Probing the Ethical Boundaries of Personalization: a Case Study of Twitter's Recommendation Algorithm

Kevin Feng
University of Washington
Seattle, Washington, USA
kjfeng@uw.edu

Michael Ibrahim
University of Washington
Seattle, Washington, USA
micibr@uw.edu

James Yoo
University of Washington
Seattle, Washington, USA
jmsy@uw.edu

ABSTRACT

Most online content platforms today are designed to maximize the time users spend engaging with their content. This engagement allows platforms to both serve advertisements and collect data from usage patterns to incorporate into their recommendation and personalization algorithms. However, personalization algorithms are often opaque; surfacing relevant and interesting content to users at best, and constructing echo chambers in which users are not exposed to a diversity of opinions or beliefs at worst. This issue is exacerbated by the fact that many content platforms enable users to fine-tune personalization algorithms without much in the way of ethical guardrails, such as "liking" a post or selecting "see less/-more". In this paper, we ask the question of whether there are ethical limits to personalization in content platforms. Twitter is a content platform with a wealth of publicly-available information surrounding its personalization algorithms, which we use as a case study for our investigation. We conduct a literature review of prior research in content personalization on social media and analyze published algorithmic personalization source code by Twitter along with related technical blog posts. We then identify ethically nuanced components of Twitter's content personalization pipeline and analyze them from the perspective of 6 ethical theories. We conclude with a discussion of how developers can more deeply engage with ethical considerations when building personalized systems.

1 INTRODUCTION

Recommendation and personalization algorithms permeate aspects of everyday life. A user's video queue on YouTube, their recommended products on Amazon, the sponsored results in their Google search results, their "For you" feed on Twitter; are all the result of algorithms that are designed to push content that maximizes engagement and user retention. Recommendation and personalization algorithms effectively shape the user experience on content platforms, having direct control over the content a user consumes with and the people they might interact with.

Despite the extent to which these algorithms control and shape the user experience on content platforms, there has been little prior work on investigating these algorithms—and their use of user data—through ethical theories. Prior work has explicitly investigated personalization systems and modelled them as a relationship between content *senders* (the content platforms) and content *receivers* (users of content platforms) [25]. This work focused on the inherent asymmetry between senders (who require huge amounts of user data to push content to receivers (who often relinquish control over their data in using content platforms).

Other prior work has applied ethical theories to the analysis of personalization for users on websites in 2004 [65], before the advent of large content networks (such as Twitter, Instagram, TikTok, etc.). This work found that customization (i.e., user-driven content curation) was ethically less questionable than personalization (i.e., system-driven content curation using continual data monitoring) [65]. However, it is not true whether these findings still hold today, or are applicable to large-scale content platforms today that leverage datasets that are magnitudes larger than they were in 2004.

To that end, we investigate the personalization algorithms employed by large content platforms through ethical theories from Quinn [55], including relativism, egoism, Kantianism, utilitarianism, social contract theory, and virtue ethics. We draw from the extensive documentation surrounding Twitter's recommendation and personalization algorithms as an exploratory study, including source code [70], literature from the engineering team [73], and explainer articles [27, 61]. Our work comprises the following contributions:

- A review of prior work focusing on the analysis of personalization and recommendation algorithms via ethical theories.
- A quantification of Twitter's recommendation systems, including features used for personalization and their weights, and descriptions of components used for visibility filtering and user segmentation.
- An analysis of Twitter's recommendation system under the lens of ethical theories, serving as first steps toward future studies that investigate other large-scale content platforms and their personalization algorithms.

2 RELATED WORK

2.1 Personalization in Social Media Recommender Systems

In 2006, Facebook launched its news feed feature as a convenient way for users to stay updated on their friends' activities without having to individually check their friends' profiles. The algorithm was simple—a reverse chronological, subscription-based model that provided a user with another user's updates if they subscribed to

CSE 581 - Computing Ethics, University of Washington, Seattle, USA,

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM…$15.00
https://doi.org/10.1145/nnnnnnn.nnnnnnn

that other user via "friending." Fast forward almost two decades, and the landscape of social media feeds has changed completely. Machine learning-powered recommendation algorithms are now ubiquitous in social media feeds [47], using a variety of techniques to elicit user preferences and values. To better align social media recommenders to these preferences and values, a growing body of work proposes techniques and broader calls to action [14, 23, 41, 50, 63, 64], drawing upon the field of value-sensitive design [9, 22], where Borning and Muller define "value" as "what a person or group of people consider important in life" [9].

Many modern recommender systems *infer* values by implicitly learning them from user information and interaction history [47, 64]. Commonly tracked attributes for implicitly eliciting values include tracking clicks [77], content dwell time [75], and affinity with other users [30, 33].

Platforms and researchers have sought to develop recommender systems with *explicit controls* [17, 23, 28, 29, 37, 42, 48, 52] in an effort to mitigate these concerns. Users have reported higher levels of satisfaction [32, 35], trust [48], and engagement [28, 42] when they were given opportunities to exert control over the system, even when the controls had no impact on the output [6, 67]. Examples of explicit controls include thumbs up/down buttons to rate recommended items [24, 76], sliders and toggles for adjusting desired content characteristics [31], drag-and-drop topic specifiers [17], keyword critique [52], and manual selection of the recommender algorithm [8, 19]. More recent work has leveraged the semantic capabilities of large language models to enable the expression of preferences conversationally through a chat interface [23] and through editing a natural language user profile [29]. Explicit controls come with their own set of challenges—users may not know that these controls exist or what they do [26, 31, 62], find them cumbersome to use and keep up-to-date [28], or do not see value in engaging with them [38]. Indeed, prior work showed that most users prefer a hybrid approach that combines implicit and explicit learning [37, 46]. One promising approach in this vein is to design controls that allow for simultaneous expression of direct feedback and less direct social signals; real-world examples include "react" options on Facebook and LinkedIn [64].

These prior works provide valuable *usability* perspectives to guide the design of personalization on social media. However, in this work, we approach personalization from an *ethical* angle to supplement existing usability-centered ones.

2.2 User Agency in HCI and Social Media Feeds

Researchers are increasingly concerned with the reduction of user agency in personalized content feeds [44, 47, 66]. A primary source of this concern is the opacity with which social media feed algorithms operate. Indeed, these algorithms are commonly referred to as "black-boxes" [4, 12, 45, 51, 56]. Prior work in explainable AI for social media has attempted to open the black-box and make algorithms more understandable to lay users [2, 18, 36, 39, 40, 58], but a fully transparent "glass-box" algorithm may still reduce agency. For one, users' decision-making abilities may be weakened via information overload relating to the complexities of "glass-box" algorithms [53]. Transparency also does not guarantee self-causality; users may observe the inner workings of the algorithm without any opportunity for control [43]. Transparency aside, widespread user behavior such as "doomscrolling" (mindless scrolling) [54] and dissociation [5] are symptomatic of users' agency loss when interacting with algorithmically-driven feeds.

Dwindling agency has led to users deriving "algorithmic folk theories" to make sense of their social media experiences [15, 16, 20, 21, 34, 59]. Examples of folk theories include that users will see more content from friends who are more similar to them in their Facebook News Feed [20] or that the TikTok For You Page algorithm prioritizes videos that feature aesthetics associated with wealthier lifestyles (e.g., large houses) [34]. Theories may emerge from both *endogenous* (originating from data within the platform [7]) as well as *exogenous* (originating from data collected outside the platform [10]) information [15, 57]. Users might apply folk theories in shaping how they interact with platforms. For example, Facebook users regularly visited profile pages of those whom they wanted to see more of and sought out more opportunities to tag them in posts, to signal their preferences to the algorithm [20]. To attempt to spread a video to other users' For You pages, TikTok users watched videos multiple times (even when they understood it perfectly the first time through), left longer comments, and hit the like/share buttons repeatedly even though they could only like or share a video once [34]. However, users remained uncertain of the efficacy of their techniques and even considered their actions to be manipulative and forced [11, 20].

Our work focuses on contributing new, ethics-focused perspectives on user agency. While much of HCI research aims to maximize user agency, some branches of ethics may push back on this ideal (e.g., those who subscribe to social contract theory may claim that maximizing user agency will actually be counterproductive to having all users adhere to the contract). On the other hand, reducing user agency has demonstrated negative effects on user experience and public perceptions of platforms more broadly. We seek to weigh these perspectives in our work.

2.3 Prior Ethical Perspectives on Personalization and Customization

Content platforms employ automated mechanisms to personalize and recommend content for users to operate at scale with billions of users and terabytes of content. However, the near-ubiquity of black-box AI recommendation systems present numerous ethical challenges [25].

One challenge arises with the data that is used to train and fine-tune AI personalization systems. Herman frames personalization systems as an interaction between content *senders* (the content platforms) and content *receivers* (the users) [25]. Senders require large amounts of data for AI personalization systems to be effective. These data may range from coarse-grained information about usage patterns to specific information about a user's age, sex, and/or ethnicity. This results in an asymmetrical dynamic between the perceived benefits to the content senders and consumers; it benefits senders to collect as much information as possible to improve their personalization systems to maximize the time users spend on their platform, while users effectively trade off their privacy for relevant content. Herman also posits that personalization systems degrade the autonomy and agency users may have in selecting the content they wish to see [25]. If a content platform offers only a personalization-based feed, then users of that platform effectively have zero autonomy and agency over what they wish to view. This irony is what Simpson et al. call a “personalization paradox” [60].

In 2004, Treiblmaier et al. conducted a study of Austrian internet users’ attitudes towards personalization on websites, with a focus on data privacy perceptions [65]. The authors then applied ethical theories including Kantianism, deontologism, social contract theory, utilitarianism, virtue ethics, and stakeholder theory to discuss the results. A key differentiation they make is one between personalization and customization—drawing from prior work, they define customization as *user-initiated* and *user-driven* adaptation of technical systems, whereas personalization is *system-initiated* and *system driven.* That is, personalization can be equated to “self-customization” [13], where the system actively monitors user behavior to adapt. Additionally, because of personalized systems’ automated nature, the user is unable to control *how* the system adapts. The authors conclude that personalization and its prerequisites fundamentally conflict with the ideals of many ethical theories, whereas customization is more ethically acceptable.

In our work, we aim to apply ethical theories to personalization in the context of social media platforms rather than general websites. Additionally, instead of surveying user attitudes, we analyze existing available information about social media recommendation algorithms, which researchers have noted is abundant [47]. Trieblamier et al.’s differentiation between personalization and customization [65] is a useful one.

3 METHOD

The goal of our study is to investigate where the ethical boundaries are to personalization in online content platforms. To accomplish this, we conducted an extensive literature review of publicly available information on the personalization algorithm for Twitter, a large content network (section 3.1) Subsequently, we embarked on an in-depth analysis of our discoveries, focusing particularly on identifying elements of the personalization algorithm that warranted ethical scrutiny or were particularly pertinent from a moral standpoint (section 3.2). Finally, we evaluated the outcomes through various established ethical frameworks to ascertain the ethical implications of our findings (section 3.3).

3.1 Understanding Twitter’s Personalization Algorithm

We focused our investigation on the Twitter content personalization algorithm, of which there exists a wealth of publicly available information about, including the source code [70], technical reports [73], and independent audits and walkthroughs [27, 61].

We extracted a set of features from these sources that are likely used as parameters in how Twitter recommends content to its users and personalizes their feeds. To increase our confidence in what features are used by Twitter’s algorithm in practice, we focused on features that appeared in a majority of the articles in our literature review. The source code [70] of the Twitter recommendation system was used as a primary oracle in cases where it was unclear whether a feature was involved in content recommendation and personalization.

3.2 Distilling the Ethical Components of the Personalization Algorithm

The distinction between binary inclusion/exclusion mechanisms, such as social proof, and the more nuanced tunable weight mechanisms in Twitter’s personalization algorithm, presents unique ethical considerations and impacts on user experience. We distill the major components of the Twitter recommendation that we consider as *ethically nuanced*—that is, the implementation choices made in these components leave room for ethical debate and shifting the implementation strategy may lead to a different ethical outcome. For example, weights used to predict engagement on a post may be ethically nuanced, while the decision to parallelize some computational processes over others in some parts of the pipeline is not. A couple more examples of ethically nuanced components in the context of the Twitter personalizaton algorithm are as follows.

The social proof mechanism [73] is a binary inclusion/exclusion criterion. If a tweet does not have a second-degree connection to the user (i.e., someone that they directly follow hasn’t engaged with the tweet or follows the author), then it will be less likely to appear in their feed.

Twitter’s personalization algorithm also takes into account the actions of users on its platform, such as following or blocking other users, “liking” or “re-tweeting” a tweet, and more. These signals are weighted [27], with the ability to be manipulated to prioritize certain content types and/or user behaviors. These components are thus suited for our analysis via ethical frameworks.

3.3 Analysis via Ethical Theories

Upon identifying ethically nuanced components of the algorithm, we will examine them through various ethical theories. We use the following 6 ethical theories from Quinn [55]: relativism (subject and cultural), egoism, Kantianism, utilitarianism (act and rule), social contract theory, and virtue ethics. The goal of this examination is to identify, for each ethically nuanced algorithm design decision, the level of acceptability of implementing that decision à la Twitter. This level of acceptability—specifically the circumstances under which a design decision may become unacceptable—is what we define to be an “ethical boundary.”

Inevitably, different ethical theories may lead to different boundaries. We account for this by mapping ethical theories to boundaries, such that each boundary corresponds to the intersection of one ethical theory and one ethically nuanced algorithm component. This resulting matrix of ethical boundaries constitutes our primary contribution from our analysis—a *framework* that answers our initial question: *what are the ethical boundaries of personalization?*

We intend for this framework to serve as a guide for researchers and practitioners to make more ethically informed decisions when designing personalization algorithms. We defer the evaluation of this framework to future work.

4 RESULTS

In this section, we begin with an overview of Twitter’s recommendation system and some of its core components related to user segmentation and visibility filtering (section 4.1). We then delve deeper into the mechanics of the system by discussing the specific features employed to tailor and recommend Tweets to individual users, as well as the significance and impact of their associated weights (section 4.2). Finally, we explore our findings in the context of ethical frameworks and use them to guide a discussion of the ethics of Twitter's algorithm and recommendation systems (section 4.3).

4.1 Understanding Twitter's Recommendation System

4.1.1 User Segmentation. Twitter's recommendation algorithm is capable of partitioning or labelling users. One example is provided in code (currently removed from the latest version of the code repository) containing labels [68] that can be applied to users. The labels range from extremely specific, describing singular users (e.g., `author_is_elon`), to extremely broad (e.g., `author_is_democrat`, `author_is_republican`, `author_is_power_user`).

These labels may conceivably be used to personalize content for users on Twitter in serving relevant content (e.g., a power user may be served ads for Twitter Blue) that may be used to increase engagement.

4.1.2 Visibility Filtering. Like many content platforms, Twitter employs a *visibility filtering* (VF) system that can be used to limit (or boost) the reach of content. For example, the VF system may be used to ensure that users do not see Tweets from accounts that they have blocked, which is a use-case that respects user agency.

However, this VF system may also be used to limit the reach of Tweets without user intervention. Twitter has a trust and safety label named `DoNotAmplify` [71], which might be used to limit the reach of a Tweet across the platform.

Twitter also appears to limit the reach of content and URLs from competitor platforms in addition to Tweets that are deemed out-of-network for a user [69]. From one perspective, filtering out irrelevant content from an out-of-network user might be beneficial to the consumer, as they may not be relevant. Filtering out content from competitor networks, however, limits users to the content available on Twitter. This is beneficial for Twitter, as they are incentivized to keep users on its platform for as long as possible. The consequences for users are less clear; one might argue that by using Twitter and not a competitor platform, they are opting out of content on other platforms. A counterargument would be that users should have access to any content, regardless of its origin.

4.2 Features and Weights for Personalization

Tweets are recommended to users in the "For you" feed via the *Heavy Ranker* [72] system, which uses at least 10 features (Figure 1) to recommend Tweets to users. Of the 10 features, two have negative weights; `report`, which is a predictor of whether a user will select "Report Tweet", and `negative_feedback_v2`, which is a predictor of whether a user might provide negative feedback, selecting "Show less often" on the Tweet, or muting or blocking the Tweet author.

Surprisingly, actions such as favouriting a Tweet, re-Tweeting, or playing a video at least halfway (`video_playback50`) are not very heavily weighted; the weights fall into a range from 0.005 to 1. Instead, the ranking system places a very high weight on features related to users interacting with content, such as clicking on a Tweet author's profile or replying to a Tweet which is in turn predicted to be replied to by the original author. With these actions, users are more likely to both engage with and create more content, which in turn may be used by Twitter to generate additional recommendations.

Type of User Engagement (Probability)	Weight
watch at least half of the video (for a tweet with video)	0.005
like the tweet	0.5
retweet the tweet	1.0
click into the conversation of a tweet and stay there for at least 2 minutes	10.0
click into the conversation of a tweet and reply or like a tweet	11.0
opens the tweet author's profile and likes or replies to a tweet	12.0
replies to the tweet	13.5
react negatively (requesting "show less often" on the tweet or author, block or mute the author)	-74.0
replies to a tweet and this reply is engaged by the tweet author	75
click report tweet	-369.0

Table 1: Types of user engagement and corresponding weights in the Heavy Ranker of Twitter's For You algorithm.

4.3 Twitter's Ranking Algorithm Through Ethical Lenses

In this section, we discuss the data practices of Twitter's algorithmic recommendation systems through the lens of six ethical theories.

4.3.1 Relativism. Practically, Twitter's data practices are not inconsistent with that of other large content platforms (e.g., Facebook, TikTok, etc. These platforms collect demographic information and have mechanisms with which they might identify user segments or limit the reach of content (visibility filtering). With this knowledge, Twitter (and other large content platforms) may be grouped into a faction with similar generally accepted practices by society. Under this characterization, Twitter's data practices are generally ethical under relativism.

However, an interesting consequence of Twitter's VF systems is that tweets for which a large amount of negative feedback has been reported are downranked to a large degree (see "click report Tweet" and "react negatively" in table 1). A scenario in which relativism would deem Twitter's VF system unethical might be when a segment of users report a tweet (as they may find it offensive or otherwise undesirable), resulting in a platform-wide downranking of the tweet. However, it may be the case that this tweet is offensive only for that given segment of users and not for others. Under relativism, it would be unethical to limit the visibility of the tweet for the entire platform.

4.3.2 Egoism. Through egoism, the data practices and recommendation algorithms of Twitter could be evaluated based on how well they align with their long-term interests and benefits. Twitter's user segmentation and visibility filtering can be considered ethical

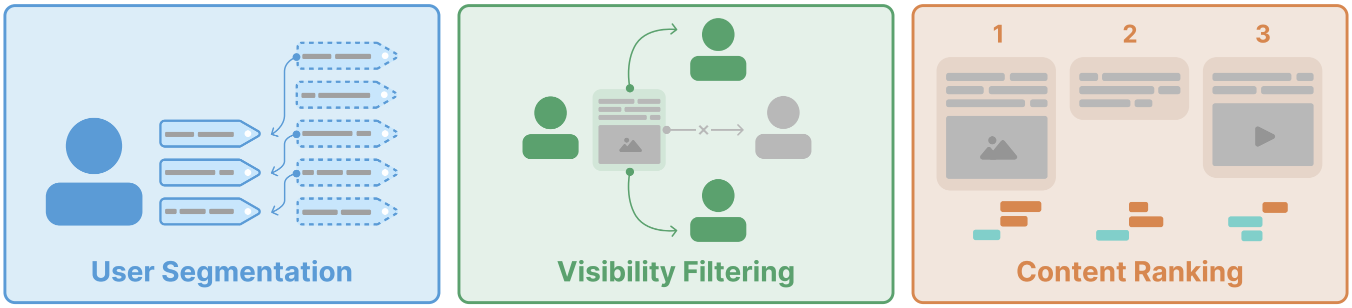

Figure 1: The components of Twitter's recommendation system that are analyzed for our study. User segmentation partitions (labels) users according to a variety of data points (e.g., inferred political views, gender and age, sexual orientation, etc.). Visibility Filtering (VF) systems are used to limit the reach of content or filter them out from feeds altogether. Content ranking systems generate engagement probability scores via model weights and features inferred from content to rank them for delivery.

from this viewpoint if they serve the platform's interests, such as increasing user engagement, maximizing revenue, or enhancing the overall user experience to retain users [73].

However, egoism also considers the long-term consequences of actions. If Twitter's practices lead to negative public perception, decreasing user trust, or regulatory scrutiny that could harm its reputation and profitability, then these practices would be deemed unethical. For example, Twitter's self-interest could ultimately be harmed by excessive use of VF to censor or manipulate content or user feeds in ways that diminish user autonomy or incites public backlash.

Moreover, Twitter's algorithmic decisions would need to balance short-term gains with long-term sustainability. If the recommendation systems prioritizes content that maximizes immediate engagement but leads to a toxic or polarized environment, the jeopardizing of Twitter's long-term viability and public image would be deemed unethical.

4.3.3 Kantianism. Under Kantianism, the data practices of Twitter—and consequently, the algorithmic systems powered by them—would be considered unethical. This aligns with the previous finding by Treiblmaier et al. which states that the collection of consumer information without consumers' consent can never be ethically justified, simply because the act as such would be ethically wrong [65].

One might argue that Twitter is collecting information with consent, as users agree to its Terms of Service when they sign up for the platform. However, this argument might be refuted by the fact that the majority of users do not read the terms in full (or at all). Under the assumption that privacy and autonomy are absolute rights, Twitter's data practices are unethical under Kantianism. For example, VF enables Twitter to artificially limit the reach of tweets and prevent users from seeing them, which violates the autonomy of both the producers and consumers of the content.

4.3.4 Utilitarianism. Under a utilitarian theory, any form of data collection might be considered ethical if the end result benefits a majority of the users in a system, even if this is at the expense of a minority of users. In analyzing Twitter's data collection practices, it is important to attempt to account for the benefits that it may provide to users. For example, Twitter's usage of social proof (which requires them to collect information about a users' followers and social circle) may mitigate irrelevant content from unknown accounts or spam from being surfaced to a user, which is arguably a benefit for the majority of users.

Twitter's user segmentation and visibility filtering (VF) practices are arguably a benefit for users under utilitarianism. For example, by collecting or inferring personal information such as age, sex, or political leanings to segment users, Twitter can better serve relevant content to the majority of users within a segment. Visibility filtering also provides a benefit for users. For example, assume a large proportion of users within a segment blocks a user or provides negative feedback for a tweet. VF may be used to ensure that the blocked user's tweets or similar tweets are surfaced less. This is an ethical action under utilitarianism; though the blocked user may not benefit from having the reach of their tweets be limited, the majority in the segment benefit from not seeing the blocked user's tweets or similar tweets. Additionally, the majority of users on popular social media platforms are consumers rather than producers [1], so the net satisfaction is likely to increase.

4.3.5 Social Contract Theory. Twitter's data practice and recommendation algorithms partially align with social contract theory, but cannot fully do so as users are not involved in the creation of their "contracts." Social contract theory asserts that 1) morality consists in the set of rules governing how members of a community will treat each other, 2) these rules are to be derived from a rational process, and 3) rational community members will collectively agree to follow these rules for mutual benefit assuming everyone else also follows them. Users on Twitter must abide by its Platform Use Guidelines [74], which one can argue is an operational contract derived through rational processes based on empirical evidence. Users also collectively agreed to follow it when they first logged onto the platform and agreed to its Terms and Conditions. However, this contract was created by the platform with little to no input from users themselves. End users also may be unaware that they agreed to the contract in the first place as few users read the Terms and Conditions in detail. Therefore, the contract exists but can hardly be considered as "social."

End users (to our knowledge) were not involved in the design of the VF system nor the selection of features and weights used in the Heavy Ranker. Both these systems shape community behavior by deciding *who* users interact with on the platform and *what* kinds of content will appear in their feeds. Users were unaware of the inner workings of these systems (at least not before they were open-sourced), let alone collectively agreeing to how they operate. Additionally, while Twitter may have engineered these systems through a rational process, there is no evidence that they did so—for example, we were unable to source any documentation for how the weights in Table 1 were determined. For Twitter's personalization techniques to fully embrace social contract theory, higher levels of algorithmic choice (such as custom algorithms on Bluesky [8]) and transparency will be needed.

4.3.6 Virtue Ethics. Twitter's data practices and personalization algorithms would be considered unethical under virtual ethics. To analyze Twitter under this theory, we consider it as an entity for which moral virtue and character are applicable characteristics. Twitter's recommendation system does not disclose whether a tweet's reach has been artificially boosted (e.g., for Twitter Blue users) or limited. This is deceptive to the users of the system who may assume that artificial boosting based on whether an author is a Twitter Blue does not occur. This is further exacerbated by the ability of Twitter Blue users to hide whether they are a Blue subscriber (e.g., remove the checkmark), meaning that it is not possible for users to easily discern whether someone may have boosted tweets.

Twitter also does not publicly disclose their user segmentation practices. Consequently, users may not be aware that they may have been placed into a specific segment for which targeted personalization algorithms may be deployed. This would be a deceptive practice considered unethical under virtue ethics.

5 THREATS TO VALIDITY, AND FUTURE WORK

5.1 Threats to Validity

The primary threat to the validity of our analysis is in the framing of each component of Twitter's personalization system in the ethical theories explored in our paper. To mitigate this threat, each author reviewed the decision made in deciding whether Twitter's personalization system was ethical (or not) under each theory.

Additionally, Twitter may currently not be operating under the systems and algorithms described in its open-source algorithm codebase [70], engineering blog post [73], and explainer articles [27, 61]. We cannot completely mitigate this threat, as most social media platforms operate under closed-source algorithms (except for a few decentralized platforms such as Bluesky or Mastodon).

5.2 Future Work

We may construct a general-purpose reasoning framework based on the six ethical theories we used to analyze Twitter's personalization systems. This framework might be applied by personalization platform developers to evaluate their design decisions in the context of ethical perspectives.

To evaluate this framework, we need to instantiate it into a tool that can actively inform design decisions during development workflows. One way to do this is via an AI policy document [3] that we then feed as knowledge to a custom OpenAI GPT [49]. Users may query this GPT about design decisions in personalization algorithms, and the GPT will ask for any necessary context (e.g., which an ethical framework the developer would like to use to inform this design decision) before providing an answer grounded in our framework.

6 DISCUSSION AND CONCLUSION

In this work, we probed the ethical boundaries of personalization—circumstances under which an algorithmic design decision of a personalization system would be considered (un)ethical—via a case study of Twitter's recommendation system in the For You feed. First, we analyzed the system's open source code, alongside technical documentation and explainers, to extract three ethically nuanced algorithmic components and their implementations—user segmentation, visibility filtering, and content ranking. We then analyzed these components through the lens of six ethical theories—relativism, egoism, Kantianism, utilitarianism, social contract theory, and virtue ethics—to determine whether and how a certain component can be considered ethical under that theory.

Our main takeaway is that there is no single standard for defining "ethical" in personalization. A particular algorithmic component and its implementation may be considered ethical under one theory but not another. For example, the set of features and weights used to compute engagement probability scores in content rankers may increase overall user satisfaction (aligning with utilitarianism) but may not encourage virtuous behavior (conflicting with virtue ethics). In addition to conflicts between theories for specific components, one theory may also yield different ethical judgements for different components within the system. For example, a content ranker may keep users more engaged on the platform and bring in more revenue (aligning under egoism), but user segmentation may harm the platform experiences of some user groups more than others and cause users to leave while negative press ensues (in conflict with egoism).

We propose three concrete steps to better grapple with the ethics of personalization systems. Developers should:

(1) Identify key ethically nuanced algorithmic components and use them as a fundamental unit of their ethical analysis. This involves asking questions such as "will this component be merely computational (e.g., database optimization) or does it have potential to make a value judgement (e.g., determine whose content is shown above others)?"
(2) Identify a set of ethical theories through which they can ground their analysis. Should the platform be shaped by universal moral rules? If so, developers should consider, compare, and contrast those (and related) theories.
(3) Identify and resolve inter-theory and inter-component conflicts within different ethical theories. How do different theories suggest implementation approaches for a particular component, and should only one theory be applied across the entire system? Developers should thoughtfully consider these questions with other team members and platform stakeholders when possible.

We hope these steps will help developers move past the nebulous notion of "ethical" and engage with theoretically-grounded approaches to determine *what exactly it means* to build for ethical personalization.

REFERENCES

[1] Alessia Antelmi, Delfina Malandrino, and Vittorio Scarano. 2019. Characterizing the Behavioral Evolution of Twitter Users and The Truth Behind the 90-9-1 Rule. In *Companion Proceedings of The 2019 World Wide Web Conference* (San Francisco, USA) *(WWW '19)*. Association for Computing Machinery, New York, NY, USA, 1035–1038. https://doi.org/10.1145/3308560.3316705

[2] Sverrir Arnórsson, Florian Abeillon, Ibrahim Al-Hazwani, Jürgen Bernard, Hanna Hauptmann, and Mennatallah El-Assady. 2023. Why am I reading this? Explaining Personalized News Recommender Systems. (2023).

[3] Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, et al. 2022. Constitutional ai: Harmlessness from ai feedback. *arXiv preprint arXiv:2212.08073* (2022).

[4] Nathan Bartley, Andres Abeliuk, Emilio Ferrara, and Kristina Lerman. 2021. Auditing Algorithmic Bias on Twitter. In *13th ACM Web Science Conference 2021* (Virtual Event, United Kingdom) *(WebSci '21)*. Association for Computing Machinery, New York, NY, USA, 65–73. https://doi.org/10.1145/3447535.3462491

[5] Amanda Baughan, Mingrui Ray Zhang, Raveena Rao, Kai Lukoff, Anastasia Schaadhardt, Lisa D. Butler, and Alexis Hiniker. 2022. "I Don't Even Remember What I Read": How Design Influences Dissociation on Social Media. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 18, 13 pages. https://doi.org/10.1145/3491102.3501899

[6] Dan Bennett, Oussama Metatla, Anne Roudaut, and Elisa Mekler. 2023. How does HCI Understand Human Autonomy and Agency? *arXiv preprint arXiv:2301.12490* (2023).

[7] Michael S. Bernstein, Eytan Bakshy, Moira Burke, and Brian Karrer. 2013. Quantifying the Invisible Audience in Social Networks. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Paris, France) *(CHI '13)*. Association for Computing Machinery, New York, NY, USA, 21–30. https://doi.org/10.1145/2470654.2470658

[8] Bluesky. 2023. Algorithmic Choice with Custom Feeds. https://blueskyweb.org/blog/7-27-2023-custom-feeds.

[9] Alan Borning and Michael Muller. 2012. Next Steps for Value Sensitive Design. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Austin, Texas, USA) *(CHI '12)*. Association for Computing Machinery, New York, NY, USA, 1125–1134. https://doi.org/10.1145/2207676.2208560

[10] Laura Burbach, Johannes Nakayama, Nils Plettenberg, Martina Ziefle, and André Calero Valdez. 2018. User Preferences in Recommendation Algorithms: The Influence of User Diversity, Trust, and Product Category on Privacy Perceptions in Recommender Algorithms. In *Proceedings of the 12th ACM Conference on Recommender Systems* (Vancouver, British Columbia, Canada) *(RecSys '18)*. Association for Computing Machinery, New York, NY, USA, 306–310. https://doi.org/10.1145/3240323.3240393

[11] Jenna Burrell, Zoe Kahn, Anne Jonas, and Daniel Griffin. 2019. When Users Control the Algorithms: Values Expressed in Practices on Twitter. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW, Article 138 (nov 2019), 20 pages. https://doi.org/10.1145/3359240

[12] Angèle Christin. 2020. The ethnographer and the algorithm: beyond the black box. *Theory and Society* 49, 5-6 (2020), 897–918.

[13] Altan Coner. 2003. Personalization and customization in financial portals. *Journal of American Academy of Business* 2, 2 (2003), 498–504.

[14] Sarah Dean and Jamie Morgenstern. 2022. Preference Dynamics Under Personalized Recommendations. *Proceedings of the 23rd ACM Conference on Economics and Computation* (2022).

[15] Michael A. DeVito, Jeremy Birnholtz, Jeffery T. Hancock, Megan French, and Sunny Liu. 2018. How People Form Folk Theories of Social Media Feeds and What It Means for How We Study Self-Presentation. 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, 1–12. https://doi.org/10.1145/3173574.3173694

[16] Michael A. DeVito, Darren Gergle, and Jeremy Birnholtz. 2017. "Algorithms Ruin Everything": #RIPTwitter, Folk Theories, and Resistance to Algorithmic Change in Social Media. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) *(CHI '17)*. Association for Computing Machinery, New York, NY, USA, 3163–3174. https://doi.org/10.1145/3025453.3025659

[17] Simon Dooms, Toon De Pessemier, and Luc Martens. 2014. Improving IMDb Movie Recommendations with Interactive Settings and Filters. In *RecSys Posters*.

[18] Upol Ehsan, Q. Vera Liao, Michael Muller, Mark O. Riedl, and Justin D. Weisz. 2021. Expanding Explainability: Towards Social Transparency in AI Systems. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 82, 19 pages. https://doi.org/10.1145/3411764.3445188

[19] Michael D. Ekstrand, Daniel Kluver, F. Maxwell Harper, and Joseph A. Konstan. 2015. Letting Users Choose Recommender Algorithms: An Experimental Study. *Proceedings of the 9th ACM Conference on Recommender Systems* (2015).

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

[21] Motahhare Eslami, Aimee Rickman, Kristen Vaccaro, Amirhossein Aleyasen, Andy Vuong, Karrie Karahalios, Kevin Hamilton, and Christian Sandvig. 2015. "I Always Assumed That I Wasn't Really That Close to [Her]": Reasoning about Invisible Algorithms in News Feeds. 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, 153–162. https://doi.org/10.1145/2702123.2702556

[22] Batya Friedman. 1996. Value-sensitive design. *interactions* 3, 6 (1996), 16–23.

[23] Luke Friedman, Sameer Ahuja, David Allen, Zhenning Tan, Hakim Sidahmed, Changbo Long, Jun Xie, Gabriel Schubiner, Ajay Patel, Harsh Lara, Brian Chu, Zexiang Chen, and Manoj Tiwari. 2023. Leveraging Large Language Models in Conversational Recommender Systems. *ArXiv* abs/2305.07961 (2023).

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

[25] Erik Hermann. 2022. Artificial intelligence and mass personalization of communication content—An ethical and literacy perspective. *New Media & Society* 24, 5 (2022), 1258–1277.

[26] Silas Hsu, Kristen Vaccaro, Yin Yue, Aimee Rickman, and Karrie Karahalios. 2020. Awareness, Navigation, and Use of Feed Control Settings Online. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '20)*. Association for Computing Machinery, New York, NY, USA, 1–13. https://doi.org/10.1145/3313831.3376583

[27] Igor Brigadir. 2023. The release of the Twitter algorithm, annotated for recsys. https://github.com/igorbrigadir/awesome-twitter-algo?tab=readme-ov-file.

[28] D. Jannach, Sidra Naveed, and Michael Jugovac. 2016. User Control in Recommender Systems: Overview and Interaction Challenges. In *International Conference on Electronic Commerce and Web Technologies*.

[29] Mahmood Jasim, Andrew McCallum, and Hamed Zamani. 2023. Editable User Profiles for Controllable Text Recommendations.

[30] Jeff Widman. 2015. EdgeRank. http://edgerank.net/.

[31] Shagun Jhaver, Alice Qian Zhang, Quanze Chen, Nikhila Natarajan, Ruotong Wang, and Amy Zhang. 2023. Personalizing Content Moderation on Social Media: User Perspectives on Moderation Choices, Interface Design, and Labor. *Proceedings of the ACM on Human-Computer Interaction* (2023).

[32] Yucheng Jin, Bruno Cardoso, and Katrien Verbert. 2017. How do different levels of user control affect cognitive load and acceptance of recommendations?. In *Jin, Y., Cardoso, B. and Verbert, K., 2017, August. How do different levels of user control affect cognitive load and acceptance of recommendations?. In Proceedings of the 4th Joint Workshop on Interfaces and Human Decision Making for Recommender Systems co-located with ACM Conference on Recommender Systems (RecSys 2017)*, Vol. 1884. CEUR Workshop Proceedings, 35–42.

[33] Hyeonsu B Kang, Rafal Kocielnik, Andrew Head, Jiangjiang Yang, Matt Latzke, Aniket Kittur, Daniel S. Weld, Doug Downey, and Jonathan Bragg. 2022. From Who You Know to What You Read: Augmenting Scientific Recommendations with Implicit Social Networks. *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (2022).

[34] Nadia Karizat, Dan Delmonaco, Motahhare Eslami, and Nazanin Andalibi. 2021. Algorithmic Folk Theories and Identity: How TikTok Users Co-Produce Knowledge of Identity and Engage in Algorithmic Resistance. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW2, Article 305 (oct 2021), 44 pages. https://doi.org/10.1145/3476046

[35] Hankyung Kim and Youn-Kyung Lim. 2023. Investigating How Users Design Everyday Intelligent Systems in Use. *Proceedings of the 2023 ACM Designing Interactive Systems Conference* (2023). https://api.semanticscholar.org/CorpusID:259376338

[36] Jan Kirchner and Christian Reuter. 2020. Countering Fake News: A Comparison of Possible Solutions Regarding User Acceptance and Effectiveness. *Proc. ACM Hum.-Comput. Interact.* 4, CSCW2, Article 140 (oct 2020), 27 pages. https://doi.org/10.1145/3415211

[37] Bart P. Knijnenburg, Niels J. M. Reijmer, and Martijn C. Willemsen. 2011. Each to his own: how different users call for different interaction methods in recommender systems. In *ACM Conference on Recommender Systems*.

[38] Pascal D. König. 2022. Challenges in enabling user control over algorithm-based services. *AI & SOCIETY* (2022).

[39] Ziyi Kou, Lanyu Shang, Yang Zhang, and Dong Wang. 2022. HC-COVID: A Hierarchical Crowdsource Knowledge Graph Approach to Explainable COVID-19 Misinformation Detection. *Proc. ACM Hum.-Comput. Interact.* 6, GROUP, Article 36 (jan 2022), 25 pages. https://doi.org/10.1145/3492855

[40] Vivian Lai, Samuel Carton, Rajat Bhatnagar, Q. Vera Liao, Yunfeng Zhang, and Chenhao Tan. 2022. Human-AI Collaboration via Conditional Delegation: A Case Study of Content Moderation. In *Proceedings of the 2022 CHI Conference*

on Human Factors in Computing Systems (New Orleans, LA, USA) *(CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 54, 18 pages. https://doi.org/10.1145/3491102.3501999

[41] Liu Leqi, Giulio Zhou, Fatma Kilincc-Karzan, Zachary Chase Lipton, and A. Montgomery. 2023. A Field Test of Bandit Algorithms for Recommendations: Understanding the Validity of Assumptions on Human Preferences in Multi-armed Bandits. *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (2023).

[42] Yu Liang, Aditya Ponnada, Paul Lamere, and Nediyana Daskalova. 2023. Enabling Goal-Focused Exploration of Podcasts in Interactive Recommender Systems. *Proceedings of the 28th International Conference on Intelligent User Interfaces* (2023).

[43] Zachary C Lipton. 2018. The mythos of model interpretability: In machine learning, the concept of interpretability is both important and slippery. *Queue* 16, 3 (2018), 31–57.

[44] Kai Lukoff, Ulrik Lyngs, Himanshu Zade, J. Vera Liao, James Choi, Kaiyue Fan, Sean A. Munson, and Alexis Hiniker. 2021. How the Design of YouTube Influences User Sense of Agency. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 368, 17 pages. https://doi.org/10.1145/3411764.3445467

[45] Caitlin Lustig, Katie Pine, Bonnie Nardi, Lilly Irani, Min Kyung Lee, Dawn Nafus, and Christian Sandvig. 2016. Algorithmic Authority: The Ethics, Politics, and Economics of Algorithms That Interpret, Decide, and Manage. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems* (San Jose, California, USA) *(CHI EA '16)*. Association for Computing Machinery, New York, NY, USA, 1057–1062. https://doi.org/10.1145/2851581.2886426

[46] Christine Mendez, Vlatko Lukarov, Christoph Greven, André Calero Valdez, Felix Dietze, Ulrik Schroeder, and Martina Ziefle. 2017. User Groups and Different Levels of Control in Recommender Systems. In *Interacción*.

[47] Arvind Narayanan. 2023. Understanding Social Media Recommendation Algorithms. https://knightcolumbia.org/content/understanding-social-media-recommendation-algorithms.

[48] Jeroen Ooge, L Dereu, and Katrien Verbert. 2023. Steering Recommendations and Visualising Its Impact: Effects on Adolescents' Trust in E-Learning Platforms. *Proceedings of the 28th International Conference on Intelligent User Interfaces* (2023).

[49] OpenAI. 2023. Introducing GPTs. https://openai.com/blog/introducing-gpts.

[50] Aviv Ovadya and Luke Thorburn. 2023. Bridging Systems: Open Problems for Countering Destructive Divisiveness across Ranking, Recommenders, and Governance. https://doi.org/10.48550/ARXIV.2301.09976

[51] Frank Pasquale. 2015. *The black box society: The secret algorithms that control money and information.* Harvard University Press.

[52] Diana Petrescu, Diego Antognini, and Boi Faltings. 2021. Multi-Step Critiquing User Interface for Recommender Systems. *Proceedings of the 15th ACM Conference on Recommender Systems* (2021).

[53] Forough Poursabzi-Sangdeh, Daniel G Goldstein, Jake M Hofman, Jennifer Wortman Wortman Vaughan, and Hanna Wallach. 2021. Manipulating and Measuring Model Interpretability. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 237, 52 pages. https://doi.org/10.1145/3411764.3445315

[54] Aditya Kumar Purohit, Kristoffer Bergram, Louis Barclay, Valéry Bezençon, and Adrian Holzer. 2023. Starving the Newsfeed for Social Media Detox: Effects of Strict and Self-Regulated Facebook Newsfeed Diets. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) *(CHI '23)*. Association for Computing Machinery, New York, NY, USA, Article 196, 16 pages. https://doi.org/10.1145/3544548.3581187

[55] Michael Jay Quinn. 2009. *Ethics for the information age.* Pearson Education Boston.

[56] Urbano Reviglio and Claudio Agosti. 2020. Thinking outside the black-box: The case for "algorithmic sovereignty" in social media. *Social Media+ Society* 6, 2 (2020), 2056305120915613.

[57] Leonid Rozenblit and Frank Keil. 2002. The misunderstood limits of folk science: An illusion of explanatory depth. *Cognitive science* 26, 5 (2002), 521–562.

[58] Ruoxi Shang, K. J. Kevin Feng, and Chirag Shah. 2022. Why Am I Not Seeing It? Understanding Users' Needs for Counterfactual Explanations in Everyday Recommendations. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency* (Seoul, Republic of Korea) *(FAccT '22)*. Association for Computing Machinery, New York, NY, USA, 1330–1340. https://doi.org/10.1145/3531146.3533189

[59] Ignacio Siles, Andrés Segura-Castillo, Ricardo Solís, and Mónica Sancho. 2020. Folk theories of algorithmic recommendations on Spotify: Enacting data assemblages in the global South. *Big Data & Society* 7, 1 (2020), 2053951720923377.

[60] Ellen Simpson, Andrew Hamann, and Bryan Semaan. 2022. How to Tame "Your" Algorithm: LGBTQ+ Users' Domestication of TikTok. *Proc. ACM Hum.-Comput. Interact.* 6, GROUP, Article 22 (jan 2022), 27 pages. https://doi.org/10.1145/3492841

[61] Sol Messing. 2023. What can we learn from 'The Algorithm,' Twitter's partial open-sourcing of it's feed-ranking recommendation system? https://solomonmg.github.io/post/twitter-the-algorithm/.

[62] Spandana Singh. 2020. Charting a Path Forward. https://www.newamerica.org/oti/reports/charting-path-forward/.

[63] Jonathan Stray. 2020. Aligning AI Optimization to Community Well-Being. *International Journal of Community Well-Being* 3 (2020), 443 – 463.

[64] Jonathan Stray, Alon Y. Halevy, Parisa Assar, Dylan Hadfield-Menell, Craig Boutilier, Amar Ashar, Lex Beattie, Michael D. Ekstrand, Claire Leibowicz, Connie Moon Sehat, Sara Johansen, Lianne Kerlin, David Vickrey, Spandana Singh, Sanne Vrijenhoek, Amy X. Zhang, McKane Andrus, Natali Helberger, Polina Proutskova, Tanushree Mitra, and Nina Vasan. 2022. Building Human Values into Recommender Systems: An Interdisciplinary Synthesis. *ArXiv* abs/2207.10192 (2022).

[65] Horst Treiblmaier, Maria Madlberger, Nicolas Knotzer, and Irene Pollach. 2004. Evaluating personalization and customization from an ethical point of view: An empirical study. In *37th Annual Hawaii International Conference on System Sciences, 2004. Proceedings of the.* IEEE, 10–pp.

[66] Tristan Harris. 2016. How Technology is Hijacking Your Mind — from a Magician and Google Design Ethicist. https://medium.com/thrive-global/how-technology-hijacks-peoples-minds-from-a-magician-and-google-s-design-ethicist-56d62ef5edf3.

[67] Kristen Vaccaro, Dylan Huang, Motahhare Eslami, Christian Sandvig, Kevin Hamilton, and Karrie Karahalios. 2018. The Illusion of Control: Placebo Effects of Control Settings. 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/10.1145/3173574.3173590

[68] X Corp. 2023. HomeTweetTypePredicates. https://github.com/twitter/the-algorithm/blob/7f90d0ca342b928b479b512ec51ac2c3821f5922/home-mixer/server/src/main/scala/com/twitter/home_mixer/functional_component/decorator/HomeTweetTypePredicates.scala#L224-L247.

[69] X Corp. 2023. OutOfNetworkCompetitorFilter.scala. https://github.com/jyoo980/the-algorithm/blob/main/home-mixer/server/src/main/scala/com/twitter/home_mixer/functional_component/filter/OutOfNetworkCompetitorFilter.scala.

[70] X Corp. 2023. Source Code for Twitter's Recommendation Algorithm. https://github.com/twitter/the-algorithm.

[71] X Corp. 2023. Trust and Safety - Do Not Amplify. https://github.com/twitter/the-algorithm/blob/7f90d0ca342b928b479b512ec51ac2c3821f5922/visibilitylib/src/main/scala/com/twitter/visibility/models/TweetSafetyLabel.scala#L55.

[72] X Corp. 2023. Twitter Heavy Ranker. https://github.com/twitter/the-algorithm-ml/tree/main/projects/home/recap#heavy-ranker.

[73] X Corp. 2023. Twitter's Recommendation Algorithm. https://blog.twitter.com/engineering/en_us/topics/open-source/2023/twitter-recommendation-algorithm.

[74] X Corp. 2024. Platform Use Guidelines. https://help.twitter.com/en/rules-and-policies.

[75] Xing Yi, Liangjie Hong, Erheng Zhong, Nathan Nan Liu, and Suju Rajan. 2014. Beyond clicks: dwell time for personalization. In *ACM Conference on Recommender Systems*.

[76] Xiaoxue Zhao, Weinan Zhang, and Jun Wang. 2013. Interactive collaborative filtering. In *Proceedings of the 22nd ACM international conference on Information & Knowledge Management.* 1411–1420.

[77] Zhe Zhao, Lichan Hong, Li Wei, Jilin Chen, Aniruddh Nath, Shawn Andrews, Aditee Kumthekar, Maheswaran Sathiamoorthy, Xinyang Yi, and Ed Chi. 2019. Recommending What Video to Watch next: A Multitask Ranking System. In *Proceedings of the 13th ACM Conference on Recommender Systems* (Copenhagen, Denmark) *(RecSys '19)*. Association for Computing Machinery, New York, NY, USA, 43–51. https://doi.org/10.1145/3298689.3346997