

Piercing the Veil: Designs to Support Information Literacy on Social Platforms

JAN WOLFF, Human-Computer Interaction Group, Leibniz University Hannover, Germany

In this position paper we approach problems concerning critical digital and information literacy with ideas to provide more digestible explanations of abstract concepts through interface design. In particular, we focus on social media platforms where we see the possibility of counteracting the spread of misinformation by providing users with more proficiency through our approaches. We argue that the omnipresent trend to abstract away and hide information from users via UI/UX design opposes their ability to self-learn. This leads us to propose a different framework in which we unify elegant and simple interfaces with nudges that promote a look behind the curtain. Such designs serve to foster a deeper understanding of employed technologies and aim to increase the critical assessment of content encountered on social platforms. Furthermore, we consider users with an intermediary skill level to be largely ignored in current approaches, as they are given no tools to broaden their knowledge without consultation of expert material. The resulting stagnation is exemplified by the tactics of misinformation campaigns, which exploit the ensuing lack of information literacy and critical thinking. We propose an approach to design that sufficiently emancipates users in both aspects by promoting a look behind the abstraction of UI/UX so that an autonomous learning process is given the chance to occur. Furthermore, we name ideas for future research within this area that take our considerations into account.

CCS Concepts: • **Social and professional topics** → **Computing literacy**; • **Human-centered computing** → *Information visualization*; *User centered design*.

Additional Key Words and Phrases: learnable interfaces, critical thinking, computer literacy, social media

ACM Reference Format:

Jan Wolff. 2021. Piercing the Veil: Designs to Support Information Literacy on Social Platforms. In *Proceedings of CHI'21 Workshop on Technologies to Support Critical Thinking in an Age of Misinformation (CTAM21)*. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

The surge of misinformation campaigns and conspiracy theories during the COVID-19 pandemic are, in part, symptomatic of a severe lack of critical digital literacy in parts of the online population. We want to propose the argument that this represents just a very visible tip of the iceberg. Ever since home computing became ubiquitous, the main target of interface design has been to initially make computers, then the internet, and by proxy a vast amount of information easily digestible. Yet, as became apparent in the debate on the term *digital native*, mere exposure to digital systems does not imply a deep understanding of the technology at hand [2]. Nevertheless, especially since the widespread introduction of smartphones and social media, computer usage and interconnectedness is on a stable path to become globally ubiquitous. In 2018 for the first time more than half of the global population were using the internet in a continuously rising trend [19]. But this development is not accompanied by a simultaneous in-depth education on digital systems. As a matter of fact, it is mainly driven by the introduction of systems that keep their internals hidden. In part to explicitly avoid a critical examination.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2021 Association for Computing Machinery.

Manuscript submitted to ACM

One visible manifestation of the problems posed by this fact is found in online spaces where users are provided with a seemingly endless influx of information that is highly personalized [5]. Within the context of the *attention economy* [6] this serves to increase the time spent lingering on each platform in order to maximize exposure to advertisements. A negative side-effect arises due to the relative ease with which content can be shared and one's own material can be circulated. The lack of information about each platforms' inner workings therefore runs the risk of leading users to draw wrong conclusions about why content recommendation systems expose them to certain media pieces [3]. Unknowingly, users are subjected to highly sensationalized content that leads to longer interaction with the platform. The ease of sharing, discussing and creating media pieces encourages one's own tendency to sensationalize, as such content is rewarded by being recommended to a larger audience. There are indications that these mechanisms play a substantial role when it comes to problems with radicalization on online platforms through reinforcement of fringe opinions within communities [18]. The personalization of presented content further causes the formation of *filter bubbles* [15], which present an overly skewed view of the global community, leading to incorrect assumptions about the prevalence and shared agreement of expressed opinions.

1.1 Problem Statement

With these previous assumptions in mind, we see the problem of lacking critical media literacy to be twofold. For one, rewarding sensationalized news and opinion pieces disincentivizes unopinionated factual statements, which then leads to a lack of fact checking as it solely reinforces and strengthens previously held viewpoints. But secondly, users are given nearly no agency to assess and explore different opinions once they are part of a filter bubble. One of the few places where exposure to other segregated communities might occur is in globally trending topics, where—once again—sensationalized content is rewarded by being exposed to a larger set of people. Such circumstances are hardly beneficial for healthy and factual discussion, because the most visible opinions are strongly influenced by the prevalent sentiments found in their respective communities and express often incompatible viewpoints. This serves to create environments in which factual evidence is given lesser weight than non-factual content, as it is surrounded by other highly charged statements if it is expressed by another social group.

This area is often neglected when it comes to solving the problem of fake news. We see a chance in giving users more agency about exploring opinions and giving more visibility to the factors that contribute to their selection of content. As a first step, we address the increasing opaqueness of user interfaces found on social media platforms. For this reason we explore examples of designs that—while user-friendly—prevent any engagement with the layers below them.

2 ADDRESSING SOCIAL PLATFORMS

Certainly the design of commercial social platforms is based most strongly on economic incentives. Some of the problems thus stem from *recommender systems*, which present users with content that is vast in scale, high in velocity¹ and automatically selected to cause long periods of interaction with the platform. Visible indications explaining the reasoning behind recommendations are often missing, which is especially concerning given that a main factor for pre-selecting content is the user's predicted interaction with it. Thus, unbeknownst to users, they are presented with controversial information in order to elicit interaction with the content [4]. As this property is not communicated by designs, users who are missing this insight have been found to assume that they are simply presented a balanced selection of the content they are subscribed to [16].

¹As one of the five *V*'s of big data, velocity refers to the rate at which new content is created.

Similarly, grasping the scale of a platform’s user base is not always possible within the bounds of the platform itself. Taking *Twitter* as an example, only few tools are provided that allow gaining an awareness of how representative one’s selection of connected profiles is in relation to the platform’s general population. The only methods available at the time of writing are a global search, the list of trending topics, and an indicator on each profile page showing the list of mutual followers. However, getting any meaningful large-scale information using only these tools is impossible, simply due to the required amount of manual work. Furthermore, knowing about the existence of certain peer groups is a prerequisite for searching them, inverting the directionality that should be provided.

When looking beyond social media, other user interfaces show similar tendencies when analysed. For example, an ongoing trend in desktop interfaces is to be increasingly less verbose about processes that are performed. This poses the risk of leading users to take several aspects presented by the GUI as uncontrollable. Thus, while simplifying the initial process of learning how to interact with such systems, any mental models and experiences are strictly coerced into the set of tasks provided by the UI. No elements indicate that deeper knowledge could be acquired, possibly engendering a mental state similar to the Dunning-Kruger effect [10, 12], as it is impossible to perform an offhand assessment of one’s own expertise. This bears the risk of contributing to factors that make users more susceptible to misinformation campaigns, especially concerning the content encountered on social platforms. Yet, other aspects of digital systems can be affected as well. For example, assumptions made about privacy, security, or malware.

2.1 Related Work

Previous works exist in which methods for increasing awareness towards the underlying data, algorithms, and applications have been explored. Common methodologies either give more context to the content found on different platforms or directly support a deeper understanding of the technological and societal mechanisms at play.

2.1.1 Opinion Space. This work by Faridani et al. [9] addresses the problem of assessing the prevalence and relationship of opinions. In their design, answers to a primary question are visually laid out in a two-dimensional space based on participants’ responses in a supplementary 5-question opinion profile. Through a principal component analysis the 5-dimensional answer space is projected onto a two-dimensional plane. Additionally, comments with a high rating stand out visually, guiding users towards them, while an overview of the diversity in opinions is maintained. In a user study, the interface has been found to increase dwell times in comment sections and to cause greater respect and agreement concerning encountered opinions when compared to a classic, time-sorted list interface. This shows the benefits of giving more agency and tools to users when dealing with discussion platforms.

2.1.2 Balancer Browser Extension. The *Balancer* extension [14] provides a nudge with the intent to diversify the selection of news sources consulted by the user. A prominently positioned browser widget displays a stick figure carrying a block in each hand, with their sizes representing the amount of either conservative or liberal media pages visited. An uneven weight distribution between the two boxes causes the figure to tilt to the corresponding side. In a supplementary study a “small but real” [14, p. 427] effect could be seen in the behaviour of participants with the visualization enabled. The work serves as an example of a method that gives unopinionated feedback concerning the user’s online behaviour and presents a design space that future approaches can take into consideration.

2.1.3 Explanations for Supporting Algorithmic Transparency. In [16] Rader et al. explore methods of providing explicit explanations for the algorithms in use by a given system. A representative imaginary student admission algorithm was used to evaluate the effectiveness of different approaches. Depending on the experiment conditions, either a detailed

visualization of the mapping between input and output (*white-box*) or the algorithm's binary result (*black-box*) was displayed. Additionally, participants could either *interactively* adjust the input parameters to learn their contribution to the result or were presented with only *static* values. A user study showed that, while participants did not self-report a heightened understanding of the system, the increased verbosity did have a positive impact when objectively measured. This fact may stem from the aforementioned Dunning-Kruger effect. People are only able to truly reflect on their knowledge once confronted with the system's internals. Thus, even if they objectively gain knowledge, their self-assessment does not change as they are made aware of further aspects that they do not know about. Such results can serve as a pointer for future research when considering similar explanatory methods.

2.1.4 Privacy Nudges. Almuhimedi et al. explore an idea in which they explicitly nudge users to review their app permissions by telling them, for example, how often their location has been requested by an app [1]. The supplementary permission manager *AppOps* further streamlines the process of permission configuration. A study could show that such nudges, in tandem with the additional management interface, cause significantly more users to review and impose limits on the data that is accessible to applications. The basic idea of these nudges can therefore serve as a basis for communicating even more detailed information to users.

3 EMANCIPATORY DESIGN

We now seek to propose ways in which the issues mentioned in subsection 1.1 can be counteracted. In what we call *emancipatory design* we envision interfaces that, while still providing a low barrier of entry, give more insight into the systems they interact with. Such interfaces can, of course, not display the entirety of each system but only a carefully chosen subset of aspects beyond their level of abstraction. This way, users are made aware of the fact that they could benefit from actively searching out more information about the technology at hand. Going further, such interface designs may directly include possible methods of knowledge acquisition. They may also allow autonomously exploring aspects of the underlying system. Designs necessarily need to coincide with the concepts of minimalism and usability in order to present a step beyond the current state-of-the-art.

Preferably, this approach is to be holistic. While individual systems are likely to benefit from users gaining a more positive disposition towards actively learning more about them, social implications could go far beyond this. At best, users are at least partially aware of all aspects concerning their everyday systems. Mental models then paint a simplified picture that fits reality closely enough to prevent situations in which uneducated guesses have to be made. We expect that, once one veil is lifted, users begin to assess other aspects more critically as well. Such developments could then help to counteract over- or underconfidence and protect users against false advertising, fake news, and other types of misinformation. We believe that the greatest benefit is to be had if such concepts are adopted in a consistent way in many different types of interfaces that users come into contact with.

We assume that several approaches need to be considered. While some interfaces provide opportunities to implement additional interactions used for introducing low-level concepts, others may not. Such cases can arise due to concepts that do not fit well into digestible visualizations or because of platforms that explicitly attempt to hide design choices. Careful analysis has to be conducted when considering the latter, as such approaches will necessarily obstruct the user's workflow within the application itself. Hence, if a positive effect is to be had, the delivery must be kept short and unobtrusive.

3.1 Examples

We will now give some examples for designs that could serve as a basis for future research. First, we present interactions that can be introduced into existing interfaces, followed by designs that are external to the system they aim to explain.

3.1.1 Exploring Filter Bubbles. As previously explained, one key factor of increased exposure to misinformation are *filter bubbles* [15], which cause an apparent amplification of otherwise fringe opinions. We therefore consider designs in which the existence of such bubbles and the partisanship within is given increased visibility. To avoid possible negative implications, care should be taken to perform visualizations as objectively as possible, because they should only serve as a tool that allows users to gain an understanding of the otherwise difficult to comprehend network graph between profiles. Furthermore, selective hiding or ranking of such bubbles should not occur, as existing research suggests exposure to opposing viewpoints to be an aide in fostering accurate beliefs [11].

One approach to study in future research could combine the design of *Opinion Space* with an automated positioning of profiles via the social graph. Faridani et al. suggested the need for an easy to understand scale in the visualization, for which they proposed displaying selected public personalities as landmarks [9]. When considering the application of this idea on *Twitter*, this design can easily be translated by selecting well-known profiles and always displaying those as landmarks in the graph. Such visualizations may be enabled manually through an action presented as *zooming out*. Users are first presented their regular, self-curated feed, from which they can zoom out to see a slightly simplified overview of the whole network using selected profiles as landmarks. Given such a view, one could imagine being able to effectively *zoom in* and experience feeds or discussions from the perspective of another bubble, which can allow people to have an insight into sentiments outside their peer group.

3.1.2 Explaining Recommender Systems. While the aforementioned idea plays strongly on social groups, another concept can be applied to recommendation systems. Considering video platforms such as *YouTube*, automated content moderation is largely based on the selection of videos the particular user has watched in the past [5]. In order to better represent the existence of clusters in such a system, users should be given the chance to understand what selection of content contributed most to each recommendation. In a possible solution, the list of recommendations and a history of content viewed in the past could be combined. To facilitate quickly assessing the information at hand, related entries in the list of past content would be grouped. Then, for each recommended item the connections to either individual entries or whole groups can be displayed.

While such visualizations and explanations can never paint a complete picture and go into deep detail, a positive effect may still be had even with some simplification. Edwards and Veale suggest that, even when given only superficial information about machine learning practices, people are sufficiently able to picture a “model-of-a-model” [8] so that meaningful conclusions about the system can be drawn.

3.1.3 Auxiliary Context Aware Learning. As already mentioned, factors can be at play that do not permit such measures to be directly embedded in an application’s interface. Hence, approaches need to be considered in which information stems from applications external to the one that is to be explained. These ideas must go hand in hand with context awareness, in order to be adequately effective and not be perceived as a nuisance. Behavioural analysis should be applied to pinpoint moments during which users are not actively performing a task, but in which they are available to interact with a knowledge gathering application for a short time. While research supporting such systems exists [7], an additional dimension to support such applications could be the users’ level of knowledge about the system. In such cases, the system could monitor whether intermediary users operate on a suboptimal plateau of performance. If they do,

new interactions or concepts could be introduced step-by-step. However, further research is required to assess if—and to what degree—such information can be gathered automatically.

One holistic approach could make use of *gamification*. Given a *Jeopardy*-style selection of topics and corresponding pieces of information, users could be motivated to gather knowledge about a wide assortment of topics. In fact, applications with comparable approaches already exist. One example would be the general knowledge gathering app *getucated*², which focuses mostly on *history*, *philosophy*, and *culture*. Combined with context awareness, such applications could put an emphasis on areas that are relevant to the specific user and in which they display a lack of knowledge. In order to further enhance the appeal of the idea, one could evaluate whether a personification of the knowledge gathering application via a digital assistant has a positive impact on user retention.

3.1.4 Judgemental interfaces. Given the fact that information given by external sources is free to emphasize negatives, more strongly worded language can be employed. One could, for example, measure the time users spend clicking through video recommendations on *YouTube*. After a while, a notification could be shown explaining that the recommended content is carefully chosen to maximise the time spent on the platform in order to expose users to as many ads as possible. Such messages should be shown regardless of the particular content on screen to prevent users from assuming a bias concerning certain media pieces. Yet, in case users become entrenched within single topics, such notifications could serve as a nudge to take a step back from the content.

Going further, one could imagine texts that take into account even more information from, for example, *Google*'s advertising profile of users. This information could then serve as a basis for messages such as: “Google knows you are in a relationship, did you?” or “Do you visit this restaurant often? Google thinks so.” This idea is related to the aforementioned *privacy nudges* [1].

4 DISCUSSION

The above-mentioned ideas reflect a subset of possible interface designs that take emancipatory approaches into account, with a focus specifically on social platforms. Therefore, mainly their potential social implications are considered, with the emancipation of users regarding technical aspects being less of a goal in this context. The particular selection of examples represents ideas which we deem feasible in their realization. In the first two examples we consider designs that would—beyond a prototyping stage—require cooperation with social platforms themselves. The other two designs allow to relax this requirement, as their basic functionality stems from locally analysing user behaviour and serving an assisting purpose. We thus assume those examples to present a valid starting point, with their prototypical evaluation being feasible for first studies and evaluations, due to their concepts being similar in scope to existing related work.

We envision that, once design approaches include the concept of presenting a digestible look behind the veil of systems they abstract away or actively offer an opportunity to seek out information, novice users are given an incentive to learn about aspects concerning the systems they use. They do not need to become experts, they just need to gain enough knowledge to develop an increased ability to critically assess digital systems. We see the chance that such developments can provide users with enough agency to overcome adversarial practices that prey on their lack of critical knowledge. For social platforms this should, at best, manifest in an awareness for the existence of social bubbles, the aim of automated recommendation systems and the impact these factors have on the selection of content. All aspects combined have the potential to counteract increasing segregation into sub-communities within platforms and thus factors that play a significant role in susceptibility to fake news. Going further, other areas of personal computing can

²<https://getucated.de>

benefit from an increased awareness as well. For example, an increased resistance of users towards data harvesting practices or premature obsolescence of otherwise functional systems.

Notably, our ideas do not directly involve automated fact checking or elements from *inoculation theory*. Rather, we see part of the reason for the success of online misinformation campaigns in a manifestation of *learned helplessness* [13] due to increasingly concise interface designs in which users are taught to act passively. We theorize that users extend newly gained critical thinking skills concerning platforms and devices towards the content presented in them. In the same manner, we argue against automated methods of confronting users with content sourced from outside their filter bubble. However, exposure to only partisan sources reinstating one's own world-view and increased polarization accompany each other [17] and fostering communication between boundaries of in-groups has been proposed as beneficial [20]. In this, we see an excellent example where increasing the visibility of such groups and providing tools for people to broaden their scope themselves is more beneficial compared to the automated counterpart, in which we see a hazard of users assuming malice. We see the possibility of passively making people aware of their bubble's scope and its true size to be a viable option in preventing social segregation in online spaces worthy of further research, whereas relying solely on automated techniques merely shifts the problem space. In particular we see aspects of inoculation theory already being used to instil beliefs in conspiracy theories, such as nurturing a distrust in reputable sources or preemptively giving non-factual counterarguments to any statement that would disprove the theory³.

We see the chance that, should such design ideas become more widespread, their positive effects transcend their respective interface. Having formed a correct mental model of one system and being made aware of non-expertise in certain areas, users are able to carry the same understanding onto other applications. This is especially important in cases where negative implications are deliberately hidden.

Finally, we want to stress that maintaining a low bar of entry for everything digital always needs to be pursued. Our ideas merely serve to prevent a stagnation concerning the level of insight people can have into their systems. Going further, measures towards preserving universal ease of use when employing such design approaches are an important cornerstone. Interfaces themselves must not become more difficult to understand or use, regardless of any factors that users may be affected by. In the same vein, the usefulness of any auxiliary information that is presented must not be lost on users.

5 FUTURE RESEARCH

The realm of users that are neither novices nor experts remains largely underexplored in current HCI research. Acquisition of deeper insight is often externalized to educational facilities or only made possible by actively searching out information. Integrating a presentation of the presence of lower levels of abstraction directly inside interface designs is a novel concept for which we assume a large body of research needs to be conducted.

To conclude, we list a number of research questions that we deem worth exploring as an initial starting point:

- Can moments be inferred in which both the application context and a user's mental state allow teaching short bits of information?
- To what extent can the functionality of lower layers be introduced via interface design?
- Are users incentivised to actively seek out information about the systems they use, if they are given a digestible subset of information beyond their skill set?
- What mental models do novice users have about their technologies?

³Instilling distrust in mainstream media and scientific sources is a cornerstone in the *QAnon* conspiracy belief. Similarly, *flat earth* advocates place a large emphasis on distributing incorrect evidence that is supposed to preemptively prove the world-view.

- Are there ways to extract mental models by analysing computer usage and inferring a user's level of knowledge to detect false assumptions?
- Do users significantly change their behaviour concerning their selection of preferred content when made aware of different filter bubbles?
- What roadblocks need to be overcome concerning the inclusivity of these approaches? Are there relevant cultural differences? To what extent do approaches need to be changed to accommodate for disabilities?

REFERENCES

- [1] Hazim Almuhammedi, Florian Schaub, Norman Sadeh, Idris Adjerid, Alessandro Acquisti, Joshua Gluck, Lorrie Faith Cranor, and Yuvraj Agarwal. 2015. Your Location Has Been Shared 5,398 Times! A Field Study on Mobile App Privacy Nudging. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (Seoul, Republic of Korea) (*CHI '15*). Association for Computing Machinery, New York, NY, USA, 787–796. <https://doi.org/10.1145/2702123.2702210>
- [2] Sue Bennett, Karl Maton, and Lisa Kervin. 2008. The 'digital natives' debate: A critical review of the evidence. *British Journal of Educational Technology* 39, 5 (2008), 775–786. <https://doi.org/10.1111/j.1467-8535.2007.00793.x>
- [3] Taina Bucher. 2017. The algorithmic imaginary: exploring the ordinary affects of Facebook algorithms. *Information, Communication & Society* 20, 1 (2017), 30–44. <https://doi.org/10.1080/1369118X.2016.1154086>
- [4] Taina Bucher, Anne Helmond, et al. 2017. The affordances of social media platforms. *The SAGE handbook of social media* (2017), 233–253.
- [5] Paul Covington, Jay Adams, and Emre Sargin. 2016. Deep Neural Networks for YouTube Recommendations. In *Proceedings of the 10th ACM Conference on Recommender Systems* (Boston, Massachusetts, USA) (*RecSys '16*). Association for Computing Machinery, New York, NY, USA, 191–198. <https://doi.org/10.1145/2959100.2959190>
- [6] Thomas H. Davenport and John C. Beck. 2001. The Attention Economy. *Ubiquity* 2001, May (May 2001), 1–es. <https://doi.org/10.1145/376625.376626>
- [7] Tilman Dingler, Markus Funk, and Florian Alt. 2015. Interaction Proxemics: Combining Physical Spaces for Seamless Gesture Interaction. In *Proceedings of the 4th International Symposium on Pervasive Displays* (Saarbruecken, Germany) (*PerDis '15*). ACM, New York, NY, USA, 107–114. <https://doi.org/10.1145/2757710.2757722>
- [8] Lilian Edwards and Michael Veale. 2017. Slave to the algorithm: Why a right to an explanation is probably not the remedy you are looking for. *Duke L. & Tech. Rev.* 16 (2017), 18. <https://doi.org/10.2139/ssrn.2972855>
- [9] Siamak Faridani, Ephrat Bitton, Kimiko Ryokai, and Ken Goldberg. 2010. *Opinion Space: A Scalable Tool for Browsing Online Comments*. Association for Computing Machinery, New York, NY, USA, 1175–1184. <https://doi.org/10.1145/1753326.1753502>
- [10] Shirley Gibbs, Kevin Moore, Gary Steel, and Alan McKinnon. 2017. The Dunning-Kruger Effect in a workplace computing setting. *Computers in Human Behavior* 72 (2017), 589 – 595. <https://doi.org/10.1016/j.chb.2016.12.084>
- [11] William Hart, Dolores Albarracín, Alice H. Eagly, Inge Brechan, Matthew J. Lindberg, and Lisa Merrill. 2009. Feeling validated versus being correct: A meta-analysis of selective exposure to information. *Psychological Bulletin* 135, 4 (2009), 555–588. <https://doi.org/10.1037/a0015701>
- [12] Khalid Mahmood. 2016. Do People Overestimate Their Information Literacy Skills? A Systematic Review of Empirical Evidence on the Dunning-Kruger Effect. *Communications in Information Literacy* 10, 2 (2016), 199–213. <https://doi.org/10.15760/comminfolit.2016.10.2.24>
- [13] Steven F. Maier and Martin E. Seligman. 1976. Learned helplessness: Theory and evidence. *Journal of Experimental Psychology: General* 105, 1 (1976), 3–46. <https://doi.org/10.1037/0096-3445.105.1.3>
- [14] Sean Munson, Stephanie Lee, and Paul Resnick. 2013. Encouraging Reading of Diverse Political Viewpoints with a Browser Widget. *Proceedings of the International AAAI Conference on Web and Social Media* 7, 1 (Jun. 2013). <https://ojs.aaai.org/index.php/ICWSM/article/view/14429>
- [15] Eli Pariser. 2011. *The Filter Bubble: What the Internet Is Hiding from You*. Penguin Books.
- [16] Emilee Rader, Kelley Cotter, and Janghee Cho. 2018. Explanations as Mechanisms for Supporting Algorithmic Transparency. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) (*CHI '18*). Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3173574.3173677>
- [17] Natalie Jomini Stroud. 2010. Polarization and Partisan Selective Exposure. *Journal of Communication* 60, 3 (08 2010), 556–576. <https://doi.org/10.1111/j.1460-2466.2010.01497.x>
- [18] Stuart A. Thompson and Charlie Warzel. 2021. *Opinion | They Used to Post Selfies. Now They're Trying to Reverse the Election*. <https://www.nytimes.com/2021/01/14/opinion/facebook-far-right.html> (Retrieved 02/05/2021).
- [19] International Telecommunication Union. 2018. ICT Indicators (Edition 2018/1). (2018). https://doi.org/11.1002/pub_series/dataset/64cb0e71-en
- [20] Sarita Yardi and Danah Boyd. 2010. Dynamic Debates: An Analysis of Group Polarization Over Time on Twitter. *Bulletin of Science, Technology & Society* 30, 5 (2010), 316–327. <https://doi.org/10.1177/0270467610380011>