The impact of computing on all aspects of life is tremendous, and artificial intelligence will have an even bigger impact. We can no longer imagine a life without computing. As usual, there are positive and negative effects.

AI for a better society

By KOEN DE BOSSCHERE

The impact of computing on society is big, and all-pervasive. Computing has both positive impacts (easier access to information, more transparency, increased productivity, …) and negative impacts (growing inequality, fake news, privacy erosion, …). The challenge is to maximize the benefits while mitigating the negative consequences. In any case, artificial intelligence (AI) will lead to change, and change always causes resistance, but we cannot (in a competitive world) turn back the clock and go back to a time when there was no internet and no AI. The global challenges of the 21st century will not be solved without the help of artificial intelligence because the problems are so complex that they cannot be solved without advanced computing. We cannot make a world of ten billion people sustainable without AI. Eventually we will have to evolve towards responsible AI.

Key insights

- Algorithms are smarter than people and today they know more about us than we know about ourselves. Their negative aspects are often a reflection of our own. We have to protect society from the negative effects such as fake news, data leaks and privacy erosion.

- Computing accelerates societal change and this creates resistance. There is, however, no way back in an open and competitive world. The only way forward is to make sure that nobody is left behind, that the systems will cope with our ethical requirements and that everybody can benefit from the changes.

- AI and the internet consume a lot of power, but also help us save energy by optimizing processes.

Key recommendations

- We cannot make a world of ten billion people sustainable without advanced computing to limit the ecological footprint of such a population. In order to save the planet, we will have to invest more than ever in computing and artificial intelligence.

- AI can be used for good and for bad. We should evolve towards responsible and ethical AI, which means that the public and the private sector agree that we should only use it for the betterment of society. This implies that they should move away from purely economic criteria when making decisions.

- Computing systems should be made loyal to their users, not to the companies that provide devices and services.

- Computing system could also protect and help users by advising and informing them about the data (both inbound and outbound) that they exchange.

When Google was founded in September 1998, and the PageRank algorithm hit the world, people fell immediately in love with the search engine that seemed to be able to guess what a user was looking for. Surprisingly, the only thing it needed was a single search box, no long list of search options, check boxes, etc. Furthermore, it never disappointed the user. If somebody entered a URL instead of a keyword, Google just displayed the webpage. If a mathematical expression was entered, Google evaluated it. It corrected spelling errors, it automatically converted currencies and units, it also checked for synonyms. Such was the powerfulness of its offering, it immediately made other indexing websites obsolete. Google became the access point for the internet. Today, we expect search engines to read our mind, and immediately show what we are looking for, be it the closest restaurant, the cheapest online shop selling a particular product, driving directions… you name it, Google finds it.

With artificial intelligence, big data analytics, deep learning, and huge computing resources, platforms like Google became almost omniscient and able to serve us almost exactly the information we wanted. The younger generations cannot imagine how much energy it took in the 20th century to find reliable and recent information. Information in books and encyclopaedias was basically obsolete shortly after they were printed. Today, even small children can find the information they need and it has made some skills like searching in an alphabetically ordered list almost obsolete. It is fair to say that search engines have completely changed the way in which we deal with information and, in the process, they have made information available to all, and made society more transparent.

The technology that makes search engines so powerful has been adopted by social media platforms to show the relevant messages on personalized timelines, by news agencies to compile a personalized digest of the latest news, by dating apps to show matches one might be interested in, by streaming platforms to show the content somebody might like (the recommender system algorithms). And on top of all this, most of these websites deliver all these services for free on condition that they can show us some adverts. But that seems harmless because printed newspapers have adverts too, don't they? Not exactly.

How users became the product

Few people fully understand internet companies' business models. Facebook is a free platform with around 2.5 billion active users. In 2019, its revenue was 70 billion USD – an average of 28 USD per user. So, that is the average value in 2019 of the seemingly worthless information we share on our Facebook accounts. Facebook's real customers are the companies and organizations paying for marketing campaigns. The goal of a marketing campaign is to change the behaviour of the target group (for example by convincing them to buy a particular product, to sign up for a service or to vote for a political party).

For companies like Facebook or Google, the users are the product, and as any other company, Facebook and Google try to optimize their product (i.e. us!) to the needs of their customers. The perfect product is a user who spends a lot of time on the platform and reacts in ways intended by the (paying) customers (i.e. buying goods and services, voting [16] and so on). The more information the platform has about its users (the queries we enter, the links we follow, the pages we spend time on), the more targeted and the more effective the marketing campaigns can be made, and the more the platform can charge for them. The longer a user spends on the platform, the more advertisements can be shown, and the bigger the revenue will be. The more features the platform offers (face recognition, language translation, video, games, and so on), the more the users will enjoy the platform, the more time they will spend on it, and the more frequently they will return. There are good reasons why Google goes to great lengths to offer a wide variety of services. They want to be a one stop shop.

Attention is a valuable resource

There is an arms race between (social) media companies for the attention of the user. Unfortunately for them, a user cannot spend more than 24 hours a day in front of a screen. All these companies are thus competing against each other to get a greater share of users' attention. Platforms deliberately use mechanisms to make them addictive, or at least habitual. These include likes, automatic notifications, clickbait and scoring. This has been called brain hacking [1].

Addicted users come back frequently, which translates into higher revenue. Finally, the number of users has to grow fast for start-up internet companies and this influences the content. On one hand, platforms try to ensure that nobody will be offended by content on the platform, so they censor all content that might be considered inappropriate to valuable groups of users. Censoring is tricky as it starts from a world view of what is acceptable and what is not, especially when it comes to political statements, religious views or sex. On the other hand, viral (including outrageous) content is welcomed because it means that more people spend more time on the platform, and hence generate extra revenue. These platforms have also become (unintended) instruments to promote the values of the large user groups (e.g. American or Chinese values).

And they are successful in gaining user attention: in the younger generations, social media has almost completely outcompeted traditional media like television and newspapers [2]. In their competition for more attention, social media platforms are also monopolizing people's time, in both their professional and private lives. Active professionals believe they have to have a presence on social media, and to amass large numbers of followers. This leads to loss of productivity and mental absence at meetings, etc. In many people's private lives, screens have replaced face-to-face interactions at home (especially in 2020), at the dining table, at the pub, in restaurants and on public transportation. This leads to a phenomenon known as "phubbing", or phone snubbing: checking your smartphone during social events instead of giving your full attention to the people who are physically there [3].
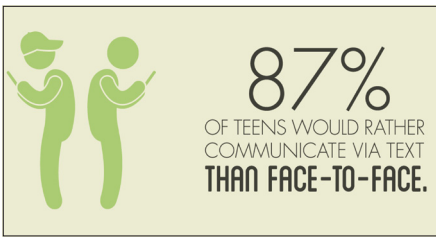
Figure 1: Phubbing (Source: stopphubbing.com)

The final frontier is competing with sleep. Studies show that millions of people suffer from sleep deprivation resulting from excessive use of smartphones and tablets [4].

Social media create echo chambers

What also sets social media apart from traditional media is that traditional media broadcast their messages publicly so that everybody can receive them and, ideally, learn about the arguments of a range of stakeholders by watching their channels. In contrast, the combination of advanced big data analytics and significant computing power hosted in large data centres has enabled social media platforms to create a personalized experience for each individual user. That means that every user gets to see a different stream of messages and that users cannot see the message streams of other users. Users can share messages in their own network, but since networks tend to be clustered, users tend to see more of the same messages rather than different points of view.

In so doing, social networks create information silos or filter bubbles and act as echo chambers which reinforce the values of the members of the network. "Wrong" posts will not garner a large number of "likes" and will quickly disappear from timelines. Hence, it is very difficult for information in one information silo to make it into another. The following figure illustrates three different communities living in Israel: pro-Palestinian, pro-Israel and religious/Muslim. There are very few links between the pro-Palestinian and pro-Israel communities. Most links are shared via the religious/Muslim community. There is little chance that messages from the pro-Israel network will ever make it into the pro-Palestinian network and vice versa.

What is worrying is that a handful of private global companies and their proprietary algorithms decide who gets to see what, in which order, and when. They can even gradually modify the user's preferences by proposing only a limited set of items and removing items that are old, in low demand or not in accordance with the ideas of the providers, for example. In the past, opinion-shaping messages came in hard copies, which were harder to remove – it was necessary to physically find them in people's house and burn them, as in *Fahrenheit 451* – than digital media on private servers and streamed to people who are not using local backups.

All this means that social media companies are in a sense helping to create a worldview per user, formed by purely business decisions – i.e. decisions that will optimize the profitability of the company – mostly unregulated by governments. The fact that traditional media such as newspapers and television news have declined in popularity among "digital natives" strengthens the impact of social media on the world view of young people. This explains to a certain extent why traditional media outlets anticipated neither Brexit nor the election of Donald Trump. They were simply unaware of messages shared in circles they did not belong to [5]. The fact that a significant number of American Trump supporters, conservatives and right-wing extremists recently moved to the Parler social networking platform is a sign that they are not interested in anything but their own messages.
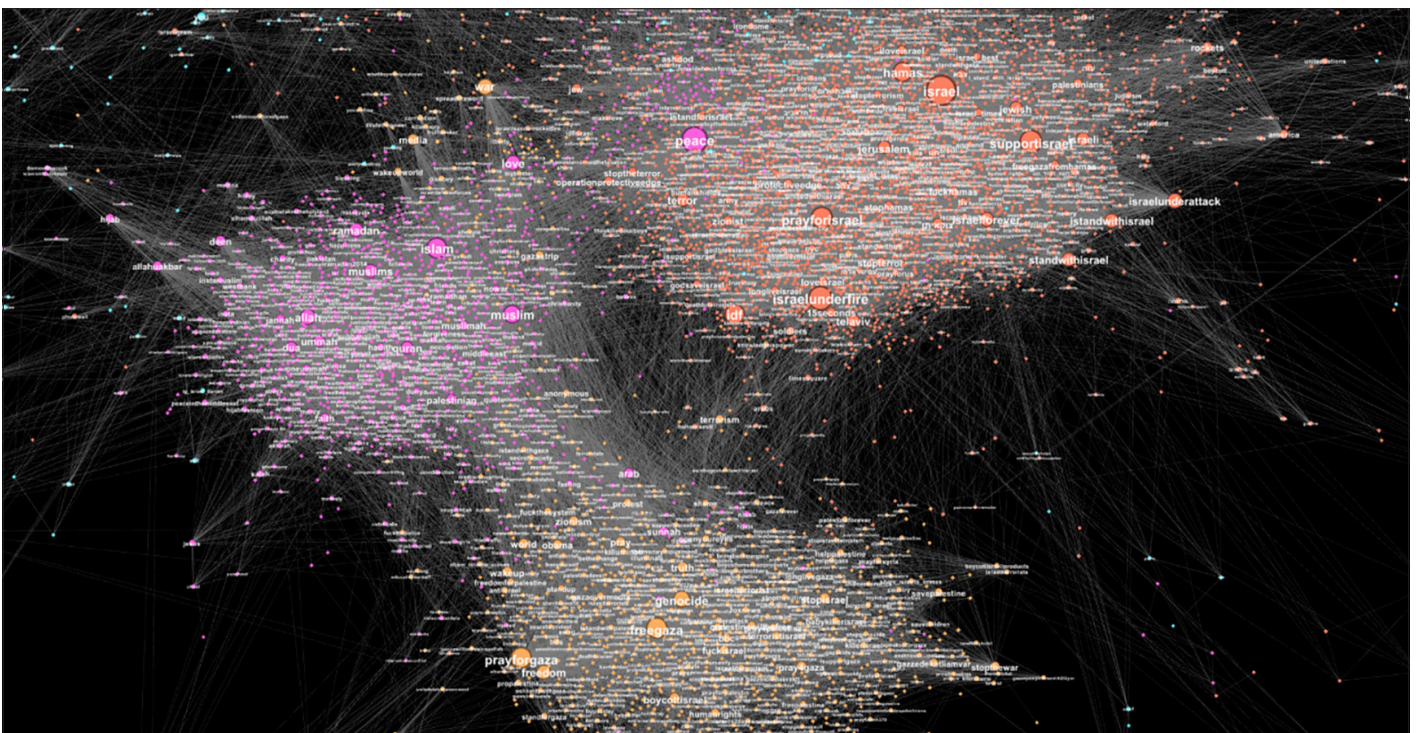

Figure 2: Israel, Gaza, War & Data – Social networks and the art of personalizing propaganda
(Source: [Gilad Lotan, "Israel, Gaza, War & Data: Social networks and the art of personalizing propaganda", 2014])

The internet leads to privacy erosion

There are multiple definitions of privacy. In the 19th century, privacy was defined as the "right to be left alone". A more modern definition is that privacy is the "control one has over the information about oneself". It is necessary for doctors to maintain medical records about their patients, but nobody expects the doctor to share this information with third parties (medical privacy) unless this were to be required for medical treatment. We expect the same behaviour from financial institutions (financial privacy), websites (internet privacy) and voting systems (political privacy). We do not expect an email service to use the content of our messages to influence the advertisements we see on websites, or a booking website to use the type of rental car we prefer to result in us seeing advertisements for that particular type of car.

Gathering information about users is crucial to the business model of internet companies. That is why many websites nudge users to complete their profiles, thereby collecting additional monetizable information. Some companies, like the now notorious Cambridge Analytica, have made a business model out of collecting information, analyzing it, and selling it to whoever is willing to pay for it.

Many people are largely unaware about the cost of convenience in terms of lost privacy; or if they are aware, they are willing to give up some of their privacy in return for convenience:

- Booking websites collect details about every single trip their users book. This is crucial marketing information for hotels, airlines and rental car companies.
- Streaming music applications have data on when and where users listen to music, as well as what their musical preferences are. The better streaming music providers can profile their users, the better suggestions they can make and the more frequently and longer people use the service.
- Companies selling e-books know the identity of every single reader of a book, when they are reading a book, which parts they actually read and so on. In a sense, they know what a buyer learned from the books they bought. The more they know,

the better suggestions they can make; it is not difficult to guess the interests of somebody buying books on classic cars, cookery, political history, or travel guides, for example. By (not) making particular suggestions, they can even steer what their users read and even think.

- Social media networks monitor all the private details users share with their most intimate friends, and use this data to infer information (for example, that the person feels depressed), in order to send them targeted advertisements they know work well (such as make-up or medication for those suffering from depression). Their aim is not to help people, but to sell and to influence. The people in social media control rooms are not medical staff; they do not have to comply with professional codes and they do not care about whether the advertised drugs are effective or safe.
- News websites track which articles users read, and adapt their content offering (news and advertisements) to their interests. They basically decide what their users will read, which might lead to a biased perception of the world. In the US, Democrats and Republicans live in two different news universes, leading to mutual demonization of the other side.
- Satellite navigation systems detect where the navigation system (and, by extension, probably its owner too) is at any time. It is comparable to being shadowed by somebody wherever you go.
- Voice-controlled devices keep track of what goes on in a house or office, and they can be hacked to eavesdrop on conversations. Few people would appreciate a stranger sitting in their house all the time.

In addition to the examples above, people are already under surveillance for a large part of the day, through access control systems in companies and hotels, numerous cameras in public places, licence-plate recognition, Google Street View filming the street, tourists taking pictures with people in the background and posting them on social media, and so on. Most people do not protest about this surveillance because they believe that it helps the government to enhance their safety and prevent terrorist attacks. Surveillance of people is an old practice, but it was limited to selected

individuals (for spying, criminals, etc), but algorithms and computing systems now allow for mass surveillance at a low cost.

Irrespective of the application, the fact is that (i) all our actions in *cyber space*, and an increasing number of actions in physical space are being recorded and stored in huge databases; (ii) that an increasing number of such databases are being linked (often through acquisition, or by linking government databases to facilitate e-government); and (iii) that that there is no guarantee that this data is only used for the purpose it was collected for. As these data collections are grown organically, where independent individuals decided to include/exclude particular information, the databases can be biased and, if used by algorithms, can come up with biased conclusions.

People are also often not really aware of the impact of the information they share; for example, a picture taken in a bar could mark you as a sociable person, but also perhaps as a drinker, which may interest your health insurance company. Computing systems could help inform people of the risks of their data exchanges.

It is clear that there is an urgent need for increased regulation. Privacy should be better protected, and there should also be more guarantees for unbiased database contents used for machine learning. People also deserve the right to be forgotten.

Fake information is part of the DNA of the internet

Whereas traditional media have built-in filters that require journalists to verify their sources, there is no such thing in social media. Anybody can post anything, and as soon as it passes social media companies' decency filters, it becomes public. The social media reviewers censor particular content (child abuse, sexual content, hate speech, …) but not fake information. The higher the number of people reading and liking the fake information, the better it is for the business results of the platform. In response to public concern over the spread of fake news and hate speech on social media, major companies such as Facebook have employed editors to monitor the content, and take it down, or add

a warning. Several baseless claims on election fraud by Donald Trump were labelled as such by social media in 2020. This is however not done with the same scrutiny for all other messages, and it can only be done if factually wrong. In other cases, it is often just a question of opinions, which cannot be kept out of social media.

Over the last few years, there has been a surge in false or misleading information such as fake news, fake science and deep fake videos. Fake information is information that is presented as a reliable piece of information, but is either completely made up or highly misleading. Such messages are like hoaxes on steroids. Popular genres are the launch of conspiracy theories (e.g. the widespread QAnon conspiracy in the United States), and the spreading of pseudo-science (such as the dangers of vaccination). The motives of people spreading such information range from making money (mostly from advertisements alongside stories that go viral) to political objectives (influencing elections, creating unrest, destabilizing societies).

The most recent technical evolution of fake information is the so-called deep fakes, a successful application of face swapping technology to video. Originally designed to put the face of celebrities on pornography actors in action, the technology has been used to create credible fake interviews [6]. For the naïve viewer, these interviews are hard to distinguish from the real thing and can thus be misleading, as well as a misrepresentation of the views or ideas of the 'interviewee'. Deep fake also applies to real-time voice substitution, leading one to believe that he or she is speaking to a known person.

The algorithms are smarter than people

The effects of digital technology on humans has been studied extensively, and there are both positive and negative effects. Customers have access to online information, they can make online appointments and buy goods and services without having to queue, physical meetings can be replaced by virtual meetings, collaboration tools allow people to work together efficiently and form the basis of the paperless office. On a personal level, it is now easier to keep in touch with friends and family members via social media. Many disabled and older people can also participate in social networks because their participation is not constrained by their limited mobility; this, in turn, helps them maintain or develop cognitive abilities. Thanks to video conference software, companies, governments and schools could continue (some of) their activities online during COVID-19 restrictions.

However, there are also some side effects. In some cases, people have become dependent on their smartphones. The smartphone does to the brain what using a lift, rather than the stairs, does to the body. Rather than memorizing information, people constantly refer to the internet, which can lead to digital amnesia [7]. Skills like mental arithmetic, memorizing numbers (mathematical constants, phone numbers), searching in a sorted list, and driving without a navigation system are disappearing in young people.

Perhaps even more disturbing is the fact that the web is full of texts that fit on just one or two screens, and that this has been linked to losing the ability of "deep reading", that is to say, the ability to focus on

HOW TO SPOT FAKE NEWS

CONSIDER THE SOURCE
Click away from the story to investigate the site, its mission and its contact info.

READ BEYOND
Headlines can be outrageous in an effort to get clicks. What's the whole story?

CHECK THE AUTHOR
Do a quick search on the author. Are they credible? Are they real?

SUPPORTING SOURCES?
Click on those links. Determine if the info given actually supports the story.

CHECK THE DATE
Reposting old news stories doesn't mean they're relevant to current events.

IS IT A JOKE?
If it is too outlandish, it might be satire. Research the site and author to be sure.

CHECK YOUR BIASES
Consider if your own beliefs could affect your judgement.

ASK THE EXPERTS
Ask a librarian, or consult a fact-checking site.

Figure 3: How to spot fake news (Source: IFLA)

a long text for an extended period of time. Research suggests that the disappearance of this skill, which is needed to read a book or to study, [8] can lead to lower academic performance.

Information technology has made sharing information so easy and cheap that it has become endemic. Many modern workers receive hundreds of messages per day; reading and responding to these messages takes up a significant part of their time, without being explicitly mentioned in their job description. Processing emails has become a struggle, putting people's bodies in fight mode for extended periods of time, and leading to exhaustion, burnout and faster ageing [9].

There is plenty of evidence that the use of technology has an impact of the amount of sleep we get. A 2015 survey showed that it was the sleep of young adults that was impacted most by technology. More recent studies show that that the problem is at least as severe in teenagers, [4] who practise late-night socializing, called vamping, which, in some extreme cases, takes place at any time of the night. Teenagers need around nine hours of sleep, but in 2015, 43% of US adolescents reported less than seven hours on most nights, which means that half of teenagers in the country are seriously sleep deprived, with 18-year-olds being the worst affected. Causes of disturbed sleep include (i) the use of social media which is both mentally and emotionally stimulating, and (ii) the blue light emitted by smartphones and tablets which simulates daylight, inhibiting the brain's production of melatonin, the hormone that regulates sleep.

Slowly, awareness about the negative effects of heavy smartphone usage is growing and even technology companies have started to offer tools to measure or restrict screen time, such as Apple's Screen Time and Google's Digital Wellbeing.

A number of former employees at the larger internet companies have started regretting what they built [10]. Some of them founded the Center for Humane Technology (http://humanetech.com) and give advice on how to take back control. The most extreme suggestion is to go "cold turkey" and delete all one's social media accounts. It has been claimed that this simple action will increase productivity, reduce stress and improve overall wellbeing. Some companies have introduced a policy not to allow their workers on the corporate network to check emails outside working hours. Sometimes it is useful to observe what insiders do; a number of high-profile executives at internet companies have admitted that they put serious restrictions on the use of social media and mobile devices by their own children.

However, at the same time, many schools are intensifying the use of technology as part of the learning process. This includes introducing "massive online open courses" (MOOCs) and flipped classroom courses, by using learning platforms that need to be used by children and students for their homework in the evening. According to an OECD study [11], the results are mixed at best. Students who use computers moderately at school tend to have somewhat better learning outcomes than students who use computers rarely. But students who use computers very frequently at school do a lot worse in most learning outcomes, even after accounting for social background and student demographics. The COVID-19 lockdown with distance learning will enable us to gauge the impact of intense computer use on learning. Time will tell whether the benefits of technology outweigh the side effects on children's development.

Computing transforms the job market

Computing, by definition, has an impact on the job market. The introduction of automation destroys jobs, creates new ones and changes the content of the remaining jobs. This has always been the case, ever since automation was invented. The key question many people have been focusing on is whether the current wave of automation fuelled by artificial intelligence and robotics will create more or fewer jobs than it destroys.

As of today, there are no signs that there are fewer jobs than, for example, twenty years ago, but the jobs have changed, and the effect of this change seems to be more

Figure 4: Cartoon by Jeff Stahler

inequality, a shrinking middle class and the emergence of a dual economy [12]. There is however a clear shift of jobs. The jobs that are most vulnerable to be destroyed are routine jobs, both manual (e.g. factory workers) and cognitive (e.g. accountants, or radiologists). The jobs that are created are non-routine jobs, i.e. jobs that require problem solving skills, creativity, entrepreneurship. In the near future, jobs in the event and hospitality sector might disappear when companies decide that some virtual events are as good as physical ones in the post-COVID-19 times, cutting down on travel. Some companies with a substantial part of the work force working from home are starting to wonder whether they still need huge office buildings. This will impact the real estate market. If people keep working a couple of days per week from home, catering and taxi services in business districts will also suffer. This is an indirect consequence of the use of technology in businesses.

Machine learning has a growing ecological cost

As the data sets used to train deep neural networks keep growing, energy consumption grows too. One study reports a 300,000-fold increase in power consumption for this purpose between 2012 and 2018. That same study reports that training one model with 175 billion parameters of a GPT-3 language model requires 28 000 GPU-days and has a carbon footprint of

85 tons. This is the emissions equivalent of one car driving 700,000 km [13]. The estimate for the power consumption of the AlphaGo Zero software is around 200 MWh, or a carbon footprint of 136 tons.

This is the cost of training one model. Fortunately, these large models (such as GPT-3) are rather effective in different use cases, decreasing the need for extra training. But if the model has to run on different platforms, different models might have to be trained. For applications in which the data set is changing regularly (like face recognition, traffic sign recognition, …) but also to fix bugs, the models might have

to be retrained regularly, fortunately incrementally if the algorithm allows it (e.g. using transfer learning). It is clear that the environmental cost of training the models is no longer negligible, but it might be offset by the benefit of the resulting application, which can be distributed in large numbers (for example in smartphones) and perhaps also used for applications that offer energy savings.

Although the energy consumption mentioned seems to be huge, one has to put it in perspective. The 85 tons of CO_2 for the GPT-3 model is equivalent to the yearly emission of 'only' four American citizens. At global scale it is the equivalent of the average yearly emission of 21 people. This is, however, not an argument to leave everything as it is. Every ton of CO_2 that can be avoided, should be avoided.

Since the carbon intensity per MWh varies wildly between countries from almost zero in countries with an abundance of hydropower or nuclear power to tenfold or more in countries with 100% fossil fuel production, the carbon footprint of a model can be reduced by running it in a place with low carbon intensity. There are also carbon intensity fluctuations during the day and the year. By avoiding the periods of the day

in which fossil fuel power plants have to support the electricity grid, the carbon footprint can be reduced. At the hardware level, low power customized accelerators should be used to reduce power consumption. And finally, at the algorithmic level, new training algorithms and models could be used like the once-for-all models that can be customized for different platforms instead of retraining them [14].

We also have to compare the computational cost to the value of the model. If a model would help reduce the energy needed for heating and cooling, or for transportation by only 1% – it is definitely worth it. Another example is agriculture, which is a big source of greenhouse gasses. We will need AI to feed the world's population while sustaining the planet. By making better use of natural resources, by creating new forms of agriculture in vertical farms, by growing meat in labs, … we will be able to reduce the ecological footprint of the world population. In these cases, the ecological benefits should of course also be bigger than the ecological costs.

According to the World Economic Forum, there are six environmental challenges where AI could be part of the solution, rather than part of the problem [15].
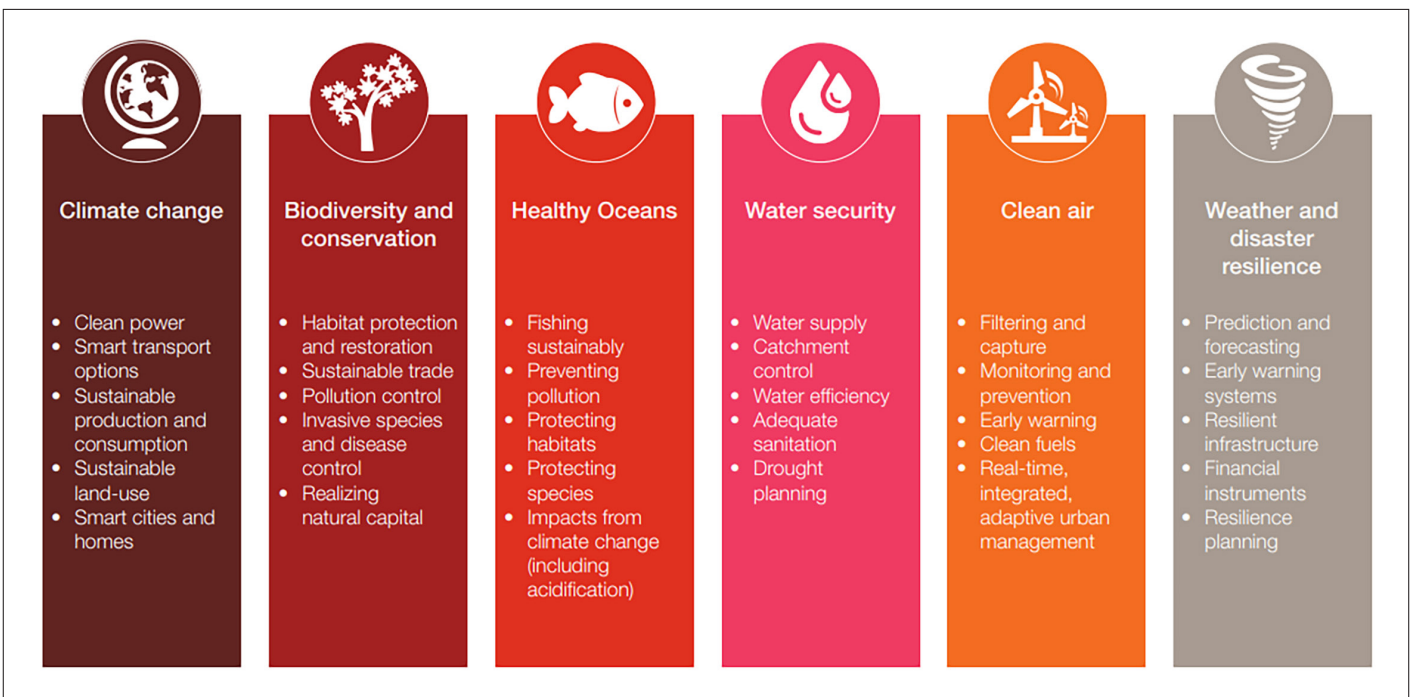

Climate change	Biodiversity and conservation	Healthy Oceans	Water security	Clean air	Weather and disaster resilience
• Clean power • Smart transport options • Sustainable production and consumption • Sustainable land-use • Smart cities and homes	• Habitat protection and restoration • Sustainable trade • Pollution control • Invasive species and disease control • Realizing natural capital	• Fishing sustainably • Preventing pollution • Protecting habitats • Protecting species • Impacts from climate change (including acidification)	• Water supply • Catchment control • Water efficiency • Adequate sanitation • Drought planning	• Filtering and capture • Monitoring and prevention • Early warning • Clean fuels • Real-time, integrated, adaptive urban management	• Prediction and forecasting • Early warning systems • Resilient infrastructure • Financial instruments • Resilience planning

Figure 5: Priority action areas for addressing Earth challenge areas [15]

Figure 6: AI risks [15]

Unfortunately, there are also six risks.

Its conclusion is that we should work on 'responsible AI' on top of the classical criteria like safety, ethics, value and governance. AI should be used not to speed up the destruction of the Earth's ecosystem by facilitating more efficient extraction of its natural resources, but to preserve the ecosystem. This will require leadership from the public and private sector. Some companies are already taking the initiative in this direction [17].

References

[1] Richard Freed, "The Tech Industry's War on Kids: How psychology is being used as a weapon against children", 2018, https://medium.com/@richardnfreed/the-tech-industrys-psychological-war-on-kids-c452870464ce

[2] Jean Twenge, Gabrielle Martin, and Brian Spitzberg, "Trends in U.S. Adolescents' Media Use, 1976-2016: The Rise of Digital Media, the Decline of TV, and the (Near) Demise of Print", 2018. Psychology of Popular Media Culture, https://www.apa.org/pubs/journals/releases/ppm-ppm0000203.pdf

[3] Kimberley Holland, "How to Identify and Manage Phubbing", 2018, https://www.healthline.com/health/phubbing

[4] Jean Twenge, "Analysis: Teens are sleeping less. Why? Smartphones", PBS, 19 Oct 2017, https://www.pbs.org/newshour/science/analysis-teens-are-sleeping-less-why-smartphones

[5] Samidh Chakrabarti, "Hard Questions: What Effect Does Social Media Have on Democracy?", https://newsroom.fb.com/news/2018/01/effect-social-media-democracy/

[6] James Vincent, "Watch Jordan Peele use AI to make Barack Obama deliver a PSA about fake news", The Verge, 17 Apr 2018, https://www.theverge.com/tldr/2018/4/17/17247334/ai-fake-news-video-barack-obama-jordan-peele-buzzfeed

[7] S. Coughlan, "Digital dependence 'eroding human memory'", BBC News, October 2015. http://www.bbc.com/news/education-34454264

[8] Nicholas Carr, "The Shallows: What the Internet Is Doing to Our Brains", 2010, Norton & Companym

[9] David Robson, "The reasons why exhaustion and burnout are so common", http://www.bbc.com/future/story/20160721-thereasons-why-exhaustion-and-burnout-are-so-common

[10] Noah Kulwin, "The Internet Apologizes", The New York magazine, 2018, http://nymag.com/selectall/2018/04/an-apology-forthe-internet-from-the-people-who-built-it.html

[11] "Students, Computers and Learning: Making the connection", OECD Publishing, 2015, https://read.oecd-ilibrary.org/education/students-computers-and-learning_9789264239555-en

[12] Peter Temin, "The Vanishing Middle Class", 2017, MIT Press

[13] Lasse F. Wolff Anthony, Benjamin Kanding, Raghavendra Selvan, "Carbontracker: Tracking and Predicting the Carbon Footprint of Training Deep Learning Models", https://arxiv.org/pdf/2007.03051.pdf

[14] Rob Matheson, "Reducing the carbon footprint of artificial intelligence", 2020, https://news.mit.edu/2020/artificial-intelligence-ai-carbon-footprint-0423

[15] "Harnessing Artificial Intelligence for the Earth", World Economic Forum, 2018, http://www3.weforum.org/docs/Harnessing_Artificial_Intelligence_for_the_Earth_report_2018.pdf

[16] Julia Carrie Wong, "The Cambridge Analytica scandal changed the world – but it didn't change Facebook", https://www.theguardian.com/technology/2019/mar/17/the-cambridge-analytica-scandal-changed-the-world-but-it-didnt-change-facebook

[17] David Hagenbuch, "The 4 Pillars of Ethical Enterprises", https://www.entrepreneur.com/article/240035

Koen De Bosschere is Professor in the Electronics department of Ghent University, Ghent, Belgium.