Intelligence and Autonomy

September 2016

The Wisdom of the Captured

Alex Rosenblat (alex@datasociety.net)
Tim Hwang (tim@datasociety.net)
Edited by Patrick Davison

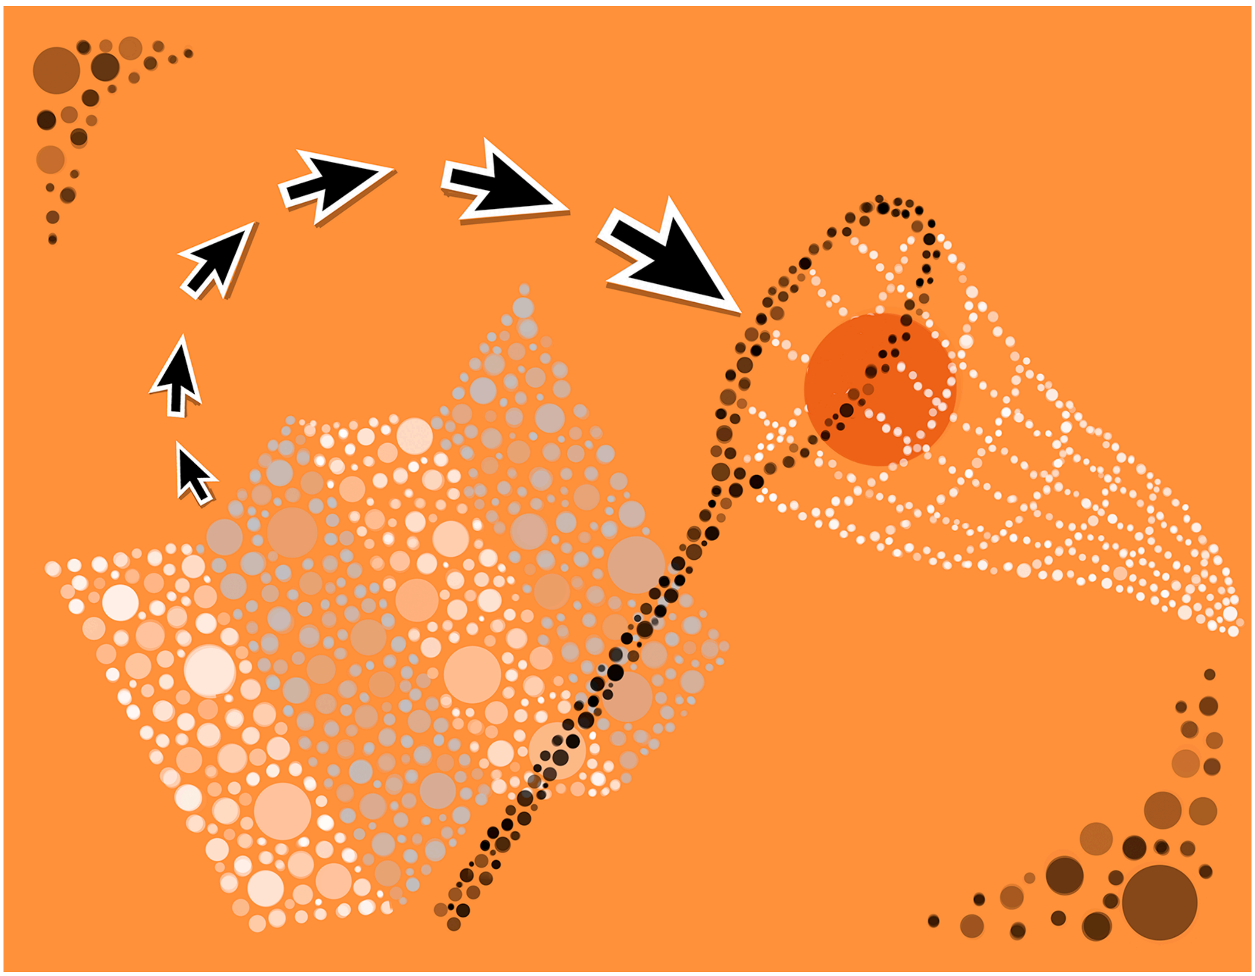

Illustration by Alexandra Mateescu

We are grateful for the insights and assistance of danah boyd, Seda Guerses, Sorelle Friedler, Surya Mattu, Karen Levy, Tom Igoe, Nick Grossman, Yan Shvartzshnaider, Alexandra Mateescu, and Angie Waller.

This project was generously supported by a grant from the John D. and Catherine T. MacArthur Foundation.

Data&Society

Part I
The Wisdom of the Captured

The "wisdom of the crowd" is a well-known conceptual description that is frequently applied to the way large-scale network services work. Projects like Wikipedia or the development of the Linux operating system rely on an updated version of the notion that "many hands make light work"—with thousands or millions of individual contributors, complicated problems can be solved with surprising speed and accuracy. Such successes have led to a general cultural confidence in large-scale collaboration networks, and so when platforms like Amazon or Netflix make recommendations based on the notion of a larger network – *users like you also purchased* – it can channel the positive associations of deriving value from just such a wise crowd. But not all large-scale networks are organized in the same way, and collecting data from a large number of users is not inherently collaborative or egalitarian.

Machine learning systems rely on algorithms that evolve in response to the data that they analyze. "Deep learning", a branch of machine learning which leverages neural network architectures, has seen many recent high-profile advances which are enabling computers to accomplish tasks – identifying objects in photographs, defeating Go masters – previously thought to be near impossible. While these accomplishments are undoubtedly impressive, machine learning is more expansive than deep learning alone: the field encompasses a vast set of different techniques that have been applied to many different problems.

Despite this variety, it is important to recognize that all machine learning systems share a common need for data. Machine learning systems form initial processes from training data, and then take in more and more data as they run, refining their models further and further. For machine learning software that identifies the subject of photos, this means inputting more and more photos. For systems that analyze language—feed it as much text as you can find.

When applied to user-facing network services, machine learning's inherent reliance on data makes the data generated by individual users of particular value. By combining machine intelligence with the collective intelligence of the crowd or "embedding crowds inside of machine learning architectures (Cheng & Bernstein, 2015, p. 1)," platforms can leverage "the strengths of both crowds and machines (Ibid., p. 2)." For services with thousands or millions of active, networked users, those users' own behavior, as captured by the system, is likely to simply be the best source of data on the system's efficacy. Capturing the data of an entire service's users can produce a type of valuable wisdom, but not from a crowd of collaborators. This is the *wisdom of the captured*.

In theory, the wisdom of the captured *does* allow individual users to benefit from data they contribute to an aggregate by receiving personalized services. Simply by using Google search, the argument goes, you are improving future Google searches for you and everyone else. But what happens when the techniques of system-improvement conflict with the delivery of services to individual users? Under what circumstances might a platform deliver sub-optimal services to individual users in order to serve some larger, system-wide goal of optimization? How might machine learning's reliance on data produce such a gap?

More broadly, how might the power dynamics of user and platform interact with the marketing surrounding these technologies to produce outcomes which are perceived as deceptive or unfair? This provocation paper assembles a set of questions on the capacity for machine learning practices to create undisclosed violations of the expectations of users – expectations often created by the platform itself -- when applied to public-facing network services. It draws on examples from consumer-facing services, namely GPS navigation services like Google Maps or Waze, and on the experiences of Uber drivers[1], in an employment context, to explore user assumptions about personalization in crowd-sourced, networked services.

1. The illustration of Uber drivers draws partly on Rosenblat's previous research with co-author, Luke Stark, in "Algorithmic Labor and Information Asymmetries: A Case Study of Uber's Drivers."

Part II
Expectations: Neutrality, Personalization, Fairness

The networked services of the contemporary internet have long been marketed in terms of personalization. Relying on responsive databases and adaptable algorithms, companies have promised users their own personalized radio stations, newspapers, television channels, love interests, and book recommendations. The two contemporary network services that we consider in this paper, GPS navigation services and Uber, base their operation on the simultaneous provision of service *to* and the collection of data *from* their users. Other services, which similarly offer users personalized services in relation to "the crowd", include Google's search services, Facebook's social media, Amazon's commercial storefront, Netflix's video streaming, Pandora's music streaming, and Uber and Lyft's ridehailing.

The cultural reputation of data-driven logic in algorithms and machine learning is that the data contains an objective truth; when companies offer users a service or recommendation based on "the data," it can elicit significant trust from users. A number of scholars and writers have critiqued the growing and complicated role of algorithms in such services and, therefore, everyday life. Functionally, algorithms can enable new business practices, removed from traditional human oversight. Rhetorically, these scholars have observed that the aura of algorithms can appeal to imagined objectivity, and claim a powerful sense of authority.

When a platform or company makes recommendations that diverge from user expectations of neutral and accurate or high-confidence results, it can foment user distrust in data-driven explanations or in the companies themselves. Yet, to improve the models of machine learning, which are gaining increasing prominence in data-centric technology services, users may be leveraged to provide streams of data in ways that they don't expect.

Both Google and Facebook, two of the most popular user-facing, networked technology companies in the U.S., have gone to great lengths to brand their algorithmic process of information-organization as objective, and their role as platforms as neutral arbiters of information. Search neutrality is a well-established principle of the algorithmic organization of search query results on Google's platform, implying a certain even-handedness in search results that does not favor particular content out of bias.

Facebook brands its platform as a place that optimizes for individuals to share and connect with everyone on a global scale. The News Feed is marketed as a personalized assortment of news that is "influenced by your connections and activities on Facebook." The News Feed is curated through a combination of algorithmic process and human signals, but many users are unaware of precisely how these two processes are combined, or that the process of curation is not inherently neutral or objective. That this lack of knowledge was widespread became evident when researchers at Facebook conducted a study by filtering the information users were shown to test whether their emotions would be influenced by the valence (such as sad or happy) of the news they were shown—essentially, using the News Feed for the purpose of an information-gathering experiment instead of surfacing the "best" or high-confidence results that users expected from the platform. News of the study prompted huge public outcry around the issue of manipulation and deception. One of the reasons for public objection was the absence of consent in a study of their emotions, and even the fact that they would never know if they were part of the study, an outcry which points to how trust in companies affects people's perceptions of their actions. Some research on the existence and importance of algorithmic cognition suggests improving users' "algorithmic awareness" may elicit user approval even if their initial reaction is negative.

While this conundrum is partly related to the complications of ethics in big data social science research, it also demonstrates the extent to which companies have established public expectations surrounding the neutrality of their platforms by obscuring their curatorial role. These experiments were deceptive insofar as the platform quietly violated established expectations of neutrality and the trust in the platform that came with this understanding, but at the same time, they drew on common practices of experimentation that are used to test and provide services.

As more services become based on machine learning, it might no longer be the occasional, human-driven experiment that violates users' expectations, but constantly updating, everyday practices directed by machines. While all software is algorithmic, there has been a recent surge in interest in a specific family of programming techniques known as machine learning. This leads us to ask: does machine learning constitute any *inherent* contradictions between individual service and system-wide efficacy? Does machine learning and the platforms they operate within present incentives to engage in activities which may be perceive as deceptive by users?

Moreover, these deceptions may not be the result of intentional, malicious behavior on the part of systems designers, but instead an emergent outcome of the optimization needs of the algorithm and the affordances of control over a user that a platform possesses in "wisdom of the captured" style architectures.

Part III
An Illustration from GPS Navigation

Data-Driven Consumer Services

One generative context in which to think through how these tensions might arise can be found in the domain of route selection (Azaria, 2014, p. 62).[2] GPS navigation services like Waze and Google Maps use sensors in smartphones to generate route recommendations based on criteria both explicit (avoid freeways, avoid tolls) and implicit (fastest or shortest).[3] Personalization is a key part of the mainstream understanding of these technologies. One recent New York Times story on GPS navigation systems is illustrative. The article discusses how these apps integrate data from multiple sources to help users avoid obstacles, like traffic, emphasizing the crowd-sourced nature of the system. At the same time, the story sums up the upshot of this architecture as simply, "Personalized Traffic Alerts from Google."

To that end, users implicitly or explicitly accept that the results generated for them involves a ranking order of those recommendations by certain criteria. They generally assume they are getting robust, high-confidence recommendations from the system, and they may feel deceived if they are instead given a low-confidence recommendation without being notified as such. GPS navigation relies on real-time data of road conditions, and those drivers already using the service to navigate are primary sources of such data. As a navigation system receives data of a slow down on the typically fastest route, it can give new users alternate directions that take advantage of this up-to-date awareness of traffic patterns. And so, as an example of an adaptive machine learning model, GPS navigation systems are able to give the best routes to the most users when the system has comprehensive data on traffic conditions.

There are different strategies and technologies for building navigation models. One such technical tool is the "multi-arm bandit algorithm" and is applied to route traffic. Hypothetically about 90% of the time, traffic will be routed on the best performing route, but the remaining 10% of the time, the algorithm might divide traffic between two versions of a route from A to B to *explore* how well they perform (Chopra, 2012). For that exploration phase, the mapping service might recommend a route to an individual driver that is under-tested, so that that the driver will unwittingly test it out and generate more data about road conditions. This may help the aggregate of users who rely on the mapping services (and therefore the perceived efficacy of the service), but if the exploring driver is given a bad route when they expect a high-confidence recommendation, that driver is effectively being deceived.

It is worth recognizing that this deception is made possible by the fact that the platform wields considerable power over the decision-making of a given driver in a "wisdom of the captured" setting. Without explicit notice, users will have no way of distinguishing between a high-confidence recommendation and a low-confidence recommendation. This is particularly challenging given that full disclosure of the exploratory intent of the system may make users behave differently or be less likely to follow its purportedly personalized route recommendations. This would parallel studies suggesting that disclosure of persuasive intent can itself erode the influence of recommendations in eliciting compliance from users (Kaptein, et. al, 2011). When the platform makes a low-confidence recommendation in order to acquire more information (exploration), there is a trade-off that produces a social welfare benefit for the users as a whole, but has ethical implications for the deception of the individual. The optimization needs of the algorithm and the powerful position of the platform in relation

2. This example forms the basis for a parallel analysis and discussion in "Exploring or Exploiting? Social and Ethical Implications of Autonomous Experimentation in AI" by Sarah Bird, Solon Barocas, Kate Crawford, Fernando Diaz, and Hanna Wallach (October 2, 2016). Workshop on Fairness, Accountability, and Transparency in Machine Learning. Available at: http://ssrn.com/abstract=2846909
3. Not all mapping services emphasize personalized recommendations. Open Street Map, for instance, is a clear example of a mapping system that crowd-sources information, but which does not emphasize recommendations.

to a driver create the opportunity for user expectations to be violated in an undisclosed and systematic way.

One of the challenges for using machine intelligence to automate human persuasion is building a system that distributes recommendations ethically (Stock et. al, 2016, p. 1; Allen et. al, 2006). The burdens of an inefficient route, for example, may also not be distributed evenly across all the drivers in the system—a platform may have an interest in disproportionately assigning inefficient, information-seeking routes to drivers who are less valuable to it. This might include imposing these costs on drivers who are the most loyal or dependent on the platform. This prompts large questions about when algorithmic services might produce inequality, and along what lines that inequality might be drawn. For instance, if 100 drivers all want to travel the same segment at the same time, a navigation platform can divide users across multiple routes to minimize time spent for the majority by giving everyone either a 20- or 25-minute trip; alternatively, it could prioritize a minority of users by directing the majority to a different route so that a few have a 15-minute trip while the rest have a 40-minute one.

Regardless of how these determinations are made (and whether or not the public will have access to them), the point remains that machine learning systems require captured data in order to function as intended. These systems cannot iterate if they have no data on which to base their iterations. GPS navigation dramatizes this, as physical space and extra minutes spent in a car make clear the kinds of trade-offs that can result, but machine learning is being applied to more than the navigation of streets. What remains to be seen is how machine learning systems for other processes might also rely on providing user services that deviate from expectations in order to provide a system-wide optimization of some value.

Part IV
Illustrations from the Ridehail Business

Data-Driven Employment Services

The relationship between personalization and machine learning is of particular interest in the case of ridehail platforms like Uber and Lyft. These platforms coordinate two different populations of users – drivers and riders – through use of a networked smartphone app that does everything from perform GPS navigation, dynamically set rates for trips, process payments, and track driver and rider ratings. Such ridehail businesses both rely on a narrative of individual user independence/personalization, as well as a business model that leverages captured data. Considering that Uber already operates under a "wisdom of the captured" style model - alongside the future opportunities it has to introduce additional machine learning and artificial intelligence techniques - it illustrates the potential for users to end up deceived as to the nature of certain network services.

Uber's business model and initial growth has made heavy use of a narrative of providing "turnkey entrepreneurship" to it drivers (Uber Newsroom, May 27, 2014), with communications like "We're always working to make Uber the best platform for partners to build a small business" (Uber Newsroom, Nov. 19, 2014). Such drivers were enticed through promises of being their own boss, or being able to work as much as they want, when and where they want. At the same time, many drivers have expressed dissatisfaction with some components of this narrative of mass entrepreneurship for individuals, particularly when their compensation structure is negatively impacted by Uber's data-driven logic for these changes (Rosenblat & Stark, p. 3764). The role of algorithms, data, and machine learning have all played a role in how drivers perceive a gap between how Uber's platform is branded for drivers (for individual optimization) and driver experiences of policy changes that hamper their earnings.

A perennial conflict between Uber and its drivers has been around the data-driven logic that Uber deploys to set and lower rates at which drivers earn their income, which Uber is unilaterally empowered to do. Uber has made several significant cuts to the base rates for trips, which inevitably produces an outcry from drivers. Uber consistently defends these cuts, however, through recourse to the data they collect through the system (data to which, crucially, individual drivers do not have access). Citing findings from data, Uber forwards slogans such as "lower rates = higher earnings." Uber has highlighted the success of various price cuts to incredulous drivers by sending them messages like, "We reduced prices in New York City last Friday and we are seeing positive results for driver-partners after only three days." Drivers have characterized such messaging as "Uber math", arguing that they have to drive longer hours, absorb additional costs, and put more wear and tear on their vehicles to earn what they made prior to rate cuts (Rosenblat & Stark, 2016, p. 3764). A BuzzFeed investigation recently validated many of their complaints when it surfaced leaked documents from Uber, indicating that the relationship between price cuts and increased driver pay through higher optimization of drivers' time is overstated and misleading.

The assumption Uber relies on when it says it "has the data" or "the numbers" is that the data contains an objective, mathematical truth, much like Facebook and Google make claims about the neutrality of their information curation. Uber asserts similar associations regarding its practice of surge pricing. While the base rates are established as a matter of contract, and drivers receive notice when changes occur, a prominent feature of drivers' compensation structure comes from "surge pricing," which produces much faster changes to the rates. The company explains that surge pricing is derived from an algorithmic estimate that demand outstrips supply by a particular metric, and base rates are magnified by a subsequent multiplier, such as 3.5x, which applies variably in a given region. For example, one area in a given neighborhood could be surging at 3.5x while an adjacent area is surging at 4.5, or not at all – the boundaries of these areas are not disclosed by the company, but it does emphasize that surge is a dynamic process – essentially a high-frequency implementation of variable rate contracting between Uber and its drivers. Drivers are unsurprisingly motivated by such inflated rates, and will sometimes relocate, but more experienced drivers caution newer drivers, in online forums, "don't chase the surge" because it's unreliable (Rosenblat & Stark, p. 3766).[4] A representative of Uber claimed that "surge pricing only kicks in in order to maximize the number of trips that happen and therefore reduce the number of people that are stranded," in 2013, but such a statement isn't exactly a precise description of the software running surge pricing decisions.

The position of platform relative to the drivers creates the opportunity and the incentive to engage in deceptive acts that contradict the public expectations Uber has established. Ultimately, there is no *real* transparency as to the means used to determine surge pricing. There is a general sense that, like driver earning numbers, surge pricing is algorithmic and data-sensitive: if the system receives a certain number of requests in an area with sufficiently few available drivers, then surge pricing automatically goes into effect. And yet, drivers are not advised whether predictive demand, which comes from historical data, is made with equal or varying confidence from real-time demand. As Rosenblat & Stark observe (p. 3768),

> The language Uber uses to describe surge pricing is often identical to the language it uses to describe predicted demand: Rhetorically, essentially predictive "guesses" about possible future demand are thus easily confused with real-time "measurements" of existing present demand. This rhetorical device is used by Uber to mobilize its workforce in a way that draws on drivers' experiences of surge pricing in real-time—with the implication that real-time measurements are made with a high degree of accuracy— although the company does not indicate whether a real-time recommendation to go to a surge zone is as accurate as predicted surge (or "high demand"), or if it is a lower-confidence recommendation.

In effect, Uber is making use of a familiar recommendation tool (nudge messaging) and corporate algorithmic practice to leverage control over how drivers choose to take a particular course of action, without alerting them to the reliability of the nudge. The intersection between data collection and data-driven recommnedations to users who are at an informational disadvantage from a centralized platform creates an (ongoing) opportunity for platform experimentation with worker-users in employment platforms just as they do in consumer-users of Facebook, for example.

As a point of future speculation, we might ask whether certain instances of surge pricing exist not to relieve rider demand, or maximize driver profits, but as a means for a machine learning algorithm to cajole its human data sensors into gathering exploratory data. This hypothetical might gain additional salience from Uber's intentions to build its own mobile maps product.

The role of Uber drivers as sensors for a total system is of growing importance as Uber begins, increasingly, to make plans for the "Uber future" of autonomous cars. Uber is already a member of a larger consortium working on the reality of such "self-driving" cars, and the data they've collected in their role of coordinating crowd-sourced ridehail

4. As Rosenblat & Stark (2016, p. 3766) observe, "Drivers also noted that they would sometimes converge en masse at a surging area, find that supply was no longer too low, and the surge would disappear. Some drivers reported experimenting with trying to game these algorithms themselves, and many developed responses to surge pricing based on their experience with its duration, reliability, and potential reward in their respective locations. It is unclear whether surge is designed equally to optimize for satisfying passenger demand or for increasing driver earnings, but Uber's stance against "surge manipulation" by drivers suggests the former."

platform is no doubt essential to their potential entry into the autonomous market. The operation of autonomous cars will necessarily inflate the importance of analyzing the function of Uber's algorithmic infrastructure: the ethics of autonomous car behavior will be programmed, and will integrate certain values and trade-offs; and public roads will essentially come under the purview of the owners of autonomous car networks, raising questions of equity and inequity. In this way, questions of deception and optimization for a particular group of users – Uber drivers- anticipates larger societal trade-offs for the governance of networked systems that become infrastructural. Perhaps even more telling is that Uber's commitment to self-driving cars, enabled in part by the data gathered by their drivers, is arguably the clearest articulation yet that Uber will make choices that benefit the system over individual drivers. Self-driving cars would directly compete with and impact the human drivers of Uber's system, effectively automating them out of a job. In essence, the technology that augments drivers as workers also puts drivers to work training the machines that will replace them.

Part V
Research Questions

This paper is meant as the provocative beginning to a conversation about deception that will only become more important as machine learning becomes more prevalent. These questions are limited by the structures of "trade secrets" and proprietary "secret sauce." There is often no practical way for critics or everyday users to gain access to the inner workings of the systems with which they interact on a daily basis. And as algorithms grow in complexity and distance from human-designed processes, it can be harder and harder to unpack how a particular system makes decisions—even with the input of the original designers. In addition, these network services are moving targets. The software that delivers Google results or Uber fares might not be the same as it was last year—it might change tomorrow, later today. There is always the capacity for familiar services to be suddenly accomplished by unknown and unfamiliar back-ends.

This is why this paper wants to start asking now what the consequences of moving toward machine learning models might be for the relationship between users and these network services. Early indications suggest that machine learning is likely to orient toward users in ways that contradict established expectations about personalization and the structure of user services. We end with three broad questions:

1. What are the conditions that would make it likely for machine learning systems to violate user expectations?

The GPS example describes "exploration" periods, where users are given data other than what they expect, for the benefit of the whole system. What other services might be similarly served by "exploration?" What distributions of power between platforms and users make these violations particularly feasible or attractive?

2. When are the decisions made by machine learning algorithms unfair or inappropriate?

Public outcry to the "experiments" of Facebook's News Feed indicate a broad public sense of propriety in algorithmic manipulation, but how do we (as researchers and society, both) divide the line between experimental and non-experimental behavior? If a system is always under recursive revision, how do we determine which methods of revision fulfill user expectations/norms, and which violate them? What drives designers of machine learning algorithms to build and implement systems in ways that commit these violations?

3. What solutions are possible to address the inequity that can result from decisions made based on machine learning models?

A well-established drive toward transparency and legibility seems an important first step, but simply being aware of a decision-making process does not account for its full impact. How platforms communicate what they do is important for maintaining good platform-user relations, but that requires an evolving interplay of understanding both from platforms, and from users. And while the legibility of deception in technology is an important factor in how we understand, accept, or reject it, the stakes can vary wildly. Reading an alternate headline for a news article (as part of A/B testing) is a far different experience than asking individual users to assume the financial costs and risks of algorithmic recommendations in an employment context. How should users and platforms work together in constructing a shared understanding of how these systems work?

References[5]

ABC7. 2016. "Uber Drivers Protest Fare Cuts at Company's Headquarters, Considering Strike." *Eyewitness News*, February 1, 2016. http://abc7ny.com/traffic/uber-drivers-protest-fare-cuts-at-companys-headquarters-considering-strike/1181656/

Adar, Eytan. 2013. "Benevolent Deception in Human Computer Interaction." *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI 2013)*. https://www.microsoft.com/en-us/research/publication/benevolent-deception-in-human-computer-interaction/

Allen, Colin, Wendell Wallach, and Iva Smit. 2006. "Why machine ethics?" *Intelligent Systems*, IEEE 21(4):12–17.

Angwin, Julia. 2016. "Make Algorithms Accountable." *The New York Times*, August 1, 2016. http://www.nytimes.com/2016/08/01/opinion/make-algorithms-accountable.html?_r=0

Azaria, Amos. 2014. "Agents for Automated Human Persuasion." Ph.D. Thesis, Bar-Ilan University, September, 2014. http://azariaa.com/Content/thesis.pdf

Barocas, Solon, Sophie Hood, and Malte Ziewitz. 2013. "Governing Algorithms: A Provocation Piece." Accessed March 30, 2016. http://papers.ssrn.com/sol3/Papers.cfm?abstract_id=2245322

Barocas, Solon, Sarah Bird, Kate Crawford, Fernando Diaz, and Hannah Wallach. 2016. "Ask Not What Your Algorithm Can Do For You: The Ethics of Autonomous Experimentation." *Microsoft Research*. Working Paper sent by email correspondence, September 17, 2016.

Bates, Daniel. 2016. "Uber Cuts Prices—and Kneecaps Drivers." *The Daily Beast*, February 1, 2016. http://www.thedailybeast.com/articles/2016/02/01/uber-cuts-prices-and-kneecaps-drivers.html

Biersdorfer, J. D. 2016. "Personalized Traffic Alerts From Google." *The New York Times*, August, 9, 2016. http://www.nytimes.com/2016/08/10/technology/personaltech/personalized-traffic-alerts-from-google.html?_r=0

boyd, danah. 2016. "Facebook Must Be Accountable to the Public." Accessed May 13, 2016. https://points.datasociety.net/facebook-must-be-accountable-to-the-public-72a6d1b0d32f#.akn6srk9u

Brewster, Signe. 2015. "How Google's New Photos App Can Tell Cats From Dogs." *Backchannel*, June 3, 2015. https://backchannel.com/how-google-s-new-photos-app-can-tell-cats-from-dogs-ffd651dfcd80#.p57bfsjkf

Brustein, Joshua. 2013. "Uber Explains Why $35 a Mile is the Right Price." *Bloomberg*, December 17, 2013. http://www.bloomberg.com/news/articles/2013-12-17/ubers-travis-kalanick-explains-the-car-services-surge-pricing

[5] Most references in this informal piece are hyperlinked, but when specific examples or quotations from an author's work are used, they are referenced in-text.

Burrell, Jenna. 2016. "How the machine 'thinks': Understanding opacity in machine learning algorithms." *Big Data & Society*. http://bds.sagepub.com/content/3/1/2053951715622512

Caplan, Robyn and danah boyd. 2016. "Who Controls the Public Sphere in the Era of Algorithms?" *Algorithms and Publics.* Accessed September 3, 2016. http://datasociety.net/pubs/ap/MediationAutomationPower_2016.pdf

Chambers, Chris. 2014. "Facebook fiasco: was Cornell's study of 'emotional contagion' an ethics breach?" *The Guardian*, July 1, 2014. https://www.theguardian.com/science/head-quarters/2014/jul/01/facebook-cornell-study-emotional-contagion-ethics-breach

Cheng, Justin and Michael S. Bernstein. 2015. "Flock: Hybrid Crowd-Machine Learning Classifiers." *CSCW 2015*, March 14-18, 2015, Vancouver, BC, Canada. http://hci.stanford.edu/publications/2015/Flock/flock_paper.pdf

Chopra, Paras. 2012. "Why Multi-armed Bandit Algorithm is Not 'Better' than A/B Testing." *VWO Blog*, June 1, 2012. https://vwo.com/blog/multi-armed-bandit-algorithm/

Edelman, Benjamin and Benjamin Lockwood. 2011. "Measuring Bias in 'Organic' Web Search." Accessed March 30, 2016. http://www.benedelman.org/searchbias/

Epstein, Robert. 2015. "How Google Could Rig the 2016 Election." *Politico Magazine*, August 19, 2015. http://www.politico.com/magazine/story/2015/08/how-google-could-rig-the-2016-election-121548

Facebook. n.d. "How News Feed Works." Accessed March 30, 2016. https://www.facebook.com/help/327131014036297/

Facebook. n.d. "Is Connectivity a Human Right?" Accessed March 30, 2016. https://fbcdn-dragon-a.akamaihd.net/hphotos-ak-xat1/t39.2365-6/12057105_1001874746531417_622371037_n.pdf

Gillespie, Tarleton. 2014. "The Relevance of Algorithms." In *Media Technologies*, ed. Tarleton Gillespie, Pablo Boczkowski, and Kirsten Foot. Cambridge, MA: MIT Press. http://www.tarletongillespie.org/essays/Gillespie%20-%20The%20Relevance%20of%20Algorithms.pdf

Grimmelmann, James. 2010. "Some Skepticism About Search Neutrality." Accessed March 30, 2016. http://james.grimmelmann.net/essays/SearchNeutrality

Hamilton, Kevin, Christian Sandvig, Karrie Karahalios, and Motahhare Eslami. 2014. "A Path to Understanding the Effects of Algorithm Awareness." *CHI 2014*, April 26-May 1, 2014, Toronto, Ontario, Canada. http://social.cs.uiuc.edu/papers/pdfs/paper188.pdf

Kaptein, Maurits, Steven Duplinsky, and Panos Markopoulos. 2011. "Means Based Adaptive Persuasive Systems." *CHI 2011*, May 7-12, 2011, Vancouver, BC, Canada. http://www.persuasion-profiling.com/wp-content/uploads/2010/04/Kaptein_MeansBased.pdf

Koch, Christof. 2016. "How the Computer Beat the Go Master." *Scientific American*, March 19, 2016. http://www.scientificamerican.com/article/how-the-computer-beat-the-go-master/

Koene, Ansgar, Elvira Perez Vallejos, Christopher J. Carter, Ramona Statache, Svenja Adolphs, Claire O'Malley, Tom Rodden, and Derek McAuley. 2011. "Ethics of Personalized Information Filtering." HORIZON Digital Economy Research, University of Nottingham, UK. http://casma.wp.horizon.ac.uk/wp-content/uploads/2015/04/ICISSGI15_EthicsOfPersonalizedInformationFilters_AKoeneEtAl_final_draft.pdf

Kramer, Adam D. I., Jamie E. Guillory, and Jeffrey T. Hancock. 2014. "Experimental Evidence of Massive-scale Emotional Contagion Through Social Networks." *Proceedings of the National Academy of Sciences of the United States of America*, vol. 111 no. 24, June 17, 2014. http://www.pnas.org/content/111/24/8788.full

Levy, Steven. 2016. "How Google is Remaking Itself as a 'Machine Learning First' Company." *Backchannel*, June 22, 2016. https://backchannel.com/how-google-is-remaking-itself-as-a-machine-learning-first-company-ada63defcb70#.2z7fnzfly

Mansour, Yishay, Vasilis Syrgkanis, and Aleksandrs Slivkins. 2015. "Bayesian Incentive-Compatible Bandit Exploration." Microsoft Research. http://128.84.21.199/pdf/1502.04147v2.pdf

Matthews, Dylan. 2014. "Facebook tried to manipulate users' emotions. But we have no idea if it succeeded." *Vox*, June 30, 2014. http://www.vox.com/2014/6/30/5856938/the-facebook-study-wasnt-just-creepy-it-was-bad-research

Morozov, Evgeny. 2011. "Don't Be Evil." *New Republic*, July 12, 2011. https://newrepublic.com/article/91916/google-schmidt-obama-gates-technocrats

Nosowitz, Dan. 2016. "Could Facebook Swing the Elecction." *NY Magazine*, April 27, 2016. http://nymag.com/selectall/2016/04/could-facebook-swing-the-election.html

O'Brien, Sara Ashley. 2016. "NYC Uber drivers protest rate cuts." *CNN Money*, February 1, 2016. http://money.cnn.com/2016/02/01/technology/uber-nyc-protest/

O'Donovan, Caroline and Jeremy Singer-Vine. 2016. "Here's What Uber Doesn't Say About Price Cuts." *BuzzFeedNews*, June 24, 2016. https://www.buzzfeed.com/carolineodonovan/uber-documents-suggest-price-cuts-dont-always-raise-driver-w?utm_term=.pqoxKdG58#.onrND2qy1

Raymond, Eric S. 2010. "The Cathedral and the Bazaar." Accessed March 30, 2016. http://www.catb.org/esr/writings/cathedral-bazaar/

Reagle, Joseph Michael Jr. 2010. *Good Faith Collaboration*. Cambridge, Massachusetts: MIT Press. https://mitpress.mit.edu/sites/default/files/titles/free_download/9780262518208_Good_Faith_Colaboration.pdf

Rosenblat, Alex. 2015. "Uber's Phantom Cars." *Motherboard*, July 27, 2015. https://motherboard.vice.com/read/ubers-phantom-cabs

Rosenblat, A., & Stark, L. (2016). Algorithmic Labor and Information Asymmetries: A Case Study of Uber's Drivers. *International Journal Of Communication, 10*, 27. Retrieved from http://ijoc.org/index.php/ijoc/article/view/4892/1739

Shepardson, David. 2016. "Google, Ford, Uber launch coalition to further self-driving cars." *Reuters*, April 26, 2016. http://www.reuters.com/article/us-autos-selfdriving-idUSKCN0XN1F1

Terrell, Rebecca. 2015. "Google's New Algorithm: Modern Version of Old-fashioned Censorship?" *The New American*, March 16, 2015. http://www.thenewamerican.com/tech/computers/item/20408-google-s-new-algorithm-modern-version-of-old-fashioned-censorship

Tufekci, Zeynep. 2015. "Algorithmic Harms Beyond Facebook and Google: Emergent Challenges of Computational Agency." Accessed March 30, 2016. http://ctlj.colorado.edu/wp-content/uploads/2015/08/Tufekci-final.pdf

Uber Newsroom. 2014a. “An Uber Impact: 20,000 Jobs Created on the Uber Platform Every Month.” https://newsroom.uber.com/an-uber-impact-20000-jobs-created-on-the-uber-platform-every-month-2/

Uber Newsroom. 2014b. “Introducing Momentum Partner Rewards.” https://newsroom.uber.com/introducing-momentum-partner-rewards/

Uber Newsroom. 2016a. “Mapping Uber’s Future.” https://newsroom.uber.com/mapping-ubers-future/

Uber Newsroom. 2016b. “Steel City’s New Wheels.” https://newsroom.uber.com/us-pennsylvania/new-wheels/

Vasserman, Shoshana, Michal Feldman, and Avinatan Hassidim. 2015. “Implementing the wisdom of waze.” *Proceedings of IJCAI’15*. https://dl.acm.org/citation.cfm?id=2832341

Weld, Daniel S., Mausam, Christopher H. Lin, and Jonathan Bragg. 2014. “Artificial Intelligence and Collective Intelligence.” https://homes.cs.washington.edu/~weld/papers/ci-chapter2014.pdf

Wikipedia.org. n.d. “Search Neutrality.” Accessed March 30, 2016. https://en.wikipedia.org/wiki/Search_neutrality

Wohlsen, Marcus. 2013. “Uber Boss Says Surging Prices Rescue People from the Snow.” *WIRED*, December 17, 2013. http://www.wired.com/2013/12/uber-surge-pricing/

Zimmer, Michael. n.d. “Research Ethics in the Big Data Era: Addressing Conceptual Gaps for Researchers and IRBs.” Accessed March 30, 2016. https://bigdata.fpf.org/papers/research-ethics-in-the-big-data-era-addressing-conceptual-gaps-for-researchers-and-irbs/