

Understanding User Beliefs About Algorithmic Curation in the Facebook News Feed

Emilee Rader and Rebecca Gray
Department of Media and Information
Michigan State University
{emilee,grayreb}@msu.edu

ABSTRACT

People are becoming increasingly reliant on online socio-technical systems that employ *algorithmic curation* to organize, select and present information. We wanted to understand how individuals make sense of the influence of algorithms, and how awareness of algorithmic curation may impact their interaction with these systems. We investigated user understanding of algorithmic curation in Facebook's News Feed, by analyzing open-ended responses to a survey question about whether respondents believe their News Feeds show them every post their Facebook Friends create. Responses included a wide range of beliefs and causal inferences, with different potential consequences for user behavior in the system. Because user behavior is both input for algorithms and constrained by them, these patterns of belief may have tangible consequences for the system as a whole.

Author Keywords

algorithms; feedback loop; intuitive theories; Facebook News Feed.

ACM Classification Keywords

H.5.m. Information Interfaces and Presentation (e.g. HCI): Miscellaneous

INTRODUCTION

People are becoming increasingly reliant on online socio-technical systems that employ *algorithmic curation*: organizing, selecting, and presenting subsets of a corpus of information for consumption. An algorithm is “a finite, discrete series of instructions that receives an input and produces an output” [16], and systems like Facebook and Google (and many, many others) use algorithms as information intermediaries that determine what information should be displayed and what should be hidden [5].

Because algorithms are automated and usually poorly understood by end users, people often assume that they are objective or impartial [5]. However, just because a system is automated does not mean it is free from the potential for bias.

Researchers are becoming increasingly concerned that algorithms enforce biases that are hard to detect, but have potentially negative outcomes [6, 14].

The Facebook News Feed is a socio-technical system composed of users, algorithms, and content. Users produce content items (posts) that become part of Facebook's corpus, which is continuously changing. At the same time, personalization algorithms select a subset of items from the corpus, rank or organize them according to a proprietary algorithm, and present them to users for consumption in their News Feeds. If a user were to contribute a post that the News Feed algorithm does not display near the top of others' News Feeds, that post becomes effectively invisible to those users. Bucher [6] calls this the “threat of invisibility,” or the potential for one's contributions to go unseen by an unknown number of people. If users become aware of this possibility, avoidance of this “threat” can then guide their behaviors and choices as they engage with others via the system.

The potential for negative consequences like the threat of invisibility are presumably the result of feedback loops: situations where the output of a process becomes an input to that same process. This happens in social media because information consumers are also producers; both explicitly via choices to post, comment or “Like”, and also implicitly via their behavioral traces recorded in system logs. What users learn about the system when they act as consumers can affect the choices they make as producers of content. Feedback loops make all users gatekeepers for each other, by using as input to the algorithm evidence collected from other users about the value (comments, Likes, shares) of the particular item in question [30]. It is extremely difficult to understand the complex, nonlinear interactions that take place in socio-technical system like the Facebook News Feed where the algorithm, the users, and the corpus of content itself are constantly interacting and evolving [27].

Because of the feedback loop characteristics of these systems, user beliefs about content filtering algorithms work are an important component of shaping the overall system behavior. To better understand the interdependence between users and algorithms, we conducted a study investigating users' beliefs about what the Facebook News Feed chooses to display, and why. We describe the kinds of evidence of algorithm behavior that users notice and respond to and the beliefs they form about algorithm selection and ranking criteria, and we present

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

CHI 2015, April 18–23 2015, Seoul, Republic of Korea

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-3145-6/15/04...\$15.00

<http://dx.doi.org/10.1145/2702123.2702174>

implications of these beliefs for their own behavior and for the system overall.

RELATED WORK

Content Streams and Recommendations

Recommender systems researchers have explored ways to use algorithms to help users connect with the content they are most interested in on social media sites. Several research teams have proposed different ways to use information about network ties, topic preferences, and characteristics of posts in the system to make recommendations. For example, Sharma and Cosley [28] proposed using information about content preferences from the Facebook profiles of users and their Facebook Friends to help generate recommendations for movies, television shows, and books. Items recommended by the algorithm variant that used Facebook Friends' profile information received the highest number of views. In a follow-up study [29], they found that such social recommendations were more persuasive when they came from people who were close friends whose interests are known to the user.

Much of this research has taken place using Twitter, which has been a more open platform to use for experimentation and evaluation of systems designed to help users with information overload in their content streams and feeds. Chen et al. [7] created a Twitter app to recommend "conversations" (a chain of @replies, or broadcasted messages directed to others via username) to users, to help them focus on the interactions that would be most interesting to them. The algorithm used conversation dimensions such as thread length, topic relevance, and tie strength to make its selections. The algorithm that performed best was a combination of topic and tie strength.

Bernstein et al. [4] also created a Twitter client to help users focus on information that would be most interesting to them. Their tool clustered tweets within a user's stream by topic, and allowed users to browse those topics. This was essentially a way to sort information in the user's feed into groups of similar items to help them focus on more relevant information. They tested a prototype with users (and users' own Twitter accounts) in the lab. They found that while users liked how it helped them read more relevant tweets when they were grouped by topic, they felt like it was harder to make sure they had seen everything. The typical chronological interface was simply "less enjoyable but more comprehensive." However, this finding may not hold on Facebook where it may be more important for relationship maintenance not to miss posts.

One consideration when designing a filtered feed is how to take relatedness of recommended content items into account. Because information appears in the feed rather than being sought out by the user, it can be jarring to encounter content that seems like it does not belong with the rest. Lv et al. [23] emphasize that relatedness is important for keeping users' attention as they browse from item to item. However, they point out that relatedness as operationalized by an algorithm is not necessarily the same as relatedness perceived by a human being, and that accounting for a user's reading interests when recommending content can be difficult.

In addition to relatedness, it is also important for the system to keep track of the timeliness of the items and not fill the user's content stream with outdated items or the same item over and over again. However, in a system where content items are turning over very quickly like Twitter or Facebook, a greater degree of personalization means the system has fewer items it can possibly recommend to the user. This makes it difficult for the algorithm to select items that are both new to the user and relevant to the user's interests [21].

Finally, several researchers have attempted to evaluate whether filter bubbles (recommendations of decreasing diversity over time) are likely to exist in content recommendation systems. Nguyen et al. [25] analyzed a MovieLens dataset and found that both the diversity of items recommended by the system and items rated by users did indeed become less diverse over time, but only slightly. They concluded that this was weak evidence for a filter-bubble-like effect. In addition, in two separate modeling and simulation studies, independent research teams found that some recommendation algorithms are susceptible to a narrowing effect under particular conditions; however, neither paper presented empirical data to verify the simulation results [11, 17].

Research projects like these use various data about users and content available in social media systems as input to algorithms that are designed to prioritize content for display, and direct user attention to information they want to see. However, few projects have explored interaction with algorithmic curation systems with large numbers of users over long periods of time. It is difficult to evaluate algorithms in isolation from their context of use—without data from real users to act upon, one cannot get an accurate picture of performance [13, 21]. Therefore, bias that might be introduced by interdependence between user and algorithm behavior is largely unknown.

The Facebook News Feed

Facebook is the largest social network site (SNS) in the world with over 1.28 billion monthly active users [9] and over 71% of American adults using the site [1]. The Facebook News Feed is a "constantly updating list of stories from people and Pages that you follow on Facebook" [10]. Maintaining relationships is the main reason most people use Facebook [22], and seeing updates and other shared content from Friends is how users stay informed about what Friends near and far are "up to" and is a source of a perpetual feeling of connection with them [18].

According to a post to the "Facebook Newsroom" on August 6, 2013, any time a user visits his or her News Feed there are on average "1500 potential stories from friends, people they follow and Pages for them to see," and the News Feed algorithm prioritizes "an average of 300 stories out of these 1500 stories to show each day" [2]. Facebook allows users to sort the News Feed in two ways: *Top Stories* or *Most Recent*. The *Top Stories* function displays Friends' posts, ranked by the News Feed algorithm. Factors used to determine the ranking include "the number of comments, who posted the story, and what type of post it is (ex: photo, video, status update)" [10]. Users can exert some control over what appears in the News

Feed if they sort by *Most Recent*, which re-orders posts by reverse chronological order within the News Feed (newest items first). Users can also *hide* people, Pages, or groups, *follow* or *unfollow* Friends, and create *Friend Lists* that when selected only display posts from members of the list.

Feedback Loops in the News Feed

The News Feed algorithm has similar goals to other recommender systems: to connect people with information they are likely to want to consume, by making some items easier to access than other items [26]. However, Facebook’s News Feed is different from other recommender systems in one important way. A feedback loop exists in the News Feed that has the potential to affect not only the information used for selection and ranking of items, but characteristics of the corpus of posts as well.

In a typical recommender system, item ratings can be susceptible to popularity bias caused by a feedback loop [31]. Items that are very popular may be recommended more often than unpopular items, which means popular items are viewed more and receive more ratings. In other words, the output of the ranking algorithm (popular items are recommended and therefore viewed more) affects the inputs to that same algorithm (items that are viewed more receive more ratings). But the corpus of movies or products is relatively stable over time, and only changes when the company revises its catalog of offerings. In the News Feed, both the item “ratings”—comments, Shares and Likes—and the posts users create can be affected by a feedback loop. This happens because users are both consumers and producers of content on Facebook, and what users learn about the system when they consume posts affects what they choose to post as producers.

The News Feed algorithm *curates* (selects, organizes and presents) a personalized stream of content, instead of offering users a set of alternatives to choose from. This approach changes the interaction between the user and the system from actively specifying preferences and selecting items, to passive consumption of items the system has chosen. For example, if a post from a user’s Friend is not highly ranked by the News Feed algorithm, it may not appear in her News Feed. Or, it may appear so far down in the list that the user is unlikely to scroll far enough to see it. If she does not visit the Friend’s Timeline, she may never know the post existed.

The posts that users see at the top of their News Feeds are the ones prioritized by the algorithm. Through repeated exposure, users may notice patterns in the characteristics of the posts and Friends they see, and begin to learn about what kinds of posts tend to receive more attention than other posts. If a user were to notice that Facebook posts that contain the word “congratulations” tend to be highly visible, she might learn that the Facebook News Feed prioritizes major life events over other kinds of posts, and subsequently believe that if she doesn’t post about celebratory occasions her posts won’t be viewed by others [32]. Even if an algorithm’s behavior is an invisible part of the system’s infrastructure, users can still form beliefs about how the system works based on their interactions with it, and these beliefs guide their behavior [12].

Research Goals

People are naturally very skilled at attending to evidence in the world and forming intuitive theories for why things happen the way they do. These theories help people structure their interactions with the world around them. However, they often contain incorrect or contradictory information, and are constantly updated as people have new experiences and learn new things [12].

Our goal in this study was to elicit users’ intuitive theories about the composition of their Facebook News Feeds, so that we can begin to understand the complex interdependencies between users and algorithms that affect system behavior [27]. Users’ beliefs about how the system works are an important component of a feedback loop that can cause systems to behave in unexpected or undesirable ways [8, 24]. The output of the News Feed algorithm (the posts that are displayed to users) affects the inputs to that same algorithm (interactions with those posts, and the characteristics new posts that are subsequently created). Understanding users’ beliefs is an important first step towards identifying effects of the feedback loop and potentially designing algorithms that are better at taking these effects into account.

METHOD AND PARTICIPANTS

This paper focuses on two questions from a larger survey that was conducted in April 2014, investigating users’ beliefs and experiences related to algorithmic curation in the Facebook News Feed. We asked respondents a closed-ended question, “Do you feel like your Facebook News Feed always shows you every post created by your Facebook friends?” to which the possible responses were Yes, Maybe, or No. We designed the closed-ended question to cause respondents to think and anticipated that it might trigger some to wonder about something they had not thought about before. We therefore included “Maybe” as a closed-ended option (in addition to “Yes” and “No”) to provide a way for respondents to express uncertainty.

Immediately following the closed-ended question, we asked an open-ended question intended to collect the reasoning behind the closed-ended responses: “Please explain your answer to the previous question.” The response field had a minimum length requirement of 150 characters (about three sentences). The median response length was 50 words (Min=26, Max=174). The survey also included questions about basic demographics, frequency of Facebook activity, and number of Facebook Friends. The questions that are the focus of this paper were asked early in the survey, and respondents had not yet been alerted via the survey to the role of the News Feed algorithm on Facebook.

We recruited respondents using Amazon’s Mechanical Turk (MTurk) for a paid incentive of \$5 for completing the entire survey. MTurk workers were exclusively from the USA and were required to have a 90% or higher approval rating after completing at least 500 HITs. Eligible respondents were at least 18 years old, had more than 20 friends on Facebook, and reported visiting Facebook at least once per week. We recruited Mechanical Turk workers for this study because we

wanted to gather beliefs and opinions from members of a population that is typically comfortable using the web and social media, and thus more likely than the general population to be aware of algorithmic curation. In other words, we were hoping to collect data from respondents who had thought about these issues before. Amazon Mechanical Turk workers may be more knowledgeable about algorithms and internet technologies in general than other Facebook users, and as such they might attend to and remember different things about their News Feeds than a random sample of Facebook users.

A total of 505 respondents completed the survey. We removed cases from the analysis when a respondent answered any of the three attention check questions incorrectly, and when responses to open-ended questions were indecipherable or virtually identical to other respondents' answers. The final number of cases for this analysis is 464.

The median time to complete all survey questions was 19 minutes, with an average of 29 minutes. The sample was 59% male, with a median age range of 26 to 34. Most respondents were white (372) and reported either attending (170) or having graduated from college (161). A large majority reported visiting Facebook several times per day (361) and posting to Facebook less than once per week (170). One hundred fourteen reported having 21-100 Facebook Friends, 208 reported 101-300 Friends, and 142 reported 301-500 Friends (Median=193, Max=2161).

ANALYSIS

People form "intuitive causal explanatory theories" [12] that help them interact with the world and accomplish their goals. These intuitive theories are generated out of each person's unique experiences, and as such are contextual and usually inaccurate in some way [19]. People are most likely to seek explanations for new evidence they encounter when they are faced with inconsistency or a problem to be solved, and as they encounter more evidence they continuously update their beliefs about what is going on [12].

One of the ways people attempt to form intuitive theories about unfamiliar situations is by making predictions and providing explanations that may or may not match the evidence at hand [12]. Our open-ended survey question asked respondents to speculate about why they believed the News Feed does or does not show every post created by their Facebook Friends. This is a very natural activity for people to engage in, and as users of Facebook they are likely already doing this.

We analyzed the open-ended responses using an iterative, inductive coding approach to identify components of respondents' intuitive theories about why they see the posts that they do in their Facebook News Feeds. In the first pass through the data, we coded for emergent similarities across responses, focusing on respondents' beliefs about who or what is in *control* of which posts appear in their News Feeds, and on the *evidence* they provided to support those beliefs.

In our second coding pass we added, removed and combined codes until we had formed a set of *control* codes and a set of *evidence* codes that represented the breadth of the responses.

In a third pass through the data we wrote definitions and summaries for each code, and in the process we found that we needed to break up some codes and combine others. We did this to ensure that our codes had descriptions that were internally consistent, and were different enough to easily distinguish them from the other codes. At this stage we also grouped codes into higher-level themes, keeping the *control* and *evidence* codes separate.

The codes and higher-level themes were *not* mutually exclusive. Most responses ended up falling under one of the higher-level *control* themes, and one or two *evidence* themes. For example, the following answer from respondent R9 was coded as *Control: News Feed*, referring to the statement that the entity the respondent believes is in control of prioritizing posts is "it". (Many responses used a pronoun to refer to the News Feed itself, which was mentioned in the question text.) This response was also coded as *Evidence: People I am Close To*, because the respondent stated that he or she mostly sees posts in the News Feed from close friends:

I usually see posts by the same people every day. This leads me to believe it might prioritize posts by closer friends. For example, with over 200 friends, I can't remember the last time I saw news updates from more than half of them. (R9)

Finally, in the last stage of the analysis we focused on drawing connections across the *control* and *evidence* themes based on code overlap within each response, to identify six intuitive theories (i.e., patterns of causal beliefs) contained in multiple responses.

Having 464 responses from different people allows us to analyze the relative prevalence of certain patterns of responses within the sampling frame we used for this study. Because of our sampling frame and the fact that our study is not a quantitative content analysis, we cannot compare absolute counts across themes to draw generalizable conclusions. However, the relative proportions can be interpreted as a rough indicator of how common one theme is compared with another.

LIMITATIONS

This study was conducted without the involvement Facebook or any of its employees. Therefore, we have no visibility into the details of the News Feed algorithm in use at the time we collected these data, so we cannot comment on the accuracy of respondents' intuitive theories about how the News Feed algorithm works. We have no way of knowing which posts each respondent's News Feed actually did or did not display, and so we have focused our findings on respondents' beliefs and the potential implications of those beliefs for the behavior of the system as a whole. Variation in respondents' perceptions about what causes their News Feeds to look the way they do may actually stem from different users experiencing different versions of Facebook or the News Feed algorithm. For example, we do not know what other research studies (e.g. [20, 3]) or A/B tests these respondents may have experienced prior to or simultaneously with our study that may have affected the composition of their News Feeds. Finally, our data are written statements from respondents and as such

contain respondents' self-reported memories of past behavior and subjective interpretation of their experiences. Certain kinds of experiences might be more memorable than others, which could bias our findings.

FINDINGS

A large majority of the responses to the closed-ended question, "Do you feel like your Facebook News Feed always shows you every post created by your Facebook friends?" indicated that most Facebook users in our sample believed they do not see every post their Friends create. Seventy-three percent (N=341) answered *No*, and only 8% (N=38) answered *Yes*. The remaining respondents answered *Maybe* (18%, N=85). The median number of Friends within each closed ended response group was: No=209 (n=341), Maybe=129 (n=85), Yes=122 (n=38). Pew data show that the US median is 200 (Fall 2013, <http://pewrsr.ch/1dm5NmJ>). Note that differences in medians across response categories are not necessarily statistically or practically meaningful, due to size differences between the categories.

Our findings below focus on patterns of beliefs—intuitive theories—expressed by respondents. The beliefs we identified are not mutually exclusive, and sometimes respondents made contradictory statements within the same response. This is normal; intuitive theories are often not well-formed or well tested, and constantly evolve as people learn new things. We briefly summarize the findings in Table 1, along with the number of responses that exhibit characteristics of each intuitive theory.

Passive Consumption

One hundred three responses contained statements that described a passive approach to experiencing the News Feed. These statements expressed few strong opinions about whether the News Feed shows every possible post. They did not mention steps respondents may have taken to ensure they are seeing the posts they want to see, nor the belief that the system is in control. It was common for *Passive Consumption* responses to speculate about the kinds of symptoms that might indicate that the user is or is not seeing all the possible posts. For example, this respondent ran through possible clues that he or she might have missed a post before concluding that he or she had not experienced any of them:

Every post I hear people talk about I have read on my Facebook feed. I have never visited someone's profile and saw a post that was never in my feed. I have never had someone tell me about a post that was not in my feed. (R161)

Another respondent mentioned how the volume of information in the News Feed might signal that posts are missing from it, and then concluded that because there are so many posts he or she must be seeing all of them:

Even though I have a relatively low amount of Facebook friends compared to most, I still see a lot of new posts to wade through no matter how often I check out the feed, so I don't see why it wouldn't be showing them all. (R317)

Evidence about News Feed volume (or lack of it) was used by other respondents to reach the opposite conclusion. Responses like this one present a belief that the News Feed would feel more full, somehow, if every post were displayed:

I don't think it shows every post created by my friends. If it did, my whole page would be filled with their posts which isn't possible because I have 80+ friends and that would flood my page making it impossible to read every post in a timely manner. (R232)

Assumptions like the above about what the News Feed would look like if it did or did not show all possible posts appear in many of the patterns of causal beliefs that respondents described. However, the group of *Passive Consumption* responses is unique in that the proportion of respondents who answered *No* (54%) is almost equal to the proportion that answered *Maybe* or *Yes* (46%). This indicates that these respondents may not have thought about this issue before and were drawing on their past experience to form plausible explanations on the fly.

Manual Control

Unlike the responses described in the previous section, statements that fall under the *Manual Control* pattern of beliefs described ways in which respondents believed themselves to be in control of what the News Feed displays. Some responses focused on explicit control of the contents of the News Feed by producers who manage their audiences using the privacy controls. Others focused on the ability of the consumer to engage the mechanisms Facebook provides to see more or fewer posts from particular individuals. In other words, these respondents believed that whether posts appear in the News Feed or not is at least sometimes up to them, not an intervention by the system.

Producer Privacy

Statements about privacy controls placed some of the responsibility over whether a post can be seen in the News Feeds of Facebook Friends in the hands of the person who created the post. This signals that some respondents were aware that the system supports a form of access control over posts. Respondents were aware that they might not see everything their Friends post because they themselves could block people from seeing their posts using Facebook settings. By definition, users cannot know if their Friends are doing this because if the producer of the post is successful the consumer will never see the post:

I know that Facebook allows you to filter who can view your post, and I'm sure that some of the people on my feed filter theirs so that only family, non-coworkers, etc can see it. I do that myself, and have talked to friends that do the same thing. (R305)

Well, I'm not sure. Some people might have chosen to elect not to show their posts to me in their privacy/post settings. I can't ever be sure if they have or haven't. (R154)

These responses illustrate that respondents are certain that access control is possible but also recognize that there is really

<i>Beliefs</i>	<i>Responses</i>	<i>Summary</i>
Passive Consumption	103	Respondents consume the contents of the News Feed without thinking too much about why they see the posts they do. Explanations range from never having experienced anything that would cause them to think some posts might not appear, to the belief (without evidence) that because there are so many posts the News Feed can't possibly show them all. <i>54% No, 46% Maybe/Yes</i>
Producer Privacy	45	Responses convey awareness that the News Feed might not show every post from Friends because users can use Facebook's audience selector to prevent specific others from seeing their posts. Respondents do not have direct evidence that they are excluded by others but believe this is possible because they do it themselves. <i>60% No, 40% Maybe/Yes</i>
Consumer Preferences	130	Respondents believe that without intervention the News Feed will not show them what they want to see. They cope with this by using the mechanisms provided by the Facebook interface for telling the system how they want it to behave, such as specifying (<i>Top Stories</i> or <i>Most Recent</i> , hiding or following people, etc. <i>74% No, 26% Maybe/Yes</i>
Missed Posts	208	Respondents explicitly implicate the News Feed as the agent that causes them to miss posts from Friends. They become aware that they have missed posts when someone they know mentions a post to them that they did not see, or when they visit Friends' Timelines to seek out posts they feel they may have missed. <i>81% No, 19% Maybe/Yes</i>
Violating Expectations	216	Respondents describe patterns or regularities in the News Feed that constitute symptoms that the News Feed may not be showing them every post. These symptoms are salient because they violate expectations about how they believe the News Feed should behave. For example, many respondents mentioned posts out of chronological order or older posts that repeatedly appear at the top of the feed as examples. <i>76% No, 24% Maybe/Yes</i>
Speculating about the Algorithm	223	Respondents indicate they believe an entity, characterized as Facebook or as an algorithm, prioritizes posts for display in the News Feed. Also, which posts they see depends on what the system knows about their preferences and characteristics, post popularity, and past interaction with other users. <i>80% No, 20% Maybe/Yes</i>

Table 1. Patterns of causal beliefs, identified through iterative inductive coding of open-ended explanations for responses to the question, "Do you feel like your Facebook News Feed always shows you every post created by your Facebook friends?" A single response may be represented in multiple rows.

no way to know, short of directly asking someone, whether one is being prevented from seeing particular posts. However, out of the 45 responses in this cluster, none expressed negative sentiments about this.

Consumer Preferences

One hundred thirty respondents expressed a belief that without manual intervention, the News Feed will not show them exactly what they want to see. This intervention was described in two ways: either directing the system *not* to display some or all of the posts from particular people, or directing the system to display *everything* from others by following, sorting, or creating Friend lists. These are the mechanisms that the system provides for users to explicitly state preferences about what they want to see. Respondents had many different names for what they called the activity: hiding (R427), blocking (R127), filtering (R390), unfollowing (R36), unsubscribing (R222), removing (R414), and ignoring (R198). Some described how they used these mechanisms for their intended purpose:

I have several friends on my list that I don't follow, which means that their posts don't show up on my news feed every time they say something on their account. I do this because I don't like what they post. (R396)

However, others had adopted these controls as workarounds. This respondent created what is essentially a secondary Friend list as an attempt to negate the influence of the News Feed algorithm:

I have a group called "Everyone" that has ALL of my Facebook friends on it. It doesn't pick and choose and

I get to see everyone's posts. When I see something I don't like or someone I might unfriend, instead, I just delete them from the group. (R246)

These are examples of people taking personal responsibility for making sure they see the things they want to see, and they don't see posts that are annoying or upsetting. Unlike the *Passive Consumption* beliefs, these responses indicate an intuitive theory involving less trust that the Facebook News Feed will make good choices on their behalf.

Symptoms of Curation

The most common pattern in respondents' answers was to provide an example of a post they did not see in their News Feeds as evidence that the News Feed does not show every post. Like the *Passive Consumption* beliefs, these statements include symptoms respondents noticed that caused them to speculate about what the News Feed is doing behind the scenes. However, unlike the *Manual Control* beliefs, these statements attribute control over this to the system rather than to other users.

Missing Posts

One form this evidence took was the description of an occasion where someone the respondent knows in real life asked if he or she had seen a particular post. Forty-three responses were like this. In the majority of these responses, the respondent did not remember seeing the post but reported being able to find it on a Friend's Timeline once they tried to do so:

Sometimes when I see a friend in real life, they will ask if I saw something that they posted, and at least half of

the time I haven't seen it. If I go directly to their profile, then I'm able to read it, but so often things don't appear on my newsfeed, which I find annoying! (R193)

One hundred seven responses described a feeling that the respondent had not seen a post from someone lately, and so they decided to find out what that person had been up to:

Sometimes I feel like I haven't heard anything about a friend for awhile. I wonder if maybe they've quit Facebook, or maybe even unfriended me. Then I go to their profile and look, and they've been posting stuff the whole time. So that means that my News Feed just hasn't been showing what they've posted. (R20)

This is a good example of a respondent recalling a past experience and incorporating it into a causal belief about how the News Feed works. It also indicates that some users have expectations for how often they should see posts from particular people or at least a background awareness of who they regularly encounter on Facebook versus who they do not. Prioritizing some posts over others is the intended behavior of the News Feed algorithm; however, respondents' background awareness of who they had heard from recently and who they had not made this behavior feel unexpected or unwelcome to them in some cases. In fact, seventy responses mentioned the feeling that there were posts missing from the News Feed that respondents really did want to see, like this one:

I miss posts from people I want to be seeing posts from all the time. Instead I see a bunch of crap that I DON'T care about: ads, comments on events I wasn't invited to, etc etc. There are some friends I see posts from so rarely that I often forget about them. (R173)

Violating Expectations

Two hundred sixteen responses included descriptions of patterns in the News Feed that could only have been noticed through repeated interactions that provided evidence over time. These patterns were incorporated into intuitive theories about the visibility of posts in the News Feed. The patterns also highlight a mismatch between expectations about how the News Feed should work, and the reality of how it does work as experienced by these respondents.

Posts out of order. Respondents seemed to have expectations about the proper order for posts to appear in. This led them to make statements about the News Feed being "out of order" or wanting to see posts in "linear" or "real" order. Signs that the News Feed was selecting some posts to prioritize over others included examples of posts that would "bump" or "bounce" to the top, "pushing" more recent posts down:

No, the timeline bounces up things from the past back up to the top. Things that may have been more recently get pushed down, if they get pushed down far enough, I may not scroll down far enough to see a friend's post. (R180)

Turnover Rate. Respondents also seemed to have expectations about the expected rate of change or turnover of posts in the News Feed. A violation of this expectation was provided in some responses as evidence that some posts might be left

out. Especially in instances of too-low turnover (i.e., repeating too many old posts), respondents wrote about feeling that seeing the same content over and over meant there were posts that were not being displayed:

I know it does not show me every post. It changes simply by refreshing the page to older posts I have already seen. (R123)

People Represented. The final pattern respondents noticed was seeing the same people represented over and over again in the News Feed. This was a sign for respondents that posts from some Friends were being repeatedly prioritized over others:

Many people I am friends with do not show up in my newsfeed even though I know that they post to Facebook. I generally only see posts from the same group of people even though I have quite a few friends. (R211)

This "same group of people" pattern was described by some respondents as only seeing posts from close friends; in other respondents' statements it was only posts from friends they are not close to, or even seemingly random friends. No respondents reported being aware that this pattern could have arisen simply because of the power-law pattern in participation: only a few of a given user's friends are likely to be very active, and the posts of these few may in reality just dominate the News Feed. However, it is interesting that respondent perceive this pattern as anomalous in some way.

Speculating about the Algorithm

The responses in this final group illustrate that some respondents were aware of a relationship between their behavior in the system and the content that the News Feed displays for them. These 223 responses contained statements about an entity—usually either "Facebook" or in 42 responses an "algorithm"—making inferences about their preferences. This is a fairly sophisticated understanding of the system, and reflects respondents' attempts to put together pieces of evidence from their own experience into a coherent story about what the News Feed is doing.

Respondents made statements illustrating that they recognize the system is attempting to show them posts that are "relevant" to them in some way, although they weren't always sure how this works:

I generally feel like Facebook shows me the posts that it 'thinks' I'd be most interested in. I feel like if I want to see all of them, I have to do so manually. (R77)

Many responses focused on the fact that the system knows how often users interact with particular Friends and posts. These responses illustrate that some users understand the system makes guesses about whose posts a user might want to see based on their past behavior. Some respondents even took this a step farther, suggesting that if one were to alter his or her browsing behavior, they might end up changing the composition of posts in the News Feed:

I know that it shows me the most commented on posts. It also shows the friends' pages that I visit the most. If

you start ‘liking’ more posts it may show more posts that I want to see. (R58)

Another group of responses recognized the popularity bias effect on the News Feed, often seen in other recommender systems (popular posts get more attention, which make them more popular). Respondents had noticed that posts with more comments and likes are prioritized over other posts:

I think that it hides some of them if they are not deemed to be popular. So what likely happens is that a popular person gets many views to their page. So people like a post. This post then gets put in the news feed for more people to like. Posts that have no likes don’t always make it to the feed. (R155)

Responses like this indicate that respondents know something is making decisions on their behalf, and that this means some posts will not be as easily accessible through the system as other posts.

Summary of Findings

Our findings describe patterns of respondents’ intuitive theories about how the Facebook News Feed works. We focused on intuitive theories because they are learned over repeated interactions with Facebook, and they guide behavior. Respondent beliefs ranged widely, from believing the News Feed shows all possible posts from their Friends, to automated filtering by an algorithm. Some respondents believed that the system could make inferences about their preferences based on which posts they read and whose Timelines they visited. Others noticed evidence of the News Feed algorithm that they interpreted as clues that the system was presenting posts “out of order”, including seeing too many or too few posts from particular Friends, or old posts displayed on too many visits. Respondents disliked missing posts from their Friends, and believed that when they did it was evidence of system intervention.

IMPLICATIONS

Algorithm designers should be aware that even if the influence of an algorithm is not made explicitly visible to users, they nevertheless can and do adapt their behavior to correspond with how they believe the system works, in order to accomplish their goals for using the system. This means that there are implications of users’ beliefs for their individual interactions with the system, and because of the feedback loop, for the behavior of the system overall. Our findings allow us to draw some conclusions below about how the interdependence of user and algorithm behavior might have system-wide consequences.

Mismatch between User and Algorithm Goals

The News Feed algorithm has (at least) two design goals¹. One goal is to “show everyone the right content at the right time so they don’t miss the stories that are important to them.” The second is that the News Feed should display posts more

¹We assume this based on our understanding of the News Feed FYI blog posts to the Facebook News Room, at <http://newsroom.fb.com/news/category/news-feed-fyi/>

prominently that will generate more interaction or “engagement” [26]. These two goals separate consumption (don’t miss important stories) from production (increased engagement).

Some respondents believed that the News Feed is like a river or a force of nature (a flood or flow metaphor was even present in some answers), rather than a road or some other kind of infrastructure that can be altered by human intervention. It was not important to these users to be too involved with manually shaping the composition of their News Feeds. This perspective of the News Feed as something that arises naturally out of Friends’ behavior rather than being constructed by the system might cause users not to attempt manual intervention. It is not clear how well the algorithm would perform for people who prefer to passively experience the News Feed, in comparison with others whose intuitive theories lend themselves to a more manual approach. If people were to accept what the system presents and not adjust their behavior to achieve a particular outcome, the impact of a feedback loop might be lessened. However, one respondent actually noticed that the less he or she interacts with others via on Facebook, the less content from friends he or she sees in the News Feed, which could be a symptom of the feedback loop:

But over the years it seems like I see their posts less often. Its very annoying actually and lately I have been considering quitting Facebook... Only about 50% of my newsfeed is posts from friends. It seems like unless you actually communicate with them often you hardly, if ever see their posts. (R150)

For users whose personal goals for using the system do not match the system’s goals, our findings suggest that unexpected outcomes like this may be possible.

Another area in which users’ and the system’s goals may not be aligned is when producers’ privacy settings restrict the visibility of posts. Greater restrictions on who can see a producer’s post would decrease the number of users who might be able to see and interact with it. When more private posts are ranked against public posts in preparation for display in the News Feed, they could appear below public posts that were visible to more people and therefore received more interaction. Privacy settings may indirectly prevent users who were in reality granted access from actually seeing the post in their News Feeds. A potential outcome like this would be extremely difficult for users to detect, as well as difficult for designers to explicitly test for. However, generating use cases where user and algorithm goals are in conflict as part of the design process is a good step towards systematic identification of edge cases that result from the interaction between user and algorithm behavior.

It may be that excluding users who do not fit the design goals of the algorithm is a side effect that is more of a feature than a bug; after all, Facebook (the company) has the ability and the right to identify and reward whatever it believes to be appropriate participation. Hallinan and Striphas [14] call this the “court of algorithmic appeal” in which ideas are judged

by system components and not human beings. However, if the operational definition of “appropriate participation” is not an intended consequence, then the algorithm may be creating categories of valid and invalid forms of contribution in ways that are contrary to the overall goals of the system. Identifying undesirable feedback loop effects that occur when the user’s goals do not match the system’s goals should be an important part of an algorithm designer’s role.

Reverse-engineering the Algorithm

Many respondents had sophisticated ideas about what kinds of interactions with content on Facebook brought about specific, repeatable consequences, and wrote about altering their behavior to work within the constraints of the algorithm or to use the mechanisms provided by the system to achieve their goals. What our respondents described is different from systematic attempts by researchers and journalists to perform experiments intended to discover details about the parameters and weighting used by the News Feed algorithm to rank posts for display. Determining the specifics of the algorithm by comparing users’ experiences is nearly impossible, because each user has a different network of Friends who serve as the sources of content [15]. However, each user is familiar with who his or her Friends are and may naturally make attributions about how interactions with the system are related to changes in what they see at a later time. Users can’t help but informally reverse-engineer the algorithm, and repeated use of the News Feed provides them with ample opportunity to form intuitive theories about how the system works.

The commonalities across respondents in the intuitive theories about how the algorithm works are important for understanding potential effects of feedback loops. For example, consider the large number of our respondents who said they had experienced missed posts. Missed posts are part of the correct operation of the News Feed algorithm, as it prioritizes a subset of posts for display. However, the focus on engagement means that people who create posts that achieve visibility in the News Feed receive more attention, leading to a form of the popularity bias. Our findings suggest that users are aware that this happens, and our respondents mentioned many symptoms of it (for example, low turnover from posts that get stuck at the top of the News Feed because of a highly active comment thread). As consumers, they can use Facebook’s mechanisms for hiding posts and Friends to try and control these effects on their News Feeds; they might also cope in other ways, such as forming a routine to visit friends’ Timelines directly to stay in touch:

I have had friends ask me about certain events or occasions in their lives and I had no clue what they were talking about because I’d never seen it in my news feed. I now make a point to visit certain people’s timelines to see what they have posted. (R28)

As producers, users can try to create posts that other users find more engaging so their posts get more attention and avoid the “threat of invisibility” [6]. Knowledge about how the algorithm works enables kind of individual experimentation, attempting to train the system through investigating what patterns of input will produce the desired output. Missed posts

seem like one symptom of algorithm behavior that is unpleasant enough to trigger users to try and figure out the algorithm. Our findings describe many other forms of expectation violations that might also serve as triggers. Therefore, we suggest that algorithm designers should consider the intuitive theories users learn about how the system works, and what might trigger them to reverse-engineer it, when trying to identify the consequences of feedback loops.

Finally, as Hamilton et al. point out, there is much debate about whether and to what extent automated system processes should be made visible to the humans who interact with them [15]. Although it may be a sign of good design when users do not detect algorithms at work, the consequences of this invisibility may exacerbate the impact of different user and algorithm goals. Our findings show that users don’t necessarily need to understand that an algorithm is shaping their experience to adapt their behavior to it. But, it may be possible for designers to leave clues for users to help them inadvertently form intuitive theories that may be more friendly to the intended operation of the algorithm—in effect, using the propensity of users to reverse-engineer the system as a mechanism for controlling some of the second-order effects created by the feedback loop. Future work should investigate the relationship between varying degrees of algorithmic literacy and system-level outcomes.

CONCLUSION

Our goal for this study was to better understand interactions between users and algorithms. While algorithms are designed to achieve particular goals, the consequences of design choices can be hard to predict. We focused on algorithmic curation in the Facebook News Feed as an example of a system in which feedback loops have the potential to affect behavior at both the individual and system level. We identified several patterns of beliefs about the News Feed that have implications for the design of algorithms for organizing, selecting, and presenting information in a complex socio-technical system. Ultimately, users vary widely in the degree to which they perceive and understand the behavior of content filtering algorithms in this online social network context, and these differences affect how they interact with and experience the systems they are using. We used intuitive theories about causal relationships as a lens for identifying implications of feedback loops that algorithm designers may not routinely consider. Algorithmic curation has the potential for generating negative outcomes that could be identified and avoided, if system dynamics like feedback loops are better understood.

ACKNOWLEDGMENTS

We thank Paul Rose for assisting with early stages of the data analysis, and the MSU BITLab research group for feedback on the survey questions. This material is based upon work supported by the National Science Foundation under Grant No. IIS-1217212.

REFERENCES

1. Social Networking Fact Sheet, 2014. <http://www.pewinternet.org/fact-sheets/social-networking-fact-sheet/>.

2. Backstrom, L. A Window Into News Feed, Aug. 2013. <https://www.facebook.com/business/news/News-Feed-FYI-A-Window-Into-News-Feed>.
3. Bakshy, E., Rosenn, I., Marlow, C., and Adamic, L. The Role of Social Networks in Information Diffusion. In *WWW 2012* (2012), 519–528.
4. Bernstein, M. S., Suh, B., Hong, L., Chen, J., Kairam, S., and Chi, E. H. Eddi: interactive topic-based browsing of social status streams. In *UIST '10*, ACM (2010), 303–312.
5. Bozdag, E. Bias in algorithmic filtering and personalization. *Ethics and Information Technology* 15, 3 (2013), 209–227.
6. Bucher, T. Want to be on the top? Algorithmic power and the threat of invisibility on Facebook. *New Media & Society* 14, 7 (Oct. 2012), 1164–1180.
7. Chen, J., Nairn, R., and Chi, E. H. Speak Little and Well: Recommending Conversations in Online Social Streams. In *CHI '11* (2011), 217–226.
8. Clegg, C. W. Sociotechnical principles for system design. *Applied Ergonomics* (2000).
9. Facebook. Company Info: Facebook Newsroom. <http://newsroom.fb.com/company-info/>, 2014.
10. Facebook. How News Feed Works. <https://www.facebook.com/help/327131014036297/>, 2014.
11. Fleder, D., and Hosanagar, K. Blockbuster Culture's Next Rise or Fall: The Impact of Recommender Systems on Sales Diversity. *Management science* 55, 5 (2009), 697–712.
12. Gelman, S. A., and Legare, C. H. Concepts and folk theories. *Annual Review of Anthropology* (2011).
13. Gillespie, T. The Relevance of Algorithms. In *Media Technologies Essays on Communication, Materiality, and Society*, P. Boczkowski and K. Foot, Eds. MIT Press, 2014.
14. Hallinan, B., and Striphas, T. Recommended for you: The Netflix Prize and the production of algorithmic culture. *New Media & Society* (2014).
15. Hamilton, K., Karahalios, K., Sandvig, C., and Eslami, M. A path to understanding the effects of algorithm awareness. In *CHI EA '14* (2014), 631–642.
16. Hogan, B. From Invisible Algorithms to Interactive Affordances: Data After the Ideology of Machine Learning. In *Roles, Trust, and Reputation in Social Media Knowledge Markets*. Springer International Publishing, Cham, 2015, 103–117.
17. Jannach, D., Lerche, L., Gedikli, F., and Bonnin, G. What recommenders recommend—an analysis of accuracy, popularity, and sales diversity effects. *User Modeling, Adaptation and Personalization* (2013), 25–37.
18. Joinson, A. N. 'Looking at', 'Looking up' or 'Keeping up with' People? In *CHI '08* (Florence, Italy, Apr. 2008), 1027–1036.
19. Jones, N. A., Ross, H., Lynam, T., Perez, P., and Leitch, A. Mental Models: An Interdisciplinary Synthesis of Theory and Methods. *Ecology And Society* 16, 1 (2011).
20. Kramer, A. D. I., Guillory, J. E., and Hancock, J. T. Experimental evidence of massive-scale emotional contagion through social networks. *PNAS* 111, 24 (2014), 8788–8790.
21. Lavie, T., Sela, M., Oppenheim, I., Inbar, O., and Meyer, J. User attitudes towards news content personalization. *International Journal of Human-Computer Studies* 68, 8 (2010), 483–495.
22. Lenhart, A. Adults and social network sites. Tech. rep., Pew Research Center, 2009.
23. Lv, Y., Moon, T., Kolari, P., Zheng, Z., Wang, X., and Chang, Y. Learning to Model Relatedness for News Recommendation. In *WWW 2011* (2011), 57–66.
24. Meadows, D. H. *Thinking in Systems: A Primer*. Chelsea Green Publishing, 2008.
25. Nguyen, T. T., Hui, P.-M., Harper, F. M., Terveen, L., and Konstan, J. A. Exploring the filter bubble: the effect of using recommender systems on content diversity. In *WWW '14* (2014), 677–686.
26. Owens, E., and Vickrey, D. News Feed FYI: Showing More Timely Stories from Friends and Pages. <http://newsroom.fb.com/news/2014/09/news-feed-fyi-showing-more-timely-stories-from-friends-and-pages/>, September 2014.
27. Rouse, W. B., and Serban, N. Understanding change in complex socio-technical systems. *Information Knowledge Systems Management* 10, 1-4 (Jan. 2011).
28. Sharma, A., and Cosley, D. Network-Centric Recommendation: Personalization with and in Social Networks. In *IEEE SocialCom 2011* (2011), 282–289.
29. Sharma, A., Gemici, M., and Cosley, D. Friends, Strangers, and the Value of Ego Networks for Recommendation. In *ICWSM '13* (2013), 721–724.
30. Singer, J. B. User-generated visibility: Secondary gatekeeping in a shared media space. *New Media & Society* 16, 1 (Jan. 2014), 55–73.
31. Steck, H. Item popularity and recommendation accuracy. In *RecSys '11* (2011), 125–132.
32. Sukumaran, A., Vezich, S., McHugh, M., and Nass, C. Normative influences on thoughtful online participation. In *CHI '11*, ACM (2011), 3401–3410.