Explaining the News Feed Algorithm: An Analysis of the "News Feed FYI" Blog

Kelley Cotter
Dept. of Media & Information
Michigan State University
cotterk6@msu.edu

Janghee Cho
Dept. of Media & Information
Michigan State University
chojangh@msu.edu

Emilee Rader
Dept. of Media & Information
Michigan State University
emilee@msu.edu

Abstract

Facebook uses algorithmic curation—automated selection
and ranking of content—to present a personalized News
Feed to each user for consumption. However, the News
Feed user interface provides little information to help users
understand how the ranking algorithm works. We analyzed
the company's "News Feed FYI" blog series to better un-
derstand the degree to which Facebook employs "how" and
"why" explanations of its News Feed algorithm. These types
of explanations have been used in other recommendation
and intelligent systems as a means of promoting user un-
derstanding and acceptance. Our findings show that the
"News Feed FYI" blog posts focus more on explanations
that justify *why* the algorithm works the way it does, and
less on explanations that describe *how* the system works.
These findings suggest that the "News Feed FYI" series
would be more helpful for increasing users' confidence in
the system, but not improving their trust in the system.

Author Keywords

algorithmic curation; Facebook News Feed; content analy-
sis; explanations

ACM Classification Keywords

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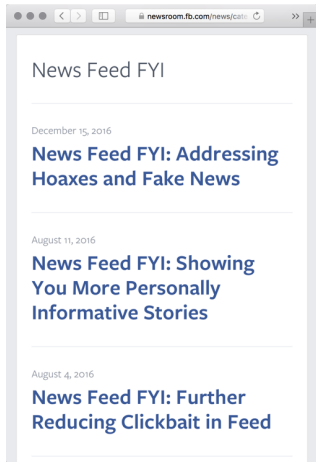Miscellaneous

Introduction

As of September 2016, Facebook had 1.18 billion active users, on average[1], making it the most widely used social media platform on the planet. Facebook uses algorithmic curation—automated selection and ranking of content—in order to present a personalized list of posts to each user for consumption. Over the years, the company has faced increased demands for greater transparency surrounding its use of algorithmic curation [1, 2, 17]. Previous research has discussed the power of algorithms to govern access to information [8]. Biased access to information online due to algorithmic curation may cause users to encounter only viewpoints that reinforce their existing attitudes, which prevents minority opinions from being expressed and deliberated [3]. Some users can develop an understanding of how the algorithm works through repeated experiences with it [5, 14]; however, the News Feed user interface provides little information that people can use to understand why they're seeing what they are seeing.

Providing an explanation—information about an intelligent system's process and objectives—results in better user understanding [4, 9, 11] and more positive perceptions of a system [7, 10, 16], as well as enhanced user performance when using a system [7, 10, 12, 15, 16]. Systems researchers have emphasized the importance of explanations as a means of influencing user acceptance of and trust in systems by increasing confidence in systems' abilities [7, 16]. In other words, explanations allow users to make inferences about a system's abilities and underlying motives, which form the basis of confidence and trust in a system.

Explanations about "how" a system produced a particular output or "why" a system produced a particular output offer a foundation for explanations [16]. "How" explanations reveal the steps a system takes to produce a particular output and connect a user's existing knowledge with the information he requires to accomplish a goal [16]. In this way, "how" explanations decrease information asymmetry between a user and a system and allow the user to weigh the costs and benefits of using the system. In contrast, "why" explanations offer rationalizations for a system's process, communicating the motivation underlying the system's design, which can establish goodwill towards the user [16]. "Why" explanations allow users to assess the degree to which a system's goals match, complement, or conflict with their goals in using the system.

In this paper we present an analysis of Facebook's "News Feed FYI" blog, which is intended to "highlight major updates to News Feed and explain the thinking behind them"[2]. We conducted this analysis to understand the extent to which Facebook employs "how" and "why" explanations of its News Feed algorithm. Reliance on either type of explanation has implications for how well the "News Feed FYI" series is able to influence confidence beliefs, which imply a sort of blind faith, and trust beliefs, which can inform a full cost-benefit analysis. Whereas trust implies comparing alternatives and deciding whether to rely on the system, confidence can occur in the absence of such a comparison [13].

Method and Analysis

We downloaded the text of all 35 of Facebook's "News Feed FYI" blog posts[3]. Blog post dates ranged from August 13, 2013 to December 15, 2016 with a new post about every couple of months. One member of the research team care-

Figure 1: A screen capture of the "News Feed FYI" blog.

---

fully read each blog post and took detailed notes, which were used to create the initial coding scheme. Four blog posts were excluded after this initial pass because the topics were out of scope or the content was a video and not text. The same member of the research team then segmented each blog post for analysis. The unit of analysis was in most instances a single paragraph from a blog post. On average, there were 8.45 segments per blog post.

Two other members of the research team independently coded the entire corpus in chronological order, and then met to resolve disagreements and produce the final codes for each segment. A description of each code, including the inter-coder reliability statistics (Cohen's κ), are presented in Table 1. Finally, two members of the research team used an inductive qualitative approach to further analyze the segments, focusing on summarizing and describing similarities within the segments for each blog post that were found to have each code present.

Findings

Our analysis revealed that "why" explanations were much more common in the blog posts than "how" explanations. While Ranking and Signal codes provided some information about "how" the system ranks News Feed content, the depth and breadth of this information is minimal and framed by "why" explanations. Segments coded as Purpose, Too Much Information, and Effectiveness exclusively provided "why" explanations.

Purpose and Too Much Information: "Why" Explanations
Nearly three quarters of the blog posts included a statement explaining at a high level an essential goal of the News Feed. The majority of these statements relay the intention to personalize the News Feed according to the user's preferences. In two blog posts, goals emphasize a

qualitative definition of the desired character of News Feed content, primarily high quality, interesting, important, or meaningful to the user. Two blog posts also underscore the importance of facilitating personal relationships as an essential goal of the News Feed. Articulation of goals are "why" explanations, as goals speak to the rationale and motivation for a system's existence. That a majority of the blog posts included at least one Purpose code demonstrates the centrality of the algorithm's goals in the blog series.

In addition to statements about the purpose of the News Feed ranking algorithm, almost one third of the blog posts make an argument for the necessity of the ranking algorithm due to the overwhelming amount of information produced by Facebook users and Pages. Evidence for this includes statements indicating that users don't want to see all the information, don't have time to read it all, or that they would be less satisfied with the News Feed if it were not personalized. Segments with the Too Much Information code also represent "why" explanations. These segments offer a specific rationale for the value of the ranking algorithm and specify the intent behind the algorithm as improving user experience. Both of these pieces of information help alleviate concerns users might have over the motivation driving Facebook's use of the ranking algorithm.

Effectiveness: Support for "Why"
One third of the blog posts offer data to illustrate the effectiveness of a change to the News Feed ranking algorithm by discussing a test or measurement intended to evaluate whether the change had the desired effect. Five blog posts reference specific numerical figures, and four blog posts specify not only the metric used to measure the success of a signal, but also a conceptual interpretation of the metric: "...in our early testing we've seen a 5% increase in people on Facebook clicking on links that take them off of Face-

Description	Cohen's κ	# Blog Posts	# Segments
Purpose: Is the purpose or goal of the News Feed stated in the text segment?	0.823	23	25
Ranking: Does the text segment mention the idea that the order or position of stories, posts, and/or ads in the News Feed are chosen or determined by the system?	0.921	21	71
Too Much Info: Does the text segment include a statement indicating that there is so much content that the News Feed is not able to show everything?	0.929	9	15
Signal: Does the text segment mention information the system uses to affect how the News Feed displays stories, posts, and/or ads?	0.712	31	103
Effectiveness: Does the text segment mention evidence from user feedback, or from other system data or metrics to support a claim that a change to how the News Feed works was effective or performing as intended?	0.742	11	18

Table 1: Descriptions, counts, and Cohen's κ for each code.

book. This is a...good sign that people are finding the remaining content in their News Feed more relevant and trustworthy" (ID 7). The text segments coded as Effectiveness justify changes to the News Feed algorithm and communicate to users what Facebook seeks to accomplish with a new signal. These codes do not help users understand how the system functions, but rather help users understand the quality and integrity of the system. In this way, Effectiveness statements represent "why" explanations.

Ranking: Both "How" and "Why"
Roughly two thirds of the blog posts describe the system's objective as determining the position or order of News Feed stories. All of these blog posts describe the placement of stories in the News Feed in directional terms, i.e. "higher" or "lower" and employ metaphors of "surfacing" or "bumping up" to help describe the dynamic placement of stories in the News Feed. The blog posts convey the idea that ranking should be conceptualized as a representation of how inter-

esting, informative, or high quality stories are at the time the user visits his or her News Feed; of how relevant the stories are to the user; and as an indicator of how much the user would want to see the stories. The segments coded as Ranking also emphasize actions taken by users as an important component of a story's ranking, such as liking, commenting on, clicking on, or sharing stories; and also properties of the content itself, like who posted them or what type of content (e.g., video or photo) the story contains. However, the blog posts also frequently explain that actions upon which ranking is based do not always accurately map to user needs and desires. The blog posts discuss Facebook's efforts to gain feedback from users about News Feed content, which have sometimes led to the discovery that the way Facebook had previously assessed a quality user experience overlooked some important considerations. For example, the blog posts often state a common refrain: "we have learned that the actions people take on

Facebook—liking, clicking, commenting or sharing a post—don't always tell us the whole story of what is most meaningful to them" (ID 27).

The Ranking code includes a mixture of "how" and "why" explanations. Most segments coded as Ranking explain that the News Feed algorithm ranks stories to determine the order in which all stories should appear in the News Feed, and a few blog posts mention general factors that contribute to ranking. However, the details of these "how" explanations remain vague and are often surrounded by rationales, goals, and/or explanations of general intent. Thus, while the blog posts do not explain how the algorithm calculates ranking—what happens inside the black box—they do describe elements of the algorithm's development process and the values that influence this process.

Signals: Both "How" and "Why"
All of the blog posts provide information about different types of *signals*, or the data that the News Feed ranking algorithm considers when it calculates the rank for a particular story. We divided the signal types into six categories:

- *Content signals* are dimensions on which stories differ: e.g., the kind of content associated with a story (e.g., link or video), the total number of likes or comments a story has received, or which friend or Page posted the story. (Blog posts: 24, segments: 42)
- *Source signals* are properties of the user or Page that created the story that are related to their production of posts; for example, over the history of the Page, how often it has posted stories with "click-bait" headlines. (Blog posts: 6, segments: 10)
- *Audience signals* also describe properties of users, but are related to consumption of stories rather than production. These signals reflect patterns in users' content consumption behavior, such as how often a

person uses the "Hide" functionality to remove stories from his or her view of the News Feed, or how often the user watches videos rather than scrolling past them. (Blog posts: 7, segments: 11)
- *Action signals* represent the digital traces of a person's behavior with respect to a particular story. For example, whether or not a given user chooses to like, click on, or share a specific story is considered to be an action signal, as is the amount of time reading a specific story or watching a video. (Blog posts: 22, segments: 29)
- *Relationship signals* are data collected about an event that involves two different people or pages. For example, how often two users interact with each other on Facebook is a Relationship signal, as is the choice by a user to unfollow a friend, Page, or Group. (Blog posts: 11, segments: 18)
- *Likelihood signals* occur when the system calculates the probability that a user will engage in some future action, such as liking or commenting on a story. These probabilities are incorporated into the ranking for a story produced by the News Feed algorithm. (Blog posts: 3, segments: 4)

Some blog posts explain that the combination of signals for a given story will be used to calculate the rank for a story, which will then result in that story being ranked higher (or lower) or seen more (or less). However, it is more common that the blog posts mention that the system takes into account a particular signal, but do not explain how the system uses the signal to calculate an overall ranking. Thus, while the blog posts do occasionally explain that the strength of a signal will be modified (i.e. made stronger or weaker), they do not explain how the system processes all of the signals collectively to calculate rank or how discrete signals may be weighted differently.

Similar to the Ranking codes, Signal codes offer a combination of "how" and "why" explanations. All Signal codes describe a discrete data point or "signal" that the News Feed system detects and uses to calculate ranking. These explanations provide valuable information about how the system ranks stories in the News Feed. This kind of information allows users to begin to draw conclusions about how the different signals may help them accomplish (or not) their goals for using Facebook. Users may also begin to make inferences about potential consequences of the signals used by the algorithm. Still, as with the discussions of Ranking, mentions of signals do not describe the system's internal process of calculating rank. Without explanations of the algorithm's internal process, users cannot fully assess the benefit of the algorithm relative to other means of communicating and satisfying information needs. Additionally, the Signal codes include information that addresses the underlying intent of the algorithm and updates made to it. Thus, much of the discussion of signals in the blog posts represents "why" explanations, not "how" explanations.

Discussion

As Gillespie [6] argued, algorithms undergo a subjective editorial process consisting of "proceduralized choices of a machine, designed by human operators to automate some proxy of human judgment or unearth patterns across collected social traces." Thus, if users are not aware of or do not understand this process of algorithmic curation, they may view their subjective reality based on what algorithms have served to them as an absolute. As more and more systems begin to use algorithms to sort and rank information for users, it is increasingly important that users are provided with information and explanations that will help them understand in what ways their access to information has been affected. Our study analyzed how Facebook uses "how" and "why" explanations of its News Feed algorithm in its "News Feed FYI" blog series, and found that "why" explanations are more common than "how" explanations.

"Why" explanations disclose the motivations, intentions, and objectives behind a system, and provide information that people may use to judge how well their goals for using a system match, complement, or conflict with that of the system's creators. On the other hand, "how" explanations provide information about a system's procedural logic. By providing information about a system's internal process, "how" explanations open the black box and instill trust by enabling users to verify that an algorithm has accurately and fairly "produc[ed] and certif[ied] knowledge" [6], and assess potential risks and evaluate possible alternatives.

"Why" explanations may suffice for users who have no interest in understanding how a system works, but merely want to know that the system will help them accomplish their goals [13]. The emphasis on these kinds of explanations in the "News Feed FYI" blog indicates that this may be one of Facebook's goals. Still, as users increasingly rely on algorithmic curation to aid not only in finding information, but also in defining the boundaries of information [6], "why" explanations may not be enough, because they do not allow users to critically analyze processes of algorithmic curation. Our future work will address user responses to different types of explanations in order to test the relationship between the explanation types and trust and confidence in algorithmic systems.

Acknowledgements

References

[1] Torie Bosch. 2015. The tyranny of algorithms: A future tense event recap. (2015). http://www.slate.com/blogs/future_tense/2015/12/16/the_tyranny_of_algorithms_a_future_tense_event_recap.html

[2] danah boyd. 2016. Untangling research and practice: What Facebook's emotional contagion study teaches us. *Research Ethics* 12, 1 (2016), 4–13. DOI:http://dx.doi.org/10.1177/1747016115583379

[3] Engin Bozdag and Jeroen van den Hoven. 2015. Breaking the filter bubble: democracy and design. *Ethics and Information Technology* 17, 4 (2015), 249–265. DOI:http://dx.doi.org/10.1007/s10676-015-9380-y

[4] Henriette Cramer, Vanessa Evers, Satyan Ramlal, Maarten van Someren, Lloyd Rutledge, Natalia Stash, Lora Aroyo, and Bob Wielinga. 2008. The effects of transparency on trust in and acceptance of a content-based art recommender. *User Modeling and User-Adapted Interaction* 5 (2008), 455–496. DOI:http://dx.doi.org/10.1007/s11257-008-9051-3

[5] Motahhare Eslami, A Rickman, Kristen Vaccaro, A Aleyasen, A Vuong, Karrie Karahalios, Kevin Hamilton, and Christian Sandvig. 2015. "I always assumed that I wasn't really that close to [her]": Reasoning about invisible algorithms in the news feed. In *CHI '15: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 153–162. DOI:http://dx.doi.org/10.1145/2702123.2702556

[6] Tarleton Gillespie. 2014. The Relevance of algorithms. In *Media Technologies*, Tarleton Gillespie, Pablo J. Boczkowski, and Kirsten A. Foot (Eds.). The MIT Press, 167–194. DOI:http://dx.doi.org/10.7551/mitpress/9780262525374.003.0009

[7] Shirley Gregor and Izak Benbasat. 1999. Explanations from intelligent systems: Theoretical foundations and implications for practice. 23, 4 (1999), 497–530.

[8] Rob Kitchin. 2016. Thinking Critically About and Researching Algorithms. *Information, Communication & Society* (2016), 1–16. DOI:http://dx.doi.org/10.1080/1369118X.2016.1154087

[9] Todd Kulesza, Simone Stumpf, Margaret Burnett, Sherry Yang, Irwin Kwan, and Weng-Keen Wong. 2013. Too much, too little, or just right? Ways explanations impact end users' mental models. In *2013 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*. IEEE, 3–10. DOI:http://dx.doi.org/10.1109/VLHCC.2013.6645235

[10] Joseph E. Mercado, Michael A. Rupp, Jessie Y. C. Chen, Michael J. Barnes, Daniel Barber, and Katelyn Procci. 2016. Intelligent agent transparency in human-agent teaming for multi-UxV management. *Human Factors: The Journal of the Human Factors and Ergonomics Society* 58, 3 (2016), 401–415. DOI:http://dx.doi.org/10.1177/0018720815621206

[11] Kenya Freeman Oduor and Eric N. Wiebe. 2008. The effects of automated decision algorithm modality and transparency on reported trust and task performance. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 52, 4 (2008), 302–306. DOI:http://dx.doi.org/10.1177/154193120805200422

[12] Alexis Papadimitriou, Panagiotis Symeonidis, and Yannis Manolopoulos. 2012. A generalized taxonomy of explanations styles for traditional and social recommender systems. *Data Mining and Knowledge Discovery* 24, 3 (2012), 555–583. DOI:http://dx.doi.org/10.1007/s10618-011-0215-0

[13] Wolter Pieters. 2011. Explanation and trust: what to tell the user in security and AI? *Ethics and Information Technology* 13, 1 (2011), 53–64. DOI:http://dx.doi.org/10.1007/s10676-010-9253-3

[14] Emilee Rader and Rebecca Gray. 2015. Under-
standing user beliefs about algorithmic curation in the
Facebook News Feed. In *CHI '15: Proceedings of the
SIGCHI Conference on Human Factors in Comput-
ing Systems*. 173–182. DOI:http://dx.doi.org/10.1145/
2702123.2702174

[15] James Schaffer, Prasanna Giridhar, Debra Jones, To-
bias Höllerer, Tarek Abdelzaher, and John O'Donovan.
2015. Getting the message? A study of explana-
tion interfaces for microblog data analysis. In *Pro-
ceedings of the 20th International Conference on
Intelligent User Interfaces* (2015). 345–356. DOI:
http://dx.doi.org/10.1145/2678025.2701406

[16] Weiquan Wang and Izak Benbasat. 2007. Recom-
mendation agents for electronic commerce: Effects
of explanation facilities on trusting beliefs. *Journal of
Management Information Systems* 23, 4 (2007), 217–
246. DOI:http://dx.doi.org/10.2753/MIS0742-1222230410

[17] Charlie Warzel. 2014. Facebook's two-way mir-
ror. (2014). http://www.buzzfeed.com/charliewarzel/
facebooks-two-way-mirror