Bits&BIPOC

A Sociotechnical Guide to Data, Algorithms, Machine Learning and AI

Lorena Almaraz De La Garza

Bits&BIPOC

A Sociotechnical Guide to Data, Algorithms, Machine Learning and AI

Lorena Almaraz De La Garza

The typefaces are Dosis (© Impallari Type) and Nunito (© Adams, Cyreal, Le Bailly).

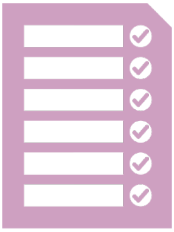

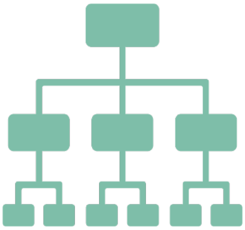

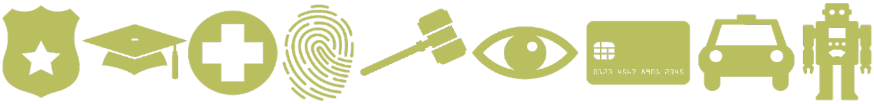
0123 4567 8901 2345

INFORMATION

turned into

DATA

feeds the

ALGORITHMS

that make up

MACHINE LEARNING

which is a kind of

AI

Table of Contents

Hello
About the author

Lorena's work stems from, and expands upon, research that reveals the relationships between power, technology, education, and justice. She received a Master of Information degree from the University of Toronto's Faculty of Information, where she focused on Human-Centred Data Science and Critical Information Policy. Prior, Lorena developed and implemented publicly funded programs in culture, education, and interdisciplinary research. As a Mexican immigrant, a woman of colour, and a settler in Canada, Lorena is committed to working in favour of equity-seeking and systemically excluded groups. She holds a Bachelor of Arts from the University of Waterloo in Fine Arts and Cognitive Science. Contact Lorena via email at *l.almaraz@mail.utoronto.ca* or LinkedIn: *linkedin.com/in/lalmaraz*.

&Welcome
Letter to the reader

Dear Tech Justice Enabler,

What do you think when you hear that data is the new oil? The phrase means that data are an untapped resource that can propel humanity forward, simplifying and bettering our lives in unforeseen ways. It is used to entice businesses, governments, and others to collect, analyze, and capitalize on data as much as possible. The metaphor is incomplete; thinking of data as a natural resource ignores all the human activity that makes data — and the use of data — possible [1]. And yet, the metaphor indirectly points at something powerful. For centuries, peoples lived on this land without the need for oil. They continue to do so. While the relentless extraction of oil has not translated into benefits for everyone, associated harms have been universal. The metaphor inadvertently contains the less-marketable realities of oil and data: unchecked practices of extraction jeopardize what we need to thrive. In the case of oil, the equilibrium of the natural world is endangered; in the case of data in the digital ecosystem, our privacy, autonomy, consent, and other similar values are threatened. The greatest insight hidden in the metaphor is that we must seek to understand not only the immediate benefits of extraction, but also the consequences across time, geographies, and ways of life.

As you read through these pages, I encourage you to question our current reliance on data and data-driven technologies. The *Reflect* and *Envision* sections can help as a start; go back and forth as you increase your knowledge to revise your thoughts. As you think through the prompts in each chapter, I kindly ask you to trust your lived experience, what you know to be true, what you value, and what you stand against. Are your perspectives represented or missing in this text? How about in the technology sector at large? Your experience is necessary if we want to answer an urgent question: how do we design technologies that more closely align with all of us?

I intended these pages for racialized youth in Canada, based on what I wish I had known when I was younger, as a Latina curious about emerging technologies but without access to technical education. I hope that through these pages we can build the confidence in our technical language and expertise to make immediate demands for the improvement — or abolition — of the technical systems that harm us. More importantly, I hope that the prompts in this text serve us, even as a small step, as we collectively envision the alternate realities that use data and data-driven technologies for justice and equity. I believe wholeheartedly that through critical engagement, self-advocacy, and civic participation, we can turn those visions into reality.

A World of Information

Information is *everywhere*

Information turned into data feeds the algorithms that make up machine learning, which is a kind of artificial intelligence. That is a mouthful! Throughout these pages, we will explore what all these buzzwords mean. We are interested in the significance of these terms both technically, as used in fields like computer and data science, and socially, as used colloquially and in media. While the terms might be familiar to many, they are grounded on a very particular and limited understanding of the world. This worldview has been dominantly crafted by, as Métis community builder Alexander Dirksen describes, "the forces of colonialism, patriarchy, and Whiteness" [1] and operationalized by systemic exclusion in the technology sector.

Calls to decolonize data and data-driven technologies are fundamental to justice. To take on this important work locally, we must understand the relevant terminology, societal impacts, and current legislative frameworks in a Canadian context. From this, we can reimagine our approach to data, algorithms, machine learning, and artificial intelligence, ensuring that technological advancement leaves no one behind. We dedicate a chapter of this guide to each of those terms and move through the text in increasing levels of complexity. For each term, we begin with a definition and a prompt for reflection based on what we already know. Then, we expand our knowledge with examples, common concerns, current regulations, and potential avenues for action. We end each chapter with an invitation to envision alternatives that uphold human values. To start, we look at the element that makes even the most complex AI system possible: data.

DATA

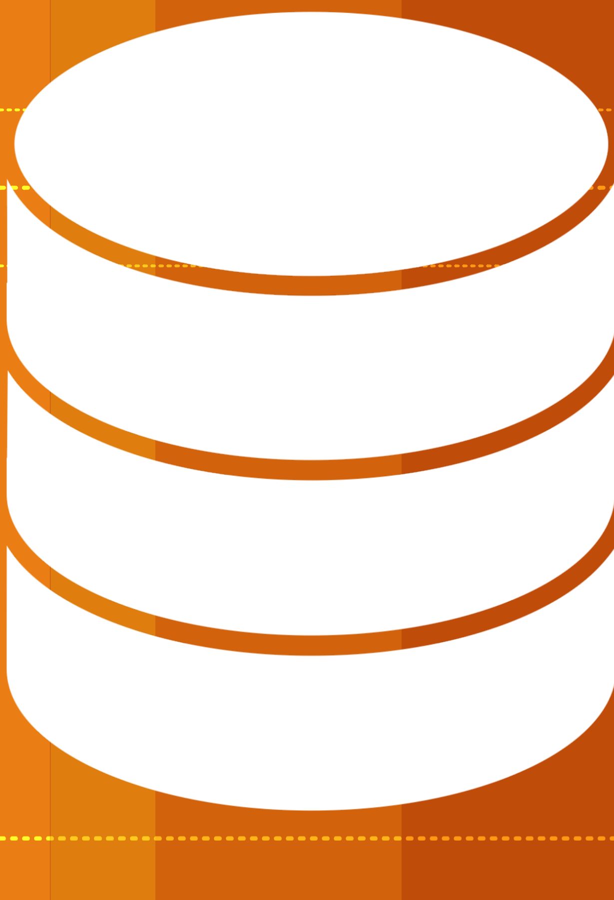

Define

What do we mean by *data*?

Data is a plural noun but is often used as singular. It refers to the grouped pieces of information that we collect, organize, measure, and analyze to make sense of the world. Each of these is a datum, although the word is not very common. In a computer, data is represented digitally in the form of binary units — bits. Data can be expressed either in numbers, as quantitative data, or in words, as qualitative data. Aside from letters and numbers, data can be images, video, or audio. A data set is a collection of information in a digital format. A data set that is character-based, for example, looks like a table in a spreadsheet made up of columns, sometimes called features or attributes, and rows, known as instances or observations. Very large or complex data sets require much more computational power than a regular home computer —sometimes thousands of servers working in parallel — and novel approaches for analysis, storage, and modelling. The field that specializes in working with these data sets is known as big data.

Reflect
What are your thoughts on *data*?

Data are fundamental to the digital economy [1]. Take a minute to think about examples of data in your life, such as the list of YouTube videos you have watched; the hours you have spent playing a video game; the categories of items you bought at a store; the voice messages you have sent on WhatsApp; the photos stored in your phone; the number of people (and bots) that follow you on social media... the list is infinite. What kinds of decisions had to be made to collect, transform, analyze, and extract value from your examples?

As you think about this, keep in mind the values of privacy, ownership and property, and informed consent. Friedman and Hendry [2] provide us with the following definitions:

PRIVACY

"Refers to a claim, an entitlement, or a right of an individual to determine what information about himself or herself can be communicated to others."

OWNERSHIP & PROPERTY

"Refers to a right to possess an object (or information), use it, manage it, derive income from it, and bequeath it."

INFORMED CONSENT

"Refers to garnering people's agreement, encompassing criteria of disclosure and comprehension (for 'informed') and voluntariness, competence, and agreement (for 'consent')."

Learn

Here is an example of *data*

Data collection and analysis are necessary for a lot of life as we know it. Every five years, for instance, Statistics Canada undertakes a massive data collection effort: the Census of Population. Through the census, every household is required to share information about their demographic, social, and economic characteristics. The data are vital to understanding changes over time and making decisions about public services across all levels of government. The raw data — or the unchanged source data — are summarized for statistical analysis through data aggregation, which presents data about groups rather than about individuals. The resulting data are published so that researchers, community organizations, businesses, and others may use them as well. Statistics Canada follows a Privacy Act and a Statistics Act which include, among technical measures like data aggregation, secrecy clauses for all employees to protect respondents' privacy [3].

Question
What could go wrong with *data*?

But of course, data collection is not limited to governments. Corporations constantly collect inordinate amounts of data. Some companies, like Meta, Google, and Amazon, even allow users to download our data archives [4][5]. The download feature is usually found under the platforms' user privacy settings, and if you use it, you will be able to download all the text, image, voice recording and even location data that these platforms have collected from you. However, the main value of data is not just my cat photos and messages to my mom. Corporations get a lot of value out of making inferences about larger groups of people, using not only my data but those of people I am similar to, or radically different from, to learn how to better target their services for us [1]. This process requires not only the data that one product or platform collects but the ability to triangulate, that is, cross-reference, several sources of information [1]. This is where big data and algorithms for computational inference come in. We will get into algorithms and machine learning in the following chapters, but for now, we will start with an example. Sociologist Zeynep Tufekci [6] explains that even if you do not have an account on Facebook or a Meta-owned platform like WhatsApp or Instagram, it is almost certain that people that have your phone number do. If they have allowed these platforms to access their contacts, Facebook can place you within the social network quite accurately. Tufekci elaborates, "Facebook even keeps 'shadow' profiles of nonusers and deploys 'tracking pixels' situated all over the web — not just on Facebook — that transmit information about your behavior to the company" [6].

Have you clicked away a pop-up banner letting you know that the website you were visiting uses cookies and similar technologies? Cookies are a clever way of using and collecting data. They are data packets that websites can store in your browser. Cookies make using certain sites more convenient, like the weather website that remembers your postal code when you open it again, or the online store that recalls the contents of your shopping cart. First-party cookies are those that the website you visited places on your browser. Third-party cookies, on the other hand, are those that are placed by the advertisers on the sites you visited. For example, if you visit a website that has five ads, those five advertisers might in turn create five cookies to track your browsing behaviour, even after you leave the site the cookies originated from. Cookies are one of the reasons you see the same ads across several unrelated websites while your friend sees different ads altogether. Cookies are a privacy concern because user profiling has gotten very sophisticated over the years: my browsing behaviour, across several websites and over time, can say much more about me than I would be comfortable sharing with any given company. More important, before the introduction of these pop-ups, no one really asked us if we were okay with being tracked in this way. The pop-ups were triggered by an effort to increase informed consent, transparency, and privacy protection measures online [7].

All healthy human relationships require consent, transparency, and respect for privacy. So how come corporations do not follow these basic principles? Because data is deemed a valuable resource for profit, values like privacy and consent are often transgressed — and not only online. While online privacy protections are crucial, offline privacy is important too. In 2020, the Privacy Commissioner of Canada published a report that stated that Cadillac Fairview — the real estate company that owns popular malls like the Eaton Centre in Toronto — had collected five million images of shoppers without meaningful consent [8][9]. A photograph of a person's face is considered sensitive biometric information because it is personally identifiable, just like a fingerprint, a video of someone's gait, or the unique sound of their voice [10]. With data from cameras on digital information kiosks, the company was using facial recognition technology to estimate shoppers' age and gender [8][9][11]. Although this is a clear transgression of human values, privacy regulators were not able to fine the company under current Canadian law [11].

Resist

What can we do about *data*?

In Canada, protections for personal data are limited, especially when compared to other jurisdictions like the State of California or the European Union. We have the Personal Information Protection and Electronic Documents Act (PIPEDA), which requires private-sector organizations to get our consent when they collect, use, or disclose our personal information for commercial purposes [12]. Alberta, British Columbia, and Quebec have similar legislation that supersedes PIPEDA, and public sector entities are covered by separate privacy acts that apply to each jurisdiction, including the federal government [13]. PIPEDA lays out ten principles that businesses must follow to protect our personal information, including establishing accountability, requiring consent, limiting collection, and other practices. It also gives us the right to access the personal information that an organization may hold about us. PIPEDA has some similarities with the General Data Protection Regulation (GDPR) in the European Union; for instance, both require companies to inform users of data breaches [14]. Key differences, however, include that GPDR applies to public bodies as well, while PIPEDA only applies to private organizations; GDPR provides a right to data erasure and PIPEDA does not; and importantly, GDPR has supervisory authorities that determine financial penalties for non-compliance, while PIPEDA does not directly allow the Office of the Privacy Commissioner (OPC) to administer fines [14]. Rather, the OPC may investigate and publish a report with recommendations that can then be used in court to determine penalties for non-compliance [14]. PIPEDA is being updated with the June 2022 introduction of Bill C-27 which reworks Bill C-11, the Digital Charter Implementation Act [15]. We will learn more about Bill C-11 and Bill C-27 in the *Resist* section in the Machine Learning and AI chapters.

Change
We can make *data* better

We must continue to improve data practices and protections in Canada. Aside from engaging in the ongoing debate on regulations for the use of data, I also encourage you to learn about the principles of Indigenous Data Sovereignty and Indigenous Data Governance, remembering that data are as much social as they are technical [16][17]. As explained in these principles, data has a dimension of power. To learn about how other groups have been impacted by data practices and about work that subverts those practices, two great books are *Dark Matters* by African & African Diaspora Studies professor Simone Browne [18] and *Data Feminism* by academic duo Catherine D'Ignazio and Lauren F. Klein [19]. If you are curious about unique careers in data, pick up a programming language like R or Python and pair it up with communication and advocacy skills to become a Data Journalist [20][21]. Data journalism uses the power of data analysis and visualization to tell stories. Newsrooms like ProPublica [22] and Al Jazeera [23] keep stellar data journalists on staff; take a look at their sites for inspiration.

&Envision
An equitable future with *data*

Think about the many ways in which data has made things easier for some but harder for others. How might we ensure equity and justice in data practices? What does an inclusive future of data look like? In your ideas, remember to uphold the values of privacy, ownership and property, and informed consent.

ALGORITHMS

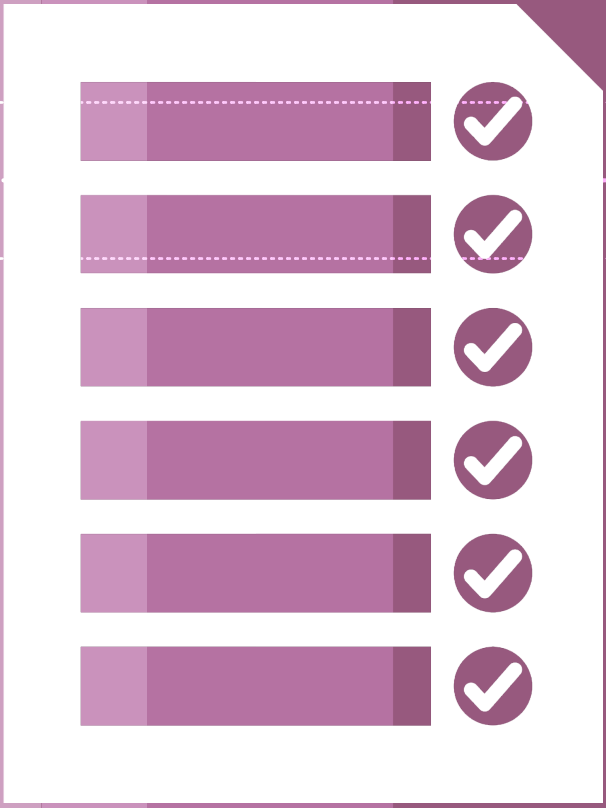

Define
What do we mean by *algorithms*?

The word “algorithm” has become popular in media with the rise of digital tools and platforms, although the concept has existed in mathematics for centuries. Put simply, we can think of an algorithm as a sequence of steps or instructions that take an input, operate upon it, and provide an output. In this way, computer algorithms can make up the simplest model and the most complex program. An algorithm exists to accomplish a task or to solve a problem, and while there are several ways of achieving the same goal, people often work towards efficiency in algorithms, whether that is in terms of time, memory, or other resources. But when we think about working with data about people, efficiency might not always be the best thing to prioritize! While things like speed are easy to measure and optimize for, working towards fairness and equity is not as simple. Therefore, algorithms that involve data that relate to people must be carefully designed.

Reflect

What are your thoughts on *algorithms*?

How has the word algorithm come up in your life recently? Do you have an example of an algorithm doing its job? Think deeply about how it might have been designed. What might it be optimizing for? As you bring this to mind, consider the values of self-direction and autonomy:

SELF-DIRECTION

"Independent thought and action — choosing, creating, exploring. Self-direction derives from organismic needs for control and mastery and interactional requirements of autonomy and independence." [1]

AUTONOMY

"Refers to people's ability to decide, plan, and act in ways that they believe will help them to achieve their goals." [2]

Learn

Here are examples of *algorithms*

A *Computer Science 101: Introduction to Algorithms* course would probably include the very basics: sorting and searching through data. Sort algorithms order the elements of a list based on different criteria. There are several kinds of sorting algorithms, like bubble sort, which compares the first two elements in a list and swaps them if needed. If you are using bubble sort to organize a list of numbers in ascending order, for instance, the algorithm will start by comparing the first and second numbers. If the first number is larger than the second, it will swap their places. It will keep going with every number in the list, which makes it neither fast nor efficient (imagine doing this with thousands of numbers). A search algorithm has a different goal, it looks for or retrieves an element from a list. A linear search algorithm, for example, will check every element in a list sequentially until it finds the one we asked for. A binary search algorithm, on the other hand, searches at intervals. It first looks for our element in the middle of the list, and if that is not it, the algorithm searches on the one half of the list closest to our element. Then, it keeps splitting the search space until it reaches its goal.

As you can imagine, algorithms can get very complicated, especially if the goal they are trying to reach involves many different sources of information and priorities. Think about dating apps like Tinder, Bumble, Hinge, and OkCupid. Do you think all their algorithms work in the same way or are they prioritizing different things? Unless we get a job at any of those companies, we cannot know how the proprietary algorithms work for sure, but we can venture a guess. OkCupid seems to calculate a match percentage based on similar preferences, perhaps checking whether the answers to the service's questionnaire are compatible between two users, or if the items that they listed as must-haves in a relationship are similar. Tinder does not have a long questionnaire, so it must be using other data. A spokesperson for the company said that the factor that most influences the outcome is how much time a user spends on the app [3]. According to her, the more the better. But of course, the more time you spend on Tinder, the more data they collect about you, which is certainly better... for them! After all, we cannot really know what each of these proprietary algorithms is optimizing for, so we have to think critically when using these platforms.

Question

What could go wrong with *algorithms*?

As often goes with human-made things, there is lots of room for error with algorithms. Critically questioning the decision-making that was coded into a computer algorithm is a fundamental skill as we navigate the digital ecosystem. Speaking of dating apps, how do you think algorithms might have a negative impact here?

Researchers at Cornell University found that dating apps can reinforce biases or sexual racism, which is the exclusion of potential partners based purely on their racial identity [4]. In their study, they found that most apps use algorithms that emphasize a person's past preferences and even match them with people with similar demographic traits. Jessie Taft, a co-author of the study, explains that minorities might be excluded as a result, since users "may find their matching results artificially limited by an algorithm that's calculated to repeat 'good' past matches without considering what 'good' future matches might be" [4].

Similarly, the recommender systems used in many platforms can artificially limit our options. Collaborative filtering is a technique used in recommenders where the goal is to predict, or filter, the interests of one person based on the interests of many other people. The assumption is that if I liked a documentary about tropical birds and other people who liked it also liked a show about life in Colombia, it is likely that I might like that show too. After watching the documentary, the show will come up next as the top suggested item for me — and perhaps even auto-play. This limits our ability to choose what we consume or to serendipitously find content; what we watch has been automagically determined for us.

Limiting choice comes with consequences, and there are many terms to define the negative impacts. “Feedback loop”, “filter bubble”, “echo chamber”, and similar terms refer to, in brief, the effect of being presented only with content that aligns with our preferences, isolating us from different ideas, perspectives, and people. This kind of isolation happens across platforms. For instance, a researcher at the University of California found that the recommended accounts on a user’s TikTok had profile pictures that matched the race, age, or facial characteristics of the people the user already followed [5]. But in a world where divisions keep increasing and discord is becoming concerningly common, perhaps being exposed to different opinions and perspectives, rather than existing in the narrow echo chambers of the digital world, is exactly what we need for empathetic, healthy debate.

Resist

What can we do about *algorithms*?

Who makes sure algorithms are not causing harm? In Canada, we do not have legislation that addresses algorithms specifically. However, as part of the Directive on Automated Decision-Making, the government has created an Algorithmic Impact Assessment (AIA) Tool to be used by government agencies or departments [6]. The tool is not binding but allows stakeholders to understand the impact level of any given automated decision system. It has 48 questions about risk and 33 about mitigation, and it provides a score based on their answers. Users are required to publish the results of the AIA on their Government of Canada websites as part of Canada’s Open Government efforts.

Change

We can make *algorithms* better

Canada needs to step up its algorithmic protection game. Now that you have the language to describe how algorithms work, and how they might fail, you can contribute to this effort by keeping abreast of proposed legislations and participating in public consultations, activism, and education. Keep learning about the impacts of algorithms with internet scholar Safiya Noble's book *Algorithms of Oppression* [7] or *Black Software* by media professor Charlton McIlwain [8]. For work that shifts power through collective action, follow the Algorithmic Justice League [9] and consider a career bridging public policy and technical skills. The job of a Technology or Technical Policy Advisor, Analyst, or Officer is to understand the complexity of technological systems and make recommendations on how best to implement them.

&Envision

An equitable future with *algorithms*

Algorithms have increased efficiency across an unimaginable number of processes, but they have also limited our choice and reduced our agency. How might we strike a balance? Take a few minutes to ideate, keeping in mind the values of self-direction and autonomy.

MACHINE LEARNING

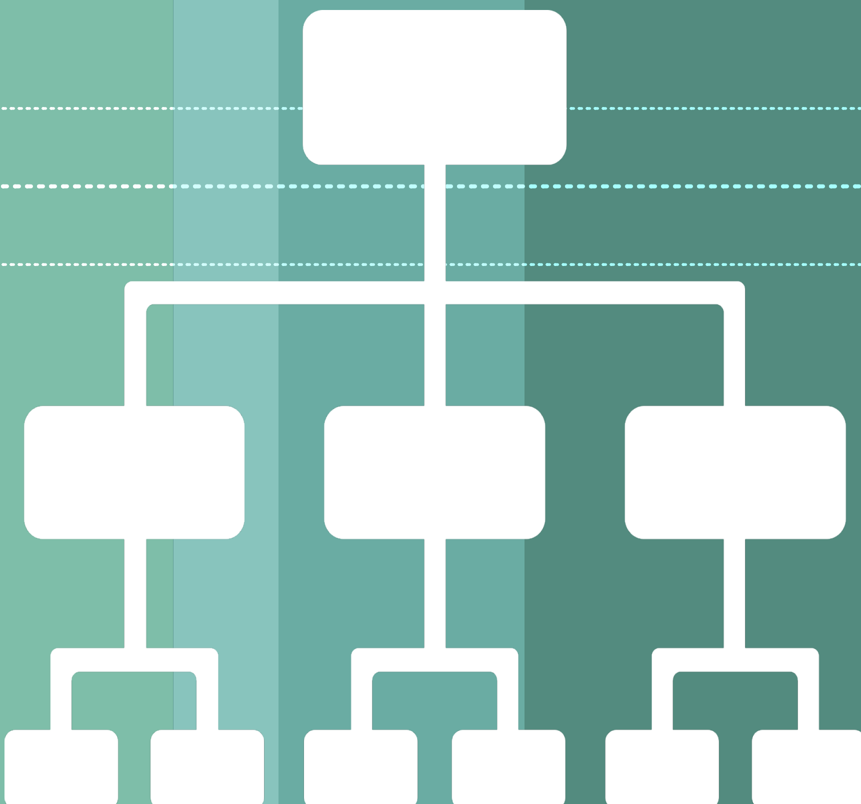

Define

What do we mean by *machine learning*?

As any other algorithm, machine learning algorithms have a task to accomplish. What makes them unique is that machine learning algorithms do not require explicit step-by-step instructions on how to solve the task. We set the goal, and with a lot of data and statistics, the algorithm generates its own instructions to solve the task. We call it "learning" because the procedure can improve its performance based on its previous iterations, just like we learn from experience. In general, there are two kinds of machine learning algorithms, supervised and unsupervised. These do not refer to whether a person is keeping an eye on the algorithm to make sure it is doing its job properly, although that is certainly important. Supervised learning means the input, or data, comes with additional information in the form of labels. For instance, if I have 1,000 pictures of mountains and label them all "mountain" and another 1,000 of lakes labelled "lake", if I ask the algorithm to classify a new picture it has never seen before, telling me whether it is a lake or a mountain, the model is supervised. Unsupervised means that there are no labels; we just have 2,000 photos and use an algorithm to, for instance, separate them into two groups based on their similarities or differences. Who knows, maybe it decides that images with a lot of blue should be separate from those with a lot of green, so it accomplishes the task without us knowing exactly how it went about it. Supervised machine learning is used for classification and numeric prediction, while unsupervised is used for clustering and pattern discovery.

An important terminology distinction is “algorithm” versus “model”. An algorithm is the set of instructions, while a machine learning model is the output of the algorithm once it has used data to “learn” something. In our mountains and lakes example, once our labelled data have been fed into the algorithm and it accomplishes its task, we will have an image classification model. We can save that model and keep using it to classify new photos.

A powerful subset of machine learning that uses slightly different data, usually messier or unstructured, is deep learning. Rather than a human determining the main features, deep learning models determine their own. Deep learning neural networks are a series of algorithms that use many layers, each based on the previous one, to make the predictions more accurate and reliable.

Reflect

What are your thoughts on *machine learning*?

Think of a task that you think machine learning would help you solve. It can be anything where you need to find a pattern, classify differences, predict numerical values, or separate a large group into smaller clusters. As you consider how machine learning could help, think about the values of accountability and trust, especially since we cannot know all the steps the machine created to solve the task. Here are some definitions of these values [1]:

ACCOUNTABILITY

"Refers to the properties that ensure that the actions of a person, people, or institution may be traced uniquely to the person, people, or institution."

TRUST

"Refers to expectations that exist between people who can experience good will, extend good will toward others, feel vulnerable, and experience betrayal."

Learn

Here is an example of *machine learning*

We talked about recommenders earlier, and although some might be using algorithms that are programmed step by step, it is safe to assume that many of them use machine learning techniques. Think about what the screen looks like when you launch Netflix, for instance. It has a list of the content you watched in the past, another with content trending in your region, and another with the content they think you specifically will enjoy. For that last one, Netflix has a proprietary, complex recommendation system that estimates the likelihood that you will watch a particular title [2]. Estimating the likelihood of something is a numeric prediction! That sounds a lot like machine learning. Some of the factors that Netflix takes into account are your ratings and viewing history, the preferences of others with similar tastes, and details about the TV show or movie, among other things. Those data are fed into an algorithm (or perhaps, several of them) which creates a model — a likelihood estimator — that is used with new users, even if their preferences are something the model has not seen before. Netflix re-trains its algorithms with new data to improve the accuracy of the prediction.

Question

What could go wrong with *machine learning*?

Machine learning is also used in dynamic pricing. Perhaps the term does not sound familiar, but if you have used a ride-hailing app, you have experienced it at work. Dynamic pricing is also called demand pricing, and it refers to, for example, the changes in cost that happen due to how many people want a service at a given time. This is exactly what Uber's surge pricing is — it is much more expensive to get an Uber on a rainy day after leaving a busy Raptors game than it is to get one on a regular weekday afternoon.

Dynamic pricing is tricky. Charging more for something that a lot of people want, or making it cheaper if there is less demand, is a common tactic for businesses and services. But when a machine learning model is involved, it is hard to know how exactly those prices are calculated. Last summer, two friends going out in Toronto each ordered an Uber from the same house to the same nightclub, at the same time [3]. To their surprise, one of them was quoted at $57.80 and the other at $32.58. That is a big difference! They shared their story on Twitter and hundreds of people commented with similar experiences.

What makes matters complicated is that the person charged a higher rate is a Black woman, while the other person is a white man. Technological systems that give advantages to some while excluding or actively harming others have been called into question by activists, scholars, and the public for decades [4][5][6][7]. The racism that unintentionally or intentionally gets embedded into technologies is a real, tangible problem, and while we do not know whether that might have played a role in this particular example, the lack of transparency does not help. According to the company, their price is determined using "many data points" but they do not provide further details [8].

An issue with machine learning algorithms is that they can be a "black box", where it is not possible to have absolute certainty of which features were prioritized while making the prediction. This means we have to not only trust the algorithm but also the data used to train the model, which might not always be a reasonable ask. Intentional transparency and clear accountability from companies and services should come first. Recently, experts have called for transparency, explainability, and interpretability of machine learning models [9][10][11][12]. Rather than simply focusing on efficiency, experts ask for openness on the model's code and goals, and for its internal logic and predictions to be understandable by humans.

Resist

What can we do about *machine learning*?

Machine learning is only addressed in our legislation as part of the broader umbrella of AI, as we will explore in the next chapter. For now, it is important to know that often legislation is designed to protect or advance consumer interests — not necessarily to uphold human rights. For instance, although the now-defunct Bill C-11: *Digital Charter Implementation Act of 2020* was meant to enhance privacy, it was not comprehensive enough [13]. While it would have afforded the OPC greater authority to enforce the Commissioner's decisions, it did not adequately address meaningful consent, transparency, nor strengthen our rights to access and correct our personal information. It also did not protect whistleblowers, who have uncovered some of the most noxious harms of surveillance-for-profit technologies. Importantly, Bill C-11 did not provide us with a right to object to how automated decision systems are used. Ultimately, Bill C-11 did not move forward. With rapid technological development, legislation struggles to keep pace and is reactionary.

Change

We can make *machine learning* better

Many of the platforms and tools we interact with in the digital ecosystem rely on machine learning. Continue learning about these, their uneven effects, and what we can do about them with the book *Digital Black Feminism* [14] by communication and technology professor Catherine Knight Steele. Level up your Python skills with machine learning libraries; you will want to pick up *pandas*, *NumPy*, and *scikit-learn* to start [15][16][17]. If you are excited about applying machine learning skills critically, consider a career as a Machine Learning Research Scientist. In this job, you would be responsible for imagining, developing, and maintaining machine learning models. Understanding their flaws and being able to communicate them prioritizing human values is a must.

&Envision

An equitable future with *machine learning*

Machine learning allows us to use massive amounts of data to make sense of the world, make better-informed decisions, and solve problems in ways we had not imagined. It also relies on the values of trust and accountability, as we cannot always explain or interpret the models' predictions. How might we harness the power of machine learning while ensuring our values are prioritized?


AI

Define

What do we mean by *AI*?

The topic of computer-facilitated artificial intelligence or AI has intrigued people since at least the 1950s [1][2]. The term has roots in computer science, philosophy, and cognitive science. In general, it refers to machines mimicking human intelligence, but AI as known today, usually refers to technologies that can be applied to improve systems through the use of large amounts of data and computational power. Machine learning is an example of AI, but AI is much more than just machine learning. AI research is broken up into several subfields, including reasoning, knowledge representation, perception, planning, and learning. AI is an exciting field that gets lots of attention, and for good reason. Not only has it revolutionized how we use computational power, but it is also aspirational. The premise is that maybe one day, we will be able to understand human intelligence so well that we will be able to program it — this is called "artificial general intelligence". Right now, AI represents the combination of technological advances in machine learning, computer vision, speech recognition, natural language processing, expert systems, and other applications.

Reflect
What are your thoughts on *AI*?

How can AI change, for better or worse, life as we know it? Thinking about your own experience, what kind of applications of AI would be positive and which would be negative? In which fields do you imagine the most promise, and in which the most risks? What kind of resources would be required to develop an AI system, including money, time, knowledge, labour, planetary, etcetera? As you think about this, consider the values of power and universalism [3]:

POWER

"Social status and prestige, control or dominance over people and resources."

UNIVERSALISM

"Understanding, appreciation, tolerance, and protection for the welfare of all people and for nature."

Learn
Here is an example of *AI*

Virtual assistants like Alexa, Siri, or Google Assistant are great examples of AI.

> *Wait* — why are they always female voices by default? Perhaps that is a question for another time, but for now, let's remember that technical systems do not exist, nor are developed, in a vacuum, and can therefore encode and uphold social stereotypes, like female bodies being the primary providers of the labour of care and service. But let's get back to AI!

These systems can recognize your voice, make sense of what you are saying, process your request, use an unimaginable number of resources to search and provide an answer, or at least continue the conversation, and they do it all in seconds. They even learn about the acoustics of the environment in which they are used [4]. A lot is going on under the hood, and yet, getting help from virtual assistants is easy and convenient. Amazon explains that Alexa gets smarter and more helpful every day. Amazon uses our requests to train its speech recognition and natural language processing systems using machine learning. And as we have learned, it is not only our requests that are taken into account but also the millions of requests of the many thousands of users of the Amazon Echo. Amazon tries to build transparency features into Alexa; for instance, if you get an odd response from Alexa that you do not understand, you can ask "Alexa, why did you do that?" for insight into the system's processes. These extraordinary applications of complex technologies are an example of one of the key features of AI: scale. AI works best with huge amounts of detailed information.

Question
What could go wrong with *AI*?

Who controls, analyses, and capitalizes on the myriad pieces of information required for AI systems? Usually, it is the corporations that develop the mechanisms to collect such data and process them — not the many people that generated the data in the first place, like the users that create the voice requests that make Alexa smarter. And how about privacy? How much information about us does Alexa need to be helpful to us and others? The many values we have discussed throughout these pages are all relevant to the applications of AI. Sometimes, the quest for personalized technologies comes at the cost of those very values.

Fortunately, researchers and activists are coming together to limit the overreach of AI technologies. In 2021, over 180 musicians and human rights organizations signed a letter to limit the application of speech-recognition technology for personalized content due to privacy and human rights concerns [5]. In this case, the company that patented the technology is Spotify. The technology was designed to analyze users' voices to be able to suggest songs based on their emotional state, gender, age, or accent. The company said it did not implement the patented technology nor plans on doing so, but the technology exists regardless.

Limiting the uses of AI to ensure human values are upheld is paramount because as we know, the use of AI is not limited to entertainment platforms. AI technologies can, and are, applied to an immense number of tasks. AI is used to interview job candidates, censor hate speech online, predict the shape of proteins to develop pharmaceuticals, change the way we conduct mental health therapy, and many other novel applications [6][7][8][9].

Because of its potentially life-changing effects, policing is a particularly problematic area for the use of AI due to the systemic inequity in law enforcement practices. In 2020, Toronto Police admitted to using a controversial facial recognition tool, Clearview AI, one month after having denied doing so [10]. The software was also used by more than 30 Canadian law enforcement agencies, including agencies in Halifax, Hamilton, and London, as well as the Ontario Provincial Police [11]. Clearview AI uses facial recognition to provide police services with a person's name, phone number, address, and other details [10][12]. The software was found to have extracted "more than three billion photos from public websites like Facebook and Instagram and used them to create a database used by more than 600 law enforcement agencies" [10] across the world.

While the promise of AI is exciting, the repercussions of faulty systems are unjustifiable. The inaccurate prediction of an AI system used by police can harm an innocent person permanently, and even an accurate prediction circumvents the processes established to ensure power is not abused by law enforcement. Clearview AI and similar facial recognition technologies have been deemed illegal by the OPC, as they are a tool for mass surveillance [12].

Resist
What can we do about *AI*?

The OPC has proposed to reform PIPEDA to include a regulatory framework for AI [13]. The framework prioritizes rights, making sure the use of AI is solely for socially beneficial or legitimate commercial purposes. It highlights privacy as a human right, but also includes the right to a meaningful explanation of automated decision-making and the right to contest automated decisions, among other priorities. Importantly, the framework proposes meaningful enforcement with real consequences for non-compliance. As we mentioned before, the OPC does not currently hold the power to levy fines or ensure compliance with its recommendations.

Internally, the Government of Canada has its own guiding principles for the responsible use of AI [15]. Government agencies and departments are expected to abide by these, which include understanding and measuring the impact of AI, being transparent about the use of AI, providing meaningful explanations, being open by sharing source code and relevant information, and providing training for employees on responsible design.

In June 2022, the Minister of Innovation, Science and Industry introduced Bill C-27, the successor of the defunct Bill C-11 we mentioned in the previous chapter, proposing a new Artificial Intelligence and Data Act (AIDA). It is not a law yet, but if it were, it would determine penalties for irresponsible or improper use of AI. It would also enhance protections by determining a system for accountability: AIDA defines requirements like the use of anonymized data, states that whoever develops an AI system is responsible for it and that systems suspected to violate the act might be audited, among other measures [14]. While it may seem like a step in the right direction, critics highlight that the bill as proposed does not include an independent oversight system — it relies solely on a politically-appointed public servant, a new AI and Data Commissioner, and not a panel of non-partisan subject-matter experts across areas like social justice, anti-racism, or intentional inclusion [15]. Hopefully, legislation will recognize the overlapping spheres of oppression to which people are subject, and that AI systems have the potential to exacerbate.

Change
We can make *AI* better

AI is an incredibly promising field. To make sure the benefits of these technologies are equitable, we must truly understand what they entail. Take a look at the Indigenous Protocol & Artificial Intelligence Working Group [17], which develops new conceptual and practical approaches to AI that centre Indigenous concerns. To learn more about the labour, materiality, planetary impact, and breadth of AI systems, the book *Atlas of AI* by AI scholar Kate Crawford [18] is an excellent option. Continue improving your technical skills with Python libraries like *Keras*, *PyTorch*, and *TensorFlow*, which allow you to create models for computer vision and natural language processing [19][20][21]. To oversee the responsible use of AI technologies, consider a career as an AI Ethicist. In this role, you would ensure that ethical use is a priority, helping draft policy and guidelines for how, when, and why AI is to be applied — or avoided altogether. AI Ethicists work directly at the companies developing AI technologies, or at the companies or organizations using AI, like banks, health organizations, retail services, research clusters, public services, or universities. Follow the work of the Montreal AI Ethics Institute [22] to learn more!

&Envision
An equitable future with *AI*

The potential applications of AI are endless. AI technologies can be applied to any aspect of our lives. But should they be? Think about the values of power and universalism as you envision the future of AI.

A World of Possibilities

Challenges and *opportunities*

The technologies we have discussed present many challenges but with a critical approach, they also have the potential of enabling equity and justice. A first step is for us to be mindful of our engagement in the digital ecosystem; as we have learned, our best interests are not always at the centre of emerging technologies. Intentionally deciding whether and how we engage allows us to reclaim our agency and power. Then, we can build our comfort and expertise in our technical language, speaking up against technologies that fail us or are unjust. If we are called to do so, we can continue our engagement through community building, advocacy, civic participation, and education — sharing what we have learned with others, so their rights are upheld.

We may also engage as creators and designers. The Design Justice Network (DJN) calls on all designers — from urban planners to software engineers — to centre systemically excluded peoples [1]. The DJN promotes design to sustain, heal, and empower communities by seeking liberation from exploitative and oppressive systems; focusing on impact over intention; sharing knowledge and tools; and honouring and uplifting traditional, Indigenous, and local knowledge and practices; and other principles [1]. Applied to emerging technologies, these practices can be transformative. As we continue to face issues of increasing complexity, like deepening inequalities or the negative impacts of climate change, drawing on sociotechnical skills expands our toolbox as we reimagine and redesign our future.

Please
A hopeful pledge

From the information contained in a single instance of a data set to the complex rules that define an AI system, humans make a lot of decisions to make technology possible. If you are able, I kindly ask you to engage in this decision-making.

Many of us are users but rarely designers, critics, advocates, educators, or regulators of digital tools and platforms. Now that you have the technical language and a critical understanding of data, algorithms, machine learning, and AI, kindly consider representing your own values and priorities, as well as those of your community and planet, as we shape the technologies of the future.

&Thank-you
Acknowledgements

These pages are the product of countless expressions of support and kindness from a wonderful community. Many thanks, sincerely, to Professor Leslie Shade for the invaluable support and guidance. To Nalini Singh, Katie Fry, and Sarah English for sharing their expertise enthusiastically. To the Social Sciences and Humanities Research Council of Canada for the Canada Graduate Scholarship (CGS-M) that enabled this research; to The eQuality Project for the production support of this guide; and to the University of Toronto's Faculty of Information for making this project possible. Thank you so much.

& to Marc, my infinite gratitude for the unwavering encouragement.

¡Muchas gracias!

References

Hello&Welcome

[1] L. Stark and A. L. Hoffmann, "Data Is the New What? Popular Metaphors & Professional Ethics in Emerging Data Culture," *J. Cult. Analytics*, 2019, DOI: 10.22148/16.036.

A World of Information

[1] A. Dirksen, "Decolonizing Digital Spaces," in *Citizenship in a Connected Canada: A Research and Policy Agenda*, E. Dubois and F. Martin-Bariteau, Eds. Ottawa: University of Ottawa Press, 2020, p. 25. [Online]. Available: https://ssrn.com/abstract=3620493

Data

[1] S. Zuboff, *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*, First Trade Paperback Edition. New York: PublicAffairs, 2020.

[2] B. Friedman and D. Hendry, *Value Sensitive Design: Shaping Technology with Moral Imagination*. Cambridge, Massachusetts: The MIT Press, 2019, p. 28.

[3] Government of Canada, "Protecting your privacy," *Government of Canada, Statistics Canada*, Nov. 12, 2021. https://www.statcan.gc.ca/en/census/census-engagement/about/privacy

[4] D. Curran, "Are you ready? This is all the data Facebook and Google have on you," *The Guardian*, Mar. 30, 2018. https://www.theguardian.com/commentisfree/2018/mar/28/all-the-data-facebook-google-has-on-you-privacy

[5] A. Smith, "Woman finds vast trove of voice recordings collected by Amazon's Alexa – and you can hear yours," *Independent*, Oct. 23, 2021. https://www.independent.co.uk/lifestyle/gadgets-and-tech/alexa-amazon-echo-voice-recordings-b1943527.html

[6] Z. Tufekci, "Think you're discreet online? Think again," *The New York Times*, Apr. 21, 2019, para. 5. https://www.nytimes.com/2019/04/21/opinion/computational-inference.html

[7] E. Stewart, "Why every website wants you to accept its cookies," *Vox*, Dec. 10, 2019. https://www.vox.com/recode/2019/12/10/18656519/what-are-cookieswebsite-tracking-gdpr-privacy

[8] Office of the Privacy Commissioner of Canada, "News release: Cadillac Fairview collected 5 million shoppers' images," *Office of the Privacy Commissioner of Canada*, Oct. 29, 2020. https://www.priv.gc.ca/en/opc-news/news-and-announcements/2020/nr-c_201029/

[9] Office of the Privacy Commissioner of Canada, "Joint investigation of the Cadillac Fairview Corporation Limited by the Privacy Commissioner of Canada, the Information and Privacy Commissioner of Alberta, and the Information and Privacy Commissioner for British Columbia," *Office of the Privacy Commissioner of Canada*, Oct. 28, 2020. https://www.priv.gc.ca/en/opc-actions-and-decisions/investigations/investigations-into-businesses/2020/pipeda-2020-004/

[10] S. A. Erdem, “Maintaining The Security In Internet Marketing: Moving From Biometrics To Behaviometrics,” *Rev. Bus. Inf. Syst. RBIS*, vol. 15, no. 3, pp. 43–48, Jul. 2011, DOI: 10.19030/rbis.v15i3.5401.
[11] C. Tunney, “Mall real estate company collected 5 million images of shoppers, say privacy watchdogs,” *CBC News*, Oct. 30, 2020. https://www.cbc.ca/news/politics/cadillac-fairview-5-million-images-1.5781735
[12] Office of the Privacy Commissioner of Canada, “PIPEDA in brief,” *Office of the Privacy Commissioner of Canada*, May 31, 2019. https://www.priv.gc.ca/en/privacy-topics/privacy-laws-in-canada/the-personal-information-protection-and-electronic-documents-act-pipeda/pipeda_brief/
[13] Office of the Privacy Commissioner of Canada, “Provincial laws that may apply instead of PIPEDA,” *Office of the Privacy Commissioner of Canada*, May 2020. https://priv.gc.ca/en/privacy-topics/privacy-laws-in-canada/the-personal-information-protection-and-electronic-documents-act-pipeda/r_o_p/prov-pipeda/
[14] I. Gaytandjieva *et al.*, “Comparing privacy laws: GDPR v. PIPEDA,” OneTrust Data Guidance. [Online]. Available: https://www.dataguidance.com/sites/default/files/gdpr_v_pipeda.pdf
[15] Government of Canada, House of Commons. Bill C-27, “An Act to enact the Consumer Privacy Protection Act, the Personal Information and Data Protection Tribunal Act and the Artificial Intelligence and Data Act and to make consequential and related amendments to other Acts”. 1st Sess, 44th Parl, 2022. [Online]. Available: https://www.parl.ca/legisinfo/en/bill/44-1/c-27
[16] Animikii, “Decolonizing Digital: Contextualizing Indigenous Data Sovereignty,” *Animikii*, Jun. 07, 2019. https://animikii.com/news/decolonizing-digital-contextualizing-indigenous-data-sovereignty
[17] Indigenous Research Support Initiative, “Indigenous Data Governance,” *Indigenous Research Support Initiative (University of British Columbia)*. https://irsi.ubc.ca/transforming-research/indigenous-data-governance
[18] S. Browne, *Dark Matters: On the Surveillance of Blackness*. Durham: Duke University Press, 2015.
[19] C. D’Ignazio and L. F. Klein, *Data Feminism*. Cambridge, Massachusetts: The MIT Press, 2020. https://data-feminism.mitpress.mit.edu
[20] R Core Team, *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing, 2018. [Online]. Available: https://www.R-project.org/
[21] G. Van Rossum and F. L. Drake, *Python 3 Reference Manual*. Scotts Valley, CA: CreateSpace, 2009.
[22] Pro Publica Inc., *ProPublica: Investigative Jounalism in the Public Interest*. 2022. https://www.propublica.org
[23] Al Jazeera Media Network, *Al Jazeera Interactives*. 2022. https://www.aljazeera.com/interactives/

Algorithms

[1] S. H. Schwartz, “Les valeurs de base de la personne : théorie, mesures et applications,” *Rev. Fr. Sociol.*, vol. 47, no. 4, p. 929, 2006, DOI: 10.3917/rfs.474.0929.

[2] B. Friedman and D. Hendry, *Value Sensitive Design: Shaping Technology with Moral Imagination*. Cambridge, Massachusetts: The MIT Press, 2019, p. 28.
[3] A. Iovine, "How do all the best dating app algorithms work?," *Mashable*, Oct. 09, 2021. https://mashable.com/article/tinder-bumble-hinge-okcupid-grindr-dating-app-algorithms
[4] B. Wong, "The Subtle Way Dating Apps Reinforce Our Racial Biases," *HuffPost*, Nov. 22, 2018, para. 8. https://www.huffpost.com/entry/dating-apps-may-reinforcesexual-racism-study_n_5bf3056ae4b0376c9e67abe1
[5] M. Mellor, "Why is TikTok creating filter bubbles based on your race?," *Wired*, Feb. 28, 2020. https://www.wired.co.uk/article/tiktok-filter-bubbles
[6] Government of Canada, "Algorithmic Impact Assessment Tool," *Government of Canada, Treasury Board Secretariat of Canada*, Mar. 10, 2022. https://www.canada.ca/en/government/system/digital-government/digital-government-innovations/responsible-use-ai/algorithmic-impact-assessment.html
[7] S. U. Noble, *Algorithms of Oppression: How Search Engines Reinforce Racism*. New York: New York University Press, 2018.
[8] C. D. McIlwain, *Black Software: The Internet & Racial Justice, from the AfroNet to Black Lives Matter*. New York, NY: Oxford University Press, 2020.
[9] Algorithmic Justice League, *Algorithmic Justice League: Unmasking AI Harms and Biases*. 2022. https://www.ajl.org

Machine Learning

[1] B. Friedman and D. Hendry, *Value Sensitive Design: Shaping Technology with Moral Imagination*. Cambridge, Massachusetts: The MIT Press, 2019, p. 28.
[2] Netflix, "How Netflix's Recommendations System Works," *Netflix Help Centre*. https://help.netflix.com/en/node/100639
[3] L. O'Neil, "Black woman quoted more than white man for the exact same Uber ride in Toronto," *blogTO*, Aug. 19, 2021. https://www.blogto.com/tech/2021/08/black-woman-charged-more-white-man-same-uber-ride-toronto/
[4] S. U. Noble, *Algorithms of Oppression: How Search Engines Reinforce Racism*. New York: New York University Press, 2018.
[5] R. Benjamin, *Race After Technology: Abolitionist Tools for the New Jim Code*. Medford, MA: Polity, 2019.
[6] V. Eubanks, *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*, First Edition. New York, NY: St. Martin's Press, 2017.
[7] C. O'Neil, *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*, First edition. New York: Crown, 2016.
[8] Uber, "Ride prices and rates - how it works." https://www.uber.com/ca/en/ride/how-it-works/upfront-pricing/
[9] L. Rainie and J. Anderson, "Theme 7: The need grows for algorithmic literacy, transparency and oversight," p. 18, Feb. 2017.
[10] V. Steyer, L. Vuarin, and S. Ben Mahmoud-Jouini, "Why we need 'explainable' AI," *Polytechnique Insights (Institute Polytechnique de Paris)*, Nov. 17, 2021. https://www.polytechnique-insights.com/en/columns/digital/why-explainable-ai-is-so-important/

[11] P. Linardatos, V. Papastefanopoulos, and S. Kotsiantis, “Explainable AI: A Review of Machine Learning Interpretability Methods,” *Entropy*, vol. 23, no. 1, p. 18, Dec. 2020, DOI: 10.3390/e23010018.
[12] N. Diakopoulos, “Transparency,” in *The Oxford Handbook of Ethics of AI*, M. D. Dubber, F. Pasquale, and S. Das, Eds. Oxford University Press, 2020, pp. 196–213. DOI: 10.1093/oxfordhb/9780190067397.013.11.
[13] M. Kenyon, “Bill C-11 Explained,” *The Citizen Lab*, Apr. 23, 2021. https://citizenlab.ca/2021/04/bill-c-11-explained/
[14] C. K. Steele, *Digital Black Feminism*. New York: New York University Press, 2021.
[15] J. Reback *et al.*, “pandas-dev/pandas: Pandas 1.4.1.” *Zenodo*, Feb. 12, 2022. DOI: 10.5281/ZENODO.6053272.
[16] C. R. Harris *et al.*, “Array programming with NumPy,” *Nature*, vol. 585, no. 7825, pp. 357–362, Sep. 2020, DOI: 10.1038/s41586-020-2649-2.
[17] F. Pedregosa *et al.*, “Scikit-learn: Machine learning in Python,” *J. Mach. Learn. Res.*, vol. 12, no. Oct, pp. 2825–2830, 2011.

Artificial Intelligence

[1] R. Anyoha, “The History of Artificial Intelligence,” *Science in the News (Harvard University)*: Special Edition on Artificial Intelligence, Aug. 28, 2017. https://sitn.hms.harvard.edu/flash/2017/history-artificial-intelligence/
[2] M. Onuoha and D. Nucera, *A People’s Guide to AI: Artificial Intelligence*. Allied Media Projects, 2018. [Online]. Available: https://alliedmedia.org/wp-content/uploads/2020/09/peoples-guide-ai.pdf
[3] S. H. Schwartz, “Les valeurs de base de la personne : théorie, mesures et applications,” *Rev. Fr. Sociol.*, vol. 47, no. 4, p. 929, 2006, DOI: 10.3917/rfs.474.0929.
[4] Amazon, “How Does Alexa Work?,” *Amazon Alexa*. https://www.amazon.com/how-does-alexa-work/
[5] Access Now, “Spotify’s speech-recognition patent tech: Global coalition says don’t spy,” May 19, 2021. https://www.accessnow.org/spotify-spy-techcoalition/
[6] S. Wall, “We tested AI interview tools. Here’s what we found.,” *MIT Technology Review*, Jul. 07, 2021. https://www.technologyreview.com/2021/07/07/1027916/we-tested-ai-interview-tools/
[7] O.-J. Campbell, “Tech Companies Want to Tackle Harassment in Gaming,” p. 8.
[8] W. D. Heaven, “AI for protein folding,” *MIT Technology Review*, Feb. 23, 2022. https://www.technologyreview.com/2022/02/23/1044957/ai-protein-folding-deepmind/
[9] C. Jee and W. D. Heaven, “The therapists using AI to make therapy better,” *MIT Technology Review*, Dec. 06, 2021. https://www.technologyreview.com/2021/12/06/1041345/ai-nlp-mental-health-better-therapists-psychology-cbt/
[10] K. Hill, “Clearview Ai’s facial recognition app called illegal in Canada,” *The New York Times*, Feb. 04, 2021, para. 3. https://www.nytimes.com/2021/02/03/technology/clearview-ai-illegal-canada.html
[11] B. Aguilar, “OPP admit to using controversial Facial Recognition Software Clearview AI,” *CTV News Toronto*, Mar. 01, 2020. https://toronto.ctvnews.ca/opp-admit-to-usingcontroversial-facial-recognition-software-clearview-ai-1.4834043
[12] Office of the Privacy Commissioner of Canada, “News release: Clearview AI’s unlawful

practices represented mass surveillance of Canadians, commissioners say," *Office of the Privacy Commissioner of Canada*, Feb. 03, 2021. https://priv.gc.ca/en/opc-news/news-and-announcements/2021/nr-c_210203

[13] Office of the Privacy Commissioner of Canada, "A Regulatory Framework for AI: Recommendations for PIPEDA Reform," *Office of the Privacy Commissioner of Canada*, Nov. 12, 2020. https://priv.gc.ca/en/about-the-opc/what-we-do/consultations/completed-consultations/consultation-ai/reg-fw_202011/

[14] A. D'Andrea, "Ottawa unveils new rules for use of personal data, AI in privacy bill," *Global News*, Jun. 16, 2022. https://globalnews.ca/news/8925156/ottawa-personal-data-ai-new-privacy-bill/

[15] E. Ifill, "Opinion: The problems with the federal data-privacy bill will disproportionately hurt marginalized Canadians," *The Globe and Mail*, Jul. 04, 2022. https://www.theglobeandmail.com/opinion/article-the-problems-with-the-federal-data-privacy-bill-will/

[16] Government of Canada, "Responsible use of artificial intelligence (AI): Our guiding principles," *Government of Canada, Treasury Board Secretariat of Canada*, Oct. 12, 2021. https://www.canada.ca/en/government/system/digital-government/digital-government-innovations/responsible-use-ai.html#toc1

[17] A. Abdilla *et al.*, "Indigenous Protocol and Artificial Intelligence Position Paper," 2020, DOI: 10.11573/SPECTRUM.LIBRARY.CONCORDIA.CA.00986506.

[18] K. Crawford, *Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence*. New Haven: Yale University Press, 2021.

[19] C. Trabelsi *et al.*, "Deep Complex Networks," 2017, DOI: 10.48550/ARXIV.1705.09792.

[20] A. Paszke *et al.*, "PyTorch: An Imperative Style, High-Performance Deep Learning Library," in *Advances in Neural Information Processing Systems 32*, Curran Associates, Inc., 2019, pp. 8024–8035. [Online]. Available: http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf

[21] Martín Abadi *et al.*, "TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems." 2015. [Online]. Available: https://www.tensorflow.org/

[22] Montreal AI Ethics Institute, *Montreal AI Ethics Institute: Democratizing AI Ethics Literacy*. 2022. https://montrealethics.ai

A World of Possibilities

[1] Design Justice Network, "Design Justice Network Principles," *Design Justice Network*, 2018. https://designjustice.org/read-the-principles