

The High Cost of Free Services: Problems with Surveillance Capitalism and Possible Alternatives for IT Infrastructure

Marvin Landwehr
University of Siegen
Siegen, Germany
marvin.landwehr@student.uni-siegen.de

Alan Borning
University of Washington
Seattle, Washington, USA
borning@cs.washington.edu

Volker Wulf
University of Siegen
Siegen, Germany
volker.wulf@uni-siegen.de

ABSTRACT

A large portion of the software side of our information technology infrastructure, including web search, email, social media, transportation information, and much more, is provided “free” to the end users, although the corporations that provide this are often enormously profitable. The business model involves customized advertising and behavior manipulation, powered by intensive gathering and cross-correlation of personal information. Significant other parts of our IT infrastructure use fees-for-service but still involve intensive information gathering and behavior manipulation. There are significant indirect costs of these business models, including loss of privacy, supporting surveillance by both corporations and the state, automated manipulations of behavior, undermining the democratic process, and consumerism with its attendant environmental costs. In a recent book, Shoshana Zuboff terms this “surveillance capitalism.” Our primary focus in this essay is how we could develop new models for providing these services. We describe some intermediate steps toward those models: education, regulation, and resistance. Following that, we discuss a partial solution, involving for-profit corporations that provide these services without tracking personal information. Finally, we describe desired characteristics for more comprehensive solutions, and outline a range of such solutions for different portions of the IT infrastructure that more truly return control to the end users. A common feature of several is the use of highly decentralized storage of information (either on the end user’s own personal devices or on small servers), a modular architecture and interface to allow for customization of what information is to be shared, and a distributed ledger mechanism for authentication.

CCS CONCEPTS

• **Social and professional topics** → **Computing / technology policy**; *Surveillance*; *Sustainability*; *Governmental regulations*.

KEYWORDS

surveillance capitalism, political manipulation, manipulation of behavior, economics, digital infrastructure, IT business models, advertising

ACM Reference Format:

Marvin Landwehr, Alan Borning, and Volker Wulf. 2019. The High Cost of Free Services: Problems with Surveillance Capitalism and Possible Alternatives for IT Infrastructure. In *Proceedings of Fifth Workshop on Computing within Limits (LIMITS’19)*. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3338103.3338106>

1 SURVEILLANCE CAPITALISM

In the aftermath of the dot-com-bubble, the surviving IT companies had to find business models that would continue providing a growing return to their venture capital funders. In order to keep increasing the user base, it was essential for free services to continue to be free of charge for the end users. Therefore it comes as no surprise that they chose to embed advertisements to generate revenue. However, the way in which they implemented this turned out to make them highly profitable at the expense of a devastating societal spillover. Instead of merely placing ads, user actions were tracked and recorded in large data bases. When computer scientists at these companies invented algorithms to effectively analyze the giant volumes of data and employed statistical methods to derive predictions from user profiles, they invented a new business model that turned out to be so profitable that even companies that are financed by fees also adopted this practice.

Based on this data, they generate revenue primarily by selling customized advertising. User data becomes much more valuable for this a) in contexts of past and planned purchases, b) when data corresponding to the same user can be connected across different contexts to generate a significantly more sophisticated user profile, c) profile data can be compared and correlated to similar profiles from the database, and d) it serves as a basis for behavioral manipulation. Since it is their predictive and manipulative capital [21], usually the data is not sold directly, as that would undermine the basis of their business model. Although smaller companies exist that employ trackers and presumably serve as additional data suppliers, these transfers of user data are typically not visible to the end user. For the big IT companies instead revenue is generated by predicting and influencing user behavior [6].

Shoshana Zuboff has named this business model “surveillance capitalism” [39, 40], and we use this same term here as well. Her recent book makes clear that its downsides are not simply issues of surveillance and loss of privacy, but also radical and ominous automated manipulation of behavior.

However, despite this very dark side of surveillance capitalism, at the same time these services have enormous utility for business, social engagement, political work, and much more. So in any potential measures to address these problems, we want to retain as much as possible the benefits.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

LIMITS’19, June 10–11, 2019, Lappeenranta, Finland

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-7281-7/19/06...\$15.00

<https://doi.org/10.1145/3338103.3338106>

To relate this analysis to work presented at LIMITS 2018 [4] and elsewhere [3], the SEED¹ project is based on the observation that quality of life for growing numbers of people on the planet is threatened by a set of integrated, systemic problems in the environment and our economic and political systems. Its goal is to form an international network of scholar/activists, advocates, and practitioners who seek to address these problems in a similarly integrated fashion. The work presented here is a step in working out how the SEED ideas play out in a particular economic sector, namely IT. (We can imagine similar efforts in the future for other sectors.)

In describing SEED, we put forth an ideal of the relationship among the environment, society, government, and the economy, firmly grounded in the natural world. Within and bounded by the natural world is human society, with a goal of prosperity for all and supporting human flourishing. Government in turn is subservient to society, and finally the economy is subservient to all three other systems. This vision is in sharp contrast to our current system, in which the economy takes priority. A major added problem of surveillance capitalism is thus: in addition to making the natural world subservient to and serving the needs of the economy, now human nature is being similarly dominated.

In the following section we discuss the processes by which surveillance capitalism impacts consumerism, threatens democracy, fuels social fragmentation, undermines our ability to tackle major environmental problems, and thereby ultimately constitutes a hurdle to living within planetary ecological limits. While a number of leaders in those IT companies were well aware of a range of negative effects on society [1, 25], in the past these effects have been underrecognized by major parts of society.

2 PROBLEMS

This business model produces a range of problems. In order to maximize the value for advertisers the companies need to capture the user for as long as possible on the website or using the application. This is achieved by applying knowledge of human psychology [22] and experimenting with different interfaces using A/B and other testing. (Another more publicized example of manipulation is the Facebook emotional manipulation study [13].)

One strategy to maximize the time spent on social media is to push toward more provocative content, both on the side of the viewer (who becomes more engaged or outraged when confronted with extreme positions, so that feeds tend to select for such content [34]), and also on the side of the contributor (since these posts then tend to receive more positive feedback, encouraging users to post more in that direction). This feedback loop is just one example based in a kind of psychological trap. Other strategies are to trigger anxieties (fear of missing out (FOMO) or how one compares with others), or exploit tendencies toward certain kinds of addiction. Overall these methods aim to result in a strong pull toward the website, platform or application.

In addition to maximizing the number of people who see an advertisement and the length of time they attend to it, the ad also becomes much more effective if seen by groups most likely to respond to it. IT companies that gather personal user information

not only identify the users as members of particular groups, they also use their elaborate algorithmic tools of statistical analysis to identify for the advertiser which target group to aim for [6, 22]. However, in order to provide these services to the advertiser, the company needs to gather and correlate more and more personal information. There are three things happening. The first is the gathering of any available possibly relevant raw information by tracking the user, including actions taken in the browser, contact addresses, and the mobile location. In addition, if the records that belong to a given person can be connected during a series of visits and between different websites and applications, a far more sophisticated user profile is generated. This is the reason these websites use trackers that are able track the entire activity during the browser session. It is also a reason that Facebook and Google are highly motivated for users to log in with their accounts for identification to other service providers. Third, the correlation of many of these sophisticated user profiles makes it possible to statistically make predictions about user behavior, and thereby understand the users possibly better than they understand themselves.

Another practice is the *manipulation* of user behavior, not just surveillance. One motivation for this is that predictions about user behavior can be made more precise if those behaviors can be nudged in particular directions, making those predictions even more valuable for the surveillance capitalist firm's customers, namely the advertisers. However, if the manipulation is no longer directed just at improving predictions to make them more valuable to advertisers, but instead (for example) to strengthening a particular political opinion or nudging a citizen toward voting a given way, the IT company has developed a tool that can wield significant power and control over civic life.

Above we noted the effect of this business model on the orientation of these IT companies. We now explore the outcomes for individuals and society, including how the intensive gathering, tracking, and correlation of personal information, along with behavioral manipulation, poses a threat to sustainable economies and democracies in multiple dimensions.

2.1 Consumerism and Excessive Consumption

The increasing precision and effectiveness of customized advertising serves the imperative for consumption and unending growth. The generation of artificial needs is valuable to the advertisers, which can serve these needs, while having a destructive influence on the individual as well as society. On the individual level the users pay twice for these artificially generated needs, while true needs remain unsatisfied: once with the money of their purchases; and the second time with the time they spend due to the provocative content and in psychological traps. On the level of society, (1) personal information bubbles amplify social fragmentation and thereby hinder effective action on global problems/challenges, while (2) such a level of consumption cannot be sustained within ecological limits. Let us now examine these two outcomes more closely.

2.2 Threats to Government and Democratic Processes

Social media platforms such as Twitter, Facebook, and Instagram offer new types of (semi-)public spaces. They play an increasingly

¹Solutions for Environment, Economy and Democracy

important role in the exchanges of ideas, visions, and convictions that are central to civic life. However, the algorithms by which content is selected are often opaque to the reader [2]. On the one hand, these public spaces can have an emancipatory effect, for example under conditions of surveilled telephone lines and censored mass media, as occurred in the early stages of the Arab Spring when discussions critical of the regime, planning for demonstrations, and the distribution of news were facilitated by these platforms [30, 38]. On the other hand, the content of these exchanges will be recorded by the platforms and can contribute to the personal profiling of the discussants. Thus, these semi-public spaces have a Janus nature. They offer emancipatory potential, but at the same time contribute to the refinement of personal profiles and opportunities for manipulation. The role of Facebook in the 2016 U.S. election and the appropriation of the Whatsapp messenger in the recent elections in Brazil shows the manipulative power that comes with the ability to create personal profiles and to distribute targeted political propaganda via social media platforms [18, 33].

A lack of privacy with regard to political communication on social media platforms can lead to less participation and to self-censorship, depriving the debate of opinions that could support political progress. Particularly due to the potentially unlimited lifetime of the data and lack of transparency with regard to what personal data was gathered and how it was used in profiling, people in public offices or running for them will have to permanently fear that unpleasant private matters from their past could be dug up. In addition, the pressures toward more and more provocative content drives extremism and social fragmentation.

Looking, for instance, at the experiences of the Arab Spring [30, 38], it is clear that authoritarian regimes can use the personal data stored in social media platforms for surveillance and propaganda. This is true for Western governments as well, as demonstrated by the Snowden papers. The platforms themselves do not need to be in hands of the government; it is sufficient for the government to gain access to the gathered data. As the case of NSA indicates, platform providers could grant the access if they are coerced by legal action or threats of losing government support or even market access. A state does not need to initially be an authoritarian regime to develop in that direction as a result of building up these types of surveillance techniques.

In China, the government is accessing different social media platforms to profile the behavior of its population by means of a point-based social credit system [9]. In the West this is often regarded as a development toward an Orwellian surveillance state [10], even though similar software architectures and personal profiling capabilities are being built up in the private sector. Although states in the West are generally not (yet) legally able to access these profiling data and match between the different platforms directly, security services seem to have these abilities.

Governments in general will not want to see any drift toward political opinions that do not support their own political mandate. If instead they can use these services as a tool to propagate their own worldviews, it could in the eyes of the government even be seen as a good thing to do so. Therefore it comes at no surprise that many countries deploy significant resources to manipulate domestic and/or foreign online public spaces [5].

Particularly in democratic countries, the targeting of voters based on their psychological profiles becomes politically charged, as the case of Cambridge Analytica shows. This company played an important role in the Brexit referendum as well as in the Trump election [7, 21]. Once the sovereignty of humans over their personal data is lost, the borders with a propaganda and surveillance state blur. Or as the Cambridge Analytica whistleblower Christopher Wylie puts it [7], “If you do not respect the agency of people, anything that you’re doing after that point is not conducive to democracy.”

2.3 Privacy Concerns

The practices of amassing personal data conflict with principles of privacy. Although people usually deliberately choose to use the services and thereby provide their personal data to the companies, they are pushed to do so, due to a lack of transparency and little understanding by most users regarding which data is tracked, how long it is held, and to whom else it is given, as well as the lack of real alternatives or options that provide end user control.

While general business conditions or national law may regulate the handling of this data, a) this regulation is only on a legal level, whereas misdemeanor can be hard to prove on a factual level, where the IT companies still control the data and b) these regulations do not apply to predictions generated out of this raw data.

2.4 Concentrations of Wealth and Power

The power that arises from control over the service and personal data is more profitably turned into revenue, not by selling the service, but instead by predicting and manipulating user behavior. This concentration of power is followed by a concentration of wealth, as evidenced by the amount of capital these companies have been able to accumulate. Furthermore, the existence of a tool to wield control provides wealthy customers with more power that can be used to accumulate even more capital. In other words, surveillance capitalism enhances the concentration of power within and outside these IT companies. In addition, the small number of employees in comparison with the revenue generated by these companies exacerbates inequality, which undermines sustainability, democracy, and much else.

2.5 Economic Growth and Sustainability

Our overall economic system is currently predicated on unending economic growth, which is on a collision course with sustainability. This requirement for unending growth interacts with surveillance capitalism in two ways. First, it leads to pressure for increasing consumption, which in the developed world means that more and more needs are generated artificially, rather than being necessary for real prosperity (also see Section 2.1). Advertising, precisely targeted to the individual consumer using the tools and data provided by surveillance capitalism, helps fuel this consumption. Second, it leads to pressure for growing corporate profits, which for firms practicing surveillance capitalism drives them to bring more and more of human experience under the domain of surveillance and conversion to data, and increasingly to engage in behavioral manipulation as well, in service of behavioral prediction markets [39, 40].

Challenging this worldview and tackling the global issue of transitioning to a truly sustainable economy requires careful thought,

deliberation, and action. However, the behavioral patterns that are encouraged by social media platforms, even if as yet under-investigated, seem to drive people toward fast emotional rewards and less focused pursuit of long term goals.

3 STEPS TOWARD SOLUTIONS

Ultimately we need new models for these services that do not rely on the logic of surveillance capitalism. We outline a partial solution in Section 4, and a number of more comprehensive and hopefully effective ones in Sections 5 and 6. Getting to any of these solutions (especially the more comprehensive ones) will be very difficult. In this section, we outline some intermediate steps that we believe will help pave the way toward making these new models feasible. These steps are grouped into three categories: education, regulation, and resistance. However, the categories are interrelated and build on each other. For example, regulations that make it easier for people to see what information about them is being stored support both education and resistance.

3.1 Education

One key step toward finding solutions is for people to understand how these services are being funded, what kinds of information is being gathered about them, how their behavior is being manipulated, and the consequences of all this. A great deal of the rhetoric around the corporations using a surveillance capitalist business model has focused on individual choice, limitless access to information, empowerment, and personalization, until recently with relatively little focus on the model's dark side of surveillance and manipulation. There were flare-ups of negative reactions, for example, in 2004 to the initial description of how Google's Gmail scans private correspondence to place targeted advertising, but subsequently this became (perhaps grudgingly) accepted as normal. In the last two years, there has been a substantial shift as more of the extent of the surveillance and manipulation has become visible, especially in light of the reports of extensive online Russian targeting of the 2016 U.S. presidential election. In addition to numerous reports on election hacking, there have been increasing numbers of editorials, articles, and books on this topic, with Zuboff's book [40] being a significant milestone in terms of presenting the depth and broad scope of the problem along with an intellectual framing.

It is essential that the education process continue, with ongoing discussion and exposure of the extent of surveillance and political and other behavior manipulation. It is also important that we do not fall into the trap of assuming such a world is now normal and acceptable. However, neither being in a state of numbness or grudging acceptance, nor being in a state of continual outrage for years, are attractive alternatives. We also need positive visions of how we can use information technology to support human flourishing without surveillance and manipulation, and the collective political will to move toward those visions.

3.2 Regulation

Another key step is stronger regulation [35] of the corporations practicing surveillance capitalism. An important milestone here is the General Data Protection Regulation (GDPR) from the European Union, which took effect in May 2018. The GDPR is certainly a

major step forward; how it plays out in the next years remains to be seen. Two obvious concerns are that it applies only in the EU; and also that the corporations most affected by it will be able to afford to hire numerous highly-skilled lawyers, lobbyists, and others to counter its impacts on their profits. Given that the profits of Google and Facebook in their current forms, for example, are almost completely linked to surveillance capitalism, these efforts will likely be formidable. (In other words, the basic business model of surveillance capitalism leads to the problems described in Sections 1 and 2 — they are not incidental byproducts. Absent different models, regulation can help mitigate but not eliminate these problems.) An in-depth discussion of regulation of this industry, including its international ramifications, is well beyond the scope of this essay. Instead, here are a few additional ideas. (Not all of these will be feasible — but we hope they will be at least provocative and help further the discussion.)

It should be possible for individuals to know just what information various corporations are storing about them and which other corporations have access to it, to challenge inaccurate information, and to demand deletions, all in easy-to-understand presentations. Regulation can help support education, for example, by making it easier for people to see what information about them is being stored. Zuboff [40, p. 482–483] describes the unsuccessful struggle of Belgian mathematician and data protection activist Paul-Olivier Dehaye to uncover all of the web pages where Facebook had tracked him, despite several years of persistent efforts. (If a determined activist will not be successful in such a pursuit, what are the chances for an average citizen?) Regulations that require companies to make this information easily available, and tools provided by activists to help us digest it, can make these abstract threats much more personal and tangible, and motivate pushing for change.

In particular, if people knew just how much data was being amassed about them under surveillance capitalism, perhaps enough of them would be sufficiently outraged to push for alternate models. This is of course likely to get even more intense pushback, but it would be difficult for a corporation to argue that the data it holds about its users should not be revealed to those users because they would be angry — the pushback will instead likely be about technical feasibility, inhibiting innovation, undermining the corporation's ability to provide highly personalized service, claims that existing privacy policies provide full protection, and so forth. (Pushback aside, there are in fact some real dangers that might arise in conjunction with such visibility, such as the danger of misuse if someone else fraudulently obtains your data or coerces you into providing it, and the problem of indirect stakeholders — other people who were tracked at the same time you were and whose information would also be exposed or made public if you did so.)

The GDPR includes a high threshold for the definition of “consent,” but when possible there should be strong privacy protections that instead entirely eliminate classes of data collection and sharing and behavior manipulation. For cases where consent is still appropriate, it is likely that additional requirements could help. Regulations could require unbundling what is being consented to, with different options for what can be gathered and how long it can be retained, and with which other entities it can be shared. Consent should be opt-in and not opt-out, and policies clearly and concisely stated — a recent *New York Times* editorial [11] notes

that “The average person would have to spend 76 working days reading all of the digital privacy policies they agree to in the span of a year. Reading Amazon’s terms and conditions alone out loud takes approximately nine hours.”

Corporations engaged in this kind of data gathering and manipulation could be subject to an impartial external audit, conducted by skilled researchers who investigate what the corporations are storing and manipulating, and publicize the results.

An action that is regularly proposed is to declare that corporations such as Google and Facebook are monopolies, and to require that they be broken up. We have concerns about such a move: the result could well be that there would be a multiplicity of surveillance capital firms, with all the same problems but harder to keep track of. Somewhat better would be to break up firms by functional capabilities, in particular to separate out the common carrier and similar functionality (network, storage, and cycle server infrastructure) from content functionality, and to restrict the data that the common carriers and their analogs can share.

Regulation is likely to result in lower profit margins for surveillance capitalist corporations [36]; arguably this should not be the goal of such regulation, but should nevertheless be an acceptable outcome. Finally, another important purpose of regulation could be to deliberately nudge the market to make it easier for other models for providing IT infrastructure to flourish — or ideally, to make surveillance capitalism models untenable.

3.3 Resistance

There are several potential goals for resistance to surveillance capitalism, including personal integrity, undermining its profitability, and raising awareness and calling people to action (i.e., education). Trying to maintain personal integrity is of course important as an end in itself, and also in helping avoid having surveillance become normalized. However, such actions, or other actions whose purpose is to undermine the profitability of the business model, seem unlikely to have sufficient impact on their own. But doing these things (and discussing doing them and the challenges of doing them) can contribute to awareness and calls to action.

Resistance can take a variety of forms. One is to simply not use certain parts of the IT infrastructure, e.g., the #DeleteFacebook movement. This certainly has merit, but can also make it difficult to participate fully in society, given the extent to which Facebook enters into many social interactions, into deliberations among members of a political movement, and so forth. It also recasts a political issue as a willpower issue [14]. And it seems simpler to delete Facebook than for example Google, given Google’s pervasiveness. (As a more extreme example, Hill describes her attempt over a period of six weeks to block the five tech giants [16].) Another important form is as art directed at the themes of surveillance and resistance [40, p. 491–492], which (among other things) can push back against such surveillance and manipulation as being considered normal. Finally, there are various kinds of technical resistance that seek to avoid being tracked, or to disrupt surveillance.

Regarding specific tools for such technical resistance, web browsers often provide a switch to block setting third-party cookies. This is only somewhat useful, since among other things it often just blocks cookie writing, not reading/sending. For example, if

a user visits Facebook directly, it would be a first party and so a cookie could be set and then subsequently used by third parties. Also, there are many other techniques for tracking besides cookies, notably browser/machine fingerprinting [27]. Web browsers may also provide a “do not track” setting — unfortunately, though, this option is effectively dead since it only works if trackers honor the request (and many do not). Simply turning off JavaScript can help as well, although doing so will also cause many sites to be unusable.

There are also a variety of ad blocker plugins and other anti-tracking browser extensions, such as uBlock Origin², Privacy Badger³, AdBlock Plus⁴, and Ghostery⁵. On the more stringent (and difficult-to-use) side, uMatrix⁶ can be set up to block all third-party requests by default, and then let the user choose which domains to enable for a particular webpage. The Firefox browser itself also includes some tracking protection [26], including in “private browsing” mode. (In other browsers, “private browsing” modes may not really protect against tracking — the goal there is more to protect the user’s web history from someone with access to the user’s device.) Panoptick⁷ from the Electronic Frontier Foundation will analyze how well the user’s browser protects against tracking.

A different approach is taken by AdNauseam⁸, built atop uBlock Origin, which simulates clicks on every blocked ad to generate a stream of meaningless data that obscures the user’s actual interests and behavior (also see [17]). Another is a Firefox add-on called Multi-Account Containers⁹, which are like normal tabs on a browser except that each container has its own preferences, advertising tracking data, and other information, which cannot be seen by the other containers, making it harder to do tracking across sites. However, they can be unintuitive for users, and it can still be difficult for users to reason about tracking since webpages often load from so many different sources.

Relevant papers in the academic literature include an early study on tracking with measurements in Summer 2011 [29], a longitudinal study of tracking 1996–present [19], and a demonstration that anyone can buy ads to track a targeted individual [37].

Stepping back, one is struck by the considerable effort that is going into these technical approaches to resistance, how complex the solutions are, and the extent to which there is a cat-and-mouse game going on between the trackers and the tracked. The economic impact on surveillance capitalism of this technical resistance is likely to be limited by its complexity. However, the main practitioners of both its development and use, such as computer science students and software engineers, are also likely the potential employees of the big IT corporations, and employees are a scarce resource, so they may have power by other means. Finally, if technically skilled users find the landscape challenging and confusing, nontechnical users must find it even more so. If one were a journalist reporting from on-the-ground in a repressive regime, one can imagine it being reasonable to require these kinds of precautions. But should ordinary citizens who just don’t want corporations tracking everything

²<https://github.com/gorhill/uBlock>

³<https://www.eff.org/privacybadger>

⁴<https://adblockplus.org>

⁵<https://www.ghostery.com>

⁶<https://addons.mozilla.org/en-US/firefox/addon/umatrix/>

⁷<https://panoptick.eff.org>

⁸<https://adnauseam.io>

⁹<https://support.mozilla.org/en-US/kb/containers>

they do online need to do this also? Ultimately, the most important role for such technical resistance may be as part of education and helping build pressure for more comprehensive change.

4 A PARTIAL SOLUTION

The problems of surveillance capitalism can to some degree be addressed by regulation, pushed by an educated population and selective resistance. A further step in this direction is to encourage the development and growth of for-profit corporations that still provide these services but without tracking personal information, both to provide an alternative for users and thereby also to nudge the existing big IT corporations. This partial solution is thus still very much capitalism, just not surveillance capitalism (or at least not surveillance to the same extent). Despite its partial nature, we believe it is worth calling out explicitly. Among other things, this is (we claim) implicitly the solution that is being proposed in approaches based on regulation alone, absent measures specifically to foster alternate business models.

Two systems to be noted in particular are Brave¹⁰ and DuckDuckGo¹¹. Brave is an open-source browser that (the company says) blocks ads and trackers, in both mobile and desktop versions. It includes a facility for giving micropayments to publishers of content being viewed using blockchain-based tokens. The DuckDuckGo search engine, according to the company, does not collect or share personal information. Its business model is still based on advertising (and also affiliate marketing). The ads shown on DuckDuckGo are based just on the keywords typed in the search box, rather than also on tracked personal information. Revenues come from Amazon and eBay affiliate programs: when users are referred to one of those sites by DuckDuckGo and then buy something, the company collects a commission. Finally, Apple among the big five tech companies has a significantly better record with respect to privacy. Apple still has other major problems — with respect to labor practices, sustainability (for example, encouraging frequent purchase of expensive new devices), and so forth — but to date they have resisted to a considerable extent adopting a surveillance capitalist business model.

In terms of why we label this as only a partial solution, Brave and DuckDuckGo are still advertising-supported — this seems potentially problematic because it leaves the companies vulnerable to the desires of the advertisers. Also, to what extent will this challenge the power of the entrenched surveillance capitalist corporations?

In Section 6 we discuss approaches for a variety of application areas. Systems that implement some of these approaches, for example for public transit information (Section 6.2.5), can be provided now without needing to use a surveillance capitalist model, making them quite compatible with this partial solution approach.

A final idea here is to nudge the market by having institutions such as libraries, universities, and others buy ad-free, no tracking versions of services for their patrons/students, either from new companies, or from existing large IT corporations if they are willing to unbundle their services to support this. (Note that it would be essential to carefully monitor the corporations to ensure they are not tracking these users [12, 28].)

¹⁰<https://brave.com>

¹¹<https://duckduckgo.com>

5 ISSUES OF FUNDING AND CONTROL

Section 4 described approaches that are still very much capitalism, just not surveillance capitalism as currently practiced. We view these as partial solutions only. In this section and the one that follows, we instead consider more far-reaching alternatives that question not just the “surveillance” part of the phrase but also the “capitalism” part (also see [24]). Two major issues are thus funding and control. We now consider some alternatives to for-profit corporations for funding and controlling IT infrastructure.

5.1 Public Funding

One alternative for funding IT services is public funding. However, if these services can be used as tools to influence public opinion and behavior (and it has been shown that they can [22]), government control of IT services brings with it the danger of manipulation. All the concerns regarding threats to the government and the democratic processes (Section 2.2) hold equally if not even more strongly if governments not only have some influence on the tools but in fact are their owners and creators. Of course, it makes a huge difference whether the government is under the control of a functioning democracy, a democracy in name only, an authoritarian government, or something in between. Democracy itself is also under huge stress at present. But for purposes of this essay, we want to at least contemplate the possibility of having a well-functioning democracy that might fund and oversee some of these services.

More radically, these considerations raise the issue of whether society is better off with the immense power that arises from the control over a communication medium that people rely on being in the hands of governments or of private corporations. Particularly if the government is not a functioning democracy, arguably neither provides a satisfactory solution, since in both cases there is a strong interest in trying to control the behaviors of individuals.

Despite these considerations, we do believe there is a substantial role for government funding. One example that provides a useful analogy is the public radio and television systems that exist in many countries; another is the subsidies to newspapers that existed in the U.S. in the 19th century via subsidized postal rates and tax policy [23]. Or there could be other programs that emphasize individual choice and responsibility. For example, “journalism vouchers” could be issued to every resident that would allow people to provide grants to investigative journalists, whose work would then appear on social media.

5.2 NGOs and Cooperatives

Another possibility is having other societal institutions that control the service. If the continued existence of such institutions is insulated from day-to-day changes in public opinion, this removes one source of pressure to engage in propaganda or surveillance. One possibility here is NGOs (e.g., the Mozilla Foundation, which is the sole shareholder in the Mozilla Corporation). However, being a nongovernmental organization does not automatically guard against conflicts of interest arising from funding, nor does being an NGO automatically mean the organization will be benevolent. Minimally, a close look at the organizational structure is needed.

Schneider and Schulz [31] advocate placing these alternatives in the hands of worker cooperatives. Using this model, more of the

relevant stakeholders would be included in the ownership model, particularly if it also includes the end users of the infrastructure. However, if one takes a closer look at the ownership structure, often the potential for conflicts of interest among different sub-groups of a coop becomes apparent. Although such a structure would be a significant advancement in the power balance, it still leaves open questions. By which mechanism would the formal owners coordinate and exercise their right to make decisions? What if one group of stakeholders, e.g., the programmers, refuse to implement the changes the majority of owners decided upon?

5.3 No Funding or Minimal Funding

Freely contributed work is another alternative. Examples such as Wikipedia and OpenStreetMap show how an enormous amount of knowledge can be contributed by volunteers, perhaps along with funding for hardware and support staff. Such a model can work well if a clear structure is provided that guides how to arrange and connect the different contributions. While the development of an entire IT service is probably not suitable for volunteers alone, such a group could very well do the maintenance, given that a clear structure to do so is provided. Such a structure can be given by a modular design, which we propose in Section 6.1.

On the other hand, to provide hardware infrastructure without funding would require freeing up idle capacity of existing infrastructure. There are working examples where the infrastructure is provided freely by the users themselves. These include file sharing and streaming services, utilizing decentralized architectures as we propose in Section 6.2.

Having freely contributed labor and infrastructure still does not address the question of who controls the data. Who is its legal owner? For social media, who designs and controls the algorithms that selects the content that is presented to the users? Such questions motivate the considerations outlined in the next section.

6 TOWARD A MORE RADICAL SOLUTION

In the previous section we discussed some general issues for alternatives to surveillance capitalism. For each of these alternatives we still have the question of who controls the service. Who is the legal owner of the data? Who (if anyone) is able to do datamining on it? Who has the ability to design the digital environments users employ and that help shape their behavior? Arguably, the power that arises from these privileges is simply too big to be placed in the hands of any entity. This gives rise to the question of whether the service itself can be structured in a way that there is no such entity. In the next section we consider some necessary properties such liberating software would need to possess.

6.1 Desired Characteristics

The first necessary condition is user control over personal data. Personal data here consists not only of the information users give about themselves, but also records of all the actions they take that are at least theoretically trackable while interacting with the service. For the latter kind of data, in many cases this will mean that they are not tracked at all. Since that depends on the program code, the source code must be visible. For the first kind, the users need to be able to choose who is allowed to see this data. Although current

service providers often offer some choice regarding which other users are allowed to see the data, there is no option to hide the information from the service provider itself. Realizing this requirement suggests not being satisfied with end-to-end-encryption, but instead saving user data locally rather than on a central server.

The second point concerns the development of features that users desire. In an economy run according to capitalistic principles, the approach to this is a corporation investigating the desired features and developing them. If these features turn out to be valuable to the users, this gives the corporation a competitive advantage. However, a corporation might have no interest in implementing a feature users desire, for example one for more privacy, if it would reduce the corporation's ability to generate revenue. If there are a few dominant players, competitors have no chance of out-competing that corporation by implementing such a feature, because network effects constitute a barrier to adoption that is larger than the benefit. In such a situation progress is stalled.

One approach to avoiding this is to design software with options to add, hide, and connect with other features, rather than monolithically with only small APIs that enable some limited connection with other services. This is also compatible with an End User Development approach [20]. Not every user will have the technical knowledge or the willingness to experiment with such options, but the lower the hurdle is to do so, the more easily user desired designs and features will be developed. As it is not required for all users to migrate to the new service, but instead new features come as optional add-ons to the existing ones, users don't lose connection to their existing peers upon testing a new feature. In such an environment two users who want to interact with each other simply have access to all of the features they share in common. For such a modular set of features to be manageable, ideally users have software that arranges these tools in a customized way that they are familiar with. Such a design also lessens their unwilling exposure to manipulation and pull toward extremism. For example, it would empower users if there were a choice among different news feeds that are not all designed to constantly draw them in and spend more and more time on the platform. Therefore a modular architecture is the second ingredient to enable users to create an online experience oriented toward their personal needs.

Modular solutions that leave individuals in charge of their data and experience could still be tools to concentrate power if there is an entity with the ability to read or deny the exchange of information via this service. Therefore, the third requirement is for truly peer-to-peer services. There are many platforms that label themselves as peer-to-peer, because peers do communicate with each other. However, if the communication uses the central server of the platform provider as a channel, this does not meet the stronger understanding of truly peer-to-peer communication as used here. This strong understanding of a peer-to-peer system requires that the entire communication, including the channel, be held by the peers. Such applications, instead of answering the question in whose hands power over the communication system might be relatively safe, solve the problem of power by not letting it manifest in the first place.

In combination these three architectural choices seem promising to solve the problems from Section 2:

- (1) User control over personal data
- (2) Modular design
- (3) Truly peer-to-peer communication

We do not see any indications that these properties are incompatible; to the contrary, user control over personal data strongly hints in the direction of peer-to-peer communication. To substantiate this, the following subsection illustrates how these properties could be realized for particular classes of IT services.

6.2 Returning Control to the Users

The outlined properties radically place the control over the services into the hands of users. They are well aligned with what Tim Berners-Lee seeks to accomplish with the SOLID project¹² at MIT. However, in order to achieve this, we propose that distributed ledger technology (DLT) play a key role in implementing truly peer-to-peer structures. (For other decentralized alternatives, also see for example [32]).

The most prominent DLT, blockchain, is not suitable for our purposes, since it requires the full data load of all members to be held by every peer. (Changing precisely this property is called sharding, which is difficult and not achieved by classical blockchains.) As a result, it neither scales up, nor does it prevent excessive data mining. Proof-of-work blockchains, such as the ones Bitcoin and Ethereum are based upon, require ecologically unsustainable energy consumption. However, there are promising blockchain alternatives that do not suffer from these weaknesses and that are potentially able to implement these desired characteristics. One of them that in particular is intended to support these properties is Holochain [15]. To give examples of how alternative services could use such distributed ledger technology to support these principles, we need to distinguish between different classes of services.

6.2.1 Social Media Services. Social media services and publishing platforms are one class of services for which currently the surveillance capitalism business model is common. Examples of this class include Facebook, Twitter, Instagram, Medium, and LinkedIn. For these applications, users create a personal profile themselves and want the content they create to be connected to that profile, visible to other users. They might want to create different categories of visibility and authenticate other users for these categories. In our proposed alternative, the user profile and its content are stored locally on the user's device. Since that will not be available all the time, the DLT stores the data with some redundancy on other devices. When somebody wants to visit the user's profile, instead of sending a request to a central server, he or she now needs to connect to one of the devices that holds the user profile data. That means that all users need to carry in addition to their own data load the backup for other users. Because mobile devices in particular are not designed to carry and upload large amounts of data, these obligations can be fulfilled by devices with a better hosting capacity and a broadband connection. It would of course be highly inefficient to require every user to own a tiny server. Instead, we imagine many hosts to fulfill these obligations, for which they can be paid using micro-payments. Technologically this is enabled by deploying an appropriate distributed ledger technology. With the

demand for tiny hosting volumes and the technical abilities to pay for these automatically, it would for example become feasible to rent out idle capacity of existing computers, resulting in a much more decentralized network than the central structure of big companies like Amazon Web Services that exist today. In addition, the capacity of public institutions such as libraries and universities could be integrated in such networks. This structure does introduce some inefficiency, since the data is stored with a higher redundancy than on a central server, where redundancy only needs to counter the rare case of hardware failure. However, it grants complete control over private data to the user, who is freed from the service provider. That service provider no longer plays any role in the communication between users and cannot easily datamine the user profiles, because there is no single database of user profiles, and because it would not be authorized to access the profiles unless the users permit it.

The second dimension of user control is over the features the service offers, especially the algorithms that selectively present content. Where today the features are determined by the service provider, with very limited APIs and algorithms that are a mystery to the users, in our proposal the requirement for a modular software structure gives users the ability to add in their own modules. Those could include modules for simple features such as a dislike-button. Using this process the community can find out whether such a module enriches or frustrates their experience. A more complex example is a content filter for the activities of friends, implementing a new algorithm. This would enable users to better control their online experience, and overcomes the network effect problem that blocks adoption of such a feature unless the infrastructure provider itself adds it. Admittedly only users who are sufficiently educated and motivated in that direction would embrace that opportunity, but the modular structure is intended to make the barrier for participation as low as possible.

6.2.2 Employment and Rental Platforms. A second class of services consists of platforms that connect different user groups with complementary needs, for example digital employment platforms. Examples for this class include Airbnb, Uber, Deliveroo, blablacar etc. These platforms are the backbone of the gig economy, and although they do not rely heavily on advertising for funding, they still involve many of the same characteristics of surveillance and behavioral manipulation [8], so we view them as also being instances of surveillance capitalism. Alternative applications that would avoid this surveillance and manipulation (and also perhaps some of the high fees) could be analogous to what we just described in the first class of examples, including different user groups. Once such an application with a license that allows free usage of the software is distributed, no fees can be imposed by the software provider, who is no longer a part of the communication.

6.2.3 Search Engines and Browsers. Two other classes of services are search engines and browsing applications. For these, users usually do not want any information to be shared with other users or corporations. There might be some interest in sharing certain customized settings or bookmarks between different devices that belong to the same user, and where a generalized structure for peer-to-peer applications is in place it might be used for this case as well. However, in general, for this class of services versions that

¹²<https://solid.mit.edu/>

do not track any personal information (Section 4), perhaps with a funding source other than advertising, seem appropriate.

6.2.4 Email and Messaging. For the class of email providers and messaging applications a serverless peer-to-peer structure would not be far fetched. Many of these applications are encrypted end-to-end already, but the provider can still try to monetize meta-data. The difference would be that the emails or messages are not stored centrally at the providers server, but instead locally on the user's device. Again some redundancy is required, so that the role of the server is replaced by the DLT, or more precisely, a small part of the peer-to-peer network of users. This backup also enables the users to synchronize between their different devices. In contrast to social media, something like a tool to micro-pay for hosting services might not even be needed when the application requires users to carry this backup for other users, because messages are small. Only for attaching large files, such as pictures, audio or video recordings, a solution similar to what we suggested for social media would need to be provided.

6.2.5 Public Transit Information and Other Specialized Applications. There are also a variety of specialized applications that are clearly appropriate for public funding and do not require the peer-to-peer approach. One example is apps for public transit information, including schedules, trip planning, and real-time arrival information. In Europe, these are often already publicly funded, but in the U.S., many transit agencies make their schedule and real-time information available via an API and perhaps a website only, so that users who want to access this information online via apps on mobile phones must do so via ones that are funded by the surveillance capitalism model (or by startups funded by venture capitalists whose business model is presumably to garner as large a user base as possible for eventual sale to the surveillance capital market). Public funding for the apps as well as the servers, with open source software that guards rider privacy, would be a relatively small additional expense and definitely worthwhile. One example of such a system is OneBusAway¹³, which also includes open source versions of the server side software as well. OneBusAway is in production use in a number of regions in the U.S., Canada, and Europe. In informal presentations and discussions with transit riders, they often do not realize the surveillance potential of for-profit transit apps, but once they do, they typically strongly support an open-source surveillance-free alternative — thus providing another example of the importance of education.

6.2.6 Wrapping Up. The example application areas presented above that rely on strongly decentralized peer-to-peer architectures are structured in a fundamentally different way from current practice, while others would need only a different funding mechanism (although that alone is a formidable barrier). The peer-to-peer applications would not only require a rebuild of a significant part of modern software infrastructure, but also a fundamental rethinking by the developers of these new services. We have presented a first examination of the design space shaped by the three properties listed in Section 6.1, and new obstacles may well arise. Let us therefore at least consider funding as one obvious source of obstacles.

¹³<https://onebusaway.org>

6.3 Revisiting Funding Options

For many of the relevant applications we could describe implementations of the peer-to-peer approach described in Subsection 6.2 embodying the three desired properties, which by design provide considerable protection against the undesired outcomes of surveillance capitalism. This provides a fresh view on the funding and control possibilities outlined in Section 5: given the modular and decentralized nature of this approach, different implementations of services, with funding from any of private, state, cooperative, or volunteer sources, should be capable of co-existing and evolving.

6.4 Open Problems

Significant numbers of open problems remain. For example, the approach described in Section 6.2 applies only to certain classes of applications, such as social media. Other classes, including the important cases of search engines and browsers, do not seem amenable to peer-to-peer structures — approaches here must focus more entirely on funding and control. There are also a set of other critical problems not addressed by these suggestions. For systems using machine learning, these include bias in training sets and explainability. For social media, these include the effects of personal information bubbles, use by extremists, and the vulnerability of the political process to targeted misinformation, trolls, bots, and the like (perhaps controlled by state actors). For employment platforms, the lack of the legal safeguards for employees, the uncertain nature of the work, and competition between workers that pushes wages to a precarious level, may result in significant economic insecurity. A last major set of issues concerns how the industry could evolve and be restructured, as well as how to overcome an imbalance between the largest companies and new entrants in funding and incentivizing good developers and designers.

7 CONCLUSIONS

We began by describing the ominous implications of the surveillance capitalism model that underlies most of the major IT corporations, which include intensive gathering and correlation of personal information and behavioral manipulation. We then described three categories of steps toward solutions: education, regulation, and resistance. Following that we discussed a partial solution that involves encouraging the development and growth of for-profit corporations that still provide similar IT services but without tracking personal information. We then outlined a range of more comprehensive solutions for different portions of the IT infrastructure that more truly return control to the end users. While many challenging problems remain to be worked out, particularly with the more comprehensive solutions, we hope that this essay can help stimulate discussion of these problems and more importantly potential solutions.

ACKNOWLEDGMENTS

Many thanks to Franzi Roesner, Anna Kornfeld Simpson, and the ever-helpful members of misc@cs.washington.edu mailing list at UW CSE for recommendations and references regarding tools for technical resistance to surveillance; to Lance Bennett, Jeff Chase, Jared Moore, Matt Ziegler, and the LIMITS referees; and to the University of Siegen for funding the SEED workshops, which started the authors on this project.

REFERENCES

- [1] Mike Allen. Sean Parker unloads on Facebook: “God only knows what it’s doing to our children’s brains”. *Axios Newsletters*, Nov 9 2017.
- [2] Renata Ávila, Juan Ortiz Freuler, and Craig Fagan. The invisible curation of content: Facebook’s news feed and our information diets. Technical report, World Wide Web Foundation, 2018. http://webfoundation.org/docs/2018/04/WF_InvisibleCurationContent_Screen_AW.pdf.
- [3] W. Lance Bennett, Alan Borning, and Deric Gruen. Solutions for environment, economy, and democracy (SEED): A manifesto for prosperity. *ACM interactions*, 25(1):74–76, December 2017.
- [4] Alan Borning. SEED: Solutions for environment, economy, and democracy. Keynote talk, 2018 Workshop on Computing Within Limits, Toronto, Canada, May 2018.
- [5] S. Bradshaw and P. Howard. Troops, trolls and troublemakers: A global inventory of organized social media manipulation. Technical report, Oxford Internet Institute, 2017. <https://ora.ox.ac.uk/objects/uuid:cef7e8d9-27bf-4ea5-9fd6-855209b3e1f6>.
- [6] Ricardo Buettner. Predicting user behavior in electronic markets based on personality-mining in large online social networks. *Electronic Markets*, 27(3):247–265, August 2017.
- [7] Carole Cadwalladr. The Cambridge Analytica files — ‘I made Steve Bannon’s psychological warfare tool’: Meet the data war whistleblower. *The Guardian*, Mar 18 2018.
- [8] Ryan Calo and Alex Rosenblat. The taking economy: Uber, information, and power. *Columbia Law Review*, 117:1623–1690, 2017. Also University of Washington School of Law Research Paper No. 2017-08. Available at SSRN: <https://ssrn.com/abstract=2929643> or <http://dx.doi.org/10.2139/ssrn.2929643>.
- [9] Yongxi Chen and Anne S. Y. Cheung. The transparent self under big data profiling: Privacy and Chinese legislation on the social credit system. *The Journal of Comparative Law*, 12(2):356–378, June 26 2017. University of Hong Kong Faculty of Law Research Paper No. 2017/011. Available at SSRN: <https://ssrn.com/abstract=2992537> or <http://dx.doi.org/10.2139/ssrn.2992537>.
- [10] Josh Chin and Gillian Wong. China’s new tool for social control: A credit rating for everything. *Wall Street Journal*, Nov 28 2018.
- [11] Editorial Board. How Silicon Valley puts the ‘con’ in consent. *New York Times*, page SR10, Feb 2 2019.
- [12] Cyrus Farivar. Former, current students sue Google over university-issued Gmail scanning. *Ars Technica*, Feb 3 2016.
- [13] Catherine Flick. Informed consent and the Facebook emotional manipulation study. *Research Ethics*, 12(1):14–28, 2016.
- [14] Anand Giridharadas. Deleting Facebook won’t fix the problem. *New York Times*, Jan 10 2019. Op-ed.
- [15] Eric Harris-Braun, Nicolas Luck, and Arthur Brock. Holochain: Scalable agent-centric distributed computing. <https://github.com/holochain/holochain-protocol/blob/master/whitepaper/holochain.pdf>, February 2018.
- [16] Kashmir Hill. Goodbye big five: Life without the tech giants. Gizmodo, February 2019. <https://gizmodo.com/tag/blocking-the-tech-giants>.
- [17] Daniel C. Howe and Helen Nissenbaum. TrackMeNot: Resisting surveillance in web search. In Ian Kerr, Valerie M. Steeves, and Carole Lucock, editors, *Lessons from the Identity Trail: Anonymity, Privacy, and Identity in a Networked Society*. Oxford University Press, New York, 2009.
- [18] Anna Jean Kaiser. The Brazilian group scanning WhatsApp for disinformation in run-up to elections. *The Guardian*, Sep 26 2018.
- [19] Adam Lerner, Anna Kornfeld Simpson, Tadayoshi Kohno, and Franziska Roesner. Internet Jones and the raiders of the lost trackers: An archaeological study of web tracking from 1996 to 2016. In *Proceedings of the 25th USENIX Conference on Security Symposium, SEC’16*, pages 997–1013, Berkeley, CA, USA, 2016. USENIX Association.
- [20] Henry Lieberman, Fabio Paternò, and Volker Wulf. End-user development: An emerging paradigm. In Henry Lieberman, Fabio Paternò, and Volker Wulf, editors, *End User Development (Human-Computer Interaction Series)*, volume 9. Springer, Berlin and Heidelberg, 2006.
- [21] Ivan Manokha. Surveillance: The DNA of platform capital – the case of Cambridge Analytica put into perspective. *Theory & Event*, 21(4):891–913, 2018. Project MUSE, <https://muse.jhu.edu/article/707015/pdf>.
- [22] S. C. Matz, M. Kosinski, G. Nave, and D. J. Stillwell. Psychological targeting as an effective approach to digital mass persuasion. *Proceedings of the National Academy of Sciences*, 114(48):12714–12719, 2017.
- [23] Robert McChesney and John Nichols. *The Death and Life of American Journalism: The Media Revolution that Will Begin the World Again*. Nation Books, Philadelphia, PA, 2010.
- [24] Evgeny Morozov. Capitalism’s new clothes. *The Baffler*, Feb 4 2019. <https://thebaffler.com/latest/capitalisms-new-clothes-morozov>.
- [25] Michael David Murphy. Transcript of excerpt from Chamath Palihapitiya’s Stanford Biz School talk. *Medium*, Dec 11 2017.
- [26] Nick Nguyen. Latest Firefox rolls out enhanced tracking protection. *The Mozilla Blog*, Oct 23 2018. <https://blog.mozilla.org/blog/2018/10/23/latest-firefox-rolls-out-enhanced-tracking-protection/>.
- [27] Nick Nikiforakis, Alexandros Kapravelos, Wouter Joosen, Christopher Kruegel, Frank Piessens, and Giovanni Vigna. Cookieless monster: Exploring the ecosystem of web-based device fingerprinting. In *Proceedings of the 2013 IEEE Symposium on Security and Privacy*, pages 541–555, Washington, DC, 2013. IEEE Computer Society.
- [28] Andrea Peterson. Google is tracking students as it sells more products to schools, privacy advocates warn. *Washington Post*, Dec 28 2015.
- [29] Franziska Roesner, Tadayoshi Kohno, and David Wetherall. Detecting and defending against third-party tracking on the web. In *Proceedings of the 9th USENIX Conference on Networked Systems Design and Implementation, NSDI’12*, pages 12–12, Berkeley, CA, USA, 2012. USENIX Association.
- [30] Markus Rohde, Konstantin Aal, Kaoru Misaki, Dave Randall, Anne Weibert, and Volker Wulf. Out of Syria: Mobile media in use at the time of civil war. *International Journal of Human-Computer Interaction*, 32(7):515–531, 2016.
- [31] Trebor Scholz and Nathan Schneider, editors. *Ours to Hack and to Own: The Rise of Platform Cooperativism, A New Vision for the Future of Work and a Fairer Internet*. OR Books, New York, 2017.
- [32] Amre Shakimov et al. Vis-à-vis: Privacy-preserving online social networking via virtual individual servers. In *Third International Conference on Communication Systems and Networks (COMSNETS 2011)*, pages 1–10. IEEE, January 2011.
- [33] Jake Swearingen. WhatsApp says it’s too late to stop far-right fake news in Brazil. *New York Magazine*, Oct 19 2018.
- [34] Zeynep Tufekci. YouTube, the great radicalizer. *New York Times*, Mar 10 2018. Opinion piece.
- [35] Moshe Y. Vardi. Are we having an ethical crisis in computing? *Communications of the ACM*, 62(1):7–7, December 2018.
- [36] Munsif Vengattil and Paresh Dave. Facebook’s grim forecast – privacy push will erode profits for years. *Reuters Business News*, July 25 2018.
- [37] Paul Vines, Franziska Roesner, and Tadayoshi Kohno. Exploring ADINT: Using ad targeting for surveillance on a budget - or - how Alice can buy ads to track Bob. In *Proceedings of the 2017 Workshop on Privacy in the Electronic Society, WPES’17*, pages 153–164, New York, 2017. ACM.
- [38] Volker Wulf, Kaoru Misaki, Meryem Atam, David Randall, and Markus Rohde. ‘On the ground’ in Sidi Bouzid: Investigating social media use during the Tunisian revolution. In *Proceedings of the 2013 Conference on Computer Supported Cooperative Work, CSCW ’13*, pages 1409–1418, New York, 2013. ACM.
- [39] Shoshana Zuboff. Big other: Surveillance capitalism and the prospects of an information civilization. *Journal of Information Technology*, 30(1):75–89, March 2015.
- [40] Shoshana Zuboff. *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. PublicAffairs Books, New York, 2019.