

Alexandra
Luccioni

Yoshua Bengio


Much of the existing research on the social and ethical impact of Artificial Intelligence has been focused on defining ethical principles and guidelines surrounding Machine Learning (ML) and other Artificial Intelligence (AI) algorithms (21). While this is extremely useful for helping define the appropriate social norms of AI, we believe that it is equally important to discuss both the potential and risks of ML and to inspire the community to use ML for beneficial objectives. In this article, which is primarily aimed at ML practitioners, we thus focus more on the latter, carrying out an overview of existing high-level ethical frameworks and guidelines, but above all proposing both conceptual and practical principles and guidelines for ML research and deployment, insisting on concrete actions that can be taken by practitioners to pursue a more ethical and moral practice of ML aimed at using AI for social good. In this way, actions can be evaluated against these principles and guidelines to serve as a kind of “social litmus test” by which others can hold ML practitioners to account.

Artificial Intelligence Leaves the Research Lab

Progress in ML in the last decade has been extraordinary and has rekindled the notion that AI systems could eventually reach human levels of performance, which was

Digital Object Identifier 10.1109/MTS.2020.2967486
Date of current version: 2 March 2020

On the Morality of Artificial Intelligence



abandoned for several decades. Even if we are still currently far from this achievement, technological progress in ML has passed a threshold that enables it to have a huge economic impact, estimated to be close to 16 trillion US dollars by 2030 (38). This contrasts with the first few decades of ML progress, when researchers had the luxury of focusing purely on the fundamental aspects of their work, not worrying too much about its potential societal impacts — an object recognition algorithm could be tested on a common dataset like MNIST

(24) or ImageNet (10), and an objective performance metric would be obtained in order to measure progress, without having to think about the messiness and complexity of deployment and social impact. Something crucial has changed in recent years, as algorithms initially developed in the lab are increasingly being improved and deployed in society in real-world applications such as healthcare, transportation, and industrial production with real-life consequences, and we are likely seeing just the tip of the iceberg in terms of social impact.

Along the way, this deployment in society has forced the realization that these algorithms have social impacts that could be positive or negative. For example, we have realized that biases hidden in data and algorithms could lead to more discrimination, in the simplest case simply because of the data imbalance: facial recognition algorithms have been found to underperform on gender and racial minorities (5). Furthermore, above and beyond hidden biases, given the high impact potential of ML research, the question stands of whether practitioners are acting with the best interests of humanity and society in mind when developing their tools and applications.

As ML researchers and engineers, we believe that we have a shared responsibility to consider both ethics and moral values when we choose what we work on, for what organization, and whether the products we contribute to directly or indirectly will be beneficial to humanity or more likely to end up hurting more than helping. Unfortunately, very few of us have been trained to think about these questions. Instead, most of us have focused from a very young age on mathematics and computer science and not so much on philosophy and other humanities. A good step towards learning about these issues is to consult the documents proposing ethical guidelines for AI, which we will cover in the next section of this article. Furthermore, in order to offer a guiding direction for such debates and soul-searching within the scope of ML, we propose the following self-directed questions:

- 1) How is the technology that I am working on going to be used?
- 2) Who will benefit or suffer from it?
- 3) How much and what social impact will it have?

- 4) How does my job fit with my values?

We are conscious that the questions listed above are subjective and the answers will depend highly on the values and ethics of the individual answering them. Nonetheless, we hope that work on some applications, such as the design and deployment of lethal autonomous weapons and automatic surveillance, will clearly be seen to contradict fundamental rights and dignity, as defined in, among other places, the UN Declaration of Human Rights (41). Other applications of ML, such as those increasing the efficiency of advertising or beating the stock market, are less clear cut in their moral value, and merit informed debate and discussion within the scientific community and society at large. As some of us become more conscious of the potential or definite social impact of ML, we have the opportunity, if not the duty, to make our voices heard. A good example of this is a recent letter, signed by numerous scientists, calling for an international treaty to ban lethal autonomous weapon systems, e.g., killer drones. Killer drones can decide to shoot at a person — without the human involvement that would make it possible to take the broad social, moral, and psychological context into account and potentially decide to abort the mission (for instance, when the target is in a school or at a family dinner surrounded by women and children).

Finally, while the legal frameworks to oversee and limit research and development violating these principles are often and unfortunately updated in a reactive rather than a proactive manner, we believe that we should not wait until all of the dots between ML and ethics are formally connected by legislation and regulation. We believe that we have a responsibility to educate

ourselves, to think ahead about potential consequences, to use our internal moral compasses, and to consciously choose the direction of the research or engineering that we practice. This is important because we believe that we are faced with a *wisdom race*, that as technology becomes more powerful, its impact can be proportionally greater, either positively or negatively.

To curb the negative impact, we need to become wiser individually (as reflected in our personal decisions) and collectively (through social norms, laws, and regulations). Unfortunately, technological progress in AI has accelerated faster than the current rate of progress of personal and social wisdom, ultimately making it possible for unwise humans or organizations, even those with good intentions and acting legally, to have large-scale, major destructive effects. This is comparable to a world in which nuclear bombs (i.e., very powerful technology) were accessible and usable by children (i.e., persons with insufficient maturity and wisdom), which could easily result in global nuclear war. This highlights the importance of the discussions still to be had by large numbers of ML practitioners about ethics and social impact, as well as the safeguards that need to be put in place to protect especially the most vulnerable members of our society. We will discuss some of the most advanced efforts to introduce these safeguards in the next section, followed by some examples of socially beneficial applications of ML.

Ethics and AI — Existing Initiatives

In recent years, there have been numerous initiatives which have taken one of two major approaches to fostering the ethical practice of AI: (1) Proposing principles guiding the socially responsible development of AI

or (2) Raising concerns about the social impact of AI. We will describe both approaches in the current section, as well as giving examples of notable initiatives and projects which have adopted either of the approaches.¹

Defining Principles for Practicing AI Responsibly

The topic of ethical research and practice in technology has been gaining momentum in different corners of the computing community in recent years, and the various initiatives that have been proposed are indicative of the interest and the

AI and the IEEE report for Ethically Aligned Design.

The Montreal Declaration for a Responsible Development of Artificial Intelligence

One of the most notable approaches to establishing guidelines for AI deployment is the Montreal Declaration for a Responsible Development of Artificial Intelligence, developed in 2017 and revised in 2018 based on public feedback. It was created under the premise that since AI will eventually affect all sectors of society, it requires principles to guide

its development to ensure its adherence to human values and social progress. The resulting Declaration has ten principles, ranging from protection of privacy to equal representation, with some principles touching responsibility and ethics directly; for instance, the principle of Prudence stipulates that *“Every person involved in AI development must exercise caution by*

anticipating, as far as possible, the adverse consequences of AIS (Artificial Intelligent Systems) use and by taking the appropriate measures to avoid them.” These principles were defined after extensive debate and dialogue between both specialists and non-specialists from different domains and parts of the world to ensure representation and cohesion. The overall aim of the declaration was to spark public debate and to encourage a progressive and inclusive orientation to the development of AI.

However, the Montreal declaration goes further than theoretical ethical principles, proposing recommendations to accomplish an ethical digital transition that includes all of the different levels of society, from researchers to policy-makers.

For instance, it includes a proposition for auditing and validating the use of AIS using concrete frameworks and certifications in order to prevent biases and discrimination. Specific steps were also proposed for ensuring the protection of democracy and reducing the environmental footprint of AI, all within the framework of a democratic and citizen-led process. This is important given that the effects of AI will permeate all levels of society, from programmers and engineers who write the code, to leaders who make laws governing its use and development, to businesses that will make products with it to be used by all. The process creating the Montreal declaration was consequently the keystone to building a way of including all of these different stakeholders in the elaboration of an ethical AI, and paves the way for subsequent work on the topic.

IEEE Ethically Aligned Design

A more recent effort, initiated by the IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems, carried out an in-depth study on the issue of the ethics surrounding the design of AI systems (40). In particular, aspects that are relevant to the topics covered in the present paper include: the usage of autonomous and intelligent systems (A/IS) in service to sustainable development for all, and more specifically for the attainment of the United Nations Sustainable Development Goals (SDGs) (42). The authors of the study specifically underline the potential of AI to contribute to resolving some of the world’s most urgent problems, such as climate change and poverty, given the necessary will and orientation towards these problems. Furthermore, they highlight the fact that despite their great potential, current AI deployment and development is currently



As researchers and engineers become more conscious of the social impacts of machine learning, we have the opportunity and duty to make our voices heard.

concern that many members share. For instance, in the United States, the Association for Computing Machinery (ACM) has proposed a Code of Ethics and Professional Conduct, to be followed by all members of the association and to guide them in their usage of computer science (16). A similar initiative was undertaken by the Royal Statistical Society (RSS) in the United Kingdom, which has created a practical guide for practitioners regarding the ethical use of mathematics (35). Here we address the two most relevant and extensive initiatives to establish ethical guidelines for AI research and practice: the Montreal Declaration for Responsible Development of

¹For a more complete overview of different global ML ethics initiatives, see a recent review in (21).

not aligned with these goals and impacts (40, p. 144), which is unsettling given the myriad of ML project and initiatives worldwide.

The IEEE report also lays down principles to guide “the ethical and values-based design, development, and implementation of autonomous and intelligent systems,” many of which are similar to those defined by the Montreal Declaration: respect of human rights, data agency, transparency, accountability, etc. They go further in proposing that “A/IS creators shall adopt increased human well-being as a primary success criterion for development” instead of focusing on isolated metrics such as accuracy, and from a deployment perspective, offering alternative metrics to quantify meaningful progress, for instance by evaluating social, economic, and environmental factors instead of profit and other common success metrics. The report also includes propositions for policymakers, legislators, and other stakeholders from the extended AI community and, as such, represents the most extensive effort of establishing ethical boundaries and guidelines for AI research to date.

In a recent survey of the various global ethics guidelines proposed around AI, the authors observed that despite a conceptual overlap between the many existing guidelines, including the two mentioned above, there are major differences regarding how the principles are interpreted (21). This underlines the complexity and nuance of applying theoretical, philosophical principles in practice, and raises questions such as: what aspects of the AI research and deployment pipeline do ethics principles affect? How would it be possible to resolve conflicts between, for instance, fairness and sustainability (i.e., training an algorithm longer and with more data — thus poten-

tially leading to more greenhouse gas emissions — to ensure that it is not discriminatory and covers all demographic groups equally well)? And, above all, how is it possible to translate ethical principles into a programming language? In any case, the bridge between theory and practice has yet to be built and there are different ways in which that can happen. This underlines the necessity of involving actors from different levels of the AI ecosystem (and neighboring ones) in order to ensure that experts in policy-making work in tandem with experts in coding and engineering to create tools and frameworks that are coherent and usable by all.

Identifying Ethical Concerns of AI Applications

There are several types of ethical concerns regarding AI applications and, in this paper, we focus more concretely on bias leading to potential discrimination. While it is true that on the one hand, AI-infused technology such as computer vision can enhance public security, for instance by identifying crime in real-time based on CCTV cameras — the trade-off is that security features can also be abused to track individuals and to establish a surveillance state where privacy is greatly threatened by those who control the technology.

On the military side, similar technology can be used to design autonomous drones that use computer vision to identify their target, representing a grave threat to global security and democracy due to the lack of human oversight. In addition to the security risk, such weapons would be moral and legal hazard: AI technology is not yet able to comprehend and represent the social and psychological context in which such a targeted attack could take place in a manner that is coherent with inter-

national laws regarding war or with human morality. Unfortunately, the most common argument brought in favor of developing lethal autonomous weapons is that they are needed as a precautionary measure (i.e., since other countries are undeniably working on them, each country needs to do the same). In reality, the weapons needed to defend against killer drones would be very different from the drones themselves, and do not need to be lethal autonomous weapons since they would be designed to destroy weapons rather than to target people, similar to the Iron Dome used by Israel.

Another common argument is that an international treaty would be useless since some countries will refuse to sign it. But we have seen in the past that even when major powers do not sign a treaty (such as the one on anti-personnel mines, signed by 133 countries, excluding the U.S., in 1997), the treaty can still be used to create a moral stigma, as well as a decline in demand; in the case of anti-personnel mines, the result has been that U.S. companies have stopped building them, even though their government never signed the treaty.

Another flawed argument is that regulating lethal autonomous weapons could threaten innovation in AI, whereas in fact AI has been developed very successfully in a civilian setting (mostly in academia and major technology companies) and its continued development does not require data or engineering from AI military development.

Another potential threat to democracy stemming from AI could come not simply from the increased ability to monitor and to target individuals, but also from the more subtle power to influence them, e.g., via AI-driven advertising, automated online trolls, and other psychological manipulations via the Internet and

social media. The recent use of AI to influence political campaigns such as the 2016 U.S. election or Brexit is just the beginning of what can be done when machines learn how to “press our buttons” in a personalized way. This is due to the fact that micro-targeting makes it possible for ads to be truly bespoke depending on your political views, networks of friends, and personal history. While we may not mind being influenced when it comes to choosing a brand of soft drinks, when the profit or power motives of a corporation or political organization go against our individual and collective interests, it becomes important to establish social norms, laws, and regulations to protect us from such psychological manipulation. But where should the line be drawn between, for example, manipulation and education? These are difficult questions but there are clues that can be used (like whether the organization that stands to profit is paying for the advertisement or social network influence). Human judgement remains key for judging the ethical aspect, e.g., in balancing different values (like autonomy vs well-being, when considering an ad campaign against cigarettes, for example). In the case of advertising, it is interesting that in addition to the moral hazard associated with psychological manipulation, it is not even clear that advertising is beneficial to society from a purely economic perspective, as it tends to favor established brands and thus slow down innovation.

Closely related to political misuse and manipulation using AI are the increasing concerns about AI-generated false images, videos, and news. Thanks to rapid progress in generative neural networks such as Generative Adversarial Networks (GANs) (15), it is becoming possible to synthesize images and sounds in a controlled way, e.g., using “deep fakes” for making a

video of a president declaring war, or with the face of a celebrity seamlessly integrated on the body and behavior of a pornography actor.

Other commonly discussed concerns of AI deployment include the effect on the job market (31) (which means that governments and communities must prepare, e.g., by adapting the education system and the social safety net, which can take decades), the potential concentration of power that it may lead to (in specific individuals, corporations and countries), and the bias and discrimination it may contribute to increasing, as we discuss next.

Identifying and Mitigating Bias

In recent years, we have been confronted numerous times with the fact that biased algorithmic systems can perpetuate injustice and discrimination, whether we are aware of it or not. There are many different ways that this kind of bias can creep into algorithms: it can be from the data itself, or the implicit bias that the creator programmed into the system, or even the way a problem is framed.² Therefore, in order to ensure that the models that we develop and the systems that they are later used in are as fair and ethical as possible, there are steps to take to identify bias and to reduce it as much as possible.

Numerical Bias

A major challenge in designing ML systems is to understand how the systems work during training and deployment, and what factors and features they use to make decisions. However, diagnosing the presence of bias in these systems is not a straightforward task, since it is not always

²For a more hands-on presentation of bias and fairness in AI, we suggest Google’s Online course designed specifically for ML practitioners: <https://developers.google.com/machine-learning/crash-course/fairness/video-lecture>.

obvious during a model’s construction what the downstream impacts of design choices may be. Therefore upstream efforts are needed to reduce this risk as much as possible. To this end, there have been several proposals to help practitioners identify and mitigate bias in ML models, some of which we will describe here.

More concretely, exploring, analyzing, and visualizing the data used for training a model is a key part of the ML process. But it is not straightforward to identify bias simply by looking at the data. Often more in-depth probing is needed to figure out what features and implicit information is present and, once a model is developed, how this will influence the model’s behavior. For instance, it was recently found that the COMPAS system, a criminal risk assessment tool developed and widely used in the United States, is often biased with respect to race (2). The bias in the COMPAS system was identified after its deployment. Once the data was made public, it was clear that this bias is an aspect of the model that should have been identified much earlier, during development and certainly before deployment. Similarly, off-the-shelf facial recognition technology used by police forces has been shown to perform much worse on racial and gender minorities, with a difference of up to 34.4% in error rate between lighter-skinned males and darker-skinned females, mostly due to the lack of reliable training data (5).

To address these types of issues, several approaches exist: for instance, researchers have recently released a tool called “What-If.” What-If is an open-source application that lets practitioners not only visualize their data, but also test the performance of their ML model in hypothetical situations. For instance researchers can modify some characteristics of data points and analyze

subsequent model behavior by measuring fairness metrics such as Equal Opportunity and Demographic Parity (47). Other approaches address bias by changing the training procedure or the structure of the ML models, for instance by transforming the raw data in a space in which discriminatory information cannot be found (49), or using a variational autoencoder to learn the latent structure from the dataset and using this structure to re-weight the importance of specific data points during model training (33). Whatever the approach chosen, using these kinds of tools during ML model development and deployment can change the life of individual people, who could go from unfairly spending decades in prison to having the chance of a better life — an immensely important difference especially when multiplied by the thousands of people whose lives can be affected by the deployment of these tools. This multiplication of bias is especially important to consider since ML is being used more and more, and therefore even edge cases and small minorities can be amplified in real-world applications.

Textual Bias

Bias is not always in numbers, it can also manifest itself in the words that we use to describe the world around us. For instance, in 2018, Reuters reported that Amazon was forced to decommission an ML-powered recruiting engine when it was discovered that it penalized any mention of female-related vocabulary, including applicants who attended all-women colleges (9). This is not surprising given the gender disparity that exists in the technology sector and since the data used to develop this tool was comprised of resumes submitted (and accepted) to Amazon over a 10-year period. It is nonetheless disturbing in terms

of algorithmic fairness, especially if algorithms such as this one make filtering or hiring decisions that can ultimately affect an entire gender's lives and careers. This can potentially create a negative feedback loop, as such a system would reduce the number of female workers and thus the number of positive role models for girls interested in technology. A similar type of gender bias was also found in pretrained word embedding models, which were found to exhibit gender stereotypes in terms of higher cosine similarity between, for instance, "woman" and "homemaker" or "receptionist" as opposed to "woman" and "doctor" or "lawyer," notably due to these biases existing in the corpus that they were trained on, which consisted of mainstream news articles (4).

In order to reduce and eventually remove gender bias in written text, researchers have proposed approaches such as identifying the gender subspace of vectors and adjusting the dimensions in a way that either neutralizes or entirely removes gender bias (4). Others have defined a formal gender bias taxonomy in order to capture gender bias and to train ML models to later identify this bias in texts (18). Debiasing the computational representation of language, notably word embedding models, is especially important because of the extent of their usage; pretrained embedding models trained on corpora such as Google News and the Common Crawl are used in a variety of applications and systems, and can therefore continue perpetuating gender bias in downstream usages in Natural Language Processing (NLP) applications such as dialogue systems. This is a challenge given the complex and subsymbolic nature of modern NLP, which makes it difficult to analyze specific features and aspects of data and identify latent connections and bias between

words and concepts. Therefore, more work is needed to explore and analyze these issues, which constitutes an interesting research direction in itself, and one that is important to pursue and to integrate into mainstream ML research.

Despite the research initiatives described above to carve appropriate social norms about AI, there remains a noticeable gap between the recommendations they make and ways to ensure that these are respected. Legislation of AI is still catching up to the progress made in research and practice, and there have not yet been any country-level laws governing AI research specifically. However, there have been, on the one hand, more high-level legislative frameworks such as the European Union (EU) General Data Protection Regulation (GDPR) (<https://gdpr-info.eu/>), which aims to ensure data privacy and protection and, on the other hand, more local initiatives such as San Francisco's Facial Recognition Software Ban. Nonetheless, more complete legal frameworks are needed to control nefarious use of AI and to ensure that the principles defined in theory are applied and enforced in practice.

AI for Good Initiatives

Whereas the profit motive is the main driver behind much of the commercial deployment of AI today, there are nonetheless many projects going on in academia, government organizations, civil society, and industry labs motivated by more noble objectives, often called AI for Social Good (AISG) projects. In addition to the specific projects being undertaken in areas such as healthcare, education, or the environment, it is interesting to highlight higher-level efforts that aim to foster and facilitate these projects. For example, the AI Commons project (<https://ai-commons.org/>) aims to construct a hub where

different kinds of actors can connect and collaborate on AISG projects, e.g., ML graduate students or engineers, problem owners in NGOs or local governments, philanthropy organizations, or startups that could deploy the ML solutions. Their interaction is to be facilitated by online tools and datasets as well as a standardized description of the status, progress, and expected impact of

In the last five years alone, groundbreaking work has been done in improving the diagnosis of diabetic retinopathy from a single visit (3), detecting breast cancer in lymph nodes (14), and large-scale discovery of diseases based on health records (32). There are also an increasing number of startups and companies working in the space, either by commercializing research

done in academia or by developing products specifically catered to the medical sector, with the most advanced applications harnessing the power of deep learning for analyzing and classifying medical imagery. Despite the many exciting advances that are being made, there are many hurdles in ML research in healthcare, starting with data privacy and control (who

owns the data? Can patients share their own data, or should the process be centralized? How to find the right balance between privacy and the lives which will be saved by applying ML on the aggregated health records from many different sources?), to the manner in which medical data should be processed (Should it contain information such as race and postal/zip code, which can impact diagnoses, be included in electronic health records, or does that open the door to discrimination and bias?), and how should such systems be deployed (human-in-the-loop or fully automated)?³ There are also often questions of responsibility and interpretability that arise, given the high stakes of deploying ML systems in situations of life and death. In order to make meaningful progress in this sector, it is therefore important to continue existing research on fair

AI for Education

and ethical usage of ML in healthcare (48) and to ensure that Hippocratic principles are a solid part of the research and development process, as well as working with stakeholders of the domain (e.g., radiologists, clinicians, patient organizations, and hospital administrators) to propose solutions to the hurdles discussed above.

The promise of using adaptive intelligent systems and agents for education has been around since the 1960s (39), but access to personalized digital education tools has yet to become a reality in most countries, especially in the developing world, where it could have the most impact to democratize education and knowledge (30). In recent years, given the increasing global shortage of qualified teachers along with the increasing number of students, the issue of access to education has become a global one, a fact highlighted by its presence among the UN SDGs. And yet, the use of ML in the education sector has been limited to specific, narrow applications such as predicting the probability of learner attrition (7), or improving learner evaluation (1). There are many reasons for this, starting from the difficulty of representing learning content in a domain-agnostic way to facilitate scalability, to overcoming cultural and linguistic barriers to deploying tutors worldwide. But the limited use of ML in education is also caused by more fundamental issues such as the lack of large-scale educational datasets and inherent technological constraints in developing countries.

Despite these hurdles, there are many new and longstanding efforts to create intelligent tutors, be it using symbolic AI approaches such as ontologies and knowledge modeling (28), educational data mining (12), or



We believe we have a shared responsibility to consider whether the products we contribute to will be beneficial to humanity.

each project. We hope that initiatives like this will help solidify and amplify the impact of AISG; in the meantime, there are also many profoundly positive uses of AI that are emerging and we would like to highlight and applaud such efforts next.

AI in Healthcare

Achieving universal health coverage is one of the seventeen UN Sustainable Development Goals (42), and although major progress has been made in numerous domains, such as maternal health as well as HIV/AIDS reduction, there are still many problems that are far from being solved. While ML is not a cure-all, there are many challenges that it can help with such as personalized medicine, diagnosis of medical imagery, and improved drug discovery (13). ML in the health sector is in fact a thriving research domain, with its own workshops at major ML conferences and research published in major medical journals read by practitioners worldwide.

owns the data? Can patients share their own data, or should the process be centralized? How to find the right balance between privacy and the lives which will be saved by applying ML on the aggregated health records from many different sources?), to the manner in which medical data should be processed (Should it contain information such as race and postal/zip code, which can impact diagnoses, be included in electronic health records, or does that open the door to discrimination and bias?), and how should such systems be deployed (human-in-the-loop or fully automated)?³ There are also often questions of responsibility and interpretability that arise, given the high stakes of deploying ML systems in situations of life and death. In order to make meaningful progress in this sector, it is therefore important to continue existing research on fair

³For a more extensive overview of the opportunities and challenges of using ML in healthcare, see (13).

more recently, ML-driven approaches (8). However, there are very high stakes in the field, since technological interventions have the potential to make a considerable, long-term impact on human livelihoods, for example lifting people out of poverty by endowing them with linguistic and numerical literacy. But these positive impacts can be hindered by bias and technological constraints. We therefore agree with recent proposals to improve and support human learning at scale and believe that ML has a key role to play in this endeavor. This can be done, for instance, by partnering with existing education initiatives and organizations in order to learn what their specific needs are and how ML can be used to meet them, or else by collaborating with Massive Open Online Course (MOOC) creators in order to gather data and make it available to the ML community, and finally by sharing learning materials and activities used in local education initiatives (e.g. university courses in Machine Learning) so that they profit learners in places where access to high-quality technical education is limited.

AI for the Environment

Climate change is, without a doubt, one of the biggest challenges humanity has faced, and we are at an important point in history when we are both aware of the issue and still have the possibility to change its course. Climate change has been described as a “wicked” problem, due to features such as the difficulty in defining the problem itself and in developing and deploying solutions to it, the lack of central authority that can solve it, the incentives for individual countries or companies to not do their share, and the cognitive biases that discount the future impacts of our actions (17), (25). Furthermore, while we do not know of any single technological silver bullet

as solution to climate change, there are nonetheless numerous technical challenges for which ML can be helpful, and which can be combined to make a significant impact on the overall issue. These challenges and the ongoing ML approaches to tackle them were presented in a recent survey paper (34). We will not go into all of these at length, but we will focus on a few examples that are particularly salient and that we hope will give an idea of both the relevance of deploying ML in environmental applications and the opportunities that this can generate.

Energy and Transportation

Together, electricity and transportation systems are estimated to produce close to half of anthropogenic greenhouse gas (GHG) emissions (20) and both sectors have their own unique challenges for decarbonization. For instance, one of the major obstacles to building and using renewable energy sources such as solar and wind is the variability of their output, which is inherently problematic since the power generated by an energy grid must equal the power used by its consumers at any given moment. Currently, this means that despite the existence of solar panels and wind turbines, these must be complemented by controllable but highly polluting energy sources such as coal and natural gas plants. ML methods that are appropriate for time-series predictions, such as Recurrent Neural Networks are particularly suited for these types of tasks (45) and can dramatically lower the barrier to entry for renewable energy globally. Furthermore, even in cases where controllable energy sources are used, demand on the energy grid will still fluctuate based on usage; in this case, ML techniques such as Reinforcement Learning and Dynamic Scheduling can be used to balance the grid in real time (43).

In transportation, reducing activity is a key part of reducing GHG emissions; however, given the highly regional nature of transportation methods (i.e., high-speed trains are only an option in Europe, whereas many major U.S. cities have limited public transportation), custom solutions are needed to make a significant impact. ML can be of particular help in estimating and predicting vehicle flow to minimize it, for example by helping to optimize the design of new roads and hubs (37) and monitoring traffic (23), as well as estimating carbon emissions in real-time (29). ML can also be used for designing more energy-efficient batteries (19), which will become an increasingly important concern as more people switch to electric vehicles. In cases of both energy and transportation, ML can be used to make systems more efficient and to improve predictions of complex phenomena based on large amounts of data; nonetheless, it remains only one part of the solution, and as tempting as it is to halt research projects once a theoretically plausible solution has been found (and a research paper published), what is key here is working with domain experts to bring projects towards deployment, where concrete impacts can be made. Transversal connections between disciplines are therefore key, and must be established and fostered for projects to flourish.

Individuals and Societies

While changes in our climate can be abstract, quantified in degrees of warming or tons of CO₂, climate change will also have very concrete impacts on society, for instance by decreasing crop yields, increasing the frequency of extreme weather events such as hurricanes and storms, and impacting biodiversity. There are a myriad of ways in which ML can help face these issues,

whether it be by analyzing real-time images and recordings of ecosystems to detect species (11) and deforestation (26), improving disaster preparation and response by generating real-time maps from satellite imagery (44), and even setting an optimal price on carbon to accelerate the transition to a low-carbon energy economy (46). Finally, while we are far from being able to predict the exact impact that increasing the carbon tax will have on the different levels of society and industry (i.e., federal and regional governments, local and international companies, and individuals), this is a worthwhile area of research and exploration, with potentially huge consequences in helping political leaders make more informed choices in addressing the climate crisis. It is therefore useful to continue gathering data and building trust between members of the political ecosystem and ML practitioners to learn from each other and to facilitate the deployment of technological solutions in setting government policies.

On an individual level, there are many reasons why individuals cannot, or will not, act on climate change, either common misconceptions regarding the fact that individuals cannot make meaningful impact on a global problem, or cognitive biases that increase an individual's psychological distance to climate change. In the first case, ML-infused tools to estimate the carbon footprint of individuals and households (22) and to model individual behavior with regards to sustainable lifestyle choices and technologies (6) can be very useful if they are sufficiently accurate and deployed on a large scale. Finally, minimizing psychological distance to the future effects of climate change is a promising way to reduce cognitive bias — in this regard, it is possible to use images generated using GANs which represent the impacts of extreme events on

locations that have personal value to the viewer (36). A crucial part of developing ML tools for individuals is, once again, working with multidisciplinary experts in psychology, scientific communication, and user design to ensure that the tools created reach the largest possible audience and maximize their positive impact.

Using AI for a Positive Impact on the World

Technology in general, and ML more specifically, carry great potential for change and disruption. While neither of these is guaranteed to make the world a better place, this potential can most definitely be used to have a positive impact on the world. We have illustrated some inspiring projects that aim to make the world a better place by using the powerful techniques and approaches that ML has brought forward. We believe that as ML researchers and practitioners, we have the responsibility to leverage our (super)powers to contribute to these efforts. This can be done by connecting with established actors from industry and policy or experts from other relevant disciplines, by learning from their past experiences, and by working together to propose innovative solutions to major problems, deployed in places they will have a positive impact.

We live in a world with many global and local challenges and issues that are in constant evolution, and it is easy to be overwhelmed by the flux of information, and to focus on a small sandbox in which we feel safe and in control in order to develop and study the aspects of ML that interest us most. But it is naive to believe that our sandbox is an isolated island that is not connected to the rest of the world, since even in the case of theoretical work, communication and cross-pollination are unavoidable. Each of us is also a citizen concerned with collective debates. Many of us also worry about the world

in which our descendants will live. We believe that there are thought processes that should take place in the head of every ML practitioner regarding the nature of the work they are doing and the potential pitfalls and impacts of this work will have on the world around them, some of which we have listed. And while we do not claim to have all the answers to these tough questions, we hope that we can start a conversation that will accompany ML research and practice throughout its infancy towards its tumultuous teenage years in the coming decades, and eventually towards mature adulthood beyond that.

Author Information

The authors are with the Department of Computer Science and Operations Research and Mila Quebec AI Institute, Université de Montréal, Montreal, Quebec, Canada. Email: luccionis@mila.quebec.

References

- (1) R.G. Abbott, "Automated expert modeling for automated student evaluation," in *Proc. Int. Conf. Intelligent Tutoring Systems*. Berlin, Heidelberg: Springer, 2006, pp. 1-10.
- (2) J. Angwin, J. Larson, S. Mattu, L. Kirchner, "Machine bias," *ProPublica*, 2016; <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>, accessed Nov. 25, 2019.
- (3) F. Arcadu, F. Benmansour, A. Maunz, J. Willis, A. Haskova, and M. Prunotto, "Deep learning algorithm predicts diabetic retinopathy progression in individual patients," *NPJ Digital Medicine*, vol. 2, no. 1, pp. 1-9, 2019.
- (4) T. Bolukbasi, K.W. Chang, J.Y. Zou, V. Saligrama, and A.T. Kalai, "Man is to computer programmer as woman is to homemaker? Debiasing word embeddings," in *Advances in Neural Information Processing Systems*, pp. 4349-4357, 2016.
- (5) J. Buolamwini and T. Gebru, "Gender shades: Intersectional accuracy disparities in commercial gender classification," in *Proc. Conf. Fairness, Accountability and Transparency*, pp. 77-91, 2018.
- (6) S. Carr-Cornish, P. Ashworth, J. Gardner, and S.J. Fraser, "Exploring the orientations which characterise the likely public acceptance of low emission energy technologies," *Climatic Change*, vol. 107, no. 3-4, pp. 549-565, 2011.
- (7) D.S. Chaplot, E. Rhim, and J. Kim, "Predicting student attrition in MOOCs using sentiment analysis and neural networks," in *AIED Workshops*, vol. 53, pp. 54-57, 2015.

- (8) C. Conati, K. Porayska-Pomsta, and K. Mavrikis, "MAI in education needs interpretable machine learning: Lessons from open learner modelling," arXiv preprint arXiv:1807.00154, 2018.
- (9) J. Dastin, "Amazon scraps secret AI recruiting tool that showed bias against women," *Reuters Business News*, Oct. 9, 2018; <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G>, accessed Nov. 25, 2019.
- (10) J. Deng, W. Dong, R. Socher, L.J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *2009 IEEE Conf. Computer Vision and Pattern Recognition*, Jun. 2009, pp. 248-255.
- (11) C. Duhart, G. Dublon, B. Mayton, and J. Paradiso, "Deep learning locally trained wildlife sensing in real acoustic wetland environment," in *Proc. Int. Symp. Signal Processing and Intelligent Recognition Systems*. Singapore: Springer, 2018, pp. 3-14.
- (12) A. Dutt, M.A. Ismail, and T.A. Herawan, "Systematic review on educational data mining," *IEEE Access*, vol. 5, pp. 15991-16005, 2017.
- (13) M. Ghassemi, T. Naumann, P. Schulam, A.L. Beam, and R. Ranganath, "Opportunities in machine learning for healthcare," arXiv preprint arXiv:1806.00588, 2018.
- (14) J.A. Golden, "Deep learning algorithms for detection of lymph node metastases from breast cancer: Helping artificial intelligence be seen," *JAMA*, vol. 318, no. 22, pp. 2184-2186, 2017.
- (15) I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, ... and Y. Bengio, "Generative adversarial nets," *Advances in Neural Information Processing Systems*, pp. 2672-2680, 2014.
- (16) D. Gotterbarn et al., "ACM code of ethics: A guide for positive action," *Commun. ACM*, vol. 61, no. 1, pp.121-128, 2018.
- (17) B.W. Head, "Wicked problems in public policy," *Public Policy*, vol. 3, no. 2, pp. 101-118, 2008.
- (18) Y. Hitti, E. Jang, I. Moreno, and C. Pelletier, "Proposed taxonomy for gender bias in text: A filtering methodology for the gender generalization subtype," in *Proc. First Workshop on Gender Bias in Natural Language Processing*, pp. 8-17, 2019.
- (19) J. Hoffmann, L. Maestrati, Y. Sawada, J. Tang, J.M. Sellier, and Y. Bengio, "Data-driven approach to encoding and decoding 3-D crystal structures," arXiv preprint arXiv:1909.00949, 2019.
- (20) IPCC, *Global Warming of 1.5°C. An IPCC Special Report on the impacts of global warming of 1.5°C above pre-industrial levels and related global greenhouse gas emission pathways, in the context of strengthening the global response to the threat of climate change, sustainable development, and efforts to eradicate poverty*, V.P. Masson-Delmotte et al. Eds. 2018.
- (21) A. Jobin, M. Ienca, and E. Vayena, "Artificial Intelligence: the global landscape of ethics guidelines," arXiv preprint arXiv:1906.11668, 2019.
- (22) C.M. Jones and D.M. Kammen, "Quantifying carbon footprint reduction opportunities for US households and communities," *Environmental Science & Technology*, vol. 45, no. 9, pp. 4088-4095, 2011.
- (23) L.H. Kaack, G. H. Chen, and M.G. Morgan, "Truck traffic monitoring with satellite images," in *Proc. Conf. on Computing & Sustainable Societies*. ACM, 2019, pp. 155-164.
- (24) Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proc. IEEE*, vol. 86, no. 11, pp. 2278-2324, 1999.
- (25) K. Levin, B. Cashore, S. Bernstein, and G. Auld, "Overcoming the tragedy of super wicked problems: Constraining our future selves to ameliorate global climate change," *Policy Sciences*, vol. 45, no. 2, pp 123-152, 2012.
- (26) N.G. McDowell, M.C. Coops, P.S. Beck, J.Q. Chambers, C. Gangogagamage, J.A. Hicke, and A.J. Meddens, "Global satellite monitoring of climate-induced vegetation disturbances," *Trends in Plant Science*, vol. 20, no. 2, pp. 114-123, 2015.
- (27) *Montreal Declaration for a Responsible Development of Artificial Intelligence*; <https://www.montrealdeclaration-responsibleai.com/>, accessed Nov. 25, 2019.
- (28) R. Nkambou, R. Mizoguchi, and J. Bourdeau, Eds., *Advances in Intelligent Tutoring Systems*, vol. 308. Springer, 2010.
- (29) S. Nocera, C. Ruiz-Alarcón-Quintero, and F. Cavallaro, "Assessing carbon emissions from road transport through traffic flow estimators," *Transportation Research Part C: Emerging Technologies*, vol. 95, pp. 125-148, 2018.
- (30) B.D. Nye, "Intelligent tutoring systems by and for the developing world: A review of trends and approaches for educational technology in a global context," *Int. J. Artificial Intelligence in Education*, vol. 25, no. 2, pp. 177-203, 2015.
- (31) I. Perisic, "How artificial intelligence is shaking up the job market," *World Economic Forum*, Sep. 17, 2018; <https://www.weforum.org/agenda/2018/09/artificial-intelligence-shaking-up-job-market/>, accessed Nov. 25, 2019.
- (32) P. Pivovarov, A.J. Perotte, E. Grave, J. Angiolillo, C.H. Wiggins, and N. Elhadad, "Learning probabilistic phenotypes from heterogeneous EHR data," *J. Biomedical Informatics*, vol. 58, pp. 156-165, 2015.
- (33) M.T. Ribeiro, S. Singh, and C. Guestrin, "Why should I trust you?: Explaining the predictions of any classifier," in *Proc. 22nd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining*. ACM, 2016, pp. 1135-1144.
- (34) D. Rolnick, P.L. Donti, L.H. Kaack, K. Kochanski, A. Lacoste, G. Sankaran, ... and A. Luccioni, *Tackling Climate Change with Machine Learning*. arXiv preprint arXiv:1906.05433, 2019.
- (35) Royal Statistical Society (RSS), *A Guide for Ethical Data Science*, 2019; <https://www.actuaries.org.uk/system/files/field/document/An%20Ethical%20Charter%20for%20Date%20Science%20WEB%20FINAL.PDF>, accessed Nov. 25, 2019.
- (36) V. Schmidt, A. Luccioni, S.K. Mukkavilli, N. Balasooriya, K. Sankaran, J. Chayes, and Y. Bengio, "Visualizing the consequences of climate change using cycle-consistent adversarial networks," arXiv preprint arXiv:1905.03709, 2019.
- (37) L.W. Sommer, T. Schuchert, and J. Beyerer, "Fast deep vehicle detection in aerial images," in *Proc. 2017 IEEE Winter Conf. on Applications of Computer Vision (WACV)*, 2017, pp. 311-319.
- (38) M. Szczepański, "Economic impacts of artificial intelligence (AI)," *European Parliamentary Research Service*, PE 637.967, July 2019; [http://europarl.europa.eu/RegData/etudes/BRIE/2019/637967/EPRS_BRI\(2019\)637967_EN.pdf](http://europarl.europa.eu/RegData/etudes/BRIE/2019/637967/EPRS_BRI(2019)637967_EN.pdf), accessed Nov. 25, 2019.
- (39) P. Suppes and M. Morningstar, "Computer-assisted instruction," *Science*, vol. 166, no. 3903, pp. 343-350, 1969.
- (40) IEEE, *The IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems. Ethically Aligned Design: A Vision for Prioritizing Human Well-being with Autonomous and Intelligent Systems*, 1st ed., 2019; <https://standards.ieee.org/content/ieee-standards/en/industry-connections/ec/autonomous-systems.html>, accessed Nov. 25, 2019.
- (41) UN General Assembly, *Universal Declaration of Human Rights*, UN General Assembly, vol. 302, no. 2, 1948.
- (42) UN High Commissioner for Refugees (UNHCR), *The Sustainable Development Goals and Addressing Statelessness*, 2017; <https://www.refworld.org/docid/58b6e3364.html>, accessed Nov. 25, 2019.
- (43) J.R. Vázquez-Canteli and Z. Nagy, "Reinforcement learning for demand response: A review of algorithms and modeling techniques," *Applied Energy*, vol. 235, pp. 1072-1089, 2019.
- (44) S. Voigt, T. Kemper, T. Riedlinger, R. Kiefl, K. Scholte, and H. Mehl, "Satellite image analysis for disaster and crisis-management support," *IEEE Trans. Geoscience and Remote Sensing*, vol. 45, no. 6, pp. 1520-1528, 2007.
- (45) C. Voyant, G. Notton, S. Kalogirou, M.L. Nivet, C. Paoli, F. Motte, and A. Fouilly, "Machine learning methods for solar radiation forecasting: A review," *Renewable Energy*, vol. 105, pp. 569-582, 2017.
- (46) S. Wei, Z. Chongchong, and S. Cuiqing, "Carbon pricing prediction based on wavelet transform and K-ELM optimized by bat optimization algorithm in China ETS: The case of Shanghai and Hubei carbon markets," *Carbon Management*, vol. 9, no. 6, pp. 605-617, 2018.
- (47) J. Wexler, M. Pushkarna, T. Bolukbasi, M. Wattenberg, F. Viégas, and J. Wilson, "The What-If Tool: Interactive probing of machine learning models," *IEEE Trans. Visualization and Computer Graphics*, 2019.
- (48) J. Wiens, S. Saria, M. Sendak, M. Ghassemi, V.X. Liu, F. Doshi-Velez ... and P.N. Ossorio, "Do no harm: A roadmap for responsible machine learning for health care," *Nature Medicine*, vol. 25, no. 9, pp. 1337-1340, 2019.
- (49) R. Zemel, Y. Wu, K. Swersky, T. Pitassi, and C. Dwork, "Learning fair representations," in *Proc. Int. Conf. Machine Learning*, 2013, pp. 325-333.

