

Facebook, surveillance capitalism, and feedback control

On May 10, a United States Senate Committee sent a [letter](#) to Facebook's Mark Zuckerberg, asking him to respond to accusations that "employees of Facebook routinely suppressed conservative political viewpoints on the social network."

Former Facebook employees reportedly said that they had manipulated the content of "Trending Topics," which displays news based on recent popularity, pages that a user has "liked," and his or her location. The committee asked whether the supposedly "neutral, objective algorithm" used by Trending topics is "in fact subjective and filtered to support or suppress particular political viewpoints." In fact, one website reported that Facebook employees asked Zuckerberg whether the company ought to help prevent Donald Trump from winning the presidential election.

The Facebook controversy highlights an increasing realization about the power of tech companies not only to collect data about reality, but to influence reality itself. The presence and depth of monitoring technologies, the achievements of data analytics, and the ubiquity of social media are combining to yield emergent new properties – some encouraging and some alarming.

These new properties are the topic of [a recent article](#) by Harvard Business School Professor Emeritus Shoshana Zuboff. She calls the phenomenon "surveillance capitalism." Reading her article, I was struck by the resonance between surveillance capitalism and what engineers call "feedback control." Specifically, surveillance capitalism is a form of "human-in-the-loop" (HiTL) feedback control. HiTL feedback

control is a useful lens through which to understand the revolution in data analytics, and I will employ its terminology in this article.

Feedback control in the US and Soviet Union

In the 1950s and 1960s, the Cold War fostered a fierce race to explore and dominate outer space. Realizing that satellites and rockets would require a degree of automatic navigation, both the US and the USSR fostered rapid advances in a discipline known as “feedback control.”

Every feedback control system has three components: sensors, controllers, and actuators. *Sensors* measure important data from the environment. *Controllers* calculate an optimal response to the data. *Actuators* put the response into practice.

Automobile cruise control is perhaps the simplest example of a feedback system. The sensor is the vehicle’s speedometer, the controller is a computer built into the car, and the “actuator” is the gas injection into the engine. If the vehicle is going too slow, the computer instructs the engine to inject more gas. If the vehicle is going too fast, it says to reduce that amount. The driver provides a target speed – say 60 miles per hour – and the system adjusts to match that speed. Similar feedback control designs have been used to orient satellites, control landings of the Mars rovers, and provide automatic guidance for precision missiles.



Cruise control is an example of a feedback system. The goal is to bring the car to 60mph. A sensor – the speedometer – measures the speed of the car. An onboard controller computes a change in the amount of gas to inject. The engine serves as an actuator, which affects the speed of the car. This changes the value that the sensor reads, and the loop repeats.

In fact, the study of feedback control does not pertain to any one discipline. Rather, feedback control is about manipulating the properties of “black boxes.” Cruise control systems do not know exactly how engines work, but they do know how to execute a general algorithm. Similarly, blood glucose control, epidemic prediction, and traffic light design all use feedback control techniques by treating underlying systems as “black boxes.” Sensors measure the behavior of the boxes, controllers analyze the data, and actuators change inputs to the boxes in order to achieve desired outputs.

Pros and cons of feedback control in web browsing and the internet of things

Sensors, controllers, and actuators are also part of the new phenomenon of Zuboff’s “surveillance capitalism,” or, as I am describing it, HiTL feedback control. Sensors in HiTL feedback control consist of search engines that track online activity, smartwatches that upload information about exercise routines, and phones that monitor app usage. Controllers are built from data analytics engines and machine learning platforms. Finally, actuators take the forms of web browsers and social media sites that modify the ways in which we view the internet and behave online.



Feedback control of humans using the internet and social media also involves sensors, controllers, and actuators. Each of these roles has been altered by new technologies in ubiquitous tracking, data analytics, and digital media.

In her article on surveillance capitalism, Zuboff points to sociotechnical changes that have revolutionized each of these components, using (and criticizing) several studies in this area published by Google Chief Economist Hal Varian. Zuboff points to developments in “new contractual forms due to better monitoring,” “personalization and customization,” and “data extraction and analysis.” In the lens of feedback control,

these fulfill the roles of actuators, controllers, and sensors, respectively. Building off of Zuboff's analysis, we can consider the role of each of these components in turn.

Actuators in HiTL feedback systems

Very recently, businesses have learned to influence human behaviors using innovative HiTL "actuators." While traditional control systems use motors or levers to drive the behavior of physical systems, corporations are increasingly using what Zuboff calls "new contractual forms" as actuators of human systems. In these new contractual forms, companies control human behavior not through legal demands, but instead through a type of automatic policy enforcement enabled by technology.

For instance, many drivers can now earn insurance discounts by maintaining safe speeds. Insurance companies install small units in cars which report how many times the vehicle has exceeded 80 mph. In this case, the commercial-technical system serves as a HiTL actuator by incentivizing good driving. This entails an interesting shift in policy enforcement from the legal to the automatic. Drive too fast and you suffer penalties as a simply automatic result.

Corporate fitness incentives are another HiTL actuator or "new contractual form." Employees in many companies wear devices which track the number of steps that they take. Some companies also send digital notices about ways to stay in shape. The more steps that employees take and the more health bulletins that they read, the more of a yearly bonus they receive. One recent headline summarized the phenomenon with the words: "As Health Incentives Rise, Many Get Paid To Work Out And Eat Kale." The trade-off is worth it for companies who want to keep their employees healthy and happy, and worth it for employees who are willing to advertise their exercise in exchange for cash. Both vehicle speed monitors and corporate fitness incentives are positive examples of actuators or "new contractual forms."

On the other hand, some HiTL actuators are dangerous. Think of one possible influence of Waze, “the world’s largest community-based traffic and navigation app.” Waze collects information from app users in order to predict the fastest driving routes by taking traffic into account. Waze users can also report accidents and the location of traffic police. Unfortunately, by telling drivers when they are approaching a cop car, it also tells them when they are *not* near one. The strategy of police who monitor traffic is probably predicated on concealing speed trap locations. Therefore, it may be that Waze notifications *increase* reckless behavior in locations where police are not present.

Controllers in HiTL feedback systems

Data analytics engines – the HiTL versions of feedback “controllers” also influence human behavior. Much of data analytics comes down to what Zuboff (again, citing Varian) studies as “personalization and customization.” Tech companies laud personalization as a beneficial service to their customers. Customers, when asked, tend to prefer to keep their privacy and discard the personalization. But to confound the matter, they do not actually put their dollars behind these statements.

In any case, this personalization actually has impacts on society as a whole. Think about what happens if you Google search the word “jaguar.” Your results depend on who Google thinks that you are.

Chances are that if you are 1) white, 2) male, and 3) earning triple-digits, then your top hit is an advertisement for luxury cars. If you do not meet these criteria, then your top result is probably instead about an animal. This type of personalization tends to deeply engrain existing social disparity. (In the language of feedback control, we call this Catch-22 a “positive feedback loop.”) Similarly, if a middle school student from the South Side of Chicago searches for

“World War II,” will he find the find the same *history.com* results as are displayed for a PhD student in Manhattan? Or will he receive advertisements for the WWII video game “Call of Duty?” What about when someone who is overweight approaches the vending machines of tomorrow? Will the machines show her soda instead of fruit juice?



You can view Google’s prediction of your interests by going to <http://www.google.com/ads/preferences/view>. Google assesses these based on your search history and activity on Google sites such as YouTube. The company uses these interests to control the ads that you receive. Google also allows users to edit their interests.

In the case of Facebook’s alleged fight against Trump, the type of HiTL feedback control would go beyond personalization. The nature of Facebook’s “Trending Topics” makes content manipulation a bad means to justify the end of getting a different candidate elected. Unlike, say, the editorial section of a newspaper, which readers know reflects the judgments of some particular individual, Trending Topics is portrayed as objective. Users do not know how its algorithm works. With around a billion active users each day, this gives Facebook the potential to deceive a lot of people.

Sensors in HiTL feedback systems

Finally, human feedback control is intensified by what Zuboff studies as “data extraction and analysis.” Other scholars describe this phenomenon as “ubiquitous sensing” or “pervasive monitoring,” emphasizing the ever-present nature of today’s tracking technologies. Human behavior is monitored not only by surveillance cameras, but also by search engines, online stores, sleep trackers, and even some refrigerators that reorder food via a touch-screen interface. Increasingly, we put on “wearable computing” devices that upload to the internet data about our physical and even medical behaviors.

Devices in “smart homes” also extract data about our electricity use and purchasing behaviors and send it to the cloud. True, sensors like these can help us to save energy. And they may be able to help us to keep a monthly budget. But since these sensors observe actions in the intimacy of our homes and bodies, we need to keep careful track of their capabilities. I wouldn’t like universities to make admissions decisions by tracking the way that students use their after-school hours. Nor would it be helpful for corporations to purchase web search histories in order to filter out “social radicals” from their payrolls.

Continuous experiments and conclusions

Certainly, many human feedback loops have positive impacts. Quick responses to medical emergencies, positive incentives to stay healthy, and municipal initiatives to save energy are all exciting implications of HiTL feedback systems. On the other hand, the intimacy of individual decision-making and access to common goods are threatened by ill-considered uses of this technology. Much of the surveillance and some of the behavioral influence is single-directional and opaque.

For example, Google’s Hal Varian lauds the tech company’s ability to carry out “continuous experiments.” Google, he says, runs about 10,000 experiments each day. “There are about 1,000 running at any one time,” he continues, “and when you access Google you are in dozens of experiments.” We already know that researchers using Facebook carried out similar types of experiments to verify that they could manipulate the moods of social network users by influencing the appearance of posts in the users’ news feeds. Exciting as this can be to academic minds, internet users ought to know whether they are the subjects of “continuous experiments.”

In sum, feedback control provides useful techniques to study “black box” systems – both physical systems and human-in-the-loop informational systems – in mathematically rigorous ways.

While insurance discounts and corporate health incentives are exciting and enabling examples of HiTL feedback loops, excessive personalization, monitoring that extends to personal behavior in the home, and implicit contracts that allow corporations to exert power over day-to-day decision making threaten important personal and social values. Feedback systems are typically oblivious to the contents of the black boxes that they control. In the case of the Facebook controversy in particular and HiTL feedback systems in general, these black boxes contain important contents: the lives of millions or billions of human beings.

—

Jeffrey Pawlick is a PhD Candidate in Electrical Engineering at the Tandon School of Engineering, New York University.

*This [article](#) by this article, visit MercatorNet.com for more.
?The views expressed by the author and MercatorNet.com are not necessarily endorsed by this organization and are simply provided as food for thought from Intellectual Takeout.???*