

DATING THROUGH THE FILTERS*

BY KARIM NADER

Abstract: In this essay, I explore ethical considerations that might arise from the use of collaborative filtering algorithms on dating apps. Collaborative filtering algorithms can predict the preferences of a target user by looking at the past behavior of similar users. By recommending products through this process, they can influence the news we read, the movies we watch, and more. They are extremely powerful and effective on platforms like Amazon and Google. Recommender systems on dating apps are likely to group people by race, since they exhibit similar patterns of behavior: users on dating platforms seem to segregate themselves based on race, exclude certain races from romantic and sexual consideration (except their own), and generally show a preference for white men and women. As collaborative filtering algorithms learn from these patterns to predict preferences and build recommendations, they can homogenize the behavior of dating app users and exacerbate biased sexual and romantic behavior.

KEY WORDS: algorithmic filtering, algorithmic bias, collaborative filtering, ethics of dating, ethics of technology

I. INTRODUCTION

In this essay, I explore ethical considerations that might arise from the use of collaborative filtering algorithms on dating apps. Collaborative filtering is used to build recommendations for users on online platforms. They learn from the preferences of other users who exhibit similar behavior in order to predict the preferences of a target user and recommend content or products that match those predictions. Collaborative filtering systems have been deployed successfully on different platforms such as Amazon and Google. While using collaborative filtering to sort through products or movies is harmless, using the same systems to power dating apps could raise some issues. These considerations are especially relevant as a majority of new couples meet online: a 2017 survey shows 39 percent of 3,510 surveyed couples in the United States met online, a higher percentage than other methods of meeting (27 percent met at a restaurant or bar, and 20 percent met through friends).¹ Another 2013 survey shows that between 2005 and

* I would like to thank Kenneth Fleischmann and the students in his Ethics of AI graduate seminar for helping me start this project, and the audience at the Arizona Feminist Philosophy Graduate Conference for their helpful comments. I would also like to thank David Schmitz, Caroline King, and an anonymous reviewer for their feedback on earlier drafts.

¹ Michael J. Rosenfeld, Reuben J. Thomas, and Sonia Hausen, "Disintermediating Your Friends: How Online Dating in the United States Displaces Other Ways of Meeting," *Proceedings of the National Academy of Sciences* 116, no. 36 (2019): 17753–58, <https://doi.org/10.1073/pnas.1908630116>.

doi:10.1017/S0265052521000133

© 2021 Social Philosophy & Policy Foundation. Printed in the USA. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted re-use, distribution, and reproduction in any medium, provided the original work is properly cited.

2013, close to 35 percent of couples in the United States met their spouse online, with about half of those meetings happening on dating sites.² Through recommender systems, dating apps are increasingly influencing whose profile users can see or match with, and so who they date and potentially marry.

I start by explaining how collaborative filtering algorithms can predict preferences to build recommendations for users. I then show that race is a confounding factor in dating app recommendation systems. And finally, I argue that collaborative filtering can affect user behavior. My goal is to establish that filtering algorithms can homogenize user behavior and deepen existing patterns of sexual and romantic bias. Since users have little control over the process, and since race plays a confounding role in how user preference is determined, this process might be worth a closer look.

II. COLLABORATIVE FILTERING

I am concerned with dating apps that use algorithms to recommend potential matches to users. For example, Tinder makes recommendations to users both in their “Top Pick” section (a collection of ten recommended profiles that a user is issued daily) and in the more general pool of user profiles by showing recommended profiles first. Another example is Hinge’s “Most Compatible” feature, which pairs two users every day based on the users’ past activity on the app and their interests.³ These apps usually show users one profile at a time, giving them two options: if they are sexually or romantically interested, they will “like” the user or “swipe right” on their profile, and if they are not, then they will “swipe left.” If two users are interested in one another, they match and can start a conversation. The data from this process is used to make future recommendations, and determine what profile is shown next. The algorithms that power recommended matches are usually inaccessible to the user and to the public, but we have strong reason to believe that those algorithms are similar to other collaborative recommender systems. I will start by looking at how collaborative filtering algorithms predict preferences to build recommendations.

With the amount of content available, recommender systems are crucial to help users choose from the abundance of movies, news articles, or products available on online platforms. Collaborative filtering algorithms filter the abundance of choices to specific recommendations that are predicted to match the user’s preference. The idea behind collaborative filtering is that if groups of users show similar patterns of preferences, the preferences of one user can be predicted from the past behavior of similar users. In other words,

² J. T. Cacioppo et al., “Marital Satisfaction and Break-Ups Differ across On-Line and Off-Line Meeting Venues,” *Proceedings of the National Academy of Sciences* 110, no. 25 (2013): 10135–40, <https://doi.org/10.1073/pnas.1222447110>.

³ Sarah Wells, “Hinge Employs New Algorithm to Find Your ‘Most Compatible’ Match” *TechCrunch*, July 11, 2018, <https://techcrunch.com/2018/07/11/hinge-employs-new-algorithm-to-find-your-most-compatible-match-for-you/>.

by collecting data on the preferences of users collectively, the algorithm predicts the preferences of an individual user, builds recommendations that match those preferences, and filters the content the user can access on the platform. Recommendations prioritize some options to the user over others, while filtering limits the choices that are available to the user. A simple example: Let's say that most online shoppers who buy chips also buy salsa. By collecting data on user shopping behavior, a filtering algorithm learns the high correlation between buying chips and buying salsa. When target users add chips to their virtual cart, they are grouped with all the previous users who bought chips and salsa. Their future interest will be predicted based on the past behavior of those users collectively; this prediction then leads to a message which should be familiar to readers who have shopped online: "You may also like" salsa. Simply put, because most people who buy chips also buy salsa, if a target user were to buy chips, the algorithm would predict that the user may respond favorably to a recommendation to buy salsa.

The sheer quantity of data available and the relative ease of creating recommender systems that are blind to content make collaborative filtering algorithms both practical and effective. First, an enormous amount of implicit data can be gathered from a simple interaction between the user and the platform. Explicit data about preference can be gathered through rating systems (a star ranking on a product or a comment left on a page). But implicit data about user preference is easier to gather and does not require users to spend time rating content or products. Implicit data includes anything from users' shopping history, to what products they look at, which link they click, and how much time they spend looking at a certain page. Second, both explicit and implicit data does not require any information about the content of the recommendation (for example, the quality of the product, the genre of movie) or any knowledge about the user (for example, demographics). Data about content and demographics is extremely hard to gather, so a recommender system that can be effective without it is preferable.

Since collaborative filtering algorithms work with abundant and reliable data, they have been deployed to filter recommendations on popular platforms. And they work! Recommendations are extremely successful in influencing user behavior. "At Netflix, 2/3 of the movies watched are recommended; at Google, news recommendations improved click-through rate (CTR) by 38%; and for Amazon, 35% of sales come from recommendations."⁴ Filtering can also be highly effective: Iyengar and Lepper show that when we are given less choice, we act faster, whether that is buying a product, watching a movie, or chatting with a match⁵: "they ran an experiment where they had two stands of jam on two different days.

⁴ Jesse Steinweg-Woods, "A Gentle Introduction to Recommender Systems with Implicit Feedback," May 30, 2016, <https://jessesw.com/Rec-System/>.

⁵ Sheena S. Iyengar and Mark R. Lepper, "When Choice Is Demotivating: Can One Desire Too Much of a Good Thing?" *Journal of Personality and Social Psychology* 79, no. 6 (December 2000): 995–1006, <https://doi.org/10.1037/0022-3514.79.6.995>.

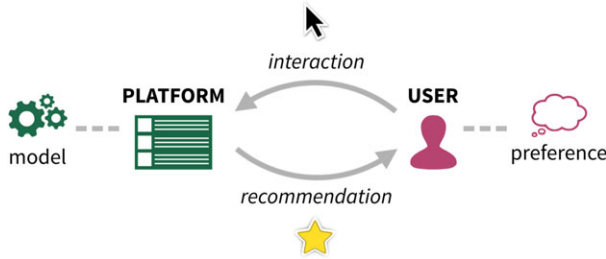


Figure 1. Recommendation feedback loop. This image is reproduced with permission of the authors of the original source. Allison J. B. Chaney, Brandon M. Stewart, and Barbara E. Engelhardt, “How Algorithmic Confounding in Recommendation Systems Increases Homogeneity and Decreases Utility,” *Proceedings of the 12th ACM Conference on Recommender Systems* (2018), 224–32, <https://doi.org/10.1145/3240323.3240370>.

One stand had 24 varieties of jam while the second had only six. The stand with 24 varieties of jam only converted 3% of the customers to a sale, while the stand with only six varieties converted 30% of the customers. This was an increase in sales of nearly ten-fold!”⁶ On the surface, recommender systems are beneficial to both users and platforms. On the one hand, users will be able to make a choice that they will be satisfied with much faster, wasting less time browsing and filtering through the results themselves. On the other, by showing recommended items first and filtering other unwanted items, a business will have both a higher customer satisfaction rate and better sale, click, or watching rates.

However, those advantages might come at a cost. By simulating a community of users interacting with items on an online platform, Chaney, Stewart, and Engelhardt show that collaborative recommender systems increase homogeneity in users’ behavior without necessarily increasing utility (see Figure 1).⁷ Since users are more likely to make the choice that is recommended to them, users will tend to make the same choice as users around them since they receive similar recommendations. The collaborative filtering system picks up on those choices, and in turn prioritizes them in its recommendations. The recommendations are then amplified through a feedback loop: users choose recommended products, and products are recommended because users choose them.

⁶ Steinweg-Woods, “A Gentle Introduction to Recommender Systems with Implicit Feedback.”

⁷ Allison J. B. Chaney, Brandon M. Stewart, and Barbara E. Engelhardt, “How Algorithmic Confounding in Recommendation Systems Increases Homogeneity and Decreases Utility,” *Proceedings of the 12th ACM Conference on Recommender Systems* (2018), 224–32, <https://doi.org/10.1145/3240323.3240370>.

Recommender systems learn users' preferences through their interaction on the platform, which leads to recommendations that impact users' interaction. This results in a feedback loop.⁸

The ethical implications of an increase in homogeneity of behavior are not necessarily obvious when we consider products on Amazon, music on Spotify, or movies on Netflix. However, other empirical research examines YouTube's recommender system⁹ to consider how our beliefs might be affected by filtering algorithms. This research is motivated by anecdotal reports of increased recommendations of conspiracy theories on the platform. The study confirms that some topics such as "natural foods" or "firearms" are likely to lead the viewer, through a series of recommendations, to videos that promote conspiracy theories. This is one example of a larger phenomenon that Alfano, Carter and Cheong call *technological seduction*, "technologically-mediated cognitive pressures and nudges that subtly but systematically induce acceptance of problematic beliefs".¹⁰ Recommendations can then turn reasonable searches into extreme recommendations because of learned correlations.

When we consider the context of dating apps, the user is not browsing through products or news but through a potential dating pool. If filtering algorithms can homogenize behavior and polarize beliefs, are they able also to affect our romantic and sexual desires? My goal is to offer the reader reasons to believe so. The effect that filtering algorithms on dating apps might have on user sexual and romantic behavior is ignored in the extensive research around designing collaborative filtering algorithms for dating apps.¹¹ I will mainly focus on race since there is established empirical research I can rely on. I suspect that dating apps can shape different kinds of preferences and behaviors, and my discussion might generalize from race to other issues.

III. RACE AND ONLINE DATING

The first step toward building a collaborative filtering algorithm is to figure out how to group similar users together. For example, on Google

⁸ Ibid.

⁹ Mark Alfano, J. Adam Carter, and Marc Cheong, "Technological Seduction and Self-Radicalization," *Journal of the American Philosophical Association* 4, no. 3 (2018): 298–322, <https://doi.org/10.1017/apa.2018.27>.

¹⁰ Ibid., 6.

¹¹ Kun Tu et al., "Online Dating Recommendations: Matching Markets and Learning Preferences," in *Proceedings of the 23rd International Conference on World Wide Web - WWW '14 Companion* (the 23rd International Conference, Seoul, Korea: ACM Press, 2014), 787–92, <https://doi.org/10.1145/2567948.2579240>; Oghenevwe Otake and Chidiebere Ugwu, "Online Matchmaking Using Collaborative Filtering an Reciprocal Recommender Systems," January 20, 2018, <https://doi.org/10.9790/1813-0702010721>; A. Krzywicki et al., "Collaborative Filtering for People-to-People Recommendation in Online Dating: Data Analysis and User Trial," *International Journal of Human-Computer Studies* 76 (2015): 50–66, <https://doi.org/10.1016/j.ijhcs.2014.12.003>; Peng Xia et al., "Reciprocal Recommendation System for Online Dating," in *Proceedings of the 2015 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining 2015 - ASONAM '15* (the 2015 IEEE/ACM International Conference, Paris, France: ACM Press, 2015), 234–41, <https://doi.org/10.1145/2808797.2809282>.

News and on YouTube, people are often grouped together on a political spectrum. This allows conservative users to receive news from conservative sources, and liberal users to receive news from liberal sources. This grouping leads to the creation of epistemic structures known as filter bubbles.¹² In filter bubbles, relevant voices are excluded by omission, amplifying how confident we are in our beliefs since they are reinforced by “echoing” testimonies. When algorithms impose epistemic filters on us, important views are excluded from the information we receive, which can in turn lead to an inflated sense of self-confidence.¹³ On dating apps, there is strong reason to believe that race is an important grouping factor, and if dating app users are grouped by race, then mechanisms similar to filter bubbles could segregate the potential dating pool of users along racial lines, reinforcing existing patterns of preferences and homogenizing behavior.

Some dating apps allow users to identify their race and the race that they would prefer in a romantic or sexual partner. If users choose to share this data, the algorithm can easily group people by race and learn their preferences from the explicit data it has access to. But even when users refuse to state any race or racial preference, the collaborative data still allows the algorithm to make predictions and recommendations that might fall along racial lines. One example is the dating app Coffee Meets Bagel. With anecdotal stories of users only receiving recommendations of their own race, even when they had no stated preferences, the app developers explained:

Currently, if you have no preference for ethnicity, our system is looking at it like you don't care about ethnicity at all (meaning you disregard this quality altogether, even so far as to send you the same every day). Consequently we will send you folks who have a high preference for [users] of your own ethnic identity, *we do so because our data shows even though users may say they have no preference, they still (subconsciously or otherwise) prefer folks who match their own ethnicity*. It does not compute “no ethnic preference” as wanting a diverse preference. I know that distinction may seem silly, but it's how the algorithm works currently.¹⁴

The upshot here is that algorithmic filtering can override individual preference, even when such preference is explicitly stated, because the preferences of users collectively might form better predictions of successful matches. In other words, the algorithm makes predictions based on implicit aggregate data rather than explicit individual data, as if it can predict your preferences better than you do.

¹² Eli Pariser, *The Filter Bubble: What the Internet Is Hiding from You* (London: Viking, 2011).

¹³ C. Thi Nguyen, “Echo Chambers and Epistemic Bubbles,” *Episteme* 17, no. 2 (2020): 141–61, <https://doi.org/10.1017/epi.2018.32>.

¹⁴ Katie Notopoulos, “The Dating App That Knows You Secretly Aren't Into Guys From Other Races,” *BuzzFeed News*, January 14, 2016, <https://www.buzzfeednews.com/article/katienotopoulos/coffee-meets-bagel-racial-preferences>.

Other dating apps do not ask their users to explicitly state their race or ethnicity. However, as mentioned, filtering algorithms can still pick up on patterns of behavior while being blind to content. Christian Rudder, founder of OkCupid, explains that “racial neutrality is only in theory” since the algorithm can easily guess the race of users based on other characteristics on their profiles. Rudder says that “one of the easiest ways to compare a black person and a white person (or any two people of any race) is to look at their match percentage,” which is OkCupid’s way to determine compatibility.¹⁵ To return to our simple example, the algorithm does not need to know about the relationship between “chips” and “salsa” in order to learn to recommend salsa to anyone who buys chips. All that is needed is a high correlation between buying one and buying the other. Similarly for dating apps, the algorithm need not know anything about the race of the users, but if people of the same race or ethnicity behave similarly, then the algorithm will be able to group them together without users stating their race on their profile.

Indeed, there is overwhelming evidence that shows that dating platform users segregate themselves along racial lines. OkCupid published, between 2009 and 2014, data about their users that reflects their racial preferences. The data shows that, overall, people show a preference for others of their own race. Men, except for black men, are significantly less likely to rate a black woman’s profile favorably compared to the profile of women of other races. Asian men and black men are subject to this same bias, except from women of their own race.¹⁶ Another survey of over six thousand heterosexual internet dating profiles shows that white men and white women are significantly less likely to be excluded from dating and sexual considerations: “Asians, blacks and Latinos are more likely to include whites as possible dates than whites are to include them.”¹⁷ Black people are ten times more likely to reach out to a white person than white people are to reach out to a black person.¹⁸ A number of other empirical studies confirms those trends: users on online dating platforms seem to segregate themselves based on race, exclude people of color from consideration (except those of their own race), and generally show a preference for white men and women.¹⁹

¹⁵ Christian Rudder, *Dataclysm: Who We Are When We Think No One’s Looking* (New York: Crown Publishers, 2014), 101.

¹⁶ Christian Rudder, “Race and Attraction, 2009–2014,” *OkTrends* (blog), accessed September 16, 2020, <https://web.archive.org/web/20140911001651/http://blog.okcupid.com/index.php/race-attraction-2009-2014/>.

¹⁷ Belinda Robnett and Cynthia Feliciano, “Patterns of Racial-Ethnic Exclusion by Internet Daters,” *Social Forces* 89, no. 3 (2011): 819.

¹⁸ Gerald A. Mendelsohn et al., “Black/White Dating Online: Interracial Courtship in the 21st Century,” *Psychology of Popular Media Culture* 3, no. 1 (2014): 2–18, <https://doi.org/10.1037/a0035357>.

¹⁹ Rudder, *Dataclysm*; Glenn T. Tsunokai, Allison R. McGrath, and Jillian K. Kavanagh, “Online Dating Preferences of Asian Americans,” *Journal of Social and Personal Relationships* 31, no. 6 (September 2014): 796–814, <https://doi.org/10.1177/0265407513505925>; Ken-Hou Lin and Jennifer Lundquist, “Mate Selection in Cyberspace: The Intersection of Race, Gender, and Education,” *American Journal of Sociology* 119, no. 1 (2013): 183–215, <https://doi.org/10.1086/673129>; Jay P. Paul, George Ayala, and Kyung-Hee Choi, “Internet Sex Ads for

This exclusionary behavior is extremely common on dating apps for gay and queer men. Since users tend to be anonymous, many state their preferences explicitly in their profiles: “no blacks, no Asians,” “white only,” and so on.²⁰

Before we move on, let me make a couple of quick points about the data. With those numbers in mind, I want to reiterate that the algorithm does not need to classify users by their race or ethnicity to make recommendations that follow racial categories. Take, for example, the profile of a heterosexual black man on an app like Tinder. Asian women will statistically rate the profile of black men lower than the profile of other men. The algorithm can learn not to recommend his profile to users who exhibit similar patterns of preferences (other Asian women), without knowing anything about the race of the users. Second, note that the racial demographic of dating apps reflects the larger demographics of Internet users in the United States. For example, on OkCupid, about 80 percent of users are white (compared to 78 percent of Internet users).²¹ And so, if we consider the larger dataset that the algorithm is learning from, it will lean toward the racial preferences of white users. Regardless of how users are grouped, race will be a strong confounding factor in their recommendations. For a great example of how such data can affect recommendations, I direct the reader to MonsterMatch.²² The website allows users to build a fictional dating app profile and swipe right and left on profiles of monsters and humanoids. It simulates the algorithm that powers dating apps and shows users exactly which profiles were left out from their dating pool and for what reasons.

To sum up, the collaborative filtering algorithms that power dating apps learn to classify users based on their race since the preferences of racial groups are usually similar enough to warrant such grouping. Racial groups on dating apps tend to segregate themselves, preferring people of their own race. Generalizing beyond racial groups, users tend to show a preference for white users, men seem to show a bias against black women, and women seem to show a bias against Asian men. Since correlations lead to recommendations through filtering, users on dating app will be recommended other users of their own race. And if they are grouped with users regardless of race, users will be recommended white users at higher rates, heterosexual men will have fewer black women in their recommendations, and heterosexual women will have fewer Asian men in their recommendations.

MSM and Partner Selection Criteria: The Potency of Race/Ethnicity Online,” *Journal of Sex Research* 47, no. 6 (November 2, 2010): 528–38, <https://doi.org/10.1080/00224490903244575>.

²⁰ Xiaofei Liu, “No Fats, Femmes, or Asians,” *Moral Philosophy and Politics* 2, no. 2 (2015), <https://doi.org/10.1515/mopp-2014-0023>.

²¹ Rudder, *Dataclism*.

²² “MonsterMatch,” MonsterMatch, accessed September 16, 2020, <https://monstermatch.hiddenswitch.com/>.

IV. SHAPING OUR SEXUAL AND ROMANTIC PREFERENCES

The last step in my argument is to establish that the recommendations that result from these filtering algorithms can affect user behavior. On platforms like Google and YouTube, affecting user behavior and preferences is exactly why filtering algorithms are deployed in the first place. We have also seen evidence that such filtering works: the power of recommendations is extremely effective in creating structures such as filter bubbles. It is not surprising that this effect could extend to the dating realm, when the same technologies are deployed to filter who we might find romantically or sexually attractive. I will end the section by raising some potential issues with the influence of algorithmic matchmaking.

First, dating apps are excluding users from others' dating pool as a result of the collaborative filtering. The effects of filtering are obvious: if you don't see someone's profile, then you cannot match or start a conversation with that person. Second, dating apps are actively suggesting some users as "good matches," and recommendations make dating app users more willing to interact with others. A research conducted by OkCupid concludes that "when we tell people they are a good match, they act as if they are [even] when they should be wrong for each other."²³ The power of recommendations is especially relevant when considering the literature on implicit bias. The imagery that we are exposed to can greatly influence the implicit biases that we hold toward groups of people.²⁴ Through their recommendations, dating apps can influence who users see as a "good match," affecting who they consider to be desirable. Mechanisms similar to Alfano, Carter, and Cheong's technological seduction²⁵ could then be at play on dating apps: pressures and nudges that subtly but systematically affect who we match with, talk to, and eventually date. As online dating platforms become increasingly popular, there is no doubt that filtering that happens outside of users' control affects their romantic and sexual behavior.

Filtering and recommendations can even ignore individual preferences and prioritize collective patterns of behavior to predict the preferences of individual users. This effectively homogenizes the behavior of those who are grouped together: they will receive the same recommendations and will tend toward matching with the same people. Not only do users have no control over what group they are placed in, the algorithm is likely to pick up on racial categories to form those groups, which ignores the preferences of users whose preferences deviate from the statistical norm. The recommender system can further amplify this process through a feedback loop:

²³ Christian Rudder, "We Experiment On Human Beings!" 07/28/2014, *OkTrends* (blog), accessed September 16, 2020, <https://web.archive.org/web/20140728200455/http://blog.okcupid.com/index.php/we-experiment-on-human-beings/>.

²⁴ Chloë FitzGerald et al., "Interventions Designed to Reduce Implicit Prejudices and Implicit Stereotypes in Real World Contexts: A Systematic Review," *BMC Psychology* 7, no. 1 (2019): 29, <https://doi.org/10.1186/s40359-019-0299-7>.

²⁵ Alfano, Carter, and Cheong, "Technological Seduction and Self-Radicalization."

if users are repeatedly recommended others of their own race, they will match with people of their own race at higher rates than with others. The algorithm can then use this as further data and additional evidence to continue with its existing pattern of recommendations. And so, if dating apps are influencing users' behavior, they do so by homogenizing this behavior through collaborative recommender systems and deepening racial biases through feedback loops.

The reader might consider that there is absolutely no harm done in the process. After all, we see no issue with recommender systems on Amazon or Spotify prioritizing some products over others. Even recommender systems that amplify filter bubbles are not obviously reprehensible, but only when they might lead to epistemically questionable practices and false beliefs. Indeed, one might think that dating apps are even more tempting now with the power of algorithmic matchmaking, even when the patterns that the algorithm learns and amplifies show deep racial biases. After all, sexual desires, and desires in general, resist moral criticism. We think of our preference for a certain body type or hair color to be out of our control and deeply personal. It would be strange for someone to praise us for being attracted to someone or blame us for our lack of attraction to someone else. Megan Mitchell and Mark Wells argue that we are morally justified from excluding certain people from our dating pool.²⁶ Xiaofei Liu also shows that there is nothing wrong with what he calls "simple looksism"—that certain physical features are "deal breakers" for our sexual or romantic consideration is perfectly okay.²⁷ If the algorithm is successfully and accurately able to determine and predict users' sexual and romantic preferences, then whatever patterns dating apps pick up on and extend should be irrelevant to a moral evaluation of the algorithmic filtering.

Yet, this comes in sharp contrast with another intuition some readers might have: to exclude everyone of a certain race from any romantic or sexual consideration seems problematic. After all, as we have seen in Part II, patterns of romantic and sexual attraction in the United States often reflect larger patterns of exclusion. Liu, for example, argues that there is a morally relevant difference between simple looksism and racial looksism. He argues that racial looksism is an overgeneralization: it assumes that people of a certain race will always look a certain way, when race does not determine the way a person will look.²⁸ Additionally, Mitchell and Wells argue that racialized sexual and romantic biases have morally relevant social meaning grounded in a history of discrimination including, for example, prohibitions on interracial marriage.²⁹ If this is true, then dating apps are contributing to

²⁶ Megan Mitchell and Mark Wells, "Race, Romantic Attraction, and Dating," *Ethical Theory and Moral Practice* 21, no. 4 (2018): 945–61, <https://doi.org/10.1007/s10677-018-9936-0>.

²⁷ Liu, "No Fats, Femmes, or Asians."

²⁸ Ibid.

²⁹ Mitchell and Wells, "Race, Romantic Attraction, and Dating," 956.

this wrong by exacerbating problematic sexual and romantic biases, as they can homogenize and deepen exclusionary preferences.

This leads to a tension in the design choice that a dating app might adopt. If we accept the obligation to resist deepening racial bias through filtering, then recommender systems ought to be designed in a way that avoids racially exclusive recommendations. But why should the algorithm resist the preferences of users who do hold such biases? At this point, we are asking dating apps to serve a function beyond the one that we started with, which was simply to learn user preferences and build recommendations based on them. An algorithm that resists biased preferences cannot do so without serving the preferences of some and not others, which is exactly the problem that we started with.

Regardless, I believe the issue lies deeper than biased recommendations: users have absolutely no control over the filtering that determines who they see on dating apps. As mentioned, stated preferences are sometimes overridden by algorithmic predictions. Using collaborative data in the context of dating apps seems to override extremely personal sexual and romantic desires. One interesting suggestion might address the tension we have encountered: Hutson et al. argue that with randomized recommendations, users can break out of the patterns that the algorithm reinforces.³⁰ This does not mean that the project of filtering is scrapped altogether. Rather, random recommendations will be part of the filtered results to allow users to explore beyond the algorithm's limit. If dating apps allow their users to branch out from what the algorithm considers a safe match, they could break patterns that the recommender system amplifies.

V. CONCLUSION

I discuss in this essay concerns that might be raised by the use of collaborative filtering on dating apps. I argue that collaborative filtering is especially effective at homogenizing behavior and amplifying existing patterns of preference. Dating app users often segregate themselves by race, showing a preference for people of their own race. Other racial biases are also at play online when we look at larger patterns of preference and behavior. Deploying collaborative filtering algorithms on dating apps can then homogenize the behavior of users of the same race, and deepen existing racial biases among online daters.

One goal of this essay is to show the extent to which recommender systems can influence user behavior. Extensive research shows how effective recommender systems are on shopping platforms or social media. But if recommender systems can affect what we buy and what we watch, and if

³⁰ Jevan Hutson et al., "Debiasing Desire: Addressing Bias and Discrimination on Intimate Platforms," *Proceedings of the ACM on Human-Computer Interaction* 2, issue CSCW (November 2018): 1–18, <https://doi.org/10.1145/3274342>.

those same systems are deployed on dating apps, then we can establish that they are also influencing who we date. Another goal of this essay was to bring attention to how collaborative filtering algorithms can learn from and amplify existing patterns of behavior. There has been recent interest in news filtering and the creation of filter bubbles: existing beliefs are echoed through news recommendations to artificially augment how confident we should be in those beliefs. Similarly, a dating app user who exhibits certain patterns of sexual or romantic preferences will have those patterns exacerbated through a feedback loop. Finally, I hope that this essay is a first step toward bringing together recent work on algorithmic justice, with the rich literature on sexual and romantic desires. It is extremely challenging to think about how our desires are shaped and if they could hold moral value. Looking at dating apps allows us to study these issues in a controlled and artificial environment; I hope, however, that my discussion does not avoid the hard questions by simplifying the reality of dating, but rather sets up a framework to address these questions. The ethics of dating cannot be divorced from discussions of online dating: as I mentioned, more new couples meet online than by any other method.

Philosophy, University of Texas, USA