


Spotlight

Let's ask more of AI



[Stefania Druga](#) from Romania teaches artificial intelligence (AI) programming to children. As a researcher, she has also studied how 450 children in

seven countries [interact with](#) and perceive connected toys and home assistants, like Amazon Alexa or Google Home.

Children can understand more than parents think, she says – including that machine learning is limited by what training data you have to work with.

The philosophy behind [the software she developed for teaching](#) is that if children are given the opportunity for agency in their relationship with “smart” technologies, they can actively decide how they would like them to behave. Children gather data and teach their computers.

This simple approach is what we urgently need to replicate in other realms of society.

In order to navigate what implications AI has for humanity, we need to understand it – and then decide what we want it to do. Use of AI is skyrocketing (for fun, as well as for governance, military and business) and not nearly enough attention is paid to the associated risks.

“Yup, it’s probably AI,” says Karen Hao’s back of the-envelope-explainer about any technologies that can listen, speak, read, move and reason. Without necessarily being aware of it, anybody who uses the internet today is already interacting with some form of AI automation.

Thought of simply, machine learning and AI technologies are just the next generation of computing. They enable more powerful automation, prediction and personalization.

These technologies represent such a fundamental shift in what is possible with networked computers that they will soon likely make even more headway into our lives.

Whether search engine results, music playlists, or map navigation routes, these processes are far from magical. Humans code “algorithms” which are basically formulas that decide how decisions should be automated based on whatever data is fed into them.

Where it begins to *feel* magical is when it makes new things possible. [This Person Does Not Exist](#) is a good example. If you visit the website and refresh the page, you will be shown an endless array of faces of people who never existed. They are images that are generated at random by a machine learning algorithm based on a database of faces that *do* exist.

Look closely, and you [will spot the errors](#) – ears that are crooked, hair that doesn’t fall naturally, backgrounds that are blurred. [This Cat Does Not Exist](#) is less convincing. The potential exists for either photo generator to improve with additional data and guidance. And the risks that such photos

could be used to misrepresent reality also exists, even for such whimsical creations.

In recognition of the dangers of malicious applications of a similar technology, researchers from [OpenAI](#) sparked a media storm by announcing they would not release the full version of an AI technology that can automatically write realistic texts, based partly on the content of 8 million web pages. “Due to our concerns about malicious applications of the technology, we are not releasing the trained model,” [they wrote](#), calling it an experiment in “responsible disclosure”.

Such recognition of the faultlines and risks for abuse of AI technologies is too rare. Over the last 10 years, the same large tech companies that control social media and e-commerce, in both the United States and China, have helped shape the AI agenda. Through their ability to gather [huge quantities of training data](#), they can develop even more powerful technology. And they do it at a breakneck pace that seems incompatible with real care for the potential harms and externalities.

Amazon, Microsoft and others have forged ahead with direct sales of facial recognition technology to law enforcement and immigration authorities, even though troubling inaccuracies and serious risks to people of color in the United States have been [rigorously documented](#) and [defended](#). Within major internet companies that develop AI technologies, including Amazon and Google, [employees have sounded alarms](#) over ethical concerns more urgently.

Company leaders deflect with confidence in their business models, hubris about their accuracy, and what appears to be ignorance or lack of care for the huge risks. Several companies, including Axxon, Salesforce, and Facebook, have sought to allay concerns over controversies by creating ethics boards that are meant to oversee decisions.

Co-founder of the research institute, [AI Now](#), Meredith Whittaker, calls this “ethics theater” and says there is no evidence that product decisions are run by them, or that they have any actual veto power. In an interview with Recode, Whittaker [asked of the companies](#), “Are you going to harm humanity and, specifically, historically marginalized populations, or are you going to sort of get your act together and make some significant structural changes to ensure that what you create is safe and not harmful?”

As it happens, Google’s announcement of an ethics board [backfired spectacularly](#) in April 2019 and was dismantled after employee protests and pub-

lic outrage about [who had \(and hadn't\) been asked to join](#). While the company has been vocal about establishing [principles for AI](#), and has engaged in [social good projects](#), it also has competing priorities across its many ventures.

What are real world ethical challenges these boards could tackle if they took Whittaker's advice? One idea would be to question an everyday function billions of people are affected by. Google's video platform, YouTube, is often said to be a "rabbit hole" – endless tunnels leading from one video to another. Though [YouTube denies it](#), research shows that content recommendation algorithms are fueling [a crisis of disinformation](#) and cultish behavior about vaccines, cancer, gender discrimination, terrorism, conspiracy theories and [add your topic].

Similarly, Pinterest and Amazon are also platforms that drive engagement by learning and suggesting new and engaging content. They experience [variations](#) of the same problem. In response to public scandals, they have each announced efforts to stop anti-vaccine content, but there is little evidence of any real [change in the basic intention](#) or function of these systems.

But it's not just technology companies that need to be interrogating the ethics of how they use AI. It's everyone, from city and government agencies to banks and insurers.

At the borders of nine European Union countries, an [AI lie detector was tested](#) to screen travelers. Systems to determine creditworthiness are being rolled out to populations in emerging markets in Africa and Asia. In the United States, [health insurers are accessing social media data](#) to help inform decisions about who should have access to what health care. AI has even been used to decide [who should and shouldn't be kept in prison](#) in the United States.

Are these implementations of AI ethical? Do they respect [human rights](#)? China, famously, has begun scoring citizens [through a social credit system](#). Chinese authorities are now also systematically [targeting an oppressed minority](#) through surveillance with facial recognition systems.

Where do we draw the line?

There are basically two distinct challenges for the world right now. We need to fix what we know we are doing wrong. And we need to decide what it even means for AI to be good.

Cutting humans out of government and business processes can make them more efficient and save costs, but sometimes too much is lost in the bargain.

Too rarely, do people ask, should we do this? Does it even work? It's worth questioning whether AI should [ever be used to make predictions](#), or whether we should so freely allow it into our homes.

Some of the most worst missteps have involved training data that is faulty or simply used with no recognition of the serious biases that influenced its collection and analysis.

For instance, some automated systems that screen job applicants consistently give women negative scores, because the data shows it's a field currently dominated by men.

"The categories of data collection matter deeply, especially when dividing people into groups," say the authors of the book [Data Feminism](#), which explores how data-driven decisions will only [amplify inequality](#) unless conscious steps are taken to mitigate the risks.

It seems that if we leave it up [to the nine big companies](#) that dominate the field of AI alone, we raise the spectre of a corporate controlled world of surveillance and conformity – especially so long as gender, ethnic and global diversity is also lacking among their ranks of employees at all levels of a company. Having engineers, ethicists and human rights experts address collaboratively how AI *should* work increases the chance for better outcomes for humanity.

We are merely at the beginning of articulating a clear and compelling narrative of the future we want.

Over the past years, a movement to better understand the challenges that AI presents to the world has begun to take root. Digital rights specialists, technologists, journalists and [researchers around the globe](#) have in different ways urged companies, governments, military and law enforcement agencies to acknowledge the ethical quandaries, inaccuracies and risks.

Each and everyone of us who cares about the health of the internet – we need to scale up our understanding of AI. It is being woven into nearly every kind of digital product and is being applied to more and more decisions that affect people around the world. For our common understanding to evolve, we need to share what we learn. In classrooms, Stefania Druga is making a small dent by working with [groups of children](#). In Finland, a grand initiative

sought to train 1% of the country's population (55,000 people) in the [elements of AI](#). What will you do?


► Further reading

- Situating Methods in the Magic of Big Data and Artificial Intelligence, danah boyd, M.C. Elish, Communication Monographs, 2017. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3040201
- AI Now 2018 Report, AI Now Institute, December 2018. https://ainowinstitute.org/AI_Now_2018_Report.pdf
- Data Feminism, Catherine D'Ignazio, Lauren Klein, MIT Press Open, January 2019. <https://bookbook.pubpub.org/data-feminism>
- Anatomy of an AI system, Kate Crawford and Vladan Joler, SHARE Lab and AI Now Institute, 2018. <https://anatomyof.ai/>

► Further listening

- RecodeDecode podcast: Meredith Whittaker and Kate Crawford: How AI could change your life, April 2019. <https://www.recode.net/podcasts/2019/4/8/18299736/artificial-intelligence-ai-meredith-whittaker-kate-crawford-kara-swisher-decode-podcast-interview>

The power of cities



When the Amazon Kindle was released, their ebooks didn't work with commonly used screen readers, making accessibility difficult for the blind community. [The National Federation of the Blind \(NFB\)](#) in the United States campaigned to change this for years, in vain. Then Amazon won a \$30 million USD contract with the New York City Department of Education in 2015 to create an ebook store for educators in 1,800 schools. City schools [delayed a final vote](#) until Amazon and the NFB came to an understanding. Since then, the Kindle now has a built in screen reader and Amazon has improved accessibility across many products.

This is an example of how cities have huge potential power to improve the health of the internet ecosystem. In this case, it was a win for children and

educators in New York, but also for people around the world. Where consumers may have a hard time persuading giant corporations to do something that they perceive as going against their business interests, a million dollar procurement contract and a commitment to serving the public interest can help.

More than half of the people in the world [now live in cities](#) and by 2050 that number is expected to rise to 68 %. Cities are where wealth and power is concentrated in most countries, and also where many technology initiatives are rolled out and tested in communities. What we may think of as local decisions today, may be of global consequence in the future.

When the Federal Communications Commission (FCC) of the United States backed away from protecting net neutrality in 2018, [a network of city mayors formed](#) to use their combined purchasing power to support internet providers who continued upholding net neutrality.

“In NYC alone, we spend over \$600 million annually to provide internet service to city employees and to offer city services. So, we convened an ad hoc coalition, starting with eight cities committed to only purchasing from broadband providers that honor net neutrality principles. Now, this coalition is over 130 cities,” says Max Sevilia, the Director of External Affairs for the NYC Mayor’s Office of the Chief Technology Officer.

This story and many others are highlighted in a publication called the [New York City Internet Health Report](#). Its creator, Meghan McDermott, adapted the format of the global Internet Health Report as part of a Mozilla fellowship project to explore among other things how cities can be strong advocates for digital rights by nurturing relationships with [civic tech](#) communities.

“The core of the digital rights agenda is to reframe how we think about and deploy technology in cities. The idea is to recapture the dignity and purpose of technology as a public good,” says McDermott, who has worked at the intersection of education and digital rights for years – formerly as director of strategy for Mozilla’s [Hive Learning Networks](#), a peer community for digital literacy.

When the internet and connected devices are applied to solving problems in cities, it tends to be referred to as a ‘smart city’ initiative. These are often projects to improve the efficiency of energy, transportation or any number of government services. For instance, it could be trash cans with sensors that alert waste management authorities when they need emptying,

or parking meters that can help people find free parking spaces in crowded streets.

Such futuristic ideas have excited city officials around the world, and the global market for ‘smart city’ technologies is worth [hundreds of billions of dollars](#) and growing. But frankly it’s also an industry where corporate interests and techno-utopianism holds high currency – where [flying taxis and autonomous helicopters](#) end up described as a solution to traffic congestion, even though they most likely won’t solve anything for people who rely on public transport.

The harshest critics say a hype about ‘smart cities’ has led to massive investments in what is essentially surveillance technology under the guise of technological progress. In both resource rich and poor cities, there are cameras, sensors, microphones, and huge multi-year procurement contracts with companies that have questionable data practices. In this way, with scant attention to data privacy, the internet has arrived to cities worldwide, for better or worse.

Where some see an opportunity [to entirely rethink how cities collect data](#) about neighborhoods to improve services, others see a lack of transparency and [a recipe for a civil rights disaster](#) spurred on by corporate interests. Where some see energy efficient LED streetlights that help gather data about pedestrians with cameras, others see [a surveillance dragnet](#) encroaching on [freedom in public space](#) and putting [vulnerable populations at risk](#). Time and again, there are design choices that could be made to minimize the risk of abuse. For instance, when could it be preferable for privacy to use a thermal sensor to collect crowd data instead of a camera?

Digital rights advocates are cast as enemies to progress in such conflicts, but it really boils down to a core difference in opinion about whose interests technology should serve, how to seed [social innovation](#), and what data should be used (or not) in the public interest.

Consider the [electronic sensors in the garbage cans](#). To some, that’s a great example of how technology can help cities operate more efficiently. To others, like Tamas Erkelens who is the program manager of data innovation in the mayor’s office of the City of Amsterdam, it’s evidence of a wasteful approach that characterizes many ‘smart city’ innovations.

“We wouldn’t need sensors in every trash can if cities could have Google Map data to see where crowds are,” says Erkelens. “Wherever people are convening is a good enough indicator of where there is likely to be more trash.

We can then use sensors just to train the models, rather than to create new data by machines with batteries that need to be changed,” he says.

Many city governments and [open data advocates worldwide](#) peer enviously at the wealth of data held by internet companies like Google, Uber, Apple and Airbnb knowing that it could help them understand crucial things about [traffic](#), housing and employment. In 2018, the Open Data Institute in the United Kingdom [published a report](#) suggesting that mapping data companies should be [compelled to share geospatial data with rival firms and the public sector](#), to stop “data monopolies” from forming and to create better opportunities for innovation.

Some companies do share aggregated data with city planners, [including Uber](#), but cities are also getting smarter about requesting things like [usage data of electric scooters](#) upfront as a condition of doing business. The city of Barcelona is one of very few cities that operates under the principle that *all* data collected in the duty of local government in public space must be [available in a data commons platform](#). Erkelens says Amsterdam is using its annual procurement budget of € 2.1 billion to help guarantee good terms for data privacy too, and that Barcelona and Amsterdam together are [experimenting](#) with partners in the European Union to develop new technologies that also give citizens more direct control over their own data.

At the [Smart Cities Expo World Congress in Barcelona](#) in November 2018, the chief technology officers of Amsterdam, Barcelona and New York together [launched](#) the [Cities Coalition for Digital Rights](#) in partnership with UN-Habitat, a United Nations program to support urban development. Cities who join the coalition agree to a declaration of just five principles that center on respect for privacy and human rights in use of the internet. They pledged to see 100 cities join in 100 days (before July) and 35 cities [have joined](#) so far. Declarations may come and go, but these cities aim to sow the seeds of a movement whereby cities decisively claim digital rights. By working together and establishing best practices they will attempt to win a race against technological progress that is not centered on principles of human dignity and inclusivity.

Despite the strong stances taken in [New York](#), [Barcelona](#) and [Amsterdam](#), people who do digital rights work at the city level describe an uphill battle of culture change within large and in some parts traditional institutions with multiple agencies and divergent interests. Creating the [policies and processes](#) by which all agencies can make better decisions about privacy, data

and transparency – and [opening up key parts of the work](#) to civil society – is a key part of the challenge.

This is where the civic tech community has blossomed in countless cities. Diverse groups of public interest startups, technical students, officials, and engaged citizens team up to hack bureaucracy and code in an attempt to make cities more responsive to their residents. They work from the inside with willing partners, and from the outside through advocacy groups, research, and [live prototypes](#) that reimagine how more responsive systems could work.

Cities worldwide are on the frontline of decisions that affect the health of the internet for all people. At the local level, whether in rural or urban communities, there are opportunities for civic engagement regarding the internet that can be more direct than at the national level. We should seize opportunities to influence how technology is used (or isn't) in our own communities, and encourage elected officials to be champions of digital rights. The more [engaged](#) we are locally, the more empowered cities will be to cast themselves as opponents to internet policies at the national or international level when they go against the interests of people.

The challenge for cities is to advance the intentional adoption of digital tools that advance values of diversity, inclusion and fairness that they already hold, rather than jumping on the latest 'smart city' trend.

When he helped facilitate conversations between Amazon and The National Federation of the Blind over ebooks, Walei Sabry in New York already worked in the Mayor's Office of People with Disabilities. Since then he has also become New York City's first official digital accessibility coordinator. About 'smart cities' he [says](#), "These initiatives can go really well or really wrong depending on who's at the table – people with disabilities must be involved at all stages of the process ... because what works for us makes products better for everyone."

Rethinking digital ads



When dozens of people fell gravely ill from eating romaine lettuce in 2018, public health authorities in the United States and Canada could not figure out where the *E. coli* contaminated leaves were farmed. The lettuce had changed hands so many times from washing, chopping, packing to shelving that they couldn't retrace the steps. The only option was to temporarily declare *all* romaine lettuce, from any source, [unsafe](#).

It's a stretch of the imagination, but let's compare that to what we are experiencing in the world of "personalised" or "targeted" digital ads.

We have absolutely no idea of the ingredients that go into the daily bread of the internet. The ads we are served as we use mobile apps and browse the Web are like lettuce leaves scattered over the planet – they can be healthy –

but information about the supply chain is muddled and we have no way to understand what is happening.

Pretty much everything we do when we interact with the internet can be tracked by someone (or something) without our knowledge. From the websites we visit, to the apps on our phones, to the things we write in emails or say to voice assistants. We have no way of knowing how this big salad of data may be combined by different companies with information that uniquely identifies us.

It appears that collecting data about everything and anything we do is of commercial interest to *someone*, whether app developers, insurance agents, [data brokers](#), hackers or scammers. The lines have been blurred between what's public and private information. Your credit card [may share a list of what you buy in stores](#) with Google. Your online dating profile has perhaps been copied and resold. Why is this?

Not all data about you is used to sell ads, but it is primarily because of the ad-driven internet economy that data has become such a hot commodity. It is why people now speak of [surveillance capitalism](#) and the attention economy. The phrase "You are the product" precedes the internet, but has gained new currency as a way to explain how so much online can be "free". Personal data may seem like a small price to pay. But the added social tax is now [mounting threats to freedom and human rights](#).

To talk about the positives: Digital ads have been a boon to the global economy. Free online services have driven the uptake of mobile internet around the world. Ads have helped publishers and startups monetize their online content and services.

For some of the most powerful companies of the internet, Google, Facebook and Baidu, ads are a [primary source of revenue](#) even as they have expanded their business into multiple directions and geographies. For Google and Facebook, especially, access to data is a source of global market power and [leverage for business negotiations](#). For the first time, in the United States, [digital ad spending is bigger than](#) for print and television.

The ad-tech industry is vast, but by some estimates Facebook and Google alone [controlled around 84% of the global digital ad market](#) in 2018 outside of China. To succeed, they have developed product design practices which are centered on [holding the user's attention](#) and [maximizing engagement](#) to drive revenue from ads.

Targeted ads for the most part promote run of the mill products and services, but these same tools can just as easily be exploited by people with crim-

inal or hateful intentions. In a few minutes, you can place content on videos in YouTube, news feeds of Twitter and Facebook and search results of Google. By selecting what demographic to target, advertisers on some platforms have been spotted excluding people of a certain race or gender [from housing or job ads](#). Or in the case of Facebook, even directly targeting “affinity groups” like “Jew Haters” (yes, really). Facebook said its categories are created by algorithms, [and when confronted](#) said it would make changes, but it begs the question of how much data should be collected and what it should ever be used for.

Your data profile is a sandwich of data that you [knowingly or unknowingly share](#), which is interpreted by secret algorithms that make use of statistical correlations. For instance, searching online for “loan payment” might say something about your finances. And if you “like” articles or join Facebook groups that could help define your affinities.

“Ads can be done in a more privacy friendly way. But publicly-traded corporations have a duty to maximize shareholder profits, which for some companies means squeezing every drop of data out of their users,” says Casey Oppenheim, the CEO of [Disconnect](#), an online privacy tool that blocks trackers and helps guard personal information from prying technologies.

The journey to a comparison with a public health crisis (remember the lettuce?) is in no small part due to the fact that the ad-tech industry, despite a focus on [“better ads”](#), has neglected privacy for years and still faces accusations of [skirting privacy and consent](#) today. Even the supposed accuracy with which the value of an ad purchase can be seen is a myth. It’s an open secret that a huge portion of the internet traffic directed to ads is actually [from bots](#) and not humans. An estimated [\\$6.5 billion USD are lost to fraud by advertisers globally](#) in 2017 because of websites that cash in from using bots to inflate numbers.

[Many advertisers are angry](#) and have demanded more transparency in the supply chain. “Silicon Valley has created a fetish around automation,” says [Rory Sutherland](#). He is the vice chairman of the advertising agency Ogilvy in the United Kingdom, and says [an obsession with measuring results](#) of targeting has led to a decline in the quality of ads compared with traditional mass media marketing. “The obsession with targeting means what you are rewarding is your algorithm’s facility at identifying a customer,” he says. He compares it to walking into a pub with a piece of paper that says, “Drink beer!” Most people are already there to drink beer, he says. “What about the people outside?”

In 2017, a [number of major marketers](#) stopped placing ads on YouTube after a slew of scandals over ads on violent and inappropriate videos. For the general global public it can be jarring to see such content monetized. It adds to the sneaking sense of discomfort that is growing among many internet users for every report of breached data, [security flaws](#), and too-far-reaching data sharing agreements with other companies. Can we really [trust these companies with our data](#)?

As internet users we may have more ‘awareness’ about privacy, but still no clear sense of what to do. We are deeply dependent on companies we [wish would protect us](#).

In a restaurant, a food and safety inspector has a checklist of things to look for that may be a danger to public health. The [Corporate Accountability Index](#) of the organization [Ranking Digital Rights](#) is a kind of checklist too – but a complex one that ranks what the biggest internet and telecom companies disclose about how they protect the privacy and freedom of expression of users. By publicly scoring companies – and none scores high – the small but influential organization creates an incentive for companies to improve year over year, and a method to track noticeable progress and setbacks over time.

Nathalie Maréchal is a senior research analyst with Ranking Digital Rights in Washington D.C. She is leading an [open consultation process](#) to create entirely new indicators for the index related to targeted advertising. “We need to decide together, what standards for disclosure and good practice should be used to hold these companies accountable,” she says. Ranking Digital Rights’ current [ideas for best practices](#) will sound familiar to many internet researchers and digital rights organizations. Among other things, they suggest companies should allow third-party oversight of the parameters for ads (eg. “affinities”) and of who is paying for them. And that companies should state rules for prohibited content and use of bots – and publish data regularly to show how they are enforced.

Such tools and practices *have* begun to emerge out of companies already. Not of their own initiative, but compelled either by regulations or public pressure. This year, Facebook says they will roll out [political ad transparency tools](#) globally by June. In 2018, Google say they [killed over two billion “bad ads”](#). And Facebook took steps to remove [5,000 ad categories](#) to prevent discrimination. Twitter [began collecting more personal data](#) in 2017, but now also [gives you to control](#) to change how they categorize you.

Data privacy regulations are improving in numerous countries and states, and courts and civil society are taking companies to task around the world on matters of data collection and consent for targeted advertising. Regulation helps!

And so does technology. To protect the security of users, most major browsers have introduced different variations of tracking protection (and sometimes also ad blocking). Total or partial ad blocking by different companies in different configurations has gone fully mainstream with hundreds of millions of users. It makes the Web faster, and batteries last longer.

Coming back to the lettuce. What would the equivalent of “farm to table” in food activism be for digital ads? Perhaps we would see who paid for ads, understand why we are targeted, and have control over who is collecting our data for what.

What really needs rethinking today is the notion that digital ads can only be effective when they are targeted, and when companies know everything about everyone. Many brands and marketers are backing away from this idea for lack of evidence. Unless internet companies are able to regain our trust by changing practices (or perhaps be legally compelled to protect our secrets and interests, like doctors and lawyers), we can invest some hope in a new generation of software initiatives that explore decentralized solutions to give people personal control over who has access to their data.

“I spent 10 years working with an environmental health organization and I have always seen parallels to the privacy world,” says Oppenheim. “Just like we can connect people to the values of the food they eat, we can also connect them to the value of their data.”

► Further reading

- A Grand Bargain to Make Tech Companies Trustworthy, Jack M. Balkin, Jonathan Zittrain, The Atlantic, 2016. <https://www.theatlantic.com/technology/archive/2016/10/information-fiduciary/502346/>
- It's time for a Bill of Data Rights, Martin Tisne, MIT Technology Review, 2018. <https://www.technologyreview.com/s/612588/its-time-for-a-bill-of-data-rights/>
- Corporate Accountability Index, Ranking Digital Rights. <https://rankingdigitalrights.org/>