"We Know You Better Than You Know Yourself": Aesthetic Taste and Recommender Systems

Kenneth L. Brewer
Montclair State University

Several years ago, walking out of a screening of *The Hangover* ("Absurd, Raunchy, Irreverent" according to Netflix), I said to my friend, "Well, that wasn't very good." "What are you talking about?" she replied. "You laughed the entire movie." While it is possible that I might have laughed a lot and still not liked the movie (perhaps the laughs were what we call "cheap laughs"), I was startled. If I had been alone, I would have walked out of *The Hangover* firmly believing that I had not enjoyed it, when apparently, I had enjoyed it tremendously. I like to think that I know what I like, but this seemingly trivial incident made me wonder how much I really understand my own taste.

Naturally, I turned to philosophy for answers to the worrisome question of whether I am in fact a fan of absurd, raunchy, irreverent bromance comedies. It turns out that my befuddlement about what I like and do not like makes me completely typical, at least according to a number of philosophers working in aesthetics who have pored over a considerable amount of data from the behavioral and brain sciences.[1] The most thorough account so far is Kevin Melchionne's "On the Old Saw 'I know nothing about art but I know what I like.'"[2] It is entirely likely, Melchionne argues, that I do not know what I like or why I actually like what I believe that I like. The unreliability of introspection seriously undermines our knowledge about our own taste, and Melchionne attacks what he views as our complacency about knowing what we like. Melchionne maintains that self-knowledge about what we like and why are necessary conditions for a valuable aesthetic life. Absent these things, we can certainly have aesthetic pleasures, but

we cannot cultivate them, as in the absence of self-knowledge "the pursuit of aesthetic satisfaction would be no more than a fishing expedition" (132). Presumably, we want more efficiency from our aesthetic experiences. We also like to understand our taste as consistent, and to see some sort of coherent narrative in our aesthetic lives.[3]

To my knowledge, Melchionne's argument has not been convincingly answered. My approach will be to accept his well-supported attack on the unreliability of introspection, and focus on one possible solution to this problem.[4] Melchionne rejects skepticism about aesthetic self-knowledge, instead recommending a "weak fallibilism" that "leave[s] room for reliable aesthetic self-knowledge while recognizing that our aesthetic responses often remain as opaque as the rest of conscious life" ("On the Old Saw…" 140). Is there any way, however, to make my aesthetic responses less opaque? I will explore a decidedly new-fangled method for better understanding one's own taste: recommender systems that rely on algorithms to predict whether or not a user will find a certain item satisfying.[5] We are familiar with these from Netflix, Amazon, Spotify, and so on. There is a vast technical literature in computer science-related fields on how to improve the performance of these systems.[6] Conversely, there is a fairly substantial amount of discussion in what we might loosely term "cultural studies" about the possible effects of such systems on individuals and societies; these effects are usually seen as pernicious.[7] Even though most of these systems are well into their second decade, philosophers have had relatively little to say about them.[8] I will discuss the possibility that technologically sophisticated recommender algorithms such as Netflix can enhance our knowledge of our own taste precisely by improving our capacity to introspect about what we like and why.

Melchionne begins by pointing out that the infallibility of introspection has pretty much been taken for granted in philosophical aesthetics, which has been more concerned with how "individuals with self-transparent taste disagree with others equally aware of their preferences" (133). Melchionne draws on a number of empirical studies to suggest "that ignorance or error about our inner states is common," such as Timothy Wilson's *Strangers to Ourselves* (134). Nor is Melchionne inclined to accept the view that while this ignorance may be true of many areas of my emotional life, aesthetic experience is special, as John Dewey argues. In fact, Melchionne notes, that we go to such elaborate lengths to avoid distractions during our aesthetic experiences (headphones, darkened movie theaters) suggests that we are well-aware of the fragility of aesthetic experience and aesthetic attention (135).

Citing John Lambie and Anthony Marcel on "emotion unawareness," Melchionne argues that there are three ways we might get knowing what we like wrong: we might have aesthetic responses without second-order awareness of them; we might have aesthetic responses with inaccurate second-order awareness; and we might give "false attributions of the causes of aesthetic response" (135).

With the first, aesthetic responses without second-order awareness, Melchionne is particularly interested in cases in which belief about my "aesthetic personality," or "biographical taste" ("I like French New Wave films") leads me to simply ignore moment-to-moment aesthetic responses that conflict with those attitudes. Thus, if I thought before *The Hangover*, "I'm not really the kind of person who enjoys raunchy comedies: I much prefer the mid-period works of Ingmar Bergman," then I would simply ignore evidence—and it was apparently considerable—that I am precisely the kind of person who likes raunchy comedies.

With the second possibility, I am aware of my affective state but may mis-categorize it. For example, with an avant-garde work of art, I misunderstand my confusion as dislike. Worse, this problem apparently cannot be solved by simply focusing more intensely on my introspective states. An interesting example of this problem was found by T. D. Wilson and J. W. Schooler, who conducted a study on pleasure in listening to music. Subjects were asked to listen to a piece of music, and either given no instructions, instructions to try to be happy, instructions to monitor their happiness, or instructions to try to be both happy and monitor their happiness.

Subjects who monitored their happiness were significantly less happy while listening to the piece than subjects in the control group who were given no instructions. Conclusion: "Vigilantly monitoring one's ongoing hedonic experience can undermine one's ability to actually gain happiness" (Wilson and Schooler 191).[9]

Finally, and perhaps most significantly, we tend to believe that taking the "aesthetic attitude" gives us more insight into what we feel. I regard aesthetic experience as a highly self- conscious activity in which I can give clear explanations for why I'm feeling what I'm feeling because I pay such close attention to the experience. In daily life, I may not know why I'm angry or sad, but with a work of art, I expect myself to be able to explain the source of these emotions. But this very effort to attribute reasons to what we feel leads, Melchionne argues, to distortions, what social psychologists term the "introspection effect." Figuring out why I like Cézanne (why looking at a Cézanne painting causes certain mental states that I classify as "pleasurable") is not at all straightforward, and I often end up relying on critical clichés I have picked up at some point in my life (Melchionne, "On the Old Saw…" 139).

So how might recommender systems deal with these three problems of reliable aesthetic self-knowledge? A couple of initial points. I am going to focus on Netflix because I believe its algorithm offers the most interesting attempts, so far, to solve these problems. At the same time, I am aware that Netflix is a profit-making enterprise. We do not know, for example, exactly how its algorithm works (it is proprietary) and this lack of transparency is obviously not ideal.[10] Moreover, there is general agreement that Netflix's algorithm promotes shows that Netflix has itself produced and favors television shows (which create longer user engagement)

over films.[11] However my aim is to explore the basic aesthetic issues these systems raise, not engage in social critique.

Briefly, it helps to have some idea of how these systems work.[12] Most of them typically use some form of *collaborative filtering*, either user-based on item-based. Based on my own ratings or behavior, I am paired up with similar users, and this similarity is used to predict my rating for a film I have not seen but similar users have seen. In an item-based system, the algorithm isolates items that tend to be consumed together (peanut butter and jelly, or Tom Cruise and Liam Neeson action movies) and makes recommendations accordingly. There are also *content-based* systems that mark features of what a user consumes and then recommends similar items based on these features: this is the "if you liked X you'll love Y" appeal present in film trailers from way back, updated into an algorithm. Most recommender systems use a combination of these and are termed *hybrid* systems, as all of them have significant limitations on their own.[13]

So the first problem Melchionne describes: I ignore my own responses because of my preconceptions about the kind of aesthetic person that I am. I do not notice the pleasure I am taking in *The Hangover* because I see myself as more of an Ingmar Bergman film sort of person. My sense of my biographical taste does not jive with my actual aesthetic experiences: my revealed preferences are quite different from my stated preferences.[14] If you go on Movielens, a research project on film recommendations based at the University of Minnesota, you are asked to rate as many movies as possible on a 1-5 scale.[15] According to the research Melchionne describes, my ratings will not tend to be very accurate. I will rate movies highly that I believe I am supposed to have liked, not even aware that I actually did not enjoy them.

While Netflix still allows a thumbs-up or a thumbs-down, this metric can be relatively crude most likely because ratings do not really matter that much for the reasons just explained. So Netflix's algorithm guru Todd Yellin has argued that I will probably give Al Gore's *An Inconvenient Truth* five stars because I think it is an important film, but it turns out that I actually watch *Paul Blart: Mall Cop 2* several more times a year than I watch the documentary.[16] As far as Netflix is concerned, it is pretty obvious which film I actually prefer, so it no longer cares to compile ratings I might give them. When Netflix shifted over to a streaming service, it began to intensively monitor user behavior with the interface, termed "data-behaviorism" by Sarah Arnold. As Arnold puts it, the user, for Netflix, is no longer an individual but a set of data points. Our behavior tells Netflix what we like, not our (suspect) ratings (Arnold 50). In rating films, we may think we know what we like, but from our behavior, Netflix actually does know what we like, and then offers a prediction: Film X is a "95% match" to our preferences. These preferences are drawn from my behavior, not from other data about me. In fact, Yellin argues, much of my biography is not relevant to my taste: my gender, class, and so on do not matter.[17]

Freed of the need to rely on users' (suspect) ratings, Netflix thus addresses the second problem of how monitoring my hedonic states actually undermines my self-knowledge about them. As Sherry Irvin notes, data from the behavioral and brain sciences indicates that if I know that I have to rate an item, my likes and dislikes typically change significantly (38). With Netflix, if all of my behavior matters, then on some level, none of it really matters in terms of affecting my behavior. Without having to monitor my hedonic state because I'm constantly thinking, "Is this film 3.5 stars or only 3?", I may be better aware of what my hedonic state actually is.

The third problem Melchionne discusses is my confident belief that I can give clear reasons for why the work of art caused certain feelings: "I felt nervous watching *Psycho* and this was due to Hitchcock's use of overhead shots." This may very well be the case, Melchionne argues; it may also just be something I learned about Hitchcock in Film Studies 101, and I'm just repeating it.[18] As Irvin explains, this is the problem of "coarse-grainedness": "Aesthetically relevant aspects of a work fail to have the expected effect on people's judgments" (39). What is really affecting me in *Psycho* is Marion Crane's moral struggle, but I overlook this, instead attributing my feelings to the usual sources.

Netflix tries to deal with the coarse-grainedness problem with "tags." Described in a hagiographic 2014 article in *The Atlantic*, Alexis C. Madrigal's "How Netflix Reverse-Engineered Hollywood," Netflix developed a 36-page manual for employees who work at tagging movies with descriptors. Netflix has broken down the dozen or so film genres we are familiar with (drama, comedy, action, horror, etc.) into 76,897 micro-genres. You may think of yourself as someone who enjoys horror, but what you actually enjoy are "Cult Evil Kid Horror Movies."

These are Kendall L. Walton's "Categories of Art" on steroids.[19] Is a movie funny? Not good enough. How is it funny? Is it witty? Gross-out? Quirky? Drily witty? And so on.

Madrigal's homage to Netflix's genius specifically makes a claim about introspection: "It's not just that Netflix can show you things you might like, but that it can tell you *what* kinds of things those are. It is, in its own weird way, a tool for introspection." Madrigal does not explain how precisely Netflix's tags function this way, but Emily Lawrence has suggested that Netflix has the capacity to make us go "Huh" about our taste, which should lead to self-reflection about what we actually like and why we like it.[20] Essentially, the tags function as a form of aesthetic education, and the next time I am enjoying a movie, I can presumably go beyond critical clichés ("I like Hitchcock because he was the master of suspense") to "I like *Rear Window* because it is cerebral and features a strong female lead."

If we lack self-knowledge about our own tastes, then we can expect inconsistency in our choices. What do recommender systems do about this problem, raised in its classic form in Ted Cohen's "On Consistency in One's

Personal Aesthetics"? Cohen describes his confusion at his wife's seemingly inconsistent tastes. In particular, he is troubled by her love the music of Béla Bartók and his dislike of the music of Aaron Copeland, specifically because she enjoys Bartók's use of folk music and *dislikes* Copeland for his use of folk music idioms (107).

Typically, Cohen argues, when we like an object x for reason R, we expect to like another object y that also features R.

What do we do? Cohen offers two basic possibilities, and I suggest that recommender systems have some relevance to both of these (122). The first is that I simply have not looked hard enough at why I like object x or dislike object y. My reasons are shallow ones. With recommender systems, tags should help alleviate this problem. I enjoy the films of Éric Rohmer, and I suspect that I will also enjoy the films of Jean-Luc Godard, a fellow French New Wave director, as they are similar in significant ways. I do not enjoy Godard's films. However, I do enjoy the films of French New Wave director Claude Chabrol. It turns out--when I examine the tags--that what I actually like are films with strong female leads, which Chabrol's and Rohmer's films have in common. If, as Melchionne argues, Cohen's wife's problem is not inconsistency but instead a vague understanding of her own taste ("opacity"), the tags should/might make her understanding less vague (Melchionne, "On the Old Saw…" 133).

A second possibility is that I do not have a "style," as Cohen puts it. My tastes are simply not consistent, and while there is educational value, he notes, in attempting to clarify my taste to myself, it may just be that there is no there there (Cohen 122). On some level, recommender systems accept inconsistency as a fact about our taste. While they need to make successful predictions to stay in business, the systems are flexible enough to accommodate what seem to be gross inconsistencies in ratings.[21] Netflix does not care if you like Bartók and dislike Copeland; it just needs to be able to successfully predict if you are going to like Phillip Glass and the system is designed to figure this out even if Cohen cannot.[22]

Melchionne concludes that "[b]y embracing anti-introspectivism, we move toward a more naturalistic view of taste, with more room for confusion, indifference, and transience in our aesthetic preferences" ("On the Old Saw…" 140). Given his earlier call for "weak fallibilism," the point is that we dethrone introspection from its position of privilege rather than rejecting it entirely as a means to understanding our own taste, though one could certainly use recommender systems to avoid introspection entirely. Netflix is confident that I would like this, so I must have liked it.[23] Dominic McIver Lopes recommends that "it is prudent to turn to data about appreciative practice gathered through empirical methods that do not rely on subjective reports" (34).[24] With its "data behaviorism," Netflix and other recommender systems seem to be carrying out, at least in part, this kind of program; these sites do not just offer opportunities for aesthetic pleasure, but also ways of thinking about aesthetic practices that are

highly relevant to recent discussions in philosophical aesthetics, and not just in the area of aesthetic self-knowledge.

Netflix is 95% confident that I would like *The Hangover*.

NOTES

1. Robson suggests that most of us, most of the time, are entirely too confident about believing that we know what we like.
2. In addition to Melchionne, see also Irvin, Lopes, and Robson for detailed discussions of the fallibility of aesthetic introspection.
3. On the importance of consistency in one's personal taste, see Cohen.
4. Though she does not discuss Melchionne's arguments, Irvin explores similar research and is a bit more optimistic about the reliability of introspection with aesthetic experience, suggesting that aesthetic "experts" are less prone to errors in this regard than laypeople (41).
5. Shriver offers a standard definition: "Recommender systems filter information to help users make decisions when lacking personal experience or knowledge or when the set of choices is overwhelmingly large. In general, they take in data about users' past preferences for items and aggregate it to predict a user's preference for unknown items, often presenting a list of the items most likely to be preferred" (1).
6. Aggarawal offers a clear and thorough overview.
7. To be fair, these are partly in response to popular media Odes to the Geniuses of Silicon Valley; Madrigal (2014) is a typical example. Lawrence offers a clearly argued, philosophically-informed account that reaches some well-earned pessimistic conclusions about the loss of personal autonomy recommender systems, especially Netflix, might encourage.
8. Melchionne notes that they must work pretty well, given user satisfaction ("On The Old Saw…" 140). In "Aesthetic Choice" he suggests that we "keep the recommendations of the streaming firms at arm's length" (298), as these firms are concerned first about their own interests, not ours. This is fairly mild stuff compared to Frankfurt School-influenced cultural studies folks, who argue that an arm's length is nowhere near far enough from these algorithms that are undermining our taste and destroying our capacity to think critically. If Melchionne is Jeff Lebowski from *The Big Lebowski*, saying to Netflix, "That's just like, your opinion, man," the cultural studies theorists are Dr. Bennell at the end of *Invasion of the Body Snatchers*, screaming that the pod people are coming for you.
9 . Melchionne discusses this research as well (see "On the Old Saw…" 138).
10. Ariellii is especially critical of Netflix's lack of transparency with its algorithm, arguing that this kind of corporate secrecy is anti-democratic (89). He also argues that Netflix aims at increasing consumption more than it aims at increasing introspection on the part of its users (79).
11. See Alexander 90.
12. Morris provides a good overview (451).
13. Bobadilla has a good discussion.
14. Lawrence 359. Arnold terms this *user expression* versus *user behavior* and notes that Netflix sees former as a poor guide to taste (58).

15. See Arielli (84) on Netflix's shift from what he agrees are the flawed star ratings that Movielens still uses, as does Amazon.

16. Arnold discusses Yellin's claims (Arnold 58).

17. Gray quotes Yellin claiming in an interview that Netflix has put traditional markers of identity "in the garbage heap" (Gary 80).

18. Lawrence: "Netflix essentially mediates the common struggle to articulate why you like what you like by providing a brief account of what features you enjoy" (360).

19. Lawrence has a brief discussion of Walton in relation to Netflix's tags (360). This is obviously an area of rich interest that I will explore further in a subsequent paper.

20. Lawrence(359. See also Arielli 78-79. In theory, of course. I am not defending the actually existing tags as necessarily good. Lawrence offers a very compelling critique of their weaknesses.

21. See Alexander's discussion of Netflix's assumption that we are "consistently inconsistent" (84).

22. Cohen also discusses a third possibility, a "deep truth" that consistent aesthetic principles are simply not possible to find (122). This is a larger issue than I can tackle in this essay. What is also interesting about Cohen's example is that it is about *dislike*. Why does his wife dislike Copeland's music when she likes music by similar artists? Here, recommender systems have run into problems. As journalist Tom Vanderbilt explains in his *You May Also Like: Taste in an Age of Endless Choice*, these algorithms have had a very difficult time explaining dislike. Your shirt tells me what kind of shirts you like, Vanderbilt notes; it does not tell me what kind of shirts you dislike (89). It seems that disliking is a lot more complicated than liking; according to Vanderbilt, the music service Pandora experimented with allowing users to give a reason for a thumbs-down on a particular piece of music, and the system was overwhelmed with explanations (111). Granted, it is perhaps because Mr. Cohen is pressing her, but his wife has much more to say in Cohen's article about why she dislikes Copeland than why she likes Bartók. This is, apparently, typical, and recommender systems, Vanderbilt notes, are very interested in cracking the problem of dislike.

23. See Ross for a discussion about fears of a loss of autonomy when we heed the advice of critics instead of thinking for ourselves. Obviously, recommender systems raise similar concerns. Ross briefly mentions recommender systems, and intriguingly suggests that we look to "public" rather than "inward" sources to preserve the authenticity of our aesthetic responses, but she does not offer any detailed analysis of how recommender systems might function in this context (367; 370). Arielli questions the mind/outside of mind model, suggesting that recommender systems operate as a kind of extended mind (78). On the extended mind, see Clark. This is another topic worth exploring that is beyond the scope of this essay.

24. He recommends "the methods of the social sciences" (Lopes 39).

WORKS CITED

Aggarwal, C. C. *Recommender Systems: The Textbook*. Springer, 2016.

Alexander, Neta. "Catered to Your Future Self: Netflix's 'Predictive Personalization' and the Mathematization of Taste." *The Netflix Effect: Technology and Entertainment in the 21st Century,* edited by Kevin McDonald and Daniel Smith-Rowsey, Bloomsbury,

2016, pp. 81-98.
Arielli, Emanuele. "Taste and the Algorithm." *Italian Journal of Aesthetics,* vol. 46, 2018, pp. 77-97.
Arnold, Sarah, (2016) "Netflix and the Myth of Choice/Participation/Autonomy." *The Netflix Effect: Technology and Entertainment in the 21st Century,* edited by Kevin McDonald and Daniel Smith-Rowsey, Bloomsbury, 2016, pp. 49-62.
Bobadilla, J., F. Ortega, A. Hernando, and A. Gutiérrez. "Recommender Systems Survey." *Knowledge-Based Systems,* vol. 46, 2013, pp. 109-32.
Clark, Andy. *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*. Oxford University Press, 2008.
Cohen, Ted. "On Consistency in One's Personal Aesthetics." *Aesthetics and Ethics: Essays at the Intersection,* edited by Jerome Levinson, Cornell University Press, 1999, pp. 106-125.
Gray, Jonathan. "Reviving Audience Studies." *Critical Studies in Media Communication,* vol. 34, no. 1, 2016, pp. 79-84.
Irvin, Sherry. "Is Aesthetic Experience Possible?" *Aesthetics and the Sciences of the Mind,* edited by Greg Currie, Matthew Kieran, Aaron Meskin, and John Robson, Oxford University Press, 2014, pp. 37-56.
Lawrence, Emily. "Everything is a Recommendation: Netflix, Altgenres and the Construction of Taste." *Knowledge Organization,* vol. 42, no. 5, 2015, pp. 358-64.
Lopes, Dominic Iver. "Feckless Reason." *Aesthetics and the Sciences of the Mind,* edited by Greg Currie, Matthew Kieran, Aaron Meskin, and John Robson, Oxford University Press, 2014, pp. 21-36.
Madrigal, Alexis C. How Netflix Reverse-Engineered Hollywood. *The Atlantic,* 2 Jan 2014, www.theatlantic.com/technology/archive/2014/01/how-netflix-reverse-engineered-hollywood/282679/.
Melchionne, Kevin. "Aesthetic Choice." *British Journal of Aesthetics,* vol. 57, no. 3, pp. 283-298.
———. "On the old saw 'I know nothing about art but I know what I like.'" *Journal of Aesthetics and Art Criticism,* vol. 68, no. 2, 2010, pp. 131-41.
Morris, Jeremy Wade. "Curation by Code: Infomediaries and the Data-Mining of Taste."*European Journal of Cultural Studies,* vol. 18, no. 4-5, 2015, pp. 446-463.
Robson, Jon. "Aesthetic Autonomy and Self-Aggrandizement." *Royal Institute of Philosophy Supplement,* vol. 75, 2014, pp. 3-28.
Ross, Stephanie. "Comparing and Sharing Taste: Reflections on Critical Advice." *Journal of Aesthetics and Art Criticism*, vol. 70, no. 4, 2012, pp. 363-71.
Shriver, David. "Assessing the Quality and Stability of Recommender Systems." *Computer Science and Engineering: Theses, Dissertations, and Student Research.* 2018, digitalcommons.unl.edu/computerscidiss/147.
Vanderbilt, Tom. *You May Also Like: Taste in an Age of Endless Choice*. Vintage, 2017.
Walton, Kendall L. (1970) "Categories of Art." *The Philosophical Review,* vol. 79, no. 3, 1970, pp. 334-367.
Wilson, T.D. and J. W. Schooler. "Thinking Too Much: Introspection Can Reduce the Quality of Preferences and Decisions." *Journal of Personality and Psychology,* vol. 60, no. 2, 1991, pp. 181-92.