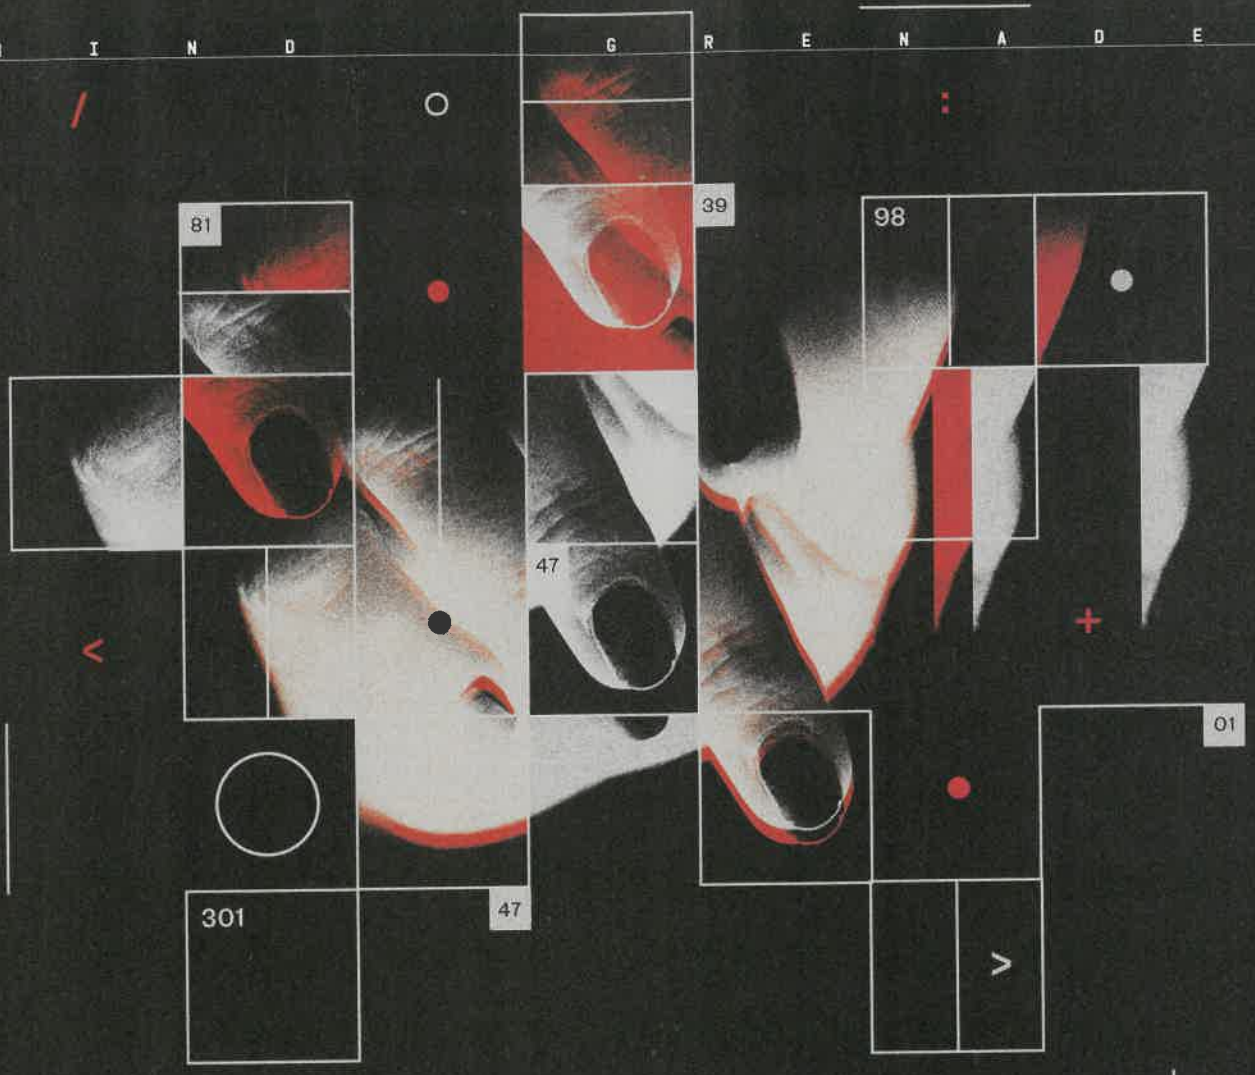

THE POWER OF SUGGESTION

How recommendation algorithms run the world.

BY ZEYNEP TUFEKCI

This March, a book that advances an outlandish conspiracy theory—a theory whose name I will not mention—soared in Amazon's sales rankings. The book's rise was helped greatly when the ecommerce giant put the book on its carousel of recommended titles, which is shown to shoppers who aren't searching for that particular book. That fueled more curiosity and sales. Which led to more recommendations. ■ The particular conspiracy theory outlined in this book holds that President Donald Trump *pretended* to collude with Russia precisely to *ensure* that he would be investigated, which would give him a chance to *secretly collaborate* with special prosecutor Robert Mueller to investigate and finally arrest former presidential candidate Hillary Clinton, who belongs to a global satanic cult of pedophiles

ILLUSTRATION / MIKE MCQUADE

with Barack Obama and George Soros. Yes, it's that unhinged. So why is Amazon recommending this book to unsuspecting shoppers? It's not because this theory has enormous persuasive power or a best-seller-size audience. Blame it on the tyranny of recommendation algorithms.

What should you watch? What should you read? What's news? What's trending? Wherever you go online, companies have come up with very particular, imperfect ways of answering these questions. Everywhere you look, recommendation engines offer striking examples of how values and judgments become embedded in algorithms and how algorithms can be gamed by strategic actors.

Cults need attention to recruit. Recommendation algorithms shouldn't be doing their bidding this easily.

Consider a common, seemingly straightforward method of making suggestions: a recommendation based on what people "like you" have read, watched, or shopped for. What exactly *is* a person like me? Which dimension of me? Is it someone of the same age, gender, race, or location? Do they share my interests? My eye color? My height? Or is their resemblance to me determined by a whole mess of "big data" (aka surveillance) crunched by a machine-learning algorithm?

Deep down, behind every "people like you" recommendation is a computational method for distilling stereotypes through data. Even when these methods work, they can help entrench the stereotypes they're mobilizing. They might easily recommend books about coding to boys and books about fashion to girls, simply by tracking the next most likely click. Of course, that creates a feedback cycle: If you keep being shown coding books, you're probably more likely to eventually check one out.

Another common method for generating recommendations is to extrapolate from patterns in how people consume things. People who watched this then watched that; shoppers who purchased this item also added that one to their shopping cart. Amazon uses this method a lot, and I admit, it's often quite useful. Buy an electric toothbrush? How nice that the correct replacement head appears in your recommendations. Congratulations on your new vacuum cleaner: Here are some bags that fit your machine.

But these recommendations can also be revealing in ways that are creepy. For a long time, Amazon recommended security keys—a safer hardware alternative to password protection—to people who bought my book about online protest movements. I do advocate that people should purchase security keys, and I guess I'm glad that some people listened to me! But now everyone who looks at my book's Amazon page gets a dose of intel about my readers and how they protect themselves.

One final method for generating recommendations is to identify what's "trending" and push that to a broader user base. But this, too, involves making a lot of judgments. For starters, there is no one way to define trending. On any given day, there are a lot of online conversations about any and all Kardashians, but they aren't trending on Twitter. That's because most trending-type recommendation algorithms employ a logic that filters out common terms as background noise and highlights those that have acceleration and velocity on their side.

This definition of trending buries ongoing conversations and amplifies sensational, new things. Ironically, this is often the trouble with traditional media as well. Chronic problems (health care, homelessness, hunger, traffic accidents) are given short shrift, while more rare events (terrorism, airplane accidents) receive sensational, saturated coverage. This isn't surprising; novelty grabs our attention. Online algorithms exploit the same societal vulnerability that editors in traditional media target.

But here's the other problem with algorithms that value acceleration over sustained signals: It's not that hard to generate a little ersatz velocity. All sorts of people game the system this way, many of them sympathetic. In fact, I first learned the details of how Twitter's trending algorithm works from democracy activists in Bahrain during the Arab Spring. They had figured out that their ongoing, large-scale discussions didn't register across Twitter even if they dominated the national conversation, but a novel hashtag that was suddenly used by lots of people would quickly trend. So the activists coordinated to create brand-new hashtags and held them in reserve until a prearranged time—and then everyone would all tweet at once. Voilà: The hashtags would quickly trend, drawing global attention to the activists' plight.

But what works for one group also works for others. I'm guessing that's how the conspiracy book got so much play on Amazon. If enough adherents opened up their wallets all at once in a coordinated effort, that would have been enough to make Amazon's algorithms take notice—and start amplifying the book even more.

This particular conspiracy theory may be batty, but it's no joke. One believer, armed with an AR-15, blocked the bridge at the Hoover Dam in an armored vehicle; two others have been tied to alleged murders. That they're organized enough to strategically game algorithms isn't a good sign. Such cults need attention to recruit, and recommendation algorithms shouldn't be doing their bidding this easily.

What's the alternative? At a minimum, there should be more transparency as to how and why certain things are recommended for us to watch or buy or read. The counterargument would be that transparency would make them easier to game. My counter-counterargument is that maybe algorithms that are so easy to game shouldn't be used at all. We should also be able to ask for different kinds of recommendations when we shop. How about books that people like me *rarely* read—but may find interesting? How about showing me a topic that's been discussed over a long period of time among a large community of people? Or how about giving me the ability to turn all these recommendations off altogether? Sometimes, less is better.

ZEYNEP TUFEKCI *(@zeynep) is a* WIRED *contributor and a professor at the University of North Carolina at Chapel Hill.*