Saving the Digital Public Sphere From Silicon Valley
Evgeny Morozov,
Inspirationsession Archive Days (www.archiefdagen.nl) Monday 20ᵗʰ September 2021

Introduction by moderator Benno van Tilburg:
Technology giants have survived the recent "techlash" only to re-emerge as the true behemoths of the post-Covid era, ever more ubiquitous in our schools and universities, museums and libraries. Some admire the innovations they introduce; others fetishize the immense cost-savings that their often free products allow. Yet, it's time to ask: are we playing with fire? Have we wasted the past two decades, having outsourced our institutional and political imagination to Silicon Valley? This talk of Evgeney Morozov will go beyond a critical analysis of today's trends and will also outline what an alternative set of public infrastructures for the digital public sphere might look like.

Thank you so much Benno for this wonderful introduction. And it's a pleasure to be addressing this audience, even if just virtually. I hope that, in the months to come, we will finally go back to the more regular meetings in real life. I've enjoyed those quite a bit, being a frequent speaker, traveling and just getting feedback from people. But for now, I guess we'll stick with zoom.

My plan for today is to speak a little bit about some of the political, ideological, but especially ideational constraints, that Big Tech and especially Silicon Valley have created for the rest of society. Some of these constraints are real and structural; others exist only in our minds. It's due to these constraints that we are not able to articulate an ambitious vision for how digital technologists can be used in order to serve the broader public good. As a result, we fall into the trap of thinking that there is no real alternative to Google and Facebook – just like a couple of decades before, we fell into the trap of thinking that there is no alternative to the market.

The sooner we can get out of this ideational and political trap, the sooner we will be able to build the right kinds of public institutions, public interventions and public infrastructures. Only thusly we would be able to build a very different public sphere, a public sphere that will not be dominated by the desire to extract and monetize user data so to maximize the number of eyeballs that see a particular article.

Before we get to the crux of my argument, it's important to have this more positive vision in mind. We should resist the temptation to join some kind of technophobic political project that resists the march of digitization. By all means, and at all costs, we have to be strategic about our choices and preferences, and we have to be very careful not to reject digitization. As such, there are many useful things that we can do with digital technologies in terms: e.g. finding connections between different content types, making recommendations to users, helping them rediscover the passion for reading and watching and listening by tapping into new multimedia formats.

None of this will happen unless we understand that that there are different types and varieties of digitization. There's digitization imposed on us from above, by the interests of not just Silicon

Valley, but also by a number of big-pocketed financial firms that stand behind it. But there is also a very different kind of democratic, bottom-up digitization that we have witnessed, partially, in cities like Barcelona, for example, where a very progressive municipal government has actually managed to build a digital strategy favorable to the interest of the citizens. However, we have been to slow to take such experiences beyond the municipalist project, failing to introduce this more emancipatory logic of digitization into other walks of life, whether it's in our bureaucracy, or in the institutions of the welfare state, or in the public sphere at large. As I near the end of my talk, I will try to articulate and at least suggest what some of this interventions might be.

Let me start by recounting how I saw the danger posed by Big Tech back in 2013, almost 10 years ago. That's when I published my second book *To Save Everything, Click Here*, which, in essence, argued that there is a new ideology in the air in Silicon Valley. I was well-positioned to sense that ideology, because I was a visiting scholar at Stanford at the time. So I would regularly stroll around Silicon Valley itself, and I would walk around Stanford campus, and young people with a lot of zeal and passion in their eyes would approach me, asking me very bizarre questions like "Would you want test the app to eradicate malaria somewhere in Africa?" And I would say, "No, I'm not interested." And they were quite hurt on hearing this. They could not actually understand how somebody could say "no" to this world-changing app that they were building.

As I was spending time in the heart of the ideological enemy, I came to understand that Silicon Valley had a very peculiar plan for the rest of the planet; it was actively building a future for the rest of us. And it was a future in which Silicon Valley itself would be key to solving whatever problem stood in the way of this or that particular social group. In the case of Google and Facebook, these problems were initially identified in the realm of communications; our facts and social graphs were inefficiently and poorly organized – so the two companies would unleash services that would take care of that. Gradually, these firms started building infrastructures that allowed them to expand far beyond their initial roles.

To me, back in 2013, it seemed obvious that it's only a matter of time before they would be present in areas such as healthcare, education, transportation, food, insurance, banking: here was hardly a domain where some startup in Silicon Valley was not building some kind of an app that would, in their language, disrupt the way things worked in the past, revolutionizing it.

The promises sounded so far: digitization would make knowledge more accessible to everybody. And it would solve many of the problems that previously occurred due to the inefficiency of the existing providers of the service, many of them state-owned or at least self-regulated entities. And it was obvious, merely by looking and listening to those promises, that there was a heavy whiff of populism about them. In essence, they were arguing that our world was full of institutions with entrenched interests that stood in the way of the greater emancipation of citizen consumers.

These institutions were highly inefficient, suspicious, and corrupt; they were utilizing resources inefficiently. All in all, they had to be removed, or, to put it more euphemistically, as many people in Silicon Valley did– and still do today – they had to be disrupted. This wave of disruption that Silicon Valley's entrepreneurs and startups actually launched was supposed to rid the world of these moribund institutions and intermediaries, replacing them with swifter and nimbler for-profit players from Silicon Valley as the key providers of those services.

When you think of a company like Uber, for example, that logic was very explicit and obvious. The bad guys with the entrenched interests and the corrupt agendas were the taxi industry; the good guys were the founders and the executives behind Uber. That digital platform was just a neutral intermediary, which allowed ordinary people who owned cars to drive for other ordinary people who needed a ride. Yes, they would take a small commission but they would nonetheless create a lot of value in the process. This became the paradigmatic example of how this wave of disruption and digitization would change the world.

A company like Uber was to result in lower costs for passengers and in more jobs for drivers; it would bring more efficiency, making the world more sustainable and greener. Given the dangers of climate change, we should just be utilizing all of the existing resources more fully and in a smarter and intelligent way. So instead of having a car that we use only three hours per day, we should think of a society with fewer cars that would be ridden in perpetuity This, at any rate, is the logic behind the "perpetual ride" – an idea with which Uber once tried to justify its activities. I can think of many other examples beyond just the world of transportation, from Airbnb in the world of accommodation to Robinhood in the world of trading.

Soon, there popped up European, and Asian and Chinese equivalents of many of these platforms, competing with the Americans for the coveted market share. What was common to all of them was this idea that, somehow, merely by digitizing thing and following this recipe of Silicon Valley, we will be able to derive tremendous efficiencies and savings from them. It was a new political-economic model of how to run society, delivered to us courtesy of start-ups and angel investors.

Looking at these tempting promises, a lot of politicians reasoned that this is exactly how the future of the welfare state and the public sphere would look like. Why should we be spending taxpayers money, of which we have less and less, on the provision of services that Silicon Valley can provide for free? A lot of governments found this tremendously appealing. And that explains this bizarre pilgrimage that a lot of European politicians have made and still make to Silicon Valley in order to meet and greet the likes of Mark Zuckerberg, shake the hands of Larry Page and Sergey Brin, to be seen next to Elon Musk.

This fascination with Silicon Valley is still there, despite all of the so-called techlash of the last few years. Silicon Valley is still seen as this magical industry that can take entrenched, corrupt, outdated, dinosaur-like institutions, disrupt them and replace them with something efficient, that saves money or even earns money.

And since this narrative was swallowed, almost in its entirety, by our politicians, our public policy in this space became such that we could not actually think of what else one culd be doing in the public sphere at large in order to enable different kinds of digitization. This is how digitization delivered on terms that were not the terms of Silicon Valley and its startups became unthinkable.

As a result, the only possible thing you could do was to encourage more startups to essentially get going. Those startups that somehow outstripped the existing legal boundaries – because, for example, they did not protect the user data or engaged in some dubious practices for foreign governments – we would strictly regulate them. The dominant approach became that of the market: you delegate everything to the startups, you build innovation hubs, incubators, startup accelerators, you do everything to make sure that there are more and more startups out there that are not only somehow filling in the holes that have emerged in the edifice of the welfare state, but that they are also contributing to the image of the country as a highly innovative place that does excellent things.

There was a crucial flaw in this theory. And it had to do with the way in which all of the initial efficiency gains and presumed savings of the likes of Uber were generated to begin with. The fairy tale that emerged from the garages of Palo Alto and Silicon Valley in general was that the reason why a service like Uber costs so little compared to an average taxi is because it's simply more efficient – and it's more efficient because it's digital. On the surface, it also made the tech industry look more more progressive, because they found a way to build things and offer services in a way that was just better and cheaper than the older taxi industry used to do. If this were the case, it certainly made the case to apply the logic of disruption to society at large.

If you look at that case closely, however, you will soon discover that behind companies like Uber stood these massive founders, e.g. sovereign wealth funds, high net worth individuals, corporations, from all sorts of industries who, after the financial crisis, were sitting on a lot of cash. They couldn't really invest in much else because the interest rates were so low. Digital platforms like Uber became magnets for these companies. Idle funds, idle pools of cash that accumulated in the system, especially cash coming from the Gulf states, like both the sovereign wealth funds of Abu Dhabi and the United Arab Emirates in general, but also those of Saudi Arabia.

What happened is that in their quest to expand and quash competition, Uber began incurring massive losses. For a very long time, Uber has been a loss-making company. There was hardly a year since its inception that it actually made a profit. And the losses that they incurred were mind-boggling and ran in the amounts of billions of euros a year. Of course, if you can incur a loss of 5 or 8 billion euros a year, it's obvious that you can also do all sorts of things that no normal company can; you can also sustain losses when others would have no choice but to shut down operations. So in the case of Uber, all these efficiencies that were celebrated as the miracle of digitization, they were mostly the result of very clever business strategy of incurring giant losses while their competitors would gradually get out of business. Politicians gave a very

different spin to this, presenting Uber's low prices as the culmination of some kind of digital innovation and inventiveness.

I'm going to such a great length through these examples, only to show you that the highly ideological story that Silicon Valley sells us about innovation and efficiency is often not based in reality. Rather, it's based on a highly misleading rhetoric which borders on propaganda. We should try to see through it, locating the rise of this industry in a very peculiar historical situation – a very particular historical conjuncture, which has to do, as I've already mentioned briefly, with the aftermath of the financial crisis.

As it hit, our own governments needed a savior. They needed a way to buy time. Before the crisis, they did it but encouraging us speculate on housing prices; for some time, that did create some kind of prosperity. They, they tried to do the same with Big Tech. Thus, now we're essentially promised that if we only put all of our savings into buying the shares of Google, Apple, Facebook, Microsoft and others, we will become millionaires and all of our financial concerns will be taken care of. And that's where we should put all of our pensions, etc. Essentially, it's a very similar proposition to the one made before the financial crisis with regards to the housing bubble, with the government expecting (rather naively) that the tech industry will be able to resolve many of the contradictions that have accumulated in our economic system in general.

For a moment, let's go back to Silicon Valley. What I saw at Stanford, in 2013, was not just confusion and a desire to help humanity; no, it was something like a coherent ideology – an ideology that I would later come to christen "solutionism." I introduced this notion to discuss the tendency by digital firms and startups to present themselves as essential to the process of solving some of the world's greatest problems. These problems could have to do with climate change, they could have to do with poverty, they could have to do with inequality: there is no limit to the ambitions of solutionists.

Eight years have now passed since I introduced that term into public discourse. And it became clear to me that there are also other readings or connotations of solutionism, which I did not account for at the beginning. So we have to think a little bit broader, about the intellectual impact of this ideology. It's not just about the centrality of these companies to resolving some of the world's greatest problems. It's certainly a huge part of it but it does not exhaust its. As an ideology, solutionism prevents non-commercial and non-capitalist actors – think about individual citizens, non-profits, public institutions –from fully utilizing the immense potential opened up by digital technologies in their own projects. Instead, it seeks to channel all this immense creative activity into for-profit, market-friendly activities.

Today, solutionism is an ideology, which insists that the only way to take full advantage of the promises of digitization is to try to fit everything into the digital straightjacket of the startup. As a result, no other actors with motivations different from those of the profit-seeking firms would be allowed in. They won't be allowed to deploy sensors, artificial intelligence, cloud computing, all sorts of other digital infrastructures for their own purposes and goals, unless they accept the

challenge of being a startup, seeking venture capital funding, hiring expensive developers, and finding a business model that will eventually pay off the investments for venture capitalists. This is happening not because of some backroom deals and maneuvering by the Silicon Valley giants. It's happening as a result of the way that our digital infrastructures have come to be owned, operated and operated.

Think about artificial intelligence, for example. Today it exists as a resource that is heavily concentrated inside the tech firms. If you want to use artificial intelligence today, you probably need to go to Amazon Web Services or Google Cloud. And you will essentially need to purchase those services from them – and it's not cheap. There is no conception whatsoever that artificial intelligence might be a public good – that it's something you could tap into to as an NGO, or as a library or as a city as a group of concerned citizens, or as any other civic association. Of course, you might apply for a waiver of the fee – and our new digital humanitarians in Silicon Valley would be happy to grant you an exception. But, from the very beginning, artificial intelligence is conceived as a commodity, or, to put it more accurately, it's conceived as a pay-for service.

The same applies to cloud computing industry, voice and face recognition, and many other types of advanced data processing: today, they are all offered as pay-for commercial services. This, in turn, imposes invisible constraints on any potential institutions that might emerge to take fuller advantage of these technologies.

So imagine if I have an ingenious idea of how people can learn a new foreign language. To turn this idea into something useful, I do require some of these advanced tools such as artificial intelligence, in order to analyze data in real time, or to translate from one language to another or to make up sentences with words that students need to practice.

Now, as somebody who has an ingenious idea that might truly have transformative implications on the world (as it will allow people to learn languages faster and more efficiently), in today's solutionist universe, I have no choice but to turn myself into a startup, go raise money, and then use the money to pay for services from Google, Amazon, Microsoft and others. And I would also be able to pay the extremely high salaries of people who know how to use those tools, because their salaries have now also become artificially inflated, due to all the idle cash that comes from the likes of Saudi Arabia and others, that pours into the tech sector.

Thus, under today's conditions, any promising idea with truly transformative implications has few chances of survival unless it conceives itself as a startup-in-waiting. Then, as you become a startup, you attract money from venture capitalists. And as you do that, it imposes further constraints on what you can and cannot do, because you need to eventually repay that money. And to repay that money, you need to have a business model. And the business model, in most cases, will not be pretty or ethical. If you really want to make money that you can repay to venture capitalists, that model will most likely not be based on subscriptions, which can give some – but not a lot – of money. Most likely, it will be based on extracting data from the users, analyzing it, selling it to advertisers or finding some other way to monetize that data.

Thus, in the absence of alternative infrastructures that we can tap into in order to realize the true potential of the digital, we end up reinforcing the solutionist ideology, which posits that the right solution to every single problem in the world is to be provided by the startups. The lesson to draw here is simple: unless we develop this alternative modes, methods and models, we would be forever trapped in the hands of Silicon Valley. And we would not be able to develop a truly digital public sphere that would feature the decentralized web of entities and organizations that are helping people to resolve the problems that they have in the most innovative and transformative manner.

What, then is the opposite of solutionism? The opposite of solutionism is a politico-economic system, whereby, regardless of your motivation, regardless of the institutional form that supports it, regardless of whether you're seeking a profit or not, you will be able to take full advantage of digital technologies. This would be a world where every creative impulse that we have – regardless of whether it's motivated by survival instinct or the feelings of altruism or solidarity – could be channeled and amplified in such a way as to produce most benefits for humanity – and not only for the app-funding venture capitalists.

The big ideological success of Silicon Valley has been in convincing all of us that we have to rely on them, we have to rely on these big technology corporations, in order to realize all the good things that come from digitizing the world and digitizing knowledge and information. If you think about something like Google Maps, clearly, there is an extremely useful dimension to it, especially as it pertains to the aggregation of real-time traffic information. Of course, you do want to know what the situation on the road is like, so that you can choose the right route, the one that makes the most sense, given your objectives and other criteria.

Of course, all of that information is not generated deliberately, because all of us want to share this information with each other. Instead, we have these passive sensors – including our smartphones – in our cars. The information that emerges from that sensory system is then aggregated and packaged in such a way as to make life better for everyone.

There is no reason for us to associate the provision of this feedback infrastructure by Google Maps with the startup, the corporation, the price system, the market form, and capitalism in the abstract. What I do in my car or with my phone often has nothing to do with the market; the way in which the sensory system operates also has nothing to do with the market. It just happens that following the Cold War, and following our fetishization of the market, companies like Google managed to move in and snap up former government contractors that had developed a lot of these maps (and funded with public money to boot). And then they integrated them into Google. Right.

The task of our post-solutionist imagination today is to disentangle all these cases, and to say: "Okay, this is a fantastic feedback infrastructure built into Google Maps. But it has nothing to do with Silicon Valley, venture capitalists, or Google. We can take this infrastructure and put it somewhere else, we can have it run as a public service, we can have it do all sorts of things."

The same logic, the same kind of imaginative questioning, we should be able to apply to many other technologists.

An example I often give is that of internet search. Here, too, we have unfortunately accepted the premise that two researchers at Stanford, Larry Page and Sergey Brin, made in the late 1990s. The initial assumption – not a very bad one – was that looking at citation patterns across websites would be a decent proxy for the quality of the content. There was, however, a second assumption, which attracted far less scrutiny, namely that the process for looking-up these citation patterns has to be heavily centralized inside a single corporate player. So the highly profitable bet made by Page and Brin was that you will have one company, that company will index everything, and then they will give you answers as you search. And the way to improve it will be by making the system more and more comprehensive. So you will be essentially introducing more and more data streams into Google's index.

I'm not at all sure that this second assumption is actually accurate. It's something that we should be able to verify empirically. It is not obvious to me that a company situated in Mountain View is best positioned to take care of my own information needs here in the local Calabrian village in the south of Italy. There's already quite a lot of local knowledge that's already concentrated here in the village; it's not really being utilized at all – or at least it's not being utilized until everyone buys into the Google universe, starts setting up their own web-sites and linking to other pages, writing Google reviews of restaurants, and what not. Is this really the best way to leverage local knowledge? I'm not so sure

Due to the heavy influence of solutionist ideology, we have internalized this idea that, somehow, the problem of internet search has been solved. So we don't have to think about alternative institutional arrangements. We don't even bother trying to imagine how a different configuration of sensors, data storage, indexing, and citation analysis might result in an information system that would not be global in its aspirations, but would serve the local needs of citizens and users much better than Google. And if we really believe that this local system might be a better option in the long term, then the goal of our innovation and digital policy should be to facilitate the emergence of such experiments, rather than to continue thinking that the problem of internet search has been solved and ll that remains for a truly radical political project is to somehow find a way for Google to pay taxes.

Here, a lot of people on the left are handicapped, because, following some of the tenets of historical materialism, they do tend to believe that capitalism does result in more efficient technologies, even if taking them to the next stage would require a socialist revolution. So, capitalism gives us Google 1.0; socialism would give us Google 2.0. But why should we be so dependent on the fruits of capitalism – and who said they are good to being with? All of this seems to require a very big leap of faith.

It's only by recovering this imaginative zeal for building alternative institutions that we will be able to show that a different kind of digitization is possible. And in order to do that, we need to free ourselves from this belief and assumption that our only saviors exist in Silicon Valley or, for

that matter, in Beijing and Shanghai. We in Europe cannot simply count on pitting them against each other, in order to see which one of them would offer us the best and cheapest deal.

What we should be doing instead is finding ways to free ourselves from the invisible shackles that our experience with Silicon Valley and its platforms has imposed on our imagination. Instead, we have to be building policies in such a way as to increase and accelerate the degree of experimentalism that is present in our society, and in our political system at large. I've been talking almost for 40 minutes now, maybe I will finish here. And I'll be more than happy to engage in dialogue and to answer any of your questions. Thank you very much.

Evgeny Morozov is the author of The Net Delusion (2011) and To Save Everything, Click Here (2013). He is also the founder and the publisher of The Syllabus, a knowledge curation initiative. He holds a PhD in History of Science from Harvard University. He has been a visiting scholar at Stanford and Georgetown and has published in The New Yorker, The New York Times, Financial Times, and many other publications. His upcoming book is Freedom as a Service (2022).