FILTERWORLD

HOW ALGORITHMS FLATTENED CULTURE

Kyle Chayka

Doubleday ▪ New York

www.doubleday.com

DOUBLEDAY and the portrayal of an anchor with a dolphin are registered trademarks of Penguin Random House LLC.

Cover design by Oliver Munday

Library of Congress Cataloging-in-Publication Data
Names: Chayka, Kyle, author.
Title: Filterworld : how algorithms flattened culture / Kyle Chayka.
Description: First edition. | New York : Doubleday, [2024] | Identifiers: LCCN 2023020397 (print) | LCCN 2023020398 (ebook) | ISBN 9780385548281 (hardcover) | ISBN 9780385548298 (ebook)
Subjects: LCSH: Culture. | Culture—Mathematical models. Classification: LCC HM621 .C488 2024 (print) | LCC HM621 (ebook) | DDC 306—dc23/eng/20231011
LC record available at https://lccn.loc.gov/2023020397
LC ebook record available at https://lccn.loc.gov/2023020398

Ebook ISBN 9780385548298

ep_prh_6.2_145871955_c0_r0

CHAPTER 1

The Rise of Algorithmic Recommendations

EARLY ALGORITHMS

Algorithm as a term simply describes an equation: any formula or set of rules that produces a desired result. The earliest examples come from ancient Babylon, in the region that is now Iraq. Cuneiform tablets, dating back to 1800–1600 BCE, record algorithms for purposes like calculating the length and width of a cistern using its depth and the volume of earth excavated for it. According to the mathematician Donald E. Knuth, the Babylonians "represented each formula by a step-by-step list of rules for its evaluation, i.e., by an algorithm for computing that formula." They had a specialized system for recording calculations, using "a 'machine language' representation of formulas instead of a symbolic language," Knuth wrote. The written explanation of each Babylonian algorithm ended with the same phrase: "This is the procedure." That line emphasizes an inherent quality of algorithms: they can be repeated, equally applicable and effective every time a given situation occurs. An acolyte of Silicon Valley today might describe them as scalable.

Algorithms are key to the history of early mathematics. Around 300 BCE, the Greek philosopher Euclid recorded in his treatise *Elements* what is called the Euclidean algorithm, a way of finding the greatest common divisor of two or more numbers. That formula and the Sieve of Eratosthenes, an algorithm from the third century BCE that identifies prime numbers within a set of numbers, are still used today, particularly in the realm of cryptography. But the actual word *algorithm* comes from a single person—or at least his birthplace.

Muhammad ibn Musa al-Khwarizmi was a Persian scholar born around 780 CE in Khwarazm, an area around present-day Turkmenistan and Uzbekistan. Though little is known about his life, Al-Khwarizmi made his way to Baghdad, which had become the intellectual center of the region after the Muslim Abassid caliphate conquered Persia in the seventh century. There he worked at the House of Wisdom, also known as the Grand Library of Baghdad, researching astrology, geography, and math. Like its predecessor the Egyptian Library of Alexandria, the House of Wisdom was an interdisciplinary center of learning where scientific study was prized and texts in Greek, Latin, Sanskrit, and Persian were translated into Arabic. Around 820, al-Khwarizmi completed *On the Calculation with Hindu Numerals,* the text that eventually introduced the numeral system we use today to Europe. He also wrote *The Rules of Restoration and Reduction,* a book on strategies for solving equations. Its Arabic name was shortened to *al-jabr* (meaning restoration, or canceling like terms on either side of an equation), which provided the source for the word and the discipline of algebra. *Restoration and Reduction* included solutions for quadratic equations and methods of calculating area and volume, with approximations of pi.

In the mid-twelfth century, an English scholar of Arabic named Robert of Chester was living in Spain, where Muslim, Jewish, and Christian cultures overlapped, at some times peacefully and at others less so. It was another moment when ideas were exchanged and disseminated, crossing between civilizations. In 1145, Robert translated *The Rules of Restoration and Reduction* into Latin. *Al-jabr* became "algeber," and al-Khwarizmi became "Algoritmi." At that time, "algorismus" referred generally to any kind of mathematical procedure using Hindu-Arabic numerals, and those who practiced such an art were called algorists. (That term was adopted by visual artists using algorithmic processes beginning in the 1960s, but it seems apt for anyone working on today's version of algorithms.) The long arc of algorithm's etymology shows that calculations are a product of human art and labor as much as repeatable scientific law.

THE INVENTION OF COMPUTER PROGRAMMING

All computers are built from series of equations performed repeatedly. Results are encoded in zeros and ones and then passed on through yet more equations to achieve an outcome. In 1822, the British inventor Charles Babbage outlined his concept for the "application of machinery to the computation of astronomical and mathematical tables"—a way to automate calculations using an assemblage of numbered wheels and gears called the Difference Engine. The machine was never fully built, but later executions look something like the inside of a piano, with wheels in long rows instead of hammers. Babbage's design would have been eight feet high and weighed four tons. His later iteration, the Analytical Engine, could, if built, take commands programmed via punch cards and perform simple programming features like loops and conditions. It was the basis for all the much more complicated modern computing that followed. As Babbage's son Henry wrote in 1888, "It is only a question of cards and time."

Ada Lovelace, the daughter of Lord Byron, is now widely regarded as the first computer programmer; she wrote algorithms for the machine as Babbage designed it, including a process for calculating Bernoulli numbers. Lovelace also realized that the repeating mechanical processes that the machine enabled could be applied to fields beyond mathematics. In 1843, Lovelace wrote that the Analytical Engine "might act upon other things besides number, were objects found whose mutual fundamental relations could be expressed by those of the abstract science of operations, and which should be also susceptible of adaptations to the action of the operating notation and mechanism of the engine." In other words, anything that can be turned into something like data—a series of numbers—could be manipulated in a formulaic way. That might include text, music, art, or even a game like chess. Lovelace imagined one form of such automation: "Supposing, for instance, that the fundamental relations of pitched sounds in the science of harmony and of musical composition were susceptible of such expression and adaptations, the engine might compose elaborate and scientific pieces of music of any degree of complexity or extent." She was envisioning something like what the composer Brian Eno created in 1995 and popularized as "generative music," a series of ambient synth compositions driven by musical software that created different tunes each time the software ran. Lovelace was envisioning how culture could be both molded and perpetuated by the new technology, the way algorithmic feeds do today.

Lovelace was early in discovering that manipulating such mechanical commands could be its own form and self-expression. In the 1990s and 2000s, computer programming began to take its place alongside basic math and science as a skill that was necessary for a child's complete education. I was introduced to it on "computer room" desktops in my high school circa 2002, where we played educational video games that resembled programming languages. But where I really learned was on the chunky plastic TI-83 calculators that we had to acquire for advanced math classes. The calculators came with the ability to code in a language called TI-BASIC, which had simple if-then loops and variable functions. At first, I made modest programs to automate formulas I needed to know for tests, but once I became more fluent in the language, I made my own versions of tic-tac-toe and Connect Four. The machine was a partner in my creativity; it felt like magic.

A century after Lovelace, during World War II, Alan Turing, a British mathematician and computer scientist, was working in code breaking for the government—he helped to decode the German Enigma cipher machine. In 1946, with the war over, Turing wrote a report for the National Physical Library proposing the development of an "Automatic Computing Engine." It was the first description of artificial intelligence as a real possibility instead of a theoretical concept. Calculating and sorting machines designed to perform specific tasks already existed, Turing wrote, but his proposal went beyond that: "Instead of repeatedly using human labor for taking material out of the machine and putting it back at the appropriate moment all this will be looked after by the machine itself."

According to Turing, the device would be able to perform any kind of calculation and at any scale without needing to be reconfigured. It had its own internal logical language that could be adapted to different ends, to solve any type of problem. "How can one expect a machine to do all this multitudinous variety of things?" Turing wrote. "The answer is that we should consider the machine to be doing something quite simple, namely carrying out orders given to it in a standard form which it is able to understand." It would execute algorithms. He hinted at the way that machine learning algorithms today evolve over time, incorporating adjustments without human decision-making.

Such a system would perform calculations much faster and at higher levels of complexity, exceeding humans. Turing wrote: "The speed of the machine is no longer limited by the speed of the human operator." He did not see such machines as utopian tools, however. Just because they would be automatic didn't mean they would always be right. "The human element of fallibility is eliminated, although it may to an extent be replaced by mechanical fallibility," Turing continued. His report predicted many now-familiar elements of personal computers, from erasable memory units to input mechanisms, conversion from binary language, and even temperature control so the machine wouldn't overheat. For Turing, however, the word *computer* referred not to the machine but to the person doing the computing, once again emphasizing that organic element.

As early as 1936, Turing conceived of what is now called a "Turing machine," which he sketched in detail in a 1948 essay called "Intelligent Machinery." The Turing machine is "an infinite tape marked out into squares, on each of which a symbol could be printed." The tape moves through a reader that scans one square at a time and performs the operation dictated by the symbol in the square, which can also be erased or overwritten. Any algorithm, in the historical sense of a mathematical process, can be calculated by such a Turing Machine. And any computational system that can compute anything that a Turing Machine can is said to be "Turing-complete." All programming languages, for example, are Turing-complete because they can model any kind of equation. (Even the spreadsheet software Excel became Turing-complete in 2021.) What Turing correctly concluded was that any computing machine would be able to do the work of any other—even Charles Babbage's nineteenth-century Analytical Engine could theoretically perform the complex tasks that our laptops do now, if given infinite scale and time.

There is something of the clash between mechanical rules and human operation within Turing's life, too. In 1952, Turing was charged with gross indecency for "homosexual expression"—the legalistic phrase for having sex with another man—during messy legal proceedings that he initiated after his own house was robbed. Homosexual sex among consenting adults remained illegal in England all the way to 1967—the law is its own kind of algorithm, deciding judgment based on an implacable set of rules. Turing eventually pled guilty to the charges and was convicted. Rather than be imprisoned, he was forced to undergo chemical castration. In June 1954, Turing, forty-one, was found dead by his housekeeper. The cause was cyanide poisoning, and his death has long been considered suicide, the suspected delivery mechanism a half-eaten apple at Turing's bedside.

When we talk about "the algorithm," it often feels like a force that began to exist only recently, in the era of social networks. But we're discussing a technology with a history and legacy that has slowly formed over centuries, long before the Internet existed. Restoring this larger picture can help us better understand the power that algorithms have today. Still, no matter how complex, an algorithm remains in its essence an equation: a method to arrive at a desired conclusion, whether it's a Sumerian diagram to divide an amount of grain equally among several men or the Facebook feed determining which post to show you first when you open the website. All algorithms are engines of automation, and, as Ada Lovelace predicted, automation has now moved into many facets of our lives beyond pure mathematics.

ALGORITHMIC DECISION-MAKING

In 1971, in Santiago, Chile, a hexagonal room designed as a kind of control room for the entire country was built in a downtown office building. Monitors and backlit displays adorned the room's wood-paneled walls, displaying data readouts with metrics like national raw material supplies and labor participation rates. Seven chairs were arrayed in a circle facing each other in the center of the room, white fiberglass wingback seats that resembled the captain's chair in a science-fictional space cruiser. Each chair had a control panel on the right-hand side to navigate the various screens, as well as an ashtray and a cupholder, perhaps for a tumbler of whiskey. The room, which was named the Operations Room under the aegis of the larger Project Cybersyn, was designed under the socialist Chilean president Salvador Allende and the consultancy of Stafford Beer, a British man who, in his own country, was known for applying the practice of "cybernetics" to business management. Beer described cybernetics as "the science of control." It involves analyzing complex systems, whether corporations or biology, and determining how they work to better model or create such intelligent, self-correcting systems. (In the United States, a similar practice of systems analysis was pioneered by the RAND Corporation in the 1950s.) Project Cybersyn was meant to provide an ideal model, aiding the Chilean government's decision-makers in real time as they sat in the room and smoked their cigarettes and drank their whiskey—another meeting of the coldly technological and the messily human. From the room, the men watched the algorithms that oversaw the nation.

Project Cybersyn's physical design, led by the German consultant Gui Bonsiepe, created an image of mid-century modernist utopianism. The monitors floated on the walls, the underlying wiring connecting them to the chairs hidden from view. The cockpit chairs themselves were sleek and uniform, smoothly curving in one molded form. The room symbolically reduced government down to the manipulation of data, like winning a video game. Project Cybersyn promised to supplant human leadership with technological oversight, the scant few screens encompassing any information that might be needed. You could sit in one of the chairs and observe everything happening in the country.

However, Project Cybersyn's technology was a facade, something akin to "design fiction"—an interactive illusion of what might be possible. What it promised was not yet feasible with computer networks at the time. Its data slideshows were created by hand, not automatically generated. It ran on a single computer, fed by telex machines that Chilean factories could use to send information over telephone lines. And finally, though the room was completed, it was never put into action. On September 11, 1973, with the assistance of the United States CIA, Allende's government was overthrown, and Augusto Pinochet took over.

There remains an undeniable appeal to the photographs of Project Cybersyn. They appear over and over in design mood boards, projecting an aesthetic that still looks like the future many decades later. Perhaps the images are so influential because we retain that dream of the raw data of reality processed and crunched into digital graphs, which are then evaluated and, from there, the correct path of action determined. Project Cybersyn exuded an air of infallibility, even though inventors like Turing knew computers couldn't work so perfectly. As the cybernetics pioneer Stafford Beer argued, we tend to use machines to automate the structures and processes that already exist, which were human creations to begin with. "We enshrine in steel, glass, and semiconductors those very limitations of hand, eye, and brain that the computer was invented precisely to transcend," Beer wrote in his 1968 book *Management Sciences,* pinpointing the paradox. As with the Mechanical Turk, the human persists within the machine.

Today we do have versions of algorithmic government and algorithmic life: banks use machine learning to dictate who receives loans; Spotify uses the data of your past actions to determine songs to recommend, those they deem most aligned with your sensibility. But the technology that accomplishes those feats doesn't look like Project Cybersyn. There are no hexagonal rooms or wingback chairs. Algorithms have become both invisible and omnipresent, contained in the apps we carry around with us on our phones even as their data are hosted physically somewhere distant, within vast air-conditioned server farms set into obscure locations in the natural landscape. Where Project Cybersyn suggested that the world run by data might be coherent and graspable, contained within a room, we now know that it is abstract and diffuse, everywhere and nowhere at once. We're encouraged to forget the presence of algorithms.

New technologies inevitably create new forms of behavior, but the behaviors are rarely those that the inventors expect. The technology has an inherent meaning of its own that eventually comes to the fore. Marshall McLuhan wrote his famous dictum "the medium is the message" in his 1964 book *Understanding Media: The Extensions of Man.* He meant that the structure of a new medium—electric light, the telephone, television—is more important than the content that travels through it. The telephone's ability to connect people exceeds any particular conversation. "The 'message' of any medium or technology is the change of scale or pace or pattern that it introduces into human affairs," McLuhan wrote. In our case, the medium is the algorithmic feed; it has scaled and sped up humanity's interconnection across the world to an unimaginable degree. Its message is that on some level, our collective consumption habits, translated into data, run together into sameness.

HOW RECOMMENDATION ALGORITHMS WORK

Algorithms are digital machines that turn a series of inputs into a particular output, like a conveyor belt in a factory. What makes one algorithm different from another is less their structure than the ingredients they are built from. All recommendation algorithms work by gathering a set of raw data. The overall term for that dataset is *signal,* the collected inputs that are fed into the machine. The signal data might include a user's past purchases on Amazon or how many other users favorited a particular song on Spotify. The data is quantitative rather than qualitative, since it must be able to be processed by the machine. So even if the data is about something as subjective as music preferences, it is translated into numbers: x number of users rated y band an average of z, or x number of users listened to y band z times. The primary signal fed into many social media recommendations is *engagement,* which describes how users are interacting with a piece of content. That might come in the form of likes, retweets, or plays—any kind of button found next to a post. High engagement means the number of likes, views, or shares is higher than the average of other posts.

The signal is fed through a *data transformer* that puts it into usable packages, set to be processed by different kinds of algorithms. Engagement data might need to be separated from ratings data, or data about the subject matter of the content itself. A *social calculator* might be used to add information about how users relate to one another within a single platform—I often engage with Instagram posts from my friend Andrew, for example, which would make a recommender system more likely to rank one of his posts highly in my personal feed.

Then comes the specific equation of an individual algorithm. In today's platforms, there is very rarely only one set algorithm—there are many. What we are experiencing is a series of different equations that consider data variables and process them in a few ways. One equation calculates a result based on engagement alone, perhaps finding the content with the highest average engagement, while another prioritizes the social context of a piece of content for a particular user. Those algorithms are also weighed against each other. *Hybrid filtering* is when multiple techniques are used. Finally, the *output* is the recommendation itself, the next song in the automated playlist or the ordered list of posts. The algorithm decides whether it should put a life update from a friend in your Facebook feed over a politics news story, for example.

An executive at the music cataloging and recommendation service Pandora once described the company's system to me as an "orchestra" of algorithms, complete with a "conductor" algorithm. Each algorithm used different strategies to come up with a recommendation, and then the conductor algorithm dictated which suggestions were used at a given moment. (The only output was the next song to play in a playlist.) Different moments called for different algorithmic recommendation techniques.

There is no single, monolithic "algorithm," because each platform works in its own way, incorporating custom-designed variables and sets of equations. It's important to remember that how the Facebook feed works is a commercial decision, the same as a food manufacturer deciding which ingredients to use. Algorithms also change over time, refining themselves using machine learning. The data they take in is used for gradual selfimprovement to encourage even more engagement; the machine adapts to users and users adapt to the machine. The differences between platforms became more prominent and more relevant moving into the mid-2010s, as social media and streaming services doubled down on algorithmic feeds and they began to dominate the user experience.

We users fundamentally do not understand how algorithmic recommendations work on a day-to-day basis. Their equations, variables, and weights are not public because technology companies have little incentive to publicize them. They are closely held trade secrets, almost like nuclear codes for how important they are to the businesses, and are rarely disclosed or hinted at. One reason for that is if the algorithms were public, users could game the system to promote their own content. Another is the fear of competition: other digital platforms could steal the secret sauce and make a better product. Yet these tools, like many digital technologies, started out in a non-commercial context.

Recommendation algorithms as a way of automatically processing and sorting information were put into practice in the 1990s. One of the first examples was a system for sorting email—to this day, an annoying chore. Even in 1992, engineers at Xerox's Palo Alto Research Center (better known as PARC) were already overwhelmed by it. They sought to solve the problem of "the increasing use of electronic mail, which is resulting in users being inundated by a huge stream of Incoming documents," David Goldberg, David Nichols, Brian M. Oki, and Douglas Terry wrote in a 1992 paper. (Little did they know the volume of digital communication we would face in the twenty-first century.) Their email filtering system, called Tapestry, used two kinds of algorithms in tandem: "content-based filtering" and "collaborative filtering." The former, which was already used in several email systems, evaluated the text of emails—say, if you wanted to prioritize everything with the word *algorithm.* But the latter, more innovative technique was based on the actions of other users. Who opened a particular email and how they responded to it would be factored into how much the system prioritized the email. As the paper described it:

> People collaborate to help one another perform filtering by recording their reactions to documents they read. Such reactions may be that a document was particularly interesting (or particularly uninteresting). These reactions, more generally called annotations, can be accessed by others' filters.

Tapestry used a "filterer" to run repeated queries over a set of documents, a "little box" that collected material that might be of interest to the user, and an "appraiser" that could prioritize and categorize documents. Conceptually, it's very similar to the algorithmic feeds we see now: Tapestry's goal was to surface the content that was most likely to be important to the user. But this system required much more up-front action on the part of the user, who had to write queries to determine what they wanted to see, based either on content or on other users' actions. The other users in the system also had to carry out very intentional actions, marking material as compelling or irrelevant in turn. Such a system required a small group of people who already knew one another and understood how their cohort interacted with email—for instance, you may need to be aware in advance that Jeff replies only to particularly important emails, so you want your filter to surface all the emails that Jeff replies to. Tapestry functioned best on a very intimate scale.

In 1995, a paper from Upendra Shardanand and Pattie Maes at the MIT Media Lab described "social information filtering," "a technique for making personalized recommendations from any type of database to a user based on similarities between the interest profile of that user and those of other users." Building on the ideas of Tapestry, it was a response to the overflow of information online: "The volume of things is considerably more than any person can possibly filter through in order to find the ones that he or she will like." Automated filters would be necessary, they concluded: "We need technology to help us wade through all the information to find the items we really want and need, and to rid us of the things we do not want to be bothered with." (Of course, this is still a huge problem online.) Shardanand and Maes argued that content-based filtering had significant drawbacks. It requires the material to be translated into data that the machine can understand, such as text; it lacks serendipity because it can filter only by the terms that the user inputs; and it does not measure inherent quality. It is unable to "distinguish a well written [and] a badly written article if the two articles use the same terms." The inability to evaluate quality brings to mind artificial intelligence: New tools like ChatGPT seem to be able to understand and generate meaningful language, but really, they only repeat patterns inherent in the preexisting data they are trained on. Quality is subjective; data alone, in the absence of human judgment, can go only so far in gauging it.

Social information filtering bypasses those problems because it is instead driven by the actions of human users, who evaluate content on their own—using judgments both quantitative and qualitative. It's more like word of mouth, the way we get advice on what to listen to or watch from friends whose preferences are similar to our own: "Items are recommended to a user based upon values assigned by other people with similar taste," according to the paper. The similarity of one user's taste to another was calculated using statistical correlation. The researchers designed a system called Ringo to make music recommendations using an email list. As a user evaluated music, rating an initial batch of 125 artists on a scale of 1 to 7, a diagram of their preferences was built. Then, by comparing that diagram to other users', the system suggested music that they were likely to enjoy—or hate, which was also an option. Ringo recommendations came with a measure of confidence, signaling how likely a suggestion was to be right and allowing the user to further consider the algorithmic choice. By September 1994, Ringo had twenty-one hundred users and five hundred emails a day evaluating music.

Ringo tested various specific algorithms to make decisions based on the music ratings. The first algorithm measured dissimilarity between users' tastes and based recommendations on the most similar users. The second algorithm measured similarity, then used positive and negative correlations with other users to make decisions. The third algorithm determined the correlation between different artists, recommending artists who were strongly correlated to those a user already liked. The fourth algorithm, and the most effective according to the researchers, matched users based on whether they rated the same things either positive or negatively. In other words, their taste matched.

Similarity was the best variable. The more users in the system, and the more input users gave up front, the better Ringo worked—some users even described it as "unnervingly accurate." Ringo's innovation was how it acknowledged that the best recommendations, or the best indications of relevance, were likely to come from other humans rather than analysis of the content itself. It represented a scaling up of human taste.

Early Internet algorithms were designed to sift through a vast body of material for whatever was important to a user, and then present it in a coherent way. Recommendations were the goal: recommending a piece of information, a song, an image, or a social media update. Algorithmic feeds are sometimes more formally and literally labeled "recommender systems," for the simple act of choosing a piece of content.

The first wholly mainstream Internet algorithm, one that almost every Internet user has encountered, was the Google Search algorithm. In 1996, while studying at Stanford University, Sergey Brin and Larry Page, the cofounders of Google, began work on what would become PageRank, a system for crawling the Internet (which at that point amounted to perhaps one hundred million documents in total) and identifying which sites and pages were more useful or informative than others. PageRank worked by measuring how many times a website was linked to by other sites, similar to the way academic papers cite key pieces of past research. The more links, the more important a page was likely to be. The metric of citation "corresponds well with people's subjective idea of importance," Brin and Page wrote in a 1998 paper, "The Anatomy of a Large-Scale Hypertextual Web Search Engine." PageRank mingled a form of collaborative filtering with content filtering. By linking various pages, human users had already formed a subjective map of recommendations that the algorithm could incorporate. It also measured factors like the number of links on a page, the relative quality of the links, and even the size of text—the larger the text, the more relevant it might be for a particular search term. Pages with a higher PageRank were more likely to appear at the top of Google's list of search results.

Page and Brin's prediction that their system would remain functional and scalable as the Internet grew were correct. Decades later, PageRank has become almost tyrannical, a system that dominates how and when websites are seen. It's vital for a business or resource to make it to that first page of Google Search results by adapting to the PageRank algorithm. In the early 2000s, I perused many successive pages of Google results to find exactly what I was looking for. More recently, I hardly ever make it to the second page, in part because Google Search now frontloads text that it gauges will be relevant, pulling it from websites and displaying it directly to the user at the top of the search page, before the actual results. Thus a query like "Can I feed my dog carrots?"—the kind of question I googled incessantly in the early days of puppy ownership—will deliver an answer without a user ever having to load another site, further consolidating Google's authority. "Knowledge itself is power," Francis Bacon wrote in the sixteenth century, but in the Internet era, sorting knowledge might be even more powerful. Information is now easy to find in abundance; making sense of it, knowing which information is useful, is much harder.

Page and Brin wanted their system to be relatively neutral, evaluating each site solely in terms of its relevance. The algorithm's directive was to prioritize the best information for the user. Catering the search to a particular site or business would ruin the results. "We expect that advertising funded search engines will be inherently biased towards the advertisers and away from the needs of the consumers," the entrepreneurs wrote in 1998. Yet, in 2000, they launched Google AdWords as the company's pilot product for advertisers. It is amusing to read their critique today, as advertising now provides the vast majority of Google's revenue—more than 80 percent in 2020. As PageRank attracted billions of users to Google Search, the company could also track what the users were searching for and could thus sell advertisers space on particular search queries. The ads a user sees were just as informed by the algorithm as the search results were. And advertising, built on the search algorithm, turned Google into a behemoth.

By the early 2000s, algorithmic filtering was already dictating our digital experiences. The Amazon website began using collaborative filtering as early as 1998 to recommend products for customers to buy. Rather than attempting to measure similar profiles of users to approximate taste, as Ringo did, the system worked by determining which items were likely to be purchased in tandem—a rattle with a baby bottle, for example. A 2017 paper cowritten by an Amazon employee described the bombardment of such suggestions on the site:

> The homepage prominently featured recommendations based on your past purchases and items browsed in the store…. The shopping cart recommended other items to add to your cart, perhaps impulse buys to bundle in at the last minute, or perhaps complements to what you were already considering. At the end of your order, more recommendations appeared, suggesting items to order later.

The algorithmic recommendations resemble the shelves stationed just before the register in a Trader Joe's, one last push of products that you may need. But in this case, what was recommended was tailored to each website user, resulting in "a store for every customer," as the paper described. Amazon found that the personalized product recommendations were much more effective in terms of click-throughs and sales than unpersonalized marketing tactics like banner advertisements and lists of bestselling products, which can't be as tightly targeted. The recommendation algorithm improved business and appeared convenient for the customer, who might find something they didn't know they needed. (Right now, my Amazon home page recommends a cordless power washer and a Japanese omelet pan.)

These early algorithms sorted individual emails, musicians (as opposed to specific songs), web pages, and commercial products. As digital platforms expanded, recommender systems moved into more complex areas of culture and operated at much faster speeds and higher volumes, sorting millions of tweets, films, user-uploaded videos, and even potential romantic partners. Filtering became the default online experience.

This history is also a reminder that recommender systems are not omniscient entities but tools built by groups of tech researchers or workers. They are fallible products. Nick Seaver is a sociologist and a professor at Tufts University who studies recommender systems. His research focuses on the human side of algorithms, how the engineers who make them think about algorithmic recommendations. In my discussions with him, Seaver always made sure to clarify the ambiguous entity of the algorithm, separating the individual equation from the corporate motives behind its design and its eventual impact on the user. “The algorithm is metonymic for companies as a whole,” he told me. “The Facebook algorithm doesn’t exist; Facebook exists. The algorithm is a way of talking about Facebook’s decisions.”

The technology is not at issue—one can no more blame an algorithm itself for bad recommendations than blame a bridge for its engineering flaws. And some degree of reordering is necessary to make the vast stores of content on digital platforms comprehensible. The negative aspects of Filterworld might have emerged because the technology has been applied too widely, without enough consideration for the experience of the user, rather than for the advertisers targeting them. The recommendations, such as they are, don’t work for us anymore; rather, we are increasingly alienated by them.

EARLY SOCIAL MEDIA

My first meaningful memories of social media come from Facebook, which I joined after I accepted an admission offer from Tufts University, where I went to college. At that time, in summer 2006, prospective users needed an official .edu email address to access the full college section of the platform. That early iteration of Facebook is nearly unrecognizable when compared to its present-day anatomy. Back then, Facebook’s reach was strictly limited; I mainly used it as a means to connect with other incoming Tufts students. If Facebook today is a frenetic highway with exits and on-ramps every few seconds, in the aughts it was more like a high school rec room where only a few people could hang out at a time. You built a profile, updated your status on the profile, and joined groups around common interests—but not much else.

Facebook was hardly the first way to socialize online. Friendster and MySpace were its predecessors. AOL’s Instant Messenger and Google’s gChat provided engrossing ways to hang out with your friends in real time. By

2006 I had already spent hundreds of hours on older forum websites discussing video games and music. But Zuckerberg's Facebook tied online identity coherently and consistently to the offline world. The platform encouraged users to use their real names rather than arcane aliases and influenced real-life plans in the small world of college: throwing parties, planning academic activities, and conducting relationships. In doing so, it paved the way for the mainstreaming of online social life for millions, and then billions, of users.

In September 2006, not long after I joined the network, Facebook implemented one of its biggest changes, a feature that would set the course for its future as the big-box everything-for-sale store of the Internet. The News Feed, a running list of updates, posts, and alerts, became the primary feature of the platform. It was unignorable, like a newly built highway that cut through a quiet village. "Now, whenever you log in, you'll get the latest headlines generated by the activity of your friends and social groups," Facebook's official update note announced.

The patent for the News Feed, filed that year, though it wasn't granted until 2012, described its purpose: "A system and method provides dynamically selected media content to someone using an electronic device in a social network environment." In other words, the News Feed was a flow of information dictated by an algorithm that determined what to show a user. Another patent claimed the ability to "generate dynamic relationship-based content personalized for the members of the web-based social network." At first the News Feed was just a stream of announcements of changed dating statuses and updated profile pictures. It wasn't particularly threatening.

The News Feed patent application's longer description suggests a system of collaborative filtering, acting on a much larger scale than the email systems of the 1990s. It's worth quoting in full because it predicts what much of life online, from social networks to streaming and e-commerce, became in the decade that followed: so many automated feeds dictated by corporations more so than users, gradually forming a more passive relationship between users and the content feed.

> Items of media content are selected for the user based on his or her relationships with one or more other users. The user's relationships with other users are reflected in the selected media content and its format. An order is assigned to the items of media content, for example, based on their anticipated importance to the user, and the items of media content are displayed to the user in the assigned order. The user may change the order of the items of media content. The user's interactions with media content available in the social network environment are monitored, and those interactions are used to select additional items of media content for the user.

All the elements of the algorithmic feed are present in this passage—a system that anticipates a piece of content's relative importance to an individual user, determined by surveillance of content they engaged with in the past, and pushes whichever content is deemed most likely to be equally engaging to the top of the list. The goal was to filter content to select what is most interesting, therefore encouraging a user to consume a higher volume of content and follow more accounts overall. Users being able to use social media more often and stay on the sites longer is what made them viable. (If our friends aren't active on Facebook, which has been the case for me in recent years, then we are likely to tail off our activity, too.)

At first, the News Feed was ordered purely chronologically, with the most recent updates first, but it gradually followed a more algorithmic logic. As Facebook grew and users added more connections, expanding from personal relationships to publications and brands, the volume of individual updates increased. Over time, the updates weren't just mundane notes from friends but messages from groups, links to news stories, and announcements of sales. Casual users couldn't hope to follow a chronological feed with such a volume and variety of posts, and if they tried, they would either be overwhelmed or fail to catch an important post, which might cause dissatisfaction with the platform. Ultimately, the scale and speed of consumption made aggressive algorithmic filtering necessary for Facebook.

Facebook's Like button, with its signature thumbs-up, was introduced in 2009, providing one form of data on how interested a user might be in a particular piece of content. User engagement, measured by likes, comments, and one account's previous interactions with another, factored into the order of the feed. That algorithmic system was called EdgeRank, and Facebook identified its principal variables as affinity score, edge weight, and time decay. "Edge" referred to any action people carry out on Facebook, which is then sent to the News Feed as an update to be listed. Affinity score represented how connected a user was to the poster and the strength of the connection (e.g., consistently commenting on friend's posts). A comment counted more than a like, and recent interactions counted more than older ones. Edge weight evaluated different categories of interactions: an update of a friend posting a new photo might be given more weight by the algorithm than posting a link to a news article or joining a new group. Time decay was the age of the action; recent actions were more likely to be at the top of the News Feed than older ones, if the other factors were equal. The EdgeRank scores were not permanently assigned once, like the outcome of a basketball game in a tournament, but changed instant to instant. And those three categories aren't simply single, neutral data points; they are collections of data packaged and interpreted in specific ways by Facebook.

It's hard to track the evolution of Facebook's algorithmic feed because it is constantly updated, and the company reveals details only intermittently. What we do know about it beyond official announcements comes down to investigative reporting from journalists and the experiences of users, who see the effects of updated algorithms long before they're made public. Familiar websites have a way of feeling different when the feed mechanism changes. On Facebook, for instance, you may notice that you see less of your friends' posts and more from groups or businesses, or that Instagram never shows you posts from a particular friend in your feed and you thus need to hunt them down using the search bar.

The algorithmic feed itself is not consistent or on a linear path toward some ultimate perfection. It changes with a company's priorities. In 2011, Facebook described the News Feed as "your own personal newspaper," suggesting its goal of mingling social updates with news stories from the outside world. In 2013, it said its algorithm worked to "detect content defined as high quality." But chasing whatever the company gauged as "high quality" was something of an absurd game over the course of the 2010s. If you wanted to get attention for your Facebook posts—a big problem for journalism publications and freelance writers—you had to guess at what kind of material was getting prioritized. The relationship was almost oppositional; only if you "gamed" the algorithm would you be heard. You could no longer rely on users who had followed or friended you seeing your posts.

At one point in my freelance journalism career, I recall a rumor going around that links to articles were no longer very highly weighted by the algorithm. So rather than posting our stories directly using a simple link, I and many other journalists added a link to the story only by commenting on the post. The trick was supposed to goose algorithmic promotion, even though it was more confusing for a reader. At another point, it became clear that writing text that resembled a marriage announcement and comments that said "congratulations" pushed posts to the top of the feed. So I began sharing my articles with fake weddings or other life milestones. These phenomena show how algorithms can warp language itself as users attempt to either game them or evade detection. More recently, on TikTok, euphemisms have emerged for terms that trigger the algorithm to block or slow down a video: "unalive" for kill, "SA" for sexual assault, "spicy eggplant" instead of vibrator, as the journalist Taylor Lorenz documented in the *Washington Post.* Such vocabulary was nicknamed "algospeak": speech molded in the image of the algorithm.

It was unclear if the tricks I used on Facebook had much of an impact, but I was willing to try anything to reach potential readers. It was like designing a website for Google search-engine optimization: journalists optimized content for the metrics of the algorithm, or at least what we perceived them to be. The process felt manipulative and at times Kafkaesque; we contended with an unseen, incomprehensible, ever-changing opponent.

Around 2015, Facebook decided to prioritize video content, so the recommendation algorithm promoted videos much more than it did previously. Media companies then "pivoted" to making videos to chase that audience, sometimes with the help of funding from Facebook itself. That effort lasted only a few years, and then Facebook deprioritized videos once more, leading to waves of layoffs at those same media companies, including BuzzFeed, Mashable, and MTV. (After the program ended, it also emerged that Facebook had lied about the traffic the videos were getting, inflating the numbers up to nine times, according to a lawsuit.) The algorithmic feed kept shifting. In 2016, Facebook added "reactions" to posts, so that viewers could respond with a range of emoticons rather than just the Like button. Posts that received many emoticon reactions got more promotion. But that change backfired, too, when incendiary content—posts that received many angry-face reactions, for example, like rage-inducing political stories—was getting too much promotion and souring the tone of the entire site. That they attracted more engagement didn't mean the posts were necessarily more worthwhile.

It wasn't only Facebook that moved from a chronological feed to an increasing volume of algorithmic recommendations. Almost every major social network followed the same path over the 2010s. Filterworld began taking shape in the middle of the decade when algorithmification intensified.

Facebook acquired Instagram in 2012, when it only had thirteen employees. In the years since, the photo-sharing app has become more like Facebook itself, moving away from a linear feed of photos uploaded by friends into a stream of videos, ads, and recommended posts. In March 2016, the Instagram feed began switching from a chronological to an algorithmic arrangement. The change was tested on small groups of users then rolled out to more and more, until it hit everyone. The increasingly out-of-order feed induced a sense of confusion and anxiety akin to the feeling of someone rearranging the furniture in your house without your knowledge. Before, by scrolling through the feed, you were moving back in time. But suddenly, a post from two days ago appeared at the top of your feed.

Early 2016 was also when Twitter became less chronological, briefly making the algorithmic feed the default when users first got on the app—a problem for a site that many people used as a real-time news ticker. (The chronological option was called “Twitter Classic,” as if it were a beloved junk-food flavor.) Later, the app would swap users over to an algorithmic feed automatically after a while and force them to opt out of it. Although Netflix’s content recommendations had long been algorithmic, 2016 was also when the streaming service began changing its home-page interface, prioritizing recommendations and individualizing it for each user.

Larger cultural consequences, unexpected by users and perhaps by the companies themselves, followed this shift—the way that damming a river changes an entire ecosystem. When feeds are algorithmic, they appear differently to different people: It’s impossible to know what someone else is seeing at a given time, and thus harder to feel a sense of community with others online, the sense of collectivity you might feel when watching a movie in a theater or sitting down for a prescheduled cable TV show. The advent of Filterworld has seen a breakdown in monoculture. It has some advantages—more than ever before, we can all consume a wider possible range of media—but it also has negative consequences. Culture is meant to be communal and requires a certain degree of consistency across audiences; without communality, it loses some of its essential impact.

Intensifying the problem of fragmentation was the fact that recommender-system updates do not roll out at the same time to all users at once across an app. For a year or two after 2016, my personal Instagram feed remained rigorously chronological, while everyone around me complained about not seeing what they wanted. Eventually my feed switched over too, and I understood what they had been complaining about. We came to rely on our feeds working in certain ways, and when those changed, how we behaved as consumers also changed. We were stuck in the algorithmic flow, driven by whichever variables it was programmed to seek.

The rise of the algorithmic feed, like the Internet itself, came slowly and then all at once. Early in the 2020s, as I’m writing, recommender systems seem unavoidable, mediating our consumption of every form of digital media. Technology often appears to belong to the distant future right up until the moment the switch flips, and the leap forward becomes totally mundane, a simple fact of daily life.

In his sprawling early twentieth-century novel *In Search of Lost Time,* Marcel Proust excavated such subtle changes in personal sensibilities against the backdrop of evolving technology. In one passage, Proust's narrator describes the telephone as "a supernatural instrument before whose miracles we used to stand amazed, and which we now employ without giving it a thought, to summon our tailor or to order an ice cream." The telephone had only been invented in the late nineteenth century, when Proust's novel is set. By 1899, there were seven thousand telephone subscribers in Paris. And yet telephones had still become banal. Even during one of his first phone calls, the narrator becomes annoyed by the device instead of awed. Proust wrote: "Habits require so short a time to divest of their mystery the sacred forces with which we are in contact, that, not having had my call at once, my immediate thought was that it was all very long and very inconvenient, and I almost decided to lodge a complaint."

In 1933, the Japanese novelist Junichiro Tanizaki memorialized another moment of technological change when he wrote *In Praise of Shadows,* a book-length essay about electric lights arriving in Tokyo. The metaphorical switch had flipped; within Tanizaki's lifetime (he was born in 1886), electric lights had gone from unknown in his country to ubiquitous, thanks to the intrusion of the West beginning in 1867, in a wave of increasing globalization and subsequent clashes of cultures. The Westerner's "quest for a brighter light never ceases," Tanizaki wrote. In the essay, Tanizaki mourned the unique forms of Japanese culture that the old dimness of candlelight had inspired, from the gleam of gold leaf on a home's interior sliding door to the murky appearance of miso soup in a darkened restaurant: "Our cooking depends upon shadows and is inseparable from darkness."

Yet Tanizaki couldn't ignore the attraction of electricity and other new devices: porcelain toilets, heaters, and neon signs. "It was not that I objected to the conveniences of modern civilization," he wrote. As he described in his fiction, the novelist loved movie theaters and modern architecture as much as he appreciated tradition. *In Praise of Shadows* tracked how technology changed, culture adapted, and personal taste shifted in turn—a pattern we see throughout Filterworld in our own time.

With new technology, the miraculous quickly becomes mundane, any glitch in its function is felt as bothersome, and finally it becomes ignorable, the miracle forsaken. We forget that life wasn't always this way, that we couldn't directly speak to people across long distances, that ceiling lights didn't make every room bright, or that we didn't have our information and media automatically filtered by machines. Such is the presence algorithmic feeds now have in our lives; the algorithm is often unconsidered, part of the furniture, noticed only when it doesn't function in the way it's supposed to, like traffic lights or running water.

CATCHING ALGORITHMIC ANXIETY

If the chess-playing Mechanical Turk was an (illusionary) encounter with miraculous technology that made decisions independent of a human hand centuries ago, we now undergo that experience dozens of times a day, in the digital spaces that we are accustomed to relying on. It's hard to overstate the ubiquity of machine influence. From what we can tell using public metrics, Facebook today has nearly three billion users. Instagram has around two billion. TikTok has over one billion. Spotify has over 500 million. Twitter has 400 million. Netflix has over 200 million. For all the people on these platforms, every interaction, every moment of passive consumption, is mediated by algorithmic recommendations. Even if some users can opt out of an algorithmic feed, their participation contributes to the data that fuels other users' recommendations. The dragnet is inescapable. Social networks and streaming services have become the primary way a significant percentage of the global population metabolizes information, whether it's music, entertainment, or art. We now live in an era of algorithmic culture.

Technology companies have long sought to achieve this massive scale. Monopolistic growth is more important to these entities than the quality of user experience and certainly more important than the equitable distribution of culture through the services' feeds. (A digital platform has none of the curatorial responsibility of, say, an art museum.) According to Silicon Valley ideology, the pursuit of scale far outweighs any negative consequence it might have, as a memo written by Andrew Bosworth, a deputy of Mark Zuckerberg's at Facebook, demonstrated in 2016:

> So we connect more people. That can be bad if they make it negative. Maybe it costs someone a life by exposing them to bullies. Maybe someone dies in a terrorist attack coordinated on our tools. And still we connect people. The ugly truth is that we believe in connecting people so deeply that anything that allows us to connect more people more often is *de-facto* good.

That statement is a stark illustration of the attitude that if people are using a platform, staying engaged and active, then it counts as successful—no matter what they are doing. That ongoing engagement is sustained by automated recommendations, delivering the next provocative news headline or hypnotic entertainment release. Today, it is difficult to think of creating a piece of culture that is separate from algorithmic feeds, because those feeds control how it will be exposed to billions of consumers in the international digital audience. Without the feeds, there is no audience—the creation would exist only for its creator and their direct connections. And it is even more difficult to think of consuming something outside of algorithmic feeds, because their recommendations inevitably influence what is shown on television, played on the radio, and published in books, even if those experiences are not contained within feeds. Filterworld spills out everywhere.

Trevor Boffone, a scholar of theater who took up work as a high school teacher, gave me an apt description of what algorithmic culture amounts to: "The films that do well are films that have TikTok followings; the Billboard Hot 100 is dictated by TikTok; you go to Barnes and Noble and you see a BookTok table," he said. (BookTok is a term for TikTok's community of literary influencers.) In other words, for a piece of culture to be commercially successful, it must already have traction on digital platforms. Boffone's career, too, has been shaped by algorithmic feeds. When he began learning TikTok dance moves with his teenage students and posting videos of them online, he quickly accrued hundreds of thousands of followers on Instagram and other platforms. He appeared on national television, briefly becoming a viral character—the dancing teacher. Following his experiences, he published an academic monograph on dance performance, a subject that had quickly become more compelling to universities and editors with its rising public popularity on TikTok. “I’ve had more interest in one month of this year in my work than in the previous ten years combined,” Boffone told me.

Boffone’s experience follows a fundamental rule of Filterworld: Under algorithmic feeds, the popular becomes more popular, and the obscure becomes even less visible. Success or failure is accelerated. “A traditional Instagram post, the life of it is dictated by the first three to five minutes of the post,” Boffone said. If a post gets engagement immediately, then it’s likely to get more, and vice versa. This dynamic can be cruel. When I post an offbeat Instagram image or an obscure tweet and it doesn’t get much action, that doesn’t stop me from checking back multiple times for more likes, even though I know I haven’t hit the algorithmic jackpot.

The absence of attention inevitably raises the question of what the feed *will* promote, tacitly encouraging safer choices, urging conformity. Who receives promotion is also a problem. It’s often not the original creators of a meme or trend who get credit, attention, and thus financial gain from its popularity in an algorithmic feed. TikTok choreography itself is an example. The TikTok influencer Charli D’Amelio became famous in 2019 for her dance videos on the platform. But one of the moves she popularized and was often credited with, called the Renegade, was actually created earlier by Jalaiah Harmon, a Black teenager from Georgia. The Renegade was a series of front-facing movements perfect for the TikTok screen, with swinging punches and hip shakes—not too difficult a sequence, but also tough to memorize and thus rewarding to re-perform.

Harmon first posted the dance on an app called Funimate as well as Instagram. But TikTok’s hyper-algorithmic feed accelerated it to mainstream fame, seeded by D’Amelio’s following, even as it helped to erase Harmon’s authorship, since D’Amelio didn’t cite her. Content creators from marginalized groups, who don’t have the same access to media and attention as, say, a white, private-school-educated, professionally trained dancer, like

D'Amelio, have a harder time benefiting from the tides of Filterworld. (Since being recognized for her work, Harmon herself has gained three million TikTok followers.)

Given that these capricious systems control so many facets of our lives, from socializing with our friends to building audiences for our creative projects, is it any wonder that social media users feel paranoid? We're encouraged to overlook algorithmic processes, but their glitches remind us of their unearned authority. The ambiguity of algorithmic influence creates a feeling that has been labeled "algorithmic anxiety." Algorithmic anxiety describes the burgeoning awareness that we must constantly contend with automated technological processes beyond our understanding and control, whether in our Facebook feeds, Google Maps driving directions, or Amazon product promotions. We are forever anticipating and second-guessing the decisions that algorithms make. Algorithmic anxiety is not hypothetical or an abstraction: It's already prevalent. Technology companies are aware of it and have been manipulating the feeling in their users for years.

In 2018, Shagun Jhaver, at the time a PhD candidate at the Georgia Institute of Technology, worked with two Airbnb employees to conduct a sociological study of the company's users. They analyzed how the platform's hosts—who rented out their homes on the service for income—interacted with and felt about Airbnb's algorithmic recommendation, search, and ratings systems, which helped renters find and book their listings. Jhaver and the other researchers coined the term *algorithmic anxiety* for the hosts' "uncertainty about how Airbnb algorithms work and a perceived lack of control," the team wrote in their findings. Hosts worried that the search algorithm was unfairly ignoring them or prioritizing other properties. Jhaver noticed that the anxiety was ascribed more to the technology than the quality of the actual homes the hosts were renting out: "It was particularly to do with the algorithm itself rather than improving their listing and property in other ways," he told me.

Airbnb forces a "double negotiation" for the hosts, the researchers wrote, because they must determine what their guests are looking for in a listing as well as which variables the algorithms are prioritizing to promote their property more often. But hosts could not tell which variables actually boosted their listing. They believed factors like the number of reviews accrued, the quality of reviews, and the number of photos available would help their chances of promotion, but they were less certain as to whether the algorithm analyzed their pricing, home amenities, or length of tenure as a host. They had little information about how the systems worked. It was all a matter of perception. As one host in the study complained: “It’s frustrating seeing the search: lots of listings that are worse than mine are in higher positions.”

Quality is subjective, of course, but the host’s sentiment speaks to how users can feel misunderstood and misjudged by algorithmic evaluations. “It’s like an exam, but you don’t know what’s going to be on this exam, or how to score well on this exam,” Jhaver explained. And it’s not just the users who don’t know what’s going on. Jhaver continued: “At the end of the day, even the people who create the algorithms cannot tell you which factor was responsible for which decision; the complexity of the algorithm is so high that disentangling different factors is just not possible.”

Failing to game the algorithm may cause an immediate drop in income for hosts, which they, as any worker, rely on remaining consistent. (The inconsistency of algorithmic promotion forces us to engage with it and stress about it even more, like repeatedly pulling a slot machine lever to hit the jackpot.) Gig-economy platforms like Airbnb have long promised flexible work and alternative ways of making or supplementing a living, but they also created a new form of labor in the need to stay up to date on changes in algorithmic priorities. Where hosts worry about Airbnb’s search algorithm, artists similarly fret about Instagram’s and musicians about Spotify’s. The hosts’ reaction to such algorithmic anxiety, the researchers found, was to develop “folk theories”—superstitious tricks that were meant to goose more algorithmic promotion and better search results—the same way I used to post my article links with fake wedding announcements. Some of the strategies included constantly updating their listings calendar, changing their profile details, and even opening the Airbnb website more often throughout the day. The tricks bring to mind a child putting a spoon under their pillow to cause a snow day and are perhaps equally as effective. As the researchers found, hosts "usually had doubts about whether such theories were true but despite their uncertainty still performed those actions in an attempt to influence the algorithm."

Algorithmic anxiety is something of a contemporary plague. It induces an OCD-ish tendency in many users toward hyperawareness and the need to repeat the same rituals, because when these rituals "work," the effect is so compelling, resulting in both a psychological dopamine rush from receiving attention and a potential economic reward if your online presence is monetized. It undergirds so many of our behaviors online: selecting the right profile picture, curating an attractive grid of photos on an Instagram account, choosing the right keywords on a marketplace listing. We worry that our posts either won't be seen by the right people or will be seen by too many if selected for virality, exposing us to strangers. There's an emotional fallout to this quest for attention: we end up both overstimulated and numb, much like a glassy-eyed slots player waiting for matching symbols to come up.

Algorithmic anxiety happens because there is a dramatically asymmetrical relationship between user and algorithm. For the individual user, trying to predict or dictate the outcome of the algorithm is like trying to control the tide. To continue the metaphor, all users can do is surf the wave that's already formed. There is little incentive for companies to assuage this anxiety because a user's confusion can be beneficial to business. When a company's product is ineffective or a user encounters difficulty, it can be blamed on the opaque entity of "the algorithm," which is perceived as external to both the users and the company itself, since they are likened to opaque "black boxes." Exploitation is disguised as an accidental glitch instead of an intentional corporate policy. In reality, a company like Facebook is wholly in control of their algorithmic systems, able to change them at will—or turn them off.

Algorithmic anxiety places the burden of action on the individual, not the business—the user must change their behavior or risk disappearing. Users sometimes complain of being "shadowbanned" when their posts or content on a platform suddenly lack the same level of engagement as before. Users often fear that their account specifically has been blocked without warning or recourse by some decision-maker; but the algorithmic priorities may simply have silently changed, and traffic is no longer flowing in their direction. The effect goes back to the Mechanical Turk; we can't always tell the difference between technology working and the *illusion* of technology working, but the perception may be just as impactful, in the end, as the reality.

In her 2019 dissertation titled *Algorithmic Anxiety in Contemporary Art,* the scholar Patricia de Vries defined algorithmic anxiety as a condition in which "the possible self is perceived to be circumscribed, bounded, and governed by algorithmic regimes." Her words feel breathtakingly accurate. The possibilities that we perceive for ourselves—our modes of expression and creation—now exist within the structures of digital platforms. The consequences of such anxiety include "algorithmic determinism, fatalism, cynicism, and nihilism," de Vries wrote. It builds to a sense that, since we users cannot control the technology, we may as well succumb to the limits of algorithmic culture and view it as inevitable. Many users have already entered such a state of despair, both dissatisfied and unable to imagine an alternative.

De Vries began observing this cultural shift as early as 2013, when she saw several museum exhibitions highlighting the work of artists who were critical of automated surveillance and data collection. While algorithmic feeds had only just begun entering mainstream experience, events like the 2010 Flash Crash, caused by algorithmic stock trading, and technology like facial recognition had implanted the word in news headlines. By the middle of the decade, when she developed her research, it became "this sort of object of our fascination," de Vries told me. The algorithm was a specter that haunted any encounter with digital platforms and their increasingly intrusive presence in our lives. That is not to say we understood what exactly algorithms were doing, per se: "Just as the fear of heights is not about heights, algorithmic anxiety is not simply about algorithms," de Vries said.

To move forward, we must disentangle the effects of algorithmic recommendations as technology from the ways that we have habitually adopted them as the primary gatekeepers of our online communication. Algorithms, after all, are inextricable from the data they run on, which has been created and is constantly refreshed by humans. Actual influence coexists with the fear of influence, which is equally manipulative. Algorithms entered our lives by promising to make decisions for us, to anticipate our thoughts and desires. Filterworld represents the establishment of the psychic world of algorithms—not just how they work, but how we users have come to rely on them, allowing them to displace our own agency, even as we come to resent their looming presence.