

DIGITAL INTERMEDIATION OF EVERYTHING: AT THE INTERSECTION OF POLITICS, TECHNOLOGY AND FINANCE

Essay by **Evgeny Morozov***

The reason why most of contemporary reflections on the digital condition fail to excite is simple: one can only understand today's digital world by seeing it as the intersection of complex logics driving the worlds of politics, technology, and finance. Grasping a phenomenon such as the rise of Uber, for example, is next to impossible without understanding where its funds — raised from sovereign wealth funds and investment powerhouses such as Goldman Sachs — come from. Likewise, Uber's ability to cheaply draw on a large pool of supposedly autonomous and independent drivers can only be understood in the context of the liberalisation of labour markets and the growing precarity of service work in general.

The conventional fairy tales that we tell ourselves about digital technology — they usually involve a bunch of hoody-wearing twenty-somethings barricaded in the proverbial garage praying to the Schumpeterian god of creative destruction — end up glorifying the entrepreneurial heroes while concealing the broader historical forces at play. Trying not to abandon various political and historical dimensions to the rise of Big Tech, this essay will attempt to elucidate five major features of today's digital society that bear some relevance for struggles against anti-democratic, extremist, and xenophobic forces that appear to be on the rise worldwide.

This essay will proceed in two parts. The first, more descriptive part will summarise the five trends and explain their broader political and cultural significance; for the sake of convenience, I will address each trend in a dedicated subsection. The second, more normative part, will discuss the kinds of cultural, artistic, and scientific interventions that could address the numerous problems raised earlier.

Data Extractivism

In the first six months of 2017, four big US tech firms — Alphabet, Amazon, Microsoft, and Facebook — have seen their stock valuations grown by an amount greater than the GDP of the oil-rich Norway. Similar trends can be observed in China, where a home-grown industry around its own giants — Tencent, Baidu, Alibaba, JD.com — has emerged to rival America's supremacy.

One can dismiss such trends as the beginning of a new bubble — perhaps, a replay of the dotcom craze of the late 1990s. There are, however, some solid reasons why such reading would not be accurate. First, unlike the late 1990s, there are some valid economic models underpinning the immense growth of such firms. Second, the immense valuations of the digital firms partly reflect their ownership of the most important resource of the 21st century: data — a digital residue of various social, economic, and cultural networks and relationships that crisscross our lives.

In fact, it would not be inappropriate to describe the logic that drives the development of this sector as “data extractivism” — in a direct parallel to the natural-resource extractivism that has driven the activities of energy firms and commodity producers across the globe.

The key premise of data extractivism is that users are valuable stocks of data; technology companies, in turn, design clever ways to have us part with that data — or at least share it with them. They need this data either to fuel their advertising-heavy business models — more and better data yields higher advertising earnings per user — or they need it in order to develop advanced modes of artificial intelligence centred around the principle of “deep learning”; here, the diversity of data inputs —and the ability to leverage millions of users to teach different behaviours to the machine — comes in handy.

Seeing the emerging digital economy through the lens of data extractivism sheds light on many phenomena that have to date remained undertheorised and poorly understood. For example, the much-discussed problem of online distraction and fatigue generated by the use of Twitter or Facebook does not require the usual moralistic and paternalistic discourse of “the Internet is making us stupid and we should be responsible enough to disconnect.”

Rather, such fatigue can be explained as a natural consequence of data extractivist models adopted by the providers of the firms: they are the ones designed their systems to offer maximum distraction as that's the way to maximise the number of clicks — and hence data about us — that we produce on their sites. They keep drilling our psyches the way oil companies drill the ground — and data keeps flowing from our emotional reservoirs.

Digital Intermediation of Everything

Data extractivism has both political and economic consequences. On the political end, we see the narrowing of political opportunities that were supposed to widen up with the growth and broader distribution of digital technologies; ambitious political projects — be it the revamp of healthcare or education systems or of public administration systems — now increasingly require some form of intermediation by the providers of digital services. On the economic end, we see immense

wealth flow to just a handful of investors who were clever and quick enough to invest into this sector; the mechanics of this industry, however, is not necessarily favourable to helping to recover global economic growth — the concentration of data and, by extension, AI services in the hands of just several firms might result in them becoming the key gatekeepers (and, potentially rent-seeking gatekeepers) in the new digital economy.

The bargain of data extractivism was too hard for most politicians to resist: the ability to have advertisers pay for the provision of telecommunications services — be it search or email or even cross-indexing of academic papers through Google Scholar — has helped to alleviate the burden on the public purse while helping to claim that a new kind of gentler, more caring digital capitalism was in the making. As data extractivism branched out from relatively trivial services into areas like health, education, insurance, and so on — these are all data-intensive fields — there immediately emerged the prospect that, somehow, many of these activities, too, can be subsidised by whoever is interested in claiming ownership upon the data produced in those areas.

This explains how Google came to striking a deal with Britain's National Health Service, which allowed the search giant to deploy its AI service to examine health data belonging to several million patients — all in an effort to find early signs of kidney diseases. Given that governments themselves suffer from the effects of self-imposed austerity measures and thus have little money to spend on welfare, while technology firms have the data and the computing infrastructure that can yield some savings while providing useful services, it's very tempting to recruit these firms and build the foundations of a new, very different kind of welfare state: a digital welfare state where most services are to be provided by data-hungry digital giants.

Thus, we need to take stock of structural factors pushing governments and other public institutions into the hands of these big tech firms; the power balance of tomorrow's politics will favour private players over public ones in a way that has not been observed since the feudal era. It's to the technology firms — not our government — that we will run in case of, say, cyber-attacks. Ironically, of course, it's also because of the flaws in the software and digital systems designed by those very firms that those cyber-attacks become possible in the first place. The growing appeal of initiatives like the Digital Geneva Convention — touted by the likes of Microsoft as the right way forward in delegating even more power to the technology sector in dealing with problems like cyber security — indicate that the very possibility of intermediary institutions that reflect public — rather than private — interests is now facing extinction.

The New Algorithmic Consensus

Nowhere is this shift towards private power more evident than in recent efforts to diagnose and eliminate fake news. Here one can see how the dominant theme of the public debate on the issue has not been the question of why it is that clearly fake and erroneous news items circulate so widely but, rather, who it is that pays for their production and how it is that we can limit their impact. The former question can only be referenced with a reference to data extractivism: fake news items — which has existed as long as news has existed — now find greater circulation online because they fit very well into the click-obsessed business models of the data extractivist giants.

In other words, these items travel so fast because this is how Facebook and Twitter make money: an item shared by just a handful of people on Facebook might even cost Facebook money — it needs to be widely shared to be profitable. Under such conditions, the right way to liberate us from fake news seems evident: we need to liberate our communication networks from their reliance on data extractivism and base them on a different set of principles not rooted in the drive to harvest data, be it for advertising or AI.

However, since the question of “data extractivism” forms part of the broader political unconscious of the contemporary world, it hardly features in policy debates. Hence the preferred way to deal with fake news: to increase even further our reliance on the tech giants and given them even more power to identify what counts as fake and what counts as genuine and true. Of course, they can only do so by means of algorithms — even if they manage to recruit partners from the journalistic world to bolster the legitimacy of their efforts.

To have so much faith in the ability of algorithms — especially given everything we already know about just how useless they are at distinguishing photos of the works of art from, say, photos of pornography — is a sure path to cultural and political disaster. The false veneer of objectivity that every news item processed by Facebook will carry forever might, in the long run, do more damage to the public sphere than the current, largely chaotic approach.

Now, the struggle to eliminate and control fake news by means of a politically enforced algorithmic consensus is part of a broader effort to recruit predictive technologies — drawing on huge swathes of data already accumulated — in the name of control and policing. This is so regardless of whether we are policing the streets — as is the case with many “smart city” initiatives — or public discourse. But there's no consensus behind the new Algorithmic Consensus: the supposedly objective data that the algorithms are supposed to draw upon in deciding who counts as a terrorist threat, who might commit a crime in the next few months, who is to be released on parole from prison — all of these predictions feed on historical data which itself reflects existing and historical power imbalances. Thanks to Alphabet, Facebook, and their peers we now have the means to make such predictions on an industrial scale, making it even harder to redress the injustices that gave rise to skewed data sets to begin with.

Predatory Emancipation

Any effort to understand why the intensification of the regime of data extractivism has failed to generate widespread discontent has to grapple with the ideological allure of Silicon Valley. Here one can also detect a certain logic at play — a logic of what I call “predatory emancipation.” The paradox at the heart of this model is that we become more and more entangled into political and economic webs spun by these firms even as they deliver on a set of earlier emancipatory promises. They do offer us a modicum of freedom — but it only comes at the cost of greater slavery.

For example, we can gain free time — thanks to Google's virtual assistants, its ability to analyse our calendar and email and set up automatic reminders and appointment dates — but only at the cost

of completely surrendering ourselves and our data to its systems. In fact, Google's entire promise is constructed in these terms: we can only enjoy it to the fullest if we surrender to the fullest.

Emancipation, as long as it is conducted on terms established by Silicon Valley, is a never-ending process because every act of emancipation creates several new types of dependencies. And the reason why a bunch of information services is seen as a path towards emancipation itself has to do with the reframing of what it means to be free in the 21st century: by and large, this is a freedom to choose in the global marketplace rather than freedom to offend and provoke those in power.

We are moving towards the model of “benevolent feudalism” — where a number of big industrial and, in our case, post-industrial grants take on the responsibilities of care and welfare — that was postulated by some analysts at the beginning of the 20th century as the future of industrial capitalism as such. It took an extra century to arrive at this vision but any sober analysis of the current situation should dispense with the “benevolent” part of the term and engage much deeper with its “feudalism” part: just because power is exercised upon us differently than in the good old days when the capitalist mode of production ruled supreme and unchallenged does not mean that we are ever more emancipated. After all, plenty of slaveholders in the American South argued that slavery, too, was a much more humane system than capitalism. The humanity of the current system is, by and large, beyond the point; the real question animating our inquiry should be whether our growing reliance on big technology firms enhances or diminishes our autonomy — not just as consumers, but also as citizens. There's little evidence that this is the case even if the resourcefulness of our communication efforts — measured by the ability to reach thousands at low cost — has undoubtedly improved.

The End of History: The Techno-Utopian Edition

One unstated but supremely important factor in the swift rise of Big Tech has been the assumption that the old conflicts — driven by class struggle or concerns about inequality of access to property or natural resources — have become a thing of the past and that a new classless world was now in the making, not least thanks to digital technology. After all, a world where the world's largest companies are also our main welfare and security providers is a world that no longer believes in either the presence of competing ideologies or in the prospect of revolutionary change.

This might have been a convenient assumption in the early 1990s when the “end of history” rhetoric was riding high. Today, however, such assumptions increasingly ring shallow; there's no denial that the levels of economic inequality are historically high, that globalisation has not benefited everyone equally, and that the world is brimming with angry people who are using their right to vote to deliver their “guilty” verdict on the establishment. However, while the rise of populist rhetoric — to be followed, undoubtedly, by populist practice — is incontrovertible, it has not in any way dented the key assumption — about the benevolence of the tech sector — on which the contemporary ideology of “technological utopianism” rests.

In fact, we could even go further: it's the dominance of this techno-utopian ideology that allows for the overall project — the one that believes that we are living in a classless society and that big

conflicts over economic resources are a thing of the past — to go on unchallenged. After all, isn't it through technology — all hail the mobile phone! — that Africa is lifting itself out of poverty, joining the middle classes of the globe? Isn't it through technology — big data and algorithms to the rescue! — that those who were previously excluded from financial services can now get a loan — and on terms that would not feature extortion-like interest rates? Isn't it through technology that those parents who cannot afford a baby sitter can still gain an hour of free time — thanks to YouTube?

As long as the dominant cultural narrative holds technology to be the weapon of the weak and the poor — rather than the weapon aimed *at* the weak and the poor — there's little hope that phenomena like data extractivism can be given their due. Here, of course, the question is not so much about technology as such but, rather, about technology as currently wielded by the data extractivist industry. Attached to a different logic, technology very well might be an ally of the weak and the poor; the reason why techno-utopian narrative always wins is because it is good at presenting every critique of the commercial and social logic attached to digital technologies as an outright conservative attack on technology — and progress! — as such. Making it harder to make such accusations would be a good first step towards having a responsible adult conversation about building a world not reliant on data extractivism.

INTERVENTIONS

In this second part of the essay, I would like to propose and discuss some specific interventions to reverse or at least slow down some of the trends outlined above.

Breaking the Intellectual Monopoly of Big Tech

First of all, we must undermine the intellectual hegemony of Big Tech on how we think about the future politics and the role that technology would play in it. We must recover the idea of citizenship that transcends the image of us being nothing more than just passive app-consumers who are pliable subjects of the global advertising empire keen to accelerate data extractivism.

To do this, we will need to break the discursive and intellectual monopoly that tech firms have held over our political imagination. Such monopoly is currently maintained through generous funding to media (via various journalism initiatives), museums (via corporate subsidies and lavish grants for digitisation of their holdings), and think-tanks (to influence research on antitrust and monopoly). But it's also maintained in much softer ways — by leveraging the immense goodwill that their brands have generated. Silicon Valley's public success derives, in part, from its ability to draw on the legacy of both 1960s hippie/counterculture *and* the Cold War science culture that many in the original counterculture were setting out to oppose — they are both Berkeley and MIT.

There's much that scientists and artists can do together to unpick specific hegemonic discourses around technology, especially many of the currently fashionable ones around AI. Given that AI and similar services have dangerous and potentially lethal consequences — especially when deployed

for military purposes — we must remember that traditional discourses about the moral responsibility of scientists have not become obsolete. Artists, on the other hand, must do better not only in playing up the dystopian theme surrounding such but also articulating a vision of the future that does not default to frictionless consumerism.

Finding Power in The Digital

It's also essential that artists — perhaps, collaborating with sociologists and historians — reveal just how empty of power struggle most of the popular techno-utopian visions are. We just need to look at many of the surrounding promises to see that underlying them is a latent assumption about the overcoming of contradictions of capitalism and the irrelevance of concepts like class. But any decent analysis of automation would suggest that class divisions will not simply fade away because we'll all get access to the same automated technologies; poor quality automation for the poor, artisanal handicraft for the rich — this seems like a far more plausible future.

In addition, civil society needs to relearn how to attack the growing normalisation of immense power that private actors — above all, corporations — hold of society, with few or trivial consequences for freedom and autonomy. That today's feudalism comes in smiley faces does not absolve us from the responsibility to think about its impact on the weakest members of society.

Beyond “Freedom as a Service”

Two related ideas worth challenging also deserve some critical attention from civil society. One is the idea that all the negative features of the digital world — from the vulnerability of our computers to cyber-attacks to the sense of fatigue and distraction that many of us feel on using social media — are just the natural consequence of our own inability to control our desires. The other one is the idea that freedom is something to be sought and purchased in the marketplace, not the result of collective struggles in the political arena.

The former is a false narrative that takes little stock of the structural effects of data extractivism. To be preaching self-control in the face of data extractivism is like preaching entrepreneurship in the face of havoc wrought by neoliberalism: it's a way to reduce a collective and political problem down to the individual, consumerism-friendly level. Rupturing the excessively moralistic nature of the digital discourse while highlighting the deeply social and political nature of constraints in which individuals find themselves is a task that artists have traditionally excelled at; this could not and should not be just the preserve of sociologically-minded intellectuals.

There's a reason why the second type of narrative — that freedom is something to be harvested and provided for by our deeper/longer use of such services — holds so strong: it draws its strength from being rooted in a much earlier discourse of consumer sovereignty which lies at the very heart of the neoliberal project. The idea that markets provide us with a much better mechanism for exercising our freedom and individuality because our every “vote” counts and because, ultimately,

companies that serve us will be punished much sooner than the political parties might be based on false premises but it does enjoy growing political support.

The challenge is to rethink and recover a very different idea of freedom and its many cognate concepts, from autonomy to privacy. A privacy delivered to us via a cute app — at just 5 Euros a month! — is a very different kind of privacy from the one delivered to us via a system of constitutional rights. The former is “privacy as a service”; the latter is “privacy as a right”. The key to the hegemonic victory of Big Tech has been its ability to blur the distinction between the two and focus just on the underlying good: freedom, autonomy, privacy —while keeping mum on the roads that lead us there. Recovering those distinctions — if necessary by means of provocation — while also revealing the rather limited conception of “freedom as a service” is a very important undertaking.

Challenging the Positivism of the Algorithmic Consensus

It’s also important to be able to reclaim algorithms as something that can cause and foster trouble — not just help preempt and eliminate it (which is how they are used today in most predictive policing systems, for example). Artists have traditionally been good at it and we need to find a way to claim algorithms as something that can also lead to random, serendipitous outcomes — and destabilisation of power relations, not just their entrenchment. The growing rationalisation of everyday life requires such kind of playful, even subversive interventions; we might not be anywhere closed to Adorno’s dark vision of the “totally administered life” but the Internet of Things might get us there faster than we think.

Likewise, the turn to Big Data — and the underlying assumption that the bigger the dataset, the more truth it would yield — has surprisingly helped to recover many simplistic premises of positivism, along with their pernicious effects on how we think about knowledge. As a consequence, the current rehabilitation and subsequent formalisation of positivism into predictive systems is likely to recast long-standing cultural, racial, and ethnic biases as objective and empirical truths, baked into algorithms, and leading to even more discrimination. We already see such positivism run amok in predictive policing, facial recognition algorithms, travel databases in airports, etc. We will see even more of it once we deputise technology firms to control the flow of “fake news”: mechanical rubber-stamping of certain concepts as “true” or “false” is not likely to enhance anyone’s ability to see through propaganda. If anything, it may only blunt our ability to think critically about information that passed the algorithmic verification process. Fake consensus, enforced by algorithms drawing on flawed data, is every bit as bad as fake news.

On a broader level, we badly need to draw sharp distinctions between algorithms and the data that feeds them; we need to show that data is the Mechanical Turk in the algorithmic machine. Since much of deep learning (the method which underpins recent progress in AI) is still powered by historical data — and data, as any product of rational techniques of administration tends to incorporate, hide, and amplify biases — such revelations can help weaken the immense trust that most of us have in such seemingly objective systems. This is a formidable task for scientists, artists, and journalists to collaborate upon.

Towards a Different Global Village

Giving the rising nationalist tide across the world, it's also important to evaluate what to do about earlier, more utopian visions of constructing the global village. While those visions failed to deliver, there's much good that ought to be preserved and recovered here — above all, the spirit of internationalism. But if there's one lesson to be learned from the 1970s and 1980s it's that touting the benefits of intercultural communication in itself would not suffice. Previous failures to build a multipolar, truly internationalist world of equal and just information flows — starting with the efforts of the Non-Aligned Movement to create the New World Information and Communication Order and culminating in the early promise of cyberspace as a third space where the emergent global civil society could converse on issues of universal importance — have to be studied much more thoroughly, if only to learn from their mistakes.

There's no harm in acknowledging that the previous conceptualisations of the global village (including all those utopian visions inspired by the likes of Marshall McLuhan and Buckminster Fuller) failed to account for both corporate power as well as geopolitical and strategic interests of governments who did not want to lose their ability to engage in surveillance activities. No such global village can be built as long as sustainable infrastructures for communication and information exchange are missing — and they would require money to build. The good news is that today's global information infrastructure is much more granular than it was in the past, allowing for a modular approach to its reconstruction: once a credible architectural plan is in place, the hard work of building such a communications system can be shared across like-minded institutions, municipalities, and citizens alike. If successful, this vision can take care of the first need — infrastructure.

But what do we do about governments and their ever-growing needs of surveillance, justified by the seemingly permanent need for exceptional powers given the terrorist threat? Well, this is where cryptography can, in fact, contribute quite a bit: there's much to be admired in the work done by hackers and privacy activists over the last few decades in building tools that seek to guarantee that each of us can communicate freely and securely without fearing that our most intimate information would be intercepted by our own governments. Here, alas, the problems are also more financial than scientific: there's an abundance of solid academic work on cryptography and anonymity but not so much independent, no-strings-attached funding that would allow to actually build the much-needed systems to enhance anonymity.

This is where we need to be able to distinguish between realistic and naive techno-utopianism: there's not much that we can learn from the latter — it has failed so many times already — but the former can, in fact, provide an alluring vision, provided we are realistic enough to know where the real bottlenecks lie (oftentimes, it's in politics — not in technology per se). Recovering the role of technology as an emancipatory force that does not default to the neoliberal role attributed to it by Silicon Valley: this is, perhaps, the biggest contribution that civil society could make to today's digital debate.

*Note: the opinions expressed in the background papers are the responsibility of the author and do not necessarily reflect the official policy of the Council of Europe.