*. 1787–1792.

[4] Yochai Benkler. 2002. Coase's Penguin, or Linux and the Nature of the Firm. *The Yale Law Journal* (2002), 369–446.

[5] Yochai Benkler and Helen Nissenbaum. 2006. Commons-based Peer Production and Virtue. *The Journal of Political Philosophy* 14, 4 (2006), 394–419.

[6] Sebastian Burgstaller-Muehlbacher, Andra Waagmeester, Elvira Mitraka, Julia Turner, Tim Putman, Justin Leong, Chinmay Naik, Paul Pavlidis, Lynn Schriml, Benjamin M Good, et al. 2016. Wikidata as a semantic framework for the Gene Wiki initiative. *Database* 2016 (2016).

[7] Jacobi Carter. 2008. ClueBot and vandalism on Wikipedia. http://www.acm.uiuc.edu/âŁijcarter11/ClueBot.pdf.

[8] Dan Cosley, Dan Frankowski, Loren Terveen, and John Riedl. 2007. SuggestBot: using intelligent task routing to help people find work in wikipedia. In

International conference on Intelligent user interfaces. 32–41.

[9] Michael D. Ekstrand, Mucun Tian, Ion Madrazo Azpiazu, Jennifer D. Ekstrand, Oghenemaro Anuyah, David McNeill, and Maria Soledad Pera. 2018. All The Cool Kids, How Do They Fit In?: Popularity and Demographic Biases in Recommender Evaluation and Effectiveness. In *1st Conference on Fairness, Accountability and Transparency.* 172–186.

[10] Motahhare Eslami. 2017. Understanding and Designing Around Users' Interaction with Hidden Algorithms in Sociotechnical Systems. In *Conference on Computer Supported Cooperative Work and Social Computing (CSCW).* 57–60.

[11] Motahhare Eslami, Amirhossein Aleyasen, Karrie Karahalios, Kevin Hamilton, and Christian Sandvig. 2015. FeedVis: A Path for Exploring News Feed Curation Algorithms. In *18th Conference Companion on Computer Supported Cooperative Work & Social Computing (CSCW).* 65–68.

[12] Motahhare Eslami, Sneha R. Krishna Kumaran, Christian Sandvig, and Karrie Karahalios. 2018. Communicating Algorithmic Process in Online Behavioral Advertising. In *Conference on Human Factors in Computing Systems (CHI).*

[13] Daniel Fleder and Kartik Hosanagar. 2009. Blockbuster culture's next rise or fall: The impact of recommender systems on sales diversity. *Management science* 55, 5 (2009), 697–712.

[14] Daniel M Fleder and Kartik Hosanagar. 2007. Recommender systems and their impact on sales diversity. In *ACM conference on Electronic commerce.* 192–199.

[15] David Gunning. 2016. Explainable Artificial Intelligence. Defense Advanced Research Projects Agency (DARPA).

[16] Aaron Halfaker, R Stuart Geiger, Jonathan T Morgan, and John Riedl. 2012. The Rise and Decline of an Open Collaboration System: How Wikipedia's Reaction to Popularity Is Causing Its Decline. *American Behavioral Scientist* (2012).

[17] Aaron Halfaker, R Stuart Geiger, and Loren G Terveen. 2014. *Snuggle: designing for efficient socialization and ideological critique.* 311–320 pages.

[18] Kevin Hamilton, Karrie Karahalios, Christian Sandvig, and Motahhare Eslami. 2014. A Path to Understanding the Effects of Algorithm Awareness. In *Extended Abstracts on Human Factors in Computing Systems (CHI).* 631–642.

[19] Xiangnan He, Tao Chen, Min-Yen Kan, and Xiao Chen. 2015. Trirank: Review-aware explainable recommendation by modeling aspects. In *CIKM.* 1661–1670.

[20] Lucie-Aimée Kaffee, Alessandro Piscopo, Pavlos Vougiouklis, Elena Simperl, Leslie Carr, and Lydia Pintscher. 2017. A Glimpse into Babel: An Analysis of Multilinguality in Wikidata. In *13th International Symposium on Open Collaboration - OpenSym.* 1–5.

[21] René F. Kizilcec. 2016. How Much Information?: Effects of Transparency on Trust in an Algorithmic Interface. In *Conference on Human Factors in Computing Systems (CHI).* 2390–2395.

[22] Ron Kohavi, Roger Longbotham, Dan Sommerfield, and Randal M. Henne. 2009. Controlled experiments on the web: survey and practical guide. *Data Mining and Knowledge Discovery* 18, 1 (2009), 140–181.

[23] Till Kohli, Renata Barreto, and Joshua A. Kroll. 2018. Translation Tutorial: A Shared Lexicon for Research and Practice in Human-Centered Software Systems. In *1st Conference on Fairness, Accountability and Transparency.* 1–7.

[24] Cliff Kuang. 2017. The Next Great Design Challenge: Make AI Comprehensible To Humans.

[25] Lyons, Kolina S. Koltai, Nhut T. Ho, Walter B. Johnson, David E. Smith, and R. Jay Shively. 2016. Engineering Trust in Complex Automated Systems. *Ergonomics in Design* (2016), 13–17.

[26] Astrid Mager. 2018. Internet governance as joint effort: (Re)ordering search engines at the intersection of global and local cultures. *New Media & Society* (2018).

[27] Claudia Müeller-Birn, Leonhard Dobusch, and James D. Herbsleb. 2013. Work-to-Rule: The Emergence of Algorithmic Governance in Wikipedia. In *C&T '13: Proceedings of the 6th International Conference on Communities and Technologies.*

[28] Cathy O'Neil. 2016. *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy.* Crown/Archetype.

[29] Richard Phillips, Kyu Hyun Chang, and Sorelle A. Friedler. 2018. Interpretable Active Learning. In *1st Conference on Fairness, Accountability and Transparency (Proceedings of Machine Learning Research).* 49–61.

[30] Emilee Rader, Kelley Cotter, and Janghee Cho. 2018. Explanations As Mechanisms for Supporting Algorithmic Transparency. In *Conference on Human Factors in Computing Systems (CHI).* 1–13.

[31] Simon Razniewski, Vevake Balaraman, and Werner Nutt. 2017. Doctoral advisor or medical condition: Towards entity-specific rankings of knowledge base properties. In *International Conference on Advanced Data Mining and Applications.* 526–540.

[32] Johan Redström and Heather Wiltse. 2015. Press Play: Acts of defining (in) fluid assemblages. *Nordes* 1, 6 (2015).

[33] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. "Why Should I Trust You?": Explaining the Predictions of Any Classifier. (2016). http://arxiv.org/abs/1602.04938

[34] Amir Sarabadani, Aaron Halfaker, and Dario Taraborelli. 2017. Building Automated Vandalism Detection Tools for Wikidata. In *International Conference on World Wide Web (WWW).*

[35] Andrew Selbst and Julia Powles. 2018. "Meaningful Information" and the Right to Explanation. In *1st Conference on Fairness, Accountability and Transparency*, Vol. 81.

[36] Phoebe Sengers, Kirsten Boehner, Shay David, and Joseph 'Jofish' Kaye. 2005. Reflective Design. In *4th Decennial Conference on Critical Computing: Between Sense and Sensibility.* 49–58.

[37] Katherine Thornton, Euan Cochrane, Thomas Ledoux, Bertrand Caron, and Carl Wilson. 2017. Modeling the Domain of Digital Preservation in Wikidata. *iPres* (2017).

[38] Peter-Paul Verbeek. 2006. *What things do: philosophical reflections on technology, agency, and design.* Penn State Press, University Park.

[39] Denny Vrandečić and Markus Krötzsch. 2014. Wikidata: A Free Collaborative Knowledgebase. *Commun. ACM* (2014).

[40] Heather Wiltse, Erik Stolterman, and Johan Redström. 2015. Wicked Interactions: (On the Necessity of) Reframing the 'Computer' in Philosophy and Design. *Techné: Research in Philosophy and Technology* 19, 1 (2015), 26–49.

[41] Ellery Wulczyn, Robert West, Leila Zia, and Jure Leskovec. 2016. Growing wikipedia across languages via recommendation. In *Proceedings of the 25th International Conference on World Wide Web.* 975–985.

[42] Cosmas Zavazava, Rati Skhirtladze, Fredrik Eriksson, Esperanza Magpantay, Lourdes Montenegro, Daniel Pokorna, Martin Schaaper, Ivan Vallejo, and David Souter. 2017. *Measuring the Information Society Report.* Technical Report 1. International Telecommunication Union. 170 pages.

[43] Yongfeng Zhang, Guokun Lai, Min Zhang, Yi Zhang, Yiqun Liu, and Shaoping Ma. 2014. Explicit factor models for explainable recommendation based on phrase-level sentiment analysis. In *SIGIR.* 83–92.