HOW RECOMMENDER SYSTEMS EXPLOIT AFFECTABILITY:

AN ESSAY ON LIBERTY IN THE FACE OF OPERANT BEHAVIOR MANIPULATION

Sofia Romão[1]

Abstract: This essay aims to investigate the relationship between human users and algorithmic recommender systems in terms of liberty, autonomy and self-determination by the assemblage of the notions of Transparency and Affectability in Denise Ferreira da Silva as well as Behavioral Variability and Operant Behavior in Burrhus Frederic Skinner. Between the optimization of the exploitation-exploration tradeoff and economic incentives underlying the use of recommender algorithms, the possibility of user's biases and preferences being exploited for profit is discussed as a matter of self-determination. At the same time, it is proposed that the promotion of algorithmic literacy may not be sufficient for true freedom of users of such technologies, taking into account the critique of the *Transparent "I"* as described by Denise Ferreira da Silva. Ultimately, regulatory measures are presented as a more interesting approach.

Keywords: Freedom, Transparency, Affectability, Human Variability, Operant Behavior

Resumo: Este ensaio tem como objetivo investigar o relacionamento entre usuários humanos e sistemas de recomendação em termos de liberdade, autonomia e autodeterminação pela assemblagem das noções de Transparência e Afetabilidade em Denise Ferreira da Silva conjuntamente de Variabilidade Comportamental e Comportamento Operante em Burrhus Frederic Skinner. Entre a otimização do custo de oportunidade 'exploitation-exploration' e incentivos econômicos sobre o uso de algoritmos de recomendação, a possibilidade de que os vieses e preferências do usuário sejam explorados para lucro é discutida enquanto questão de autodeterminação. Ao mesmo tempo, propõe-se que a proposta de alfabetismo algorítmico pode não ser suficiente para uma verdadeira liberdade de usuários dessas tecnologias, levando em consideração a crítica ao *"Eu" Transparente* como descrito por Denise Ferreira da Silva. Por fim, medidas de regulamentação são apresentadas como uma abordagem mais interessante.

Palavras-chave: Liberdade, Transparência, Afetabilidade, Variabilidade Humana, Comportamento Operante

[1] Mestranda em Filosofia pelo PPGFIL/UNB.

Recommender algorithms serve a crucial purpose as the informational output and recording grows with the integration allowed by the internet. The sorting of information is what protects the massive volumes of data from becoming utterly useless, as accessibility shrinks in accordance to this growth. These programs, then, make an evaluation based on certain parameters determining which information is to be shown. As a foundational cornerstone of a massive market, this technology has transformed in ways that optimize these evaluations and, unsurprisingly, adapted to economic incentives. Not only would recommender algorithms be vital to the use of the internet, but they would also be vital to a new paradigm for advertising.

As market and technology develop, recommender algorithms start to take into account more and more data collected from their users as a means of refining its parameters, becoming more precise on what could be considered relevant information for the specific user, which enables targeted advertisements. Along with IP address, gender, age and other data, these programs take into account user behavior as relevant information. Examples of parameters informed by human behavior will be discussed later.

The very nature of recommender systems gives rise to many fruitful discussions. The ethics regarding what data should be collected, what uses of such information are to be accepted, for how long they should be held among many others come instantly to mind. This work is interested in the predictive nature of these programs. For a recommender system to be effective, it needs to be able to predict if the information to be shown is useful to the user or not, or to what degree it is. As the user's behavior is evaluated in order to maximize the predictive power of the tool and, due to the economic model around it, profit, this essay's discussion arises. With the existence of behavioral psychology as a field and the reported use of it in influencing user behavior, what does it mean for one's self-determinacy, liberty and dignity to be subject to these practices?

In the next section, the discussion will center on how recommender systems interact with human users, presenting the explore-exploit tradeoff and the concept of human variability. Then, a brief examination of the state of the art on psychology-aware recommender systems is made to consider the possibility of exploiting and strengthening user biases. Next, liberty as self-determination is discussed through the works of Denise Ferreira da Silva and Burrhus Frederic Skinner. Finally, algorithmic literacy and regulation are discussed as possible remedies to the prospect of digital vulnerability.

The extraction of data and its uses for the purpose of recommendation is affected by what is called the exploration-exploitation tradeoff, a particular choice that developers have to consider when designing their recommender programs as they exist currently. This consists of the evaluation of whether the program should prioritize showing recommendations which have higher user ratings (safe options in terms of predictability) or novelty in information (options that are more likely to be unknown to the user but are nonetheless within the user's preference landscape) (YANG, 2023, p. 4). In their work, Yang et al. propose a theoretical model capable of representing the tradeoff in terms of the metrics of *predictive accuracy* and *predictive probability*. As it is presented, the algorithm which prioritizes the exploitation side of the tradeoff should present high recommendation accuracy and low predictive accuracy, and the reverse is true for prioritizing exploration (YANG, 2023, p. 7). Both these metrics are influenced by the user's base rate, which is the user's "inherent bias to find any recommendation relevant" (YANG, 2023, p. 3). They go on to elaborate that a user with an extreme base rate will probably find almost any item relevant or not in the sense that this metric can mask both the recommendation and predictive accuracy of the program. With a short range of responses (either totally relevant or totally irrelevant), it is easier to predict than the responses of users with other base rates (YANG, 2023, p. 6).

With such parameters, Yang et al.'s model was able to study the extent of this tradeoff in terms of optimizing recommendations in order to maximize both metrics together, considering different levels of complexity of the decision function and base rate variation (YANG, 2023, p. 29). Of particular interest to the current discussion is the finding that varying base rates are responsible for strengthening "the rate at which the highest achievable recommendation accuracy dampens with increasing complexity" (YANG, 2023, p. 28). Within their work, said complexity (of a user's decision function) is related to how quickly the recommender system learns and is able to predict the user's preference (YANG, 2023, p. 3). Although the authors recognize that their results can also be interpreted as the result of biased base rates (YANG, 2023, p. 28), the understanding that user behavior can be categorized as more or less obstructive to the severeness of the tradeoff and, therefore, the optimization of such algorithms, leads to a realization: human variability is a limiter to predictive and recommendation accuracy.

While this is to be expected, the existence of an incentive to maximize these metrics leads to the conclusion that manipulating such a threshold, if possible, is under consideration. Such an incentive, as was alluded to before, is economic.

Zuboff, in her work on Surveillance Capitalism, describes the history of the development of the market around behavioral surplus and its conversion into predictive products, built around both extraction and prediction economic imperatives. The first consists of the incentive to record human behavior in terms of scale (ZUBOFF, p. 126), while the latter focuses on the quality of prediction products, in terms of depth of the information gained from behavior (ZUBOFF, p. 192-193). The second is achieved through algorithms defined under the umbrella "Psychology-informed Recommender Systems". These can have several strategies. One is to build algorithms around models from cognitive psychology (Cognition-inspired Recommender Systems). Examples of what these models can take into account are cognitive processes, memory, attention, reasoning and competence (LEX et al., 2021, p. 18). Another strategy is implemented on "Personality-aware Recommender Systems", which takes into account persistent user characteristics correlating to preferences (LEX et al., 2021, p. 31,32). Lastly, there is the strategy used in "Affect-aware Recommender Systems". These focus on the user's mood and emotions (LEX et al., 2021, p. 40).

As algorithms become psychology-informed, Zuboff argues, economic imperatives act in unison with them in order to produce behavioral modification. This is not to say that behavioral norms are imposed, but that economically beneficial behavior is produced (ZUBOFF, p. 195). The review on the state of the art concerning psychology-informed recommender systems by Lex et al. gives the caveat, though, that current technology typically cannot control decision biases. On the other hand, they themselves recommend that the user be aware of such biases in order to make more informed decisions (LEX, 2021, p. 54). This leads to the understanding that exploiting the decision biases the user already has is possible.

Consumer vulnerability is a well explored field of study. As Helberger et al. argue, the personalization of persuasion strategies enabled by digitalization exacerbate vulnerability, even reaching persons that wouldn't be considered vulnerable in other markets, precisely because of the possibility of discovering and exploiting individual biases, weaknesses, preferences and needs (HELBERGER, 2021, p. 180).

Facing such technologies and practices, one ought to ask themselves what is to be done. On the one hand, to be made to act in accordance with a market's interest seems to violate liberty. Isn't an individual supposed to determine their own actions? On the other hand, if being made to act in such a way doesn't contradict one's already held beliefs, biases and preferences, what is this violating exactly?

In the next section, the discussion will center on the concept of self-determination as liberty. For that end, the assemblage of some works of Denise Ferreira da Silva and Burrhus Frederic Skinner will be deployed to put into consideration the role of affectability on the matter.

> The struggle for freedom and dignity has been formulated as a defense of autonomous man rather than as a revision of the contingencies of reinforcement under which people live (SKINNER, 1973, p. 124).
>
> Treating a person on the basis of a presumed membership of a "group" (stereotyping) or as an object of knowledge [...] is an assault on their *uniqueness* because it expresses a disregard for interiority. More importantly, under the interpretive moment, because the gesture obviates Subjectivity, it robs the person of the qualities that tell that the Human is singular among other existing things, that is, the qualities of *liberty* and *dignity* it acquires in its ethical figuring as the *transparent* (self-determined because self-knowing) *I* (SILVA, 2022, p. 102).

Both of these passages are highlighted on the grounds that each of them, in their authors' respective fields and theoretical frameworks, go on to revise the commonly held presuppositions regarding matters of freedom and self-actualization. While Silva goes on to describe the history of the notion of the *transparent I* and ultimately rejects it as a means of fighting racial subjugation, Skinner separates the figure of the *autonomous man* from the struggle for freedom and dignity.

While they are not the most usual pairing for a theoretical framework, their works on black feminism and behaviorism have this convergence in critiquing the unaffectable, transparent and autonomous man, although through very different strategies, goals and, ultimately, conclusions.

Silva's description of the *transparent I* places it as a construct of the post-Enlightenment episteme. In the turning point from nineteenth to the twentieth century paradigms, the Human is rescued from the possibility of being fully uncovered, and therefore, determined, by scientific reason. In this sense, the human body and its organic functions are taken to not be the causes of human (mental) attributes (SILVA, 2022, p. 90).

To be transparent is to be capable of self-determination. This is in contrast to the affectable, which is determined by external influences. As an example, Silva posits how, in gender studies, the concept of patriarchy is unable to guide action in a political position to be inhabited by the dispossessed black mother. That is because, while the former assumes a woman

capable of deciding and enacting her own desires, the latter is an outer-determined subject, due to her biological, cultural and social position (SILVA, 2007, p. 265-266).

Alternatively, Skinner describes how, traditionally, liberty is defined as the absence of aversive control, with an emphasis on one's state of mind associated with pursuing one's wants (SKINNER, 1973, p. 36). The behaviorist readily points to the shortcomings of such an understanding: would-be controllers' use of non-aversive control. A person may feel free and get rewarded (positively reinforced) in a system that ultimately is aversive to them (SKINNER, 1973, p. 37).

In that sense, freedom is then defined not as a feeling, but as a matter of contingencies of reinforcement (SKINNER, 1973, p. 42). This distinction is important because one's feelings towards being controlled is not the same as knowing if said control is aversive or not. To the benefit of clarity, it is also important to understand contingencies of reinforcement as they are described by Skinner and their role in controlling behavior.

Skinner proposes a kind of probabilistic determinism in regard to operant behavior, as he understands that "seldom does any environmental condition 'elicit' behavior in the all or-nothing fashion of a reflex; it simply makes a bit of behavior more likely to occur" (SKINNER, 1973, p. 97). To change one's behavior, then, is to alter the likelihood of it occurring, and this means that the contingencies of reinforcement were altered.

Through this understanding, the author makes a criticism of traditional defenses of freedom in the sense that changing another's mind (in the sense of persuasiveness) is condoned because it is ineffective in changing behavior (SKINNER, 1973, p. 97). His argument is that, by supposedly leaving the control to the individual (to decide, in this case), we are shifting from one mode of control to another. That is because contingencies are still affecting that person and, by masking them, "autonomous man survives to be credited with his achievements and blamed for his mistakes" (SKINNER, 1973, p. 97).

This point seems to converge with Silva's understanding of the *transparent I* and, though Skinner posits his criticism of the autonomous man arguably strictly through the lenses of scientific reason, it is in the interest of this work to evaluate the degree of compatibility between the concepts here assembled for analysis.

Silva's criticism of gender studies could be presented through this reading of Skinner as such: the assumption of transparency leads to the masking of the contingencies of reinforcement related to the double affectability to which the economically dispossessed black

woman's desires are mediated by while her body is made available to a transparent male desire. She is not made to act as a white woman would as a subject separated from the public domain- but rather as an object of public policy (SILVA, 2007, p. 266).

The racial, argues Silva, is necessary to the preservation of the *transparent I* to the extent that transparency is created and placed specifically in post-enlightenment Europe, in which supposedly subjects reach rational thought that determines oneself, as opposed to the affectable consciousness which is determined by its surroundings (SILVA, 2007, p. 116).

Through this questioning of the correlation made between the *transparent I*, autonomous man, self-determination and, therefore, freedom, often discussed remedies related to the problem of exploitation and vulnerability in digital spaces begin to carry some nuance. Getting back to the question previously left open: what exactly is violated when one is made to act according to another's economic interest without contradicting one's already held beliefs, biases and preferences? Given the framework proposed, the user's liberty is what is at stake in the sense that they are put (intentionally) in a situation where the probability of economically advantageous behavior is higher. This is because even though they are acting according to their conditions of reinforcement, if we don't presuppose one's transparency, odds are, they couldn't be expected to act in any other way.

The exploitation is not grounded in some vague idea of mind control. It is the creation of a state of vulnerability designed specifically for the user at hand. Skinner discusses an example of how non-aversive measures can have the same effect as coercion:

> In the 1930s it seemed necessary to cut agricultural production. The Agricultural Adjustment Act authorized the Secretary of Agriculture to make 'rental or benefit payments' to farmers who agreed to produce less - to pay the farmers, in fact, what they would have made on the food they agreed not to produce. It would have been unconstitutional to compel them to reduce production, but the government argued that it was merely inviting them to do so. But the Supreme Court recognized that positive inducement can be as irresistible as aversive measures when it ruled that 'the power to confer or withhold unlimited benefit is the power to coerce or destroy'. The decision was later reversed, however, when the Court ruled that 'to hold that motive or temptation is equivalent to coercion is to plunge the law into endless difficulties'. We are considering some of these difficulties" (SKINNER, 1973, p. 42).

Even if the fact that what is being manipulated is the probability of behavior means that the person *could* be expected to act otherwise, it is still pertinent to consider what it

means to purposely weigh these odds. Complete control of the contingencies surrounding one's behavior is not available to any control tactic. Therefore, to dismiss this kind of exploitation because there is no absolute certainty of its result is to also dismiss the role of violent coercion by the same logic.

Next, the often-considered solutions of algorithmic literacy and regulation are discussed through the lens of the proposed assemblage of conceptual tools.

While it is useful to point out the effects of algorithmic literacy as a means for the individual to protect himself from bad actors, such a solution (by itself) seems to have most of the weaknesses just outlined as a means of understanding and preserving freedom. Being aware of how algorithms function and the related companies' intents helps the individual to make better decisions, but that in itself leaves them susceptible to other contingencies. Also, it is naive to consider the possibility of erasing all contingencies of reinforcement for there is not an objective rational standard to achieve in doing so. On top of all that, there is the problem of maintaining the divide between those with such literacy and those dispossessed of it, indicating this solution's inability to engage with the racial.

But to just say that defeats the purpose of actually engaging with the practice. Considering just these limitations, it is difficult to see algorithmic literacy as an all-encompassing solution. It does require something else. What it brings to the table, though, is the prospect of what a good relationship with recommender systems looks like.

In the aforementioned work on the state of the art concerning Recommender Systems by Lex *et al.*, the concept of Nudging is presented. Defined explicitly as the act of influencing human behavior via suggestions towards choices that are: (a) in the users' and societies' interest; (b) don't limit any options or; (c) don't significantly change people's economic incentives (LEX, 2021, p. 52). Studies vary as some introduce techniques that exploit user biases to incentivize beneficial behavior such as using environmentally friendly transportation or and eating healthy foods (LEX, 2021, p. 53). Other works focus on detecting such biases in recommendation scenarios (LEX, 2021, p. 54). Finally, some discuss ethical concerns around the very concept of nudging, out of which grew the competing framework of *boosting*: the employment of systems that improve a user's competencies in decision making (thus altering cognitive biases) instead of nudging them (LEX, 2021, p. 55).

When considering the proposition of algorithmic literacy, implementing it through recommender systems themselves is not the only option, though the discussions around nudging

are revealing of the technology's shortcomings in this task. Considering a person's contingencies of reinforcement, to incentivize them to depend on things such as popularity bias to do the 'correct' behavior is to maintain a state of vulnerability. These practices can, to some extent, serve as a way to introduce and educate the user on what are beneficial practices, but they don't rely on the user understanding the biases exploited in doing so. Boosting, having such a sensitivity, has the possible benefit of equipping the person with countermeasures to the attempt of exploitation in other digital settings, which would arguably prevent aversive consequences.

Some studies on news literacy applied to digital settings give further insight to both the vulnerabilities people are subjected to, and also on possible tools to protect oneself.

While interviewing young people from a diverse set of countries, Nee observed their use of messaging services and visual platforms as news sources. She found that most respondents report using both social media and mass media for news, while those who said they get news from Instagram do so mostly by following mass media organizations (NEE, 2019, p. 7). On the same study, it was observed that participants evaluated the credibility of a post and decided to read it based on the comments of other users and who posted or shared the content, while being more likely to find the post believable if it aligns with their own political opinions (NEE, 2019, p. 10).

On a similar note, Swart set out to investigate how young people access, evaluate and engage with news on social media. In an attempt to deviate from the predominant focus on college and university students the field tends to have, Swart chose a focus group of Dutch students at schools for vocational education. The interviews found that news would reach the users incidentally, through face-to-face conversations, peers in social media and algorithmic recommendations (SWART, 2023, p. 510). Another finding is that there are considerable gaps in algorithmic literacy, which is presented as knowledge on the programs' functioning, as some interviewees were unaware that these platforms select content based on their behavior (SWART, 2023, p. 512).

Lastly, Jylhä, Hirvonen and Haider also dwell on youth's information practices but have a focus group of 15- and 16-year-olds and set out specifically to understand the mediation of algorithmic recommendations in these practices (JYLHÄ, 2024, p. 30). One of their findings is how the algorithms invite the user to delegate thinking and decision making to its recommendations, while also reporting how the interviewees framed the experience of actively

searching as ‘wasting time’ (JYLHÄ, 2024, p.32). An interesting consequence of the ease of recommendations is that, even if those can get repetitive to the user, the search for more heterogeneous recommendations has to be active and requires more of the user (JYLHÄ, 2024, p.36). The incentives at play seem to narrow user behavior by facilitating the passive acceptance of recommended content, shaping the previously discussed interest in lowering human variability for optimization of the predictive power of the platform.

Similarly to Nee, they also found that the participants tend to base their trust on the contents of an information based on who posts it but, more specifically, that the first search result of a query is to be found reliable because of the trust put on the recommender system in itself (JYLHÄ, 2024, p. 33-34).

All of these findings focus on attainment of information going through the means in which information gets to the participants and how it is perceived, showing some user biases as to the tendency of finding search results reliable or truthful. Also of note is the awareness the users have of this whole process. It is not a given that the user understands: (a) the systems in use for collecting their data; (b) the contents of such data (although many governments protect one’s right to information) and; (c) how their data is used to influence their behavior. This takes the conversation back to what seems to be a good relationship between users and algorithms, which, by understanding how these computers programs function and their owner’s intent, users can engage with the technology in the terms that they’re comfortable with, preserving their agency and, therefore, consent on the matter.

Again, while it is important to know how people can have their vulnerabilities exploited, even vast educational campaigns will leave people behind. And this is reflective of how the issue at hand is usually framed. The abuse of these corporations is thought consistently as a given, as if there’s nothing to do but to protect oneself as an individual. That is, in its core, the most damning limitation that algorithmic literacy has as a remedy to the misuse of recommender systems. It is to be used, but its shortcomings cannot go understated.

With that in mind comes another (and thankfully not competing) solution to the problem: Regulation. Having ever growing evidence of how these technologies can (and do) exacerbate vulnerabilities, leaving each person to fend for themselves is to either blindly trust a sense of transparency (and thus not recognize the problem) or to accept corporate control.

Collecting various studies on regulatory practices around what they call the ‘Attention Market’, Michel and Gandon not only describe other aversive effects caused by

recommender systems but also propose general principles which can inform future research and policy. These can be: (a) more traditional methods of incentives and penalties to reach desirable behavior (from companies) (MICHEL, 2024, p. 6); (b) usage regulation, which can prevent, for example, what they call the infinite feed trap, or having an opt-in principle for notifications and pop-ups as a default (MICHEL, 2024, p. 6); (c) content recommendation monitoring in which, for example, the platform is to be required to monitor and limit the exposure of recommendations associated to negative emotions, as these are abused for the spread of fake news (MICHEL, 2024, p. 7); among others.

These are enumerated as a means of grasping the extent to which change is possible and that many of the misuses observed are not inherent to the technology itself but are brought to the picture nonetheless because of the economic incentives. The economic agents also act in accordance to their contingencies of reinforcement. Through regulation it is possible to alter these contingencies, while choosing not to engage with them is to let them act by their economic incentives as other contingencies of reinforcements. In that sense, if there is no regulation, the companies couldn't be expected to act in any other way than how they currently do.

The research behind this essay began with the query of whether having your biases reinforced externally would hamper one's prospects of wanting to become someone else. It was in a sense naive because it couldn't begin to comprehend the extent to which a person is affectable. But an important concept which was only a little touched on and yet remains central to this discussion is human variability. Through some of the works presented it was marked as a feature that hampers the algorithms' optimization in regard to the exploitation-exploration trade-off. It remains as a point of interest that can guide future research on not only variability itself but also change.

The discussion surrounding this work centered on the interaction between user and recommender systems and the possibility of exploitation of a person's biases, needs and emotions. These abuses were analyzed through the assemblage of some works from Denise Ferreira da Silva and Burrhus Frederic Skinner as a means of understanding the implications on freedom and self-determination of such practices. Then the use of algorithmic literacy is explored as a means of solving the issue while ultimately it is concluded that it needs to be implemented alongside regulative measures to the market surrounding recommender systems. That is because the promotion of algorithmic literacy by itself allows for the persistence of abuse to the already dispossessed.

BIBLIOGRAPHY

HELBERGER, N., SAX, M., STRYCHARZ, J. *et al*. *Choice Architectures in the Digital Economy: Towards a New Understanding of Digital Vulnerability*. **Journal of Consumer Policy** Vol. 45, Dez, 2021. https://doi.org/10.1007/s10603-021-09500-5

JYLHÄ, Ville, HIRVONEN, Noora, HAIDER, Jutta. *Algorithmic Recommendations enabling and constraining information practices among young people*. **Journal of Documentation**, Vol. 80: N°7, Jan, 2024. p. 25-42. https://doi.org/10.1108/JD-05-2023-0102

SILVA, Denise Ferreira da. **Unpayable Debt**. London, Sternberg Press, 2022

SILVA, Denise Ferreira da. **Toward a Global Idea of Race**. Mineapolis, University of Minnesota Press, 2007

LEX, Elisabeth, KOWALD, Dominik, SEITLINGER, Paul; TRAN, Thi Ngoc Trang, FELFERNIG, Alexander, SCHEDL, Markus. *Psychology-informed recommender systems*. Boston, Delft, **Foundations and Trends in Information Retrieval**, Vol. 15: N° 2, Jul, 2021. http://dx.doi.org/10.1561/1500000090

MICHEL, Franck, GANDON, Fabien. **Pay Attention: Call to Regulate the Attention Market and Prevent Algorithmic Emotional Governance**. *arXiv*, Fev, 2024.

https://doi.org/10.48550/arXiv.2402.16670

NEE, Rebecca C.. *Youthquakes in a Post-Truth Era: Exploring Social Media News Use and Information Verification Actions Among Global Teens and Young Adults*. **Journalism & Mass Communication Educator**, Vol. 74: N° 2, Jan, 2019. https://doi.org/10.1177/1077695818825215

SKINNER, B. F.. **Science and Human Behavior**. Cambridge: The B. F. Skinner Foundation, 2014.

SKINNER, B. F.. **Beyond Freedom & Dignity**. Bungay: Pelican Books, 1973

SWART, Joelle. *Tactics of news literacy: How young people access, evaluate, and engage with news on social media*. **New Media & Society**, Vol. 25: N° 3, Mai, 2023. p. 505-521. https://doi.org/10.1177/14614448211011447

YANG, Scott Cheng-Hsin, RANK, Chirag, WHRITNER, Jake A., NASRAOUI, Olfa, SHAFTO, Patrick. *Human Variability and the Explore-Exploit Trade-Off in Recommendation*. **Cognitive Science** Vol. 47, Abr, 2023. https://doi.org/10.1111/cogs.13279

ZUBOFF, Shoshana - **The age of surveillance capitalism: the fight for a human future at the new frontier of power**. New York: PublicAffairs, 2019.