< >
Montréal Declaration
Responsible AI_
</ >

PART 5

SUMMARY REPORT OF ONLINE SURVEYS AND PROPOSALS RECEIVED FOR THE MONTRÉAL RESPONSIBLE AI DECLARATION

This document is part of the 2018 **MONTRÉAL DECLARATION FOR A RESPONSIBLE DEVELOPMENT OF ARTIFICIAL INTELLIGENCE.** You can find the complete report HERE.

TABLE OF CONTENTS

WRITTEN BY

MARTIN GIBERT, Ethics Counsellor at IVADO and researcher in Centre de recherche en éthique (CRÉ)

1. INTRODUCTION

In November 2017, the first step of the *Montréal Declaration for a Responsible Development of Artificial Intelligence* was launched following a convention organized by the Université de Montréal at the Palais des congrès in Montréal. The preliminary version of this Declaration, articulated around seven principles, would serve as the basis for a co-construction phase from which a new version would be created. Although the discussion workshops helped reach citizens and experts, there were also other ways to join the collective reflection: 1) by filling out an online survey accessible through the Declaration's website (www.declarationmontreal-iaresponsable.com), and 2) by sending in a proposal on one or more aspects of the Declaration. This report presents a summary of the proposals received and the answers to the survey. The report on the co-construction workshops is also available on the Declaration's website.

2. ONLINE SURVEY

The online survey consisted of 35 questions, five for each principle. A total of 83 people answered the survey, 17 of whom were anglophones. As the summary reveals, many had advanced knowledge of AI and the ethical and social issues raised by its development.

Questions are presented per order in the questionnaire, which was based on the preliminary Declaration plan. Since the revised Declaration is more complete (it is made up of 10 principles), the relevant new principles were added to those from the preliminary version.

WELL-BEING (ENVIRONMENT, CAUTION)

1. HOW CAN AI CONTRIBUTE TO WELL-BEING?

This was a general question that sparked many answers, and many varied ones at that. One recurring hope was for healthcare and assistance for the elderly or disabled. AI also seems to hold promise for reducing environmental impacts, though it was noted that "AI development has an environmental footprint (and thus a direct impact on well-being) that is often neglected, even though it is significant." Many pointed out that AI could replace humans for dangerous tasks. The aspect of "decision-making assistance" especially in the form of a personal assistant that could also assist in information searches, was also mentioned numerous times.

We expect AI to improve productivity and free us from repetitive and routine tasks as well. IT could also anticipate our needs and expectations, or simply do the vacuuming for us. One important provision: AI will improve our well-being "as long as we live in a true democracy, where it serves everyone, not only a privileged few".

SELECTED EXCERPTS

"The AI or any technology will create a lot more value for the rural population than the urban population. A single smartphone can provide immense value, and anything that can collect data is a breeding ground for AI: Better education, better farming technology (e.g. crop analysis, robot farming)."

2. CAN AN AUTONOMOUS WEAPON BE USED TO KILL A HUMAN BEING? AN ANIMAL?

An overwhelming majority of people answered "no" to this question, often very emphatically and with numerous exclamation marks. Reasons included that "killing must remain in the hands of humans, who must be fully aware of their actions". The idea of legally banning autonomous weapon systems was also mentioned many times. One survey respondent also pointed out the risk of an arms race and possible programming errors. Some respondents made a distinction: "no' for humans, "yes" for animals ("for population control"). There seems to be exceptions in some cases: a machine killing a death row inmate or a "tiger that breaks free from its cage and threatens the general population". In each case, it appears that AI should only be a tool used for killing and that in the end humans should be held responsible. One respondent, however, offered a more critical point of view and raised a valid question: if such a weapon can make a better decision than a human being, why not?"

We also noticed that this question depends largely on context: "It would be acceptable for an autonomous weapon to kill a human being or an animal in any circumstances where it would be acceptable for a human or other creature to kill a human or animal."

SELECTED EXCERPTS

"Autonomous weapons shouldn't exist, they should be banned just like chemical weapons. Humans should always be in control of a weapon; they would therefore be morally responsible for their actions."

"No! (...) a horrific scenario could ensue from an unethical manufacturer or rogue programmer who perhaps, unbeknownst to either the weapons' company or weapon purchaser, may secretly design, code & program autonomous weapons which reflect their secret views & biases as a neo-Nazi or K.K.K. supporters, for example."

"Why would you think a HUMAN should be able to kill somebody? If you have a reason, then why doesn't it apply to AI? There's no reason humans should always occupy a privileged position with respect to killing other humans. Obviously 'AI' at the moment is not even ready for consideration for this, but that's unlikely to be permanent (assuming you think anything or anybody should be killing whoever or whatever you're thinking about killing). The interesting question may be how you'll know when it's changed, and how you manage the transition."

3. SHOULD AI BE USED TO CONTROL A SLAUGHTERHOUSE?

As in the previous question, the vast majority answered "no" (though some were more in favour). The argument around normalizing violence through psychological distance, which was evident in the previous question, comes up once again: "This would distance man even more from the action of killing the animal," or even: "we must not offer humans a new level of cowardice to hide behind by delegating a morally reprehensible task to a robot". The environmental argument was also cited: "This is not the direction to take for the future of the planet and humans who live on it. I would rather see AI control greenhouses and zero CO2 emission and zero-waste buildings."

There were a few arguments in favour of such a project: avoid mistreating and creating stress for the animals, as well as improving hygiene. (One respondent nonetheless explains that a human should always oversee slaughtering operations precisely to avoid cruelty...) Some positions were mixed: AI could control the cutting and wrapping, but not the killing. Many wonder, however, if slaughterhouses are acceptable, even without AI.

SELECTED EXCERPTS

"I understand it would relieve those who are responsible for these morbid tasks. However it would be absolutely unethical."

"Slaughtering conditions may be slightly improved, but the practice itself would last longer, because it would be easier to look away."

"Yes, to the extent that the AI follows humane (and human) protocol."

"Interesting question. One side of me says 'yes', the other 'no'. What we are doing to other animals that we raise for food already has some serious ethical issues. When I read about the life of the average chicken raised for food, I was shocked. Totally automating the process of raising food, including having AI do the killing would just put the fate of these animals even more out of sight, out of mind. So, on balance, I think I am against an AI-controlled abattoir."

"Slaughterhouses already exist and won't stop existing anytime soon. AI can make sure that the method of slaughter is ethical and is done in the most humane way possible. This can also strictly ensure and maintain safety standards."

4. SHOULD WE ENTRUST AI WITH MANAGING A LAKE, A FOREST OR THE EARTH'S ATMOSPHERE?

This question elicits significant skepticism of AI, but also hope for solutions to the environmental crisis. Once more, the main idea that emerges is that the AI taking care of the environment should be "configured by responsible human beings who take conservation to heart". Some are even hopeful, especially for the climate, but from the perspective of a human/AI collaboration, rather than delegating the problem to AI. Many respondents hope for AI that cannot be corrupted or seek profit at all costs.

The risk of a malicious hijacking of AI entrusted with such a mission was also raised. There was a hint of cynicism too: "humans have destroyed nearly every natural environment they have come into contact with, so it can't really get much worse..." The theme of replacing humans also comes up: "We could have AI do everything, but we need to ask ourselves if we want humans to be assisted with reaching their full potential." There is also a democratic principle: no human or artificial entity alone should be able to decide how the environment is managed—this should be based on cooperation between all humans.

SELECTED EXCERPTS

"No, because we are sorely lacking the knowledge to be able to judge the long-term repercussions of actions taken by AI."

"It depends on what the instructions given to the IA are and how much absolute control it holds. I think AI trained on environmental systems and with ability to monitor and consider big (environmental) data could make much better decisions than any group of individuals, effectively helping to protect the environment and regenerate those that may have been affected by industry, etc."

"It could be done with the assistance of AI; but AI itself doesn't know what is good for the lake, the forest or the atmosphere. It is more efficient from the point of view of instrument rationality, but cannot determine its own goals by itself."

"At present, AI would be most useful in the collection and analysis of data."

"Eventually, machines may be more competent than people to make almost all decisions. But, if we give the machine control and stop monitoring what and how it does what it does, the ability of human beings to manage our affairs will pass out of the living memory of humans, and we will be entirely dependent upon machines. This does not seem to me to be a good future for human beings."

5. SHOULD WE DEVELOP AI CAPABLE OF EXPERIENCING WELL-BEING?

Participants hesitated on this question and answers were contradictory. Is it even possible, from a technical standpoint? Some mentioned that sentience could allow humans to control or punish AI. Others see the interest in AI being able to better understand humans (and other sentient beings) and empathize with them. Emotional intelligence also appears to be a requirement to make good moral judgments. However, simulated empathy appears sufficient, because some foresee danger in sentient AI looking after their own well-being over the functions it was assigned by humans. "AI must remain a tool that serves humans, not a quasi-human." And many wonder: "what's the point?" One respondent worries for AI: "I'd rather AI understand well-being than experience it, especially because in the notion of well-being, there's also the notion of being unwell."

SELECTED EXCERPTS

"I think it makes most sense to approach the development of general AI as the development of a calculator/tool. Developing a personified, sentient AI may be bringing new life to the world. I'm sure it would be treated fairly or with rights."

"This is a very complex question. If it experiences a sense of well-being, it will feel the need to maximize it. This is useful when rewarding learning, but will it balance its machine well-being with that of humans and other living creatures?"

"Yes, but only proportional to it accomplishing the tasks it was assigned. You could develop AI which, thanks to the satisfaction of a job well done, constantly seeks to improve, but only in the specific field in which they operate."

AUTONOMY

1. HOW CAN AI CONTRIBUTE TO HUMAN AUTONOMY?

AI presents us with an ambiguous relationship to autonomy: it makes us rely on it ("we could no longer dissociate from AI), while freeing humans from certain alienating cognitive tasks (e.g. driving a car, administrative functions), and even the need to work. What mostly shines through in the comments, however, is the positive and liberating aspect to it. The partnership model is an option, as is having AI as

a simple assistant. And the sky seems to be the limit: AI could improve the human condition, especially for people suffering from a disability, and could lead to less invasive medical care, which would also help the elderly who are losing their autonomy.

SELECTED EXCERPTS

"AI should be used to restore autonomy (physical or mental) to people with disabilities. Only a person who can entirely control the configuration of the algorithms for an AI system could gain autonomy, everyone else would lose some because they rely on decisions made by someone else."

"No system will ever help (or currently helps) human autonomy if it comes from a private company. Regulations and involvement from the public sector are essential to maintain balance."

"By freeing them from the tasks they don't want to do, and by improving their emotional state and their understanding of the world."

"HUMANs will deliberately develop AI to force other humans to follow their values or act according to their interests. And those humans will see themselves as benevolent in doing that which is the really scary part."

"AI can automate most of the trivial things that we spend a lot of time doing. Almost everything that we do without actively thinking about it can in a way be simplified or made more convenient using AI. But this also has to ensure that humans don't become too dependent on the technology, which would then handicap their life instead of providing more autonomy."

2. SHOULD WE FIGHT THE ATTENTION-CAPTURING PHENOMENON THAT COMES WITH AI BREAKTHROUGHS?

Participants were highly skeptical of this phenomenon ("I need more information"). But many highlight the risk of "technological hypnosis", "especially among teenagers". One respondent sums it up this way: "we must not become slaves to our technologies". Another suggests treating AI with AI: "we'd have to know why our attention is being captured. Income-generating goals, which drive applications such as Facebook should be blockable through other AI applications, countermeasures of sorts made available to users to fight against intrusions."

Evidently, drawing people's attention to ethical problems seems to be a good idea: "it's another way of making sure these conversations take place". And one comment that makes a lot of sense: "Yes, businesses should be prevented from manipulating people's attention in ways that people don't control or understand. That's not intrinsically related to AI; it's just that AI is a convenient, powerful and, therefore, dangerous tool for it."

SELECTED EXCERPTS

"Yes. First by educating people, then by passing legislation to impose an operational framework that reflects humanist values (truth, justice, kindness, respect, etc.)"

"All technologies, from radio frequencies, nuclear energy to

cryptography must live within a regulatory framework. Attention-seeking AI could be classified as addictive entertainment, like gambling."

"One way to combat this is awareness about the problem, the fact that this is happening is not known to many (hypothesis). And give the user proper tools to combat this: nudge the user to actually learn the skill using small dopamine hits until the user doesn't need it anymore."

3. SHOULD WE BE WORRIED IF HUMANS PREFER THE COMPANY OF AI TO THE COMPANY OF OTHER HUMANS OR ANIMALS?

No clear trend emerges from the answers. Certainly, there is concern that technology separates or isolates humans: "human beings must remain social beings" and we must ensure that humans do not forget social skills such as empathy. But from this point of view, AI "would be no worse than video games". Psychological studies are likely required to evaluate the risks of a new type of addiction. But there is also a potential benefit for people who are alone, or for certain psychological profiles: autistic children, for example, may find it easier to communicate with AI than with a human. It should remain a marginal thing, however, because "if a human no longer wants any contact with other humans, then humanity disappears". But what about paternalism? If it causes no harm to others, why prevent strong relationships between humans and AI? After all, we generally accept that certain people prefer the company of animals to that of humans.

SELECTED EXCERPTS

"If AI agents have no awareness or feelings, they can't be good company. Less so than animals even. For the moment, they are things. Machines."

"If a person has no other option, it can be a good thing. Otherwise, we're going to start having a hard time living as a community."

"No, many human beings are already trapped in relationships with objects or fictional characters (television, soap operas, social network friends). AI companionship would at least have the advantage of presenting a certain degree of interaction that could prove especially beneficial to people who are elderly or alone."

"This is a legitimate concern. It can be compared to the preference for texting as a substitute for direct human-to-human interaction."

"If you care about people's autonomy, then LET THEM MAKE THEIR OWN DECISIONS. It doesn't matter whether you're 'worried', because it's purely none of your business, full stop."

"Even if technologies like VR [virtual reality] are developed to an almost realistic level, it would only increase social isolation, and it would be detrimental in the long

run. Social security, i.e. the fact that there are people to support you and will be there with you in your time of need, is invaluable!"

4. CAN WE GIVE OUR INFORMED CONSENT WHEN FACED WITH INCREASINGLY COMPLEX AUTONOMOUS TECHNOLOGIES?

Many respondents felt this would be difficult for two reasons: the complexity of machines and the complexity of legal clauses. No one reads the terms of consent for apps or platforms when they are too complex (legalese): when people "accept" (do they really have a choice?), this consent cannot be considered truly informed. "How often do we sign off on online agreements saying we read them when we didn't?"

Systems that are secure and that inspire trust therefore still need to be created. The lack of digital literacy was also highlighted, as was the need to remedy the situation through education. That also highlights the "importance of establishing a code of ethics on which AI is built". One solution could come from AI itself: it should be able to answer our questions to help us make informed decisions. But another danger lurks: "Information presented to humans will naturally inform (and bias) decision-making. Humans are quick to assume that algorithms or information provided by statistical analysis is somehow void of bias."

SELECTED EXCERPTS

"It would be a good idea to provide a legal framework for this notion when it comes to companies and public organizations doing business in Quebec."

"It's impossible. The only thing to do is establish or restore trust with those who build or own these technologies through social or political control that satisfies the greatest number of users, by reducing the harmful use and misappropriation that the owners and designers of these technologies could be tempted to carry out."

"Probably not. I think it's already impossible to provide informed consent for digital technologies that aren't even based on AI. For example, how can we be sure the software we buy isn't spying on us?"

"For decently complex systems, the user has to be fully made aware of how the data being generated can and might be used, along with theoretical guarantees or open code base proving their claims. But for very complex systems, here, even the creator wouldn't know how the data might be used completely. But, even in the worst of the cases, rigorous proof of claims and possible benefits, analysis on a test group can help earn the trust of the user and allow the person to give consent."

"As technology advances, the demands for our consent will increase exponentially. Under those conditions, the unaided human will not be able to give truly informed consent in many of the cases where it is demanded. The proof is that we

already have become conditioned to signing off on agreements that we have not actually read or understood. The demands are only going to increase. The solution, if there is one, would involve 'loyal' AI agents assisting us."

5. SHOULD WE LIMIT THE AUTONOMY OF SMART COMPUTER SYSTEMS? SHOULD A HUMAN BEING ALWAYS HAVE THE FINAL SAY?

Many responses were positive. Human beings must always be at the helm. AI is a tool to help with decision-making. Diverging points of view are nonetheless interesting: "AI is potentially more accurate, less biased and soon more creative than humans. Let's take advantage of it!" Along the same lines: "Humans can be corrupted. AI could have a stricter moral code than humans." Restrictions could also exist when an urgent decision had to be made.

Context is obviously everything: making croissants or launching an attack are obviously not the same thing. Humans should at the very least be able to make the decision to shut down an autonomous system. And that does not appear negotiable "in the case of complex decisions that include an ethical dimension involving responsibility".

SELECTED EXCERPTS

"There may come a point in system development where we will be able demonstrate that a human being can no longer make better decisions than a computer."

"Not the final decision, because the advantage of AI is to make a decision instantly based on a series of parameters that no human could ever analyze so quickly. But decision-making responsibility should always be assumed by a human being."

"Yes and yes, computer systems are there to assist us with decision-making, and that's how they must remain. Why give a cyborg control over us?"

"The fundamental decisions must be human and be as consensus-based as possible."

"You always want to have the option of an off switch. And we need to build systems in such a way that we can come to an understanding of how the machine is making the decision."

"Obviously with the current state of the technology, you can't let it have total control over too many things. That is unlikely to be true forever; eventually the AI is probably going to be smarter than the human ... and possibly more benevolent than the human, which is where you should really be putting your energy. At some point the question may be whether the human should even get any input into certain decisions, especially into decisions that affected more than just that human."

"If the human does NOT always make the final decision, then there needs to be a transparent interface so that users can correct the

decision-making computer system when it makes mistakes. (For example, with Google translate, you can provide a better translation.)"

JUSTICE (EQUITY, SOLIDARITY, DIVERSITY)

1. HOW CAN WE ENSURE THAT EVERYONE HAS ACCESS TO THE BENEFITS OF AI?

By making it affordable (or free), through open source and clearly exposing which decisions AI will be making for us (transparency). But is it possible in the capitalist system we know? "The private sector should not be able to exploit an annuity for its sole profit, at the expense of the rest of humanity." We could even tax companies that get excessively rich thanks to AI (would that harm innovation?).

Education could play a role in combating the digital divide. The role of governments (or even the UN) is to redistribute the benefits in equitable fashion and ensure that the AI values are aligned with human values. A basic income, a call for political realism tempers expectations: "let's not be utopian, AI isn't the one creating inequalities, humans are". One respondent also noted that information technologies make participative democracy possible. Another brought up the issue of basic income.

SELECTED EXCERPTS

"Build AI for the common good rather than private property. Regulate it to force advanced forms to adopt a GNU free licence for example, and promote information sharing."

"We cannot leave AI entirely at the mercy of the private sector."

"Possible AI advances, the discovery of a new protein for example, must be for the collective good."

"General quality of life for everyone should be improved with AI. Legal system seems to be one that will be greatly affected and see a lot of change, for the better."

"The digital divide could depend on whether or not large private companies get their hands on the data generated by the population."

"We must review international patent laws extensively. AI development will only truly progress if the information at its core is in the public domain. Special interest groups (corporations, army, governments) should not be able to appropriate this technology, otherwise it will inevitably be hijacked to serve their interests rather than those of the citizens."

"Make equity a central pillar. Include researchers and community groups that collaborate on designing equitable solutions. Take a look at the work done by the Support Unit (SRAP) and the mobilization and citizen involvement section of Alliance santé Québec."

"Give free Wi-Fi to the poor for starters."

"This is a very complex question. One could argue that everyone already benefits from AI through 'free' products like Facebook and Google Maps. What is missing is an understanding of the market value of someone's data relative to the machine's ability to build a more powerful model. Governments at all levels need to be using AI with the data they currently manage as another part of their policy-making tool set."

2. SHOULD WE FIGHT AGAINST THE CONCENTRATION OF WEALTH AND POWER ENJOYED BY ONLY A FEW AI COMPANIES?

Answers are clearly positive. By promoting open source and GNU free licences. Because it's the State rather than the private sector (GAFA) that citizens trust. The concerns are real: "Would democracy survive if predominant AI power fell into the wrong hands?" How could we even do this? We couldn't even manage to get open-source software to replace proprietary software. Nationalize to remain "masters of our own domain"? Regardless, AI should be seen as a common good that does not serve a minority. One respondent highlighted the need for an antitrust organization to break up certain monopolies. However, some preferred a more competitive model: "If some companies manage to carve a niche that brings them wealth and power, more power to them. But knowledge must remain in the public domain to cultivate competition." One respondent suggested that individuals own their own data and use an AI personal assistant that is loyal to them.

SELECTED EXCERPTS

"Obviously, we must fight the concentration of power, period."

"Yes, it seems there will be a lot of power available to those who control AI systems. New legislation/law will be required to monitor this, along with taxes on automation, etc."

"What's most important is that the basic programs are universal and built for the common good. Otherwise they will only be robots serving those who already maliciously rule the world with their own interests in mind, so either nothing changes, or the inequalities, violence, conflicts, etc. all get worse."

"The hands of a small number of AI companies or the hands of a small number human entities (i.e. the 1%) should not have more power and wealth than the 99% of human beings on earth. Powerful entities should adopt socially responsible behaviours at all time, especially when in presence of the public. (...) The democratization of AI should definitely empower the 99% of human beings."

3. WHAT KIND OF DISCRIMINATION COULD AI CREATE OR EXACERBATE?

Every "classic" form of discrimination seems to be exacerbated by AI, especially "social, racial, economic", but also "linguistic and cultural" discrimination, not only among people, but also among groups or states. A dystopian scenario looms: one where a new class of ultra-rich people (the 1%?) use AI to perpetuate socioeconomic inequalities. Participants also mentioned that access

to technology can be exclusive and excluding, especially for older people.

One respondent specified the type of mechanism that could encourage AI: "AI could be the perfect scapegoat under the guise of a BLACK BOX: Why didn't I get that line of credit, Mr. Bank Manager? Ah, I'm sorry, the system gave us that result." It also appears clear for respondents that humans, whether as individuals or as groups (e.g. systemic racism) that are responsible for this discrimination—not AI.

SELECTED EXCERPTS

"[We have to beware] of a 'caste' of AI experts emerging, whether known or secret, that holds all the knowledge, and therefore all the power. [We must also beware] of discrimination based on a health condition (flirting with eugenics), racial or sexual discrimination, towards the elderly, towards women, etc. [Lastly, also beware] of economic discrimination, increasing poverty for the majority and the power the rich hold over decision-makers."

"There are too many... That's precisely the problem. We have a hard time establishing what discrimination is or whether we're already doing it. How can AI determine this for us without discriminating exactly same way with the data we give it?"

"Social networks are already a source of stereotypes and racist, sexist, stigmatizing content. We can consider filtering that, offloading the problem. These filters could also be unduly discriminating."

"Algorithms must be developed by multidisciplinary and multicultural teams to avoid perpetuating stereotypes based on gender, wealth, race, etc."

"If AI contributes to well-being and autonomy, the people who need it, but can't access it, are worse off."

"If AI is deployed by special interest groups (armies, governments, corporations), it will only serve their momentary interests at the expense of the population."

"See weapons of math destruction. AI models with labelled training data that is discriminatory will simply perpetuate and reinforce these discriminations."

"It's going to be hard to deal with that, because in order to admit that the AI is going to find a regularity, you have to admit that the regularity exists. You have to be willing to say, 'Yes, XXX people *are* more likely to default on loans, but we want to ignore that anyway.' After that, it's a relatively simple technical problem to make AI implement your wishes. Short-term AI, anyhow."

4. SHOULD AI DEVELOPMENT BE NEUTRAL, OR SHOULD IT SEEK TO REDUCE SOCIAL AND ECONOMIC INEQUALITIES?

Most respondents are in favour of AI that would actively contribute to reducing social and economic inequalities. Many even consider this a priority. One optimistic respondent thinks that reducing inequalities will be the automatic result of AI development. Another wants it to mostly promote equal opportunities. However, a few skeptics would prefer it to remain neutral: "Who decides which inequalities to reduce?" And the more pessimistic maintain that there will always be inequalities … which shouldn't prevent trying to reduce them. Finally, one respondent suggests that AI remain neutral on social and economic inequality for commercial use, but that non-commercial use should aim for more equality.

SELECTED EXCERPTS

"Yes, that should always be its goal, along with reducing environmental impacts."

"AI cannot be neutral, so we might as well guide it in a direction that benefits everyone."

"Why are we developing AI? Reducing inequalities does not appear to be the primary reason; that does not mean, however, that AI development should be neutral: social and economic inequalities could serve as a 'constrained site', so the development doesn't occur at the expense of important values".

"AI models should be applied within a policy framework. No information system is neutral and any architect or policymakers must embrace the ethical challenges and opportunities when applying AI. In this context, reducing existing inequalities is a moral imperative. Machine learning models need to be conceived inside a larger pipeline that can mitigate regressions and provides recourse for error."

"But we should make sure that by doing so we are not actively causing friction between different groups or trying to homogenize them. The effect, in that manner, should be neutral."

"It should be neutral in commercial settings, otherwise the technology might never be adopted at all—leading to no benefit to the society. But it should also reduce socioeconomic inequalities in a non-commercial setting by giving everyone access to the same tools and opportunities."

5. **WHAT TYPE OF LEGAL DECISIONS COULD BE DELEGATED TO AI?**

The consensus was that no important decision should be delegated to AI. AI must simply serve as a tool to assist in decision-making. It could therefore "accelerate case processing", and even "make easy decisions after analyzing proof", such as decisions tied to paying tickets.

AI could be beneficial in other aspects of justice: "Detecting a lie or a false memory. Detecting risks of relapse." If a general artificial intelligence were developed, then AI that replaces judges could be envisioned; but this option is far from unanimous, even if it has been proven that human judges are

often biased in their rulings and subjected to various pressures. Perhaps the entire legal institution needs to be overhauled from top to bottom to make "artificial rulings" possible. Regardless, lowering costs and making justice more democratic would be good news and AI could certainly contribute to that, namely by making access to jurisprudence easier.

SELECTED EXCERPTS

"AI could replace a stenographer."

"AI would be fairer because it isn't subject to emotions or pressure from the media or other lobby groups. The only thing that would eventually need to be revised would be the Criminal Code, given the differences observed between human and artificial rulings."

"I don't believe any final decisions should be made by the AI. Seems the legal aid/technician and data processing could be best managed by AI."

"Decisions that involve complex practical rulings (jurisprudence) should be reserved for humans. Justice is also a social process. Let's not forget that.

"AI could do research for the general population (as well as jurists), by having access to all jurisprudence. This would make access to justice more democratic, since most of the costs for citizens go to jurists doing this kind of research."

"Current and near-future AI aren't going to be able to comprehend the law or apply it other than in cases so mechanical that you don't really need 'AI' at all. I suspect that any real legal decisions will take a truly general intelligence."

"AI predictive technology can be used to help judges make better decisions. The idea is not to replace judges."

PRIVACY (INTIMACY)

1. HOW CAN AI GUARANTEE RESPECT OF PRIVACY?

Many respondents wondered whether this question was relevant: How can AI ensure respect? The impression is rather that it violates it, repeatedly, without user consent. It also appears contradictory, since AI needs our data in order to develop.

But options may exist: "encrypt everything", not be invasive when requesting personal data. Someone remarked: "Respect of privacy is guaranteed if the person isn't exposed to AI by default." Users also have a responsibility: "It's up to each of us to control our exposure: shop in independent stores and pay in cash, rather than buying off the Internet."

Some were wary of the private sector: "Nothing is guaranteed if it's solely managed by the private sector." That is why we call the State and legislators to the rescue: Quebec's laws regarding privacy must be respected and improved. "It's a major challenge," because isn't it too late already? Our Facebook data, for example, may have been siphoned long ago by Cambridge Analytica or any other such company. And that's not even mentioning "hackers".

SELECTED EXCERPTS

"I believe the information economy, based on traceability, can create more information sharing, but at the same time more transparency in its use, and so having your information shared won't have such negative consequences if those who see it are also traced."

"Let's face facts, there are, realistically speaking, no truly reliable guarantees that AI can respect people's privacy. Health records & private accounts are hacked all the time despite the best security upgrades that technology has to offer. Google reads our private emails, doesn't it?"

"Differential privacy—the idea that you can give away information about yourself without ever having it trace back to you as the source. But, if such a practice is possible and can be made prevalent then I believe that informed consent is possible. The user has to be fully made aware of how the data being generated can and might be used, along with theoretical guarantees or open code base proving their claims."

"Make people's private data truly their private property."

2. DOES OUR PERSONAL DATA BELONG TO US, AND SHOULD WE HAVE THE RIGHT TO ERASE IT?

Agreement from participants was overwhelmingly positive for both questions. Someone specified, "and it should be very easy to do so, so everyone can do it". One respondent disagreed with the idea that our data belongs to us, but that that shouldn't prevent us from having "the right to examine its use". Although most respondents implicitly admitted that individuals should own their data, some see it rather as a collective good.

Erasing data should not obstruct justice (or healthcare services), which may need to access older data, nor should it harm others.

SELECTED EXCERPTS

"Yes, every citizen should own their personal data, just as artists own their creations."

"No, but data should be considered a national good, like libraries or nature reserves."

"Absolutely, and unequivocally. Only data required for good government functioning should be kept: demography, income, health, legal. All other data should be controlled by the user."

"As long as companies own and licence IP, individuals should have a right to all data they create."

"Generally yes. But I have a very broad definition of what should be considered personal data (and should be private property). Within this larger view even our criminal records would be personal data that we own (though not without

controls). It would be a category of personal data that we should not be able to delete—at least not whenever we choose."

3. SHOULD WE KNOW WHO OUR PERSONAL DATA IS SHARED WITH AND, GENERALLY, WHO IS USING IT?

The answer was a unanimous "yes"! Someone specified: "Just like we need to know who comes into our home, we need to know who can access our personal data." Another said: "Yes, [and we should know] who, how and for what." One respondent mentioned that we might grow tired of knowing who is using our data, and may quickly lose interest. But that obviously should not stop us from having the right to know.

SELECTED EXCERPTS

"Yes, I think I should even have a portal where I control 100% of the data I am sharing."

"Our data should never be shared without a clear and concise request to do so first being made. No 20-page contracts in small print where we have to guess that lifetime permission has been granted. If we subscribe to a service, information should never be able to be used in any other way than for the service requested."

"Absolutely and they should be required to ask permission to do so on a regular basis. Permission is not granted in perpetuity."

"Absolutely. Personal data should be private property. We should defend it and allow the owner to control who can access it and to what extent they can access it. The current default—wherein we cede our data to others—is bad for citizens and bad for democracy. There is another option."

4. DOES IT RUN COUNTER TO ETHICS AND ETIQUETTE TO HAVE AI ANSWER E-MAILS FOR YOU?

There were contradictory answers to this question. Many remarked that this kind of service already exists, or that certain people have human assistants that answer e-mails for them. One option would be to have AI prepare a response, but to have it validated by a human being (thereby giving them the last word). One respondent specified that "what's important, in my opinion, is that the person using this service has the necessary understanding and trust in the service". Another request that the process be transparent, meaning that the correspondent knows that the reply to their e-mail was written by AI. There may not be a generic answer to this question: it depends on the type of question ("Are you available for this meeting?" vs. "Do you think we should hire this person?").

SELECTED EXCERPTS

"No, as long as it is clearly stated that the response was written by AI rather than the individual concerned. If that person chooses to have AI answer for them, that's their responsibility … as long as the Internet service provider lets you activate or deactivate this feature. Obviously, it's not about imposing this service.

"It's useful for those who have to manage a high volume of similar messages with low complexity."

"That depends, if you always answer the same way to the same questions, it's not going to make much of a difference for you."

"If it's a question of customer service, of fulfilling a human need to satisfy that involves responsibility for others, I expect a human being to answer."

"Yes. Human intent is a critical component to our society's framework. We can delegate to AI, but human dignity claims that you should know if you are interacting with a machine."

"Similarly, if an organization has a bot deal with people, it should always identify itself as a bot. People should always know if they are dealing with a human or a machine. And the organization that has bots dealing with people should always be held responsible for any actions the bot takes on the organization's behalf."

5. WHAT COULD AI DO ON YOUR BEHALF?

An open question that elicited very different answers, ranging from "nothing" to "everything" (as long as consent was provided). Between the two extremes: book an appointment manage my finances, my schedule, file my taxes and other administrative tasks, vote (!). But I should always be held responsible for the consequences of what AI does on my behalf. (Many respondents confused this question with "What could AI do for you?", e.g. vacuum.)

SELECTED EXCERPTS

"Everything that I have approved beforehand."

"Any task that does not commit to any future engagement."

"Nothing serious that could have legal or emotional repercussions."

"Book appointments respond with numerical data that is already in the public domain, check on the well-being of family pets."

"That depends on the AI. I wouldn't trust any *present* AI to do anything that I couldn't countermand or that people would interpret as a direct application of my personal judgment."

"My recommendation is to adopt a paradigm in which each citizen owns private, 'loyal' AI tools (agent) that can help protect, manage, analyze and use a citizen's private data (stored in a protected online profile) to help that citizen at their behest and only their behest. (...) Some people might say they can do simple repetitive tasks, perhaps review email. Others might allow their AI agent to browse the web to plan online shopping. Others might let the agent actually make purchases autonomously. Others might allow the AI agent to perform investment transactions for them. In an advanced future, some prefer

to trust their AI to participate in a family vote about 'pulling the plug', given on its intimate access to its owner's private data, which could analyze a variety information types taken from a personal profile, allowing it to use predictive analysis to help decide what the citizen might want if they were able to speak."

KNOWLEDGE (PUBLICITY, CAUTION)

1. COULD AI DEVELOPMENT JEOPARDIZE CRITICAL THINKING?

Answers varied, but leaned towards "no". On the "risk" side, many fears are raised: loss of sense of curiosity, publicity, standardizing thought and dismissing marginal viewpoints. AI might also speak on behalf of humans and appear too reliable: "The machine can't be wrong; everything has been said, there's nothing left to add".

On the plus side, many noted that the time gained through automation could be invested in critical thinking and the fact that AI and information technologies make information more accessible, or that we could even program AI to perform critical thinking— the idea that AI could be more neutral than humans was also brought up. Finally, AI could be viewed as a wonderful opportunity—or need—for humans to exercise critical thinking.

SELECTED EXCERPTS

"Yes, but not if it is used to make people's lives easier thus leaving them with more time to educate themselves and develop their critical thinking."

"No, quite the opposite. The sum total of human knowledge is growing at an exponential rate, to the point where it has become impossible to know all the ins and outs of a problem. AI, with its ability to summarize, allows humans to filter redundant information and focus on what's essential."

"I believe it certainly could compromise humans quest for knowledge & need to problem solve & therefore seriously impair our critical thinking & problem solving capacities & increase depression in people who may in future, have no motivation to use their god-given gifts & intelligence because they have been replaced by AI."

"It would definitely be more of a crutch than a tool if we become overly reliant on it. Instead the development and the products that are created using AI tech should be such that it aids critical thinking, aids skill development and indirectly making life easier."

2. HOW CAN WE STOP FAKE NEWS OR FALSE INFORMATION FROM SPREADING?

This question was open-ended and generated a wide range of potential solutions: providing financial support to media (local, traditional) that fact-check information, investing in quality journalism (with multiple information sources), educating people, using AI to fact-check information, punish those who spread false information, erase it, impose regulations for platforms (such as Facebook) that spread these

fake news. Our collective dependency on "free" (one-way only) news was also highlighted.

Should fake news be censored? One respondent took a stand: "We should circulate fact-checking articles as much as possible instead, because censorship is counterproductive (it can feed conspiracy theories, for example)". One pessimistic point of view: "It may become impossible as AI advances so too will its ability to mimic voices and fabricate images and video."

SELECTED EXCERPTS

"Redefine the journalism profession. Develop an accreditation system for information sources. Recognize communication experts in various sectors of human activity."

"There will always be fake news, we must develop critical thinking and educate youth on the matter."

"Censorship must not come directly from AI. However AI can become a tool to help predict the likelihood of a news item being fake."

"Teach people how to develop critical thinking, search for credible information and open their minds."

3. SHOULD THE RESULTS (POSITIVE OR NEGATIVE) OF STUDIES ON AI BE MADE AVAILABLE AND ACCESSIBLE?

The answer was positive beyond a shadow of a doubt. Many respondents felt that this should be the case for study results in all fields. These results should be open source, according to other respondents (it should be noted that a vast majority of them already are).

SELECTED EXCERPTS

"Absolutely. And as much as possible, break down these results to make them accessible to all. No opaque results, with incomprehensible terms..."

"This question has more to do with research than AI. Publicly funded research, with few exceptions, should be made available as a Social Good."

"Yes. I know people who think really powerful results should be kept from the 'bad guys'. That is a total pipe dream. All you'll do by trying is to disadvantage the 'good guys'. Your best bet is to be open."

"YES!! Especially negative results. They would provide as much information, if not more about a particular problem."

4. IS IT OKAY NOT TO BE INFORMED WHEN MEDICAL OR LEGAL ADVICE IS DISPENSED BY A CHATBOT?

For our survey respondents, the answer was predominantly "no". Their answers were influenced by two concerns: transparency and privacy: "Advice dispensed by a chatbot may be taken into consideration differently if the person knows they're speaking with a chatbot, or believe they're speaking to a human. A chatbot cannot know all the variables for a given situation." Many mentioned that it was easy to let a person know that they are speaking with a chatbot.

SELECTED EXCERPTS

"Eventually, yes. No passenger on a plane asks the M/C whether it's the pilot or the autopilot who is steering the plane."

"The source of such advice being often critical to a person's well-being, one should be aware of the source of this information."

"No, all information should be presented along with the source exactly as it is, along with the analysis of how accurate or biased the information/advice might be. It may happen that the person may rely on that information even after realizing that it is from a chatbot, as it would get good results. And that is the kind of relationship we'd like to foster."

5. TO WHAT EXTENT SHOULD ALGORITHMS BE TRANSPARENT ABOUT THEIR DECISION-MAKING PROCESS?

This question left many respondents uncertain. The most popular response was "as much as possible", while acknowledging the technical difficulties in play (the "black box" problem). Although some believe that AI should simply not make any decisions, others seemed to agree that AI can make decisions, on the condition that there is access to a "justification that can be understood by a human". Transparency may not be desirable in certain contexts. Many mentioned that transparency is important in building trust in AI. One respondent suggested giving a reliability rating for decisions made by AI.

They also noted that transparency involves knowing which data (or which type of data) an AI makes its decision and the values (or interests) guiding its decision.

One participant suggested instead that we not ask any more from AI than we would from a human.

SELECTED EXCERPTS

"A description of the algorithm's decision-making process should be included with the purchase of an AI product, like an instruction manual or the manufacturer's warranty that comes with regular purchases."

"If AI creators cannot precisely define the reach and the limits of an AI's decision-making ability, then that AI shouldn't be marketed and sold."

"The scale of values used to make their decision. See the relative values for different decision-making elements. For example: cat vs. dog, collective vs. individual, etc."

"Completely transparent. How can you trust something if you don't know what principles they are basing themselves on to conduct their analysis? Just like understanding the methodology used by researchers is always relevant."

"You should be able to ask AI why it made a choice then if you find its reasons lacking you should be able to make it change its behaviour."

"We may be able to infer decision-making processes but we should not assume that there is any internal motive or intent in an algorithm."

DEMOCRACY (PUBLICITY, DIVERSITY)

1. SHOULD INSTITUTIONS CONTROL AI RESEARCH AND APPLICATIONS?

The response was positive overall, especially for AI applications (freedom of scientific research is an important value). An "office of the AI Ombudsman" was suggested, along with AI ethics committees or some sort of Hippocratic oath. Participants also noted that "the subject is too intensely political and social to be left in the hands of the private sector". This control, however, should not impede innovation (as long as it is compatible with the common good and human rights). One inherent difficulty for institutional control stems from international politics: how can countries with competing interests agree on common institutions?

SELECTED EXCERPTS

"Yes, on the condition that we develop a participative democracy and that governments are in the service of the majority, not of money."

"No, but establishing boundaries is essential."

"Yes but good luck getting China or Russia to follow along."

"Controlling AI research is simply not possible. The research itself should continue, but a broader communication framework explaining what AI can and cannot do is critical. Sensitizing researchers to the ethical ramifications of their work is also important (e.g. the Hippocratic oath).

2. IN WHICH FIELD IS THIS THE MOST PERTINENT?

The question was open-ended. Many answered, "in all fields". Healthcare easily takes the lead in the fields listed, followed by weapons, justice, environment, food, surveillance, privacy, finance, safety, education and government, respectively. The following were also mentioned: economy, industry, epigenetics, journalism, transportation, municipal services, research on a super-IA (AGI), self-driving cars and targeted advertisements.

SELECTED EXCERPTS

"In all fields related to life (biology) and living in society."

3. WHO SHOULD DECIDE—AND WHAT SHOULD THE TERMS BE—ON THE STANDARDS AND VALUES DETERMINING THIS CONTROL?

Respondents were often unsure how to answer to this, and hesitated between various options: Parliament, public consultations, the overall population (referendum, random draw), a multidisciplinary committee (experts, elected officials, citizens), a science and technology ethics commission, an advisory committee, an international institution (UN-style). The idea that this decision-making body must be independent (from political and economic power) was raised many times, along with the concern that this body must represent the diversity of citizens.

SELECTED EXCERPTS

"I don't know... A joint, multidisciplinary, academic, public and impartial committee."

"All of us, by developing information resources, consultation and decision-making methods that involve as many people as possible from all walks of life. Not the current "democracy".

"A lot of committees. They could establish rules, values, etc., tied to each institution where there would be one of these committees. They could thereby establish some sort of "charter" for the institution and make recommendations... That obviously shouldn't be left to gather dust on a shelf!

"In Quebec, the science and technology ethics commission already produced a document on smart cities outlining the issues to consider. Other AI projects could be analyzed by this body or other government bodies specializing in the field. An ombudsman could be named to certify AI projects and receive flags about the Montreal AI Declaration principles not being respected."

"Since AI affects every field (law, health care, science, society, arts), specialists from each of these fields must be represented within the organization. The government must fund this organization properly, but cannot intervene in how it operates. Furthermore, the government should not have the power to eliminate the organization or interfere with its work."

"Canadians from all groups, backgrounds & beliefs."

"This should function like an IRB as in the drug development and testing industry."

4. WHO SHOULD CHOOSE THE "MORAL SETTINGS" FOR SELF-DRIVING CARS?

There were a number of different answers to this question: Parliament, a government agency, the State, provincial powers, the State in collaboration with the industry, an ethics expert committee, the SAAQ, the car manufacturer, a software certification authority, a user committee, Supreme Court judges, a U.N.-like international organization. The user could also have the choice of certain options. It should be noted that many respondents distrusted self-driving cars ("they should be banned").

SELECTED EXCERPTS

"Again, it could be citizen committees. We'd need a representative for pedestrians, one for seniors, another for youth 16 and under, another for bikes, etc. Each one could have a say on the moral settings for self-driving cars."

'Certainly not the companies building them!"

"An ethics commissioner and the Bureau du Coroner in Quebec."

"It should be a multilateral decision (after thorough public discussion)."

"Judges/supreme court, whoever decides and upholds the existing ethical guidelines should have a major role to play in the decision. But along with them, community participation, transport businesses and authorities, AI researchers and developers."

5. SHOULD WE DEVELOP ONE OR MORE "ETHICAL LABELS" FOR AI, WEBSITES AND COMPANIES THAT RESPECT CERTAIN STANDARDS?

The vast majority of participants agreed, saying it was a "good idea", a "good start". It could be similar to an ISO standard. One respondent wondered, however, why all companies and websites did not have to respect these standards. Another specified: "yes, case-by-case with a standard chart". This also raised some skepticism: Will these certifications be respected? Don't they risk being corrupted?

SELECTED EXCERPTS

"Certifications that would eventually be subject to a vigilant review to adapt to a given situation."

"Communities are different, people are different. (...) We should make sure that by doing so we are not actively causing friction between different groups or trying to homogenize them. The effect should be neutral."

"Definitely, at least three major ones should be developed: corporate, government and individual ethical labels."

RESPONSIBILITY (CAUTION)

1. WHO ARE THE STAKEHOLDERS RESPONSIBLE FOR THE CONSEQUENCES OF AI DEVELOPMENT?

Respondents identified numerous stakeholders: universities, researchers, companies, ethicists, politicians, those who market the apps, the government, the economic decision-makers, those who hold a financial stake, elected representatives, society, users, every one of us. But it was probably "the developers/creators, the companies and the government" that came up most often. Some drew a parallel with pets or children: the owners/guardians are responsible. In the case of AI, it could be the owners, or even those who test the AI, who authorize its deployment.

SELECTED EXCERPTS

"The people who build them, the people who distribute them, and, if we can nab them, the people who use them maliciously to harm, injure, kill or dominate others (including animals), or harm the environment."

"Every member of the supply chain: from the graduate researcher to the multinational firm, including local, regional and national regulatory bodies."

"Companies offering services must be accountable and responsible, but especially company stakeholders."

"Whoever provides the results/ predictions of the AI decision-making. For example, Google is responsible for Google Translate."

"Researchers developing models are partially responsible. However the application of AI ultimately rests with the owner/operators."

2. HOW CAN WE DEFINE PROGRESSIVE OR CONSERVATIVE AI DEVELOPMENT?

Participants had no definitive answer to this question. Progressive development is synonymous with the collective, transparency, smaller wage gap. Conservative development goes hand in hand with caution: there's no point rushing in, better to go gradually. Someone remarked that it seemed easier to adapt legislation around AI than adapt AI around legislation because progress moves quickly and shows no signs of stopping. Another mentioned that progressive development should "Foster alternative research". And a sentiment shared by many was: Let's go, let's go, can we do things differently?"

SELECTED EXCERPTS

"By holdings forums on the subject! :-) The more we talk about it, and the more inclusively, the more progressive AI development will be, in a good way. Also through education. The more our society is educated, the more informed it will be, the more informed its decisions will be."

"For common good vs. private property."

"It is progressive when it is maximizing freedom and agency. It is conservative when it is carefully monitored and cultivated as to insure safety."

"Conservative development: Checking, testing at each and every step. First in isolation, then within an isolated test group, and gradually deploy the AI."

3. HOW CAN WE RESPOND TO PREDICTABLE OUTCOMES FOR THE WORKPLACE?

Many ideas came back over and over: a solid social net or basic income, a tax reform with a tax on robots, or a better distribution of wealth. Access to education and training is the preferred route; however, people will have to adapt, which requires more ongoing training. The transition will certainly have to be gradual and transparent: people must be kept informed. But not everyone is worried: The workplace has always been evolving and will continue to do so." Incidentally, many seem to hope to free themselves from work.

SELECTED EXCERPTS

"Offer a guaranteed salary in exchange for participating in the creation of digital commons."

"Work is not humanity's ideal, nor is it its goal. The free time obtained and the productivity gains generated should be pooled to allow everyone to work less without sacrificing standards of living."

"By redirecting people towards other types of employment that are more involved in building social cohesion."

"Need to slow down the pace; we must first define priorities that aim to develop what can serve humans before what can replace humans."

"AI tax, job displacement compensation, basic living wage, and research/development of new jobs."

"The real cost of the introduction of AI technology is not just the money some people pay for it. It is the social, political, and economic costs—to everybody in society that need to be considered."

4. IS IT ACCEPTABLE TO ENTRUST AI WITH CARING FOR A VULNERABLE PERSON? (FOR EXAMPLE, WITH A "ROBOT NANNY")

Respondents were very torn on this question: "to entertain, but not to heal", "not sure". There seemed to be a fear of humans in healthcare disappearing. The importance of "human warmth", especially for vulnerable people, was brought up. Of course, it still seemed better than nothing: Yes, if there's no other choice." There is also the fact that it may provide better access to care, especially when human resources are scarce. Many highlight the risk of shirking our duties toward these people by entrusting them to AI. The subject is a sensitive one and such robots should certainly be guided and supervised.

SELECTED EXCERPTS

"Not completely. The robot nanny should always be there as a complement to human staff."

"Yes, if you can program AI correctly so it doesn't bypass certain more sensitive skills."

"Up to the vulnerable person to decide."

"No. The result could be disastrous as it has not been studied for decades to determine the social, psychological, mental & physical effects it would have on our children. It could also possibly make our children emotionally unable to connect & bond with their parents, siblings & other humans."

"Of course ... consider how television is sometimes referred to as a babysitter."

5. CAN AN ARTIFICIAL AGENT SUCH AS TAY, MICROSOFT'S "RACIST" CHATBOT BE MORALLY BLAMABLE AND RESPONSIBLE?

The question drew mostly negative responses. The chatbot is not defined as racist "because it doesn't understand anything", and the blame is rather placed on its designers (Microsoft). Nonetheless, the "consequences of its declarations" could have very real impacts. Most respondents therefore agreed that it is unacceptable. One respondent brought up the legal aspect by considering placing AI under guardianship (like children or animals), while another considered them simply as objects for whom responsibility falls on its owner.

SELECTED EXCERPTS

"No, I think we should consider artificial intelligence products as if they were children. Giving them the title of a person without a complete legal personality would be a good idea. That way, each artificially intelligent product would have a human guardian that would be responsible for its actions."

"In the end, it's just a program. And we know to what extent some programs can be bugged, faulty and poorly made."

"Not for the moment, responsibility comes with sentience, if AI isn't sentient, it can't be blamed."

"It's the programmer's responsibility to make sure its robot isn't racist and to make any required changes as quickly as possible."

"We should accept that machine learning algorithms are non-deterministic and empower operators to explore their utility while being responsible operators."

"The responsibility (until proven that the being is actually sentient, if that's even possible) should be taken by: People who gave permission to deploy them > People who tested them > People who developed them. In that order."

"Humans are not good examples for AI agents. AI agents will be more efficiently learning from other AI agents than from human activities."

"No. I think it is always people who must be held responsible. I am against giving machines any kind of legal status similar to people. You cannot punish or hold responsible a machine. So, people must always be responsible."

3. SUMMARY OF SUBMISSIONS RECEIVED

Over 15 documents were submitted following the call for proposals published on the Montréal Declaration's website in November 2017 (with a deadline at the end of April 2018). The objective of this was to contribute to the Declaration's content, either by discussing the seven principles in the preliminary version, or by suggesting concrete recommendations. These documents range from summary reports of collective discussions to individual opinion pieces. They are written in French and English, and can be read on the Declaration's website (this summary obviously cannot do justice to the rich content of the submissions received).

The following abbreviations are used to indicate the documents from the following people or organizations:

AQT
for Association québécoise des technologies

CAIQ
for Commission d'accès à l'information du Québec

MAIEM
for the Montreal AI Ethics Meetup group

OIQ
for Ordre des ingénieurs du Québec

SRAD
for the evening of reflection around the Declaration which was held at UQAM

Hernandez
for Annick, Guillaume and Raphaël Hernandez

McNally
for John McNally

Musseau
for Pierre Musseau-Milesi

Parent
for Lise Parent

Quintal *et al.*
for Ariane Quintal, Matthew Sample and Eric Racine

Ravet
for Jean-Claude Ravet

Robert
for Bruno Robert

Wark
for Grant Wark

PRIVACY

PROPOSED PRINCIPLE

"AI development should guarantee the respect of privacy and allow people who use it to access their personal data as well as the kind of information involved in the algorithm."

GENERAL OBSERVATIONS

The privacy principle has probably been commented on the most in the submissions received. The Commission de l'accès à l'information du Québec (CAIQ) in particular, but also the Montreal AI Ethics Meetup (MAIEM) group, the discussion session on the Declaration held at UQAM (SRAD), the Ordre des ingénieurs du Québec (OIQ), Lise Parent (Parent), Annick, Guillaume and Raphaël Hernandez (Hernandez), Grant Wark (Wark), Quintal, Sample and Racine (Quintal et al.) all suggested recommendations explicitly linked to privacy.

As the CAIQ mentions, in Quebec, there are already well-established principles for the protection of personal information (RLRQ, A-2.1 ; la Loi sur l'accès, as well as RLRQ, P-39.1 ; la Loi sur le privé) that AI development will have to respect: for example, the organizations collecting data must determine ahead of time the reason they are collecting this data and advise the people concerned. Once more the principles of necessity, consent, confidentiality, destruction, transparency, access and responsibility (see CAIQ appendix) can be noted.

Regarding new practices, at least two types of regulation can be considered: one coercive, which focuses on penalties in the event the legal framework is not respected, and the other preventive, which aims for greater flexibility in adapting to change. In Quebec, the CAIQ suggests the second approach and insists on evaluating the risks beforehand, using parameters with the strictest possible default settings, using technology to improve confidentiality, designating a person in each organization who is responsible for the protection of personal information and held accountable as well as "transparency, working for citizens". We have to wonder, however, if the balance of power with major digital multinationals will not also entail more coercive than preventive measures.

This position perhaps echoes that of the OIQ (and Parent) which promotes privacy through design and suggests drawing inspiration from existing best practices, such as the General Data Protection Regulation (GDPR) which recently came into effect across Europe.

This concern for the respect of privacy is often accompanied by a concern for transparency. The MAIEM group suggests, therefore, expanding on the privacy principle by specifying that transparency is essential—an analysis also made by the CAIQ and the SRAD. The close relationship between the issues of protecting certain information (personal data) and being able to know who holds what (access to information) is evident, two elements which are likely to be expanded in the Declaration. We also note the tension that sometimes surfaces between these two elements: when transparency applies to personal information that we would rather keep confidential. Mediation between these two notions may prove necessary.

As well, consenus on this mediation may not be reached, because as the MAIEM highlights, privacy preferences can "vary considerably according to cultures, generations and individuals". One idea for which there is certainly consensus it that we must "preserve citizen control over their personal information and the management of their consent" (CAIQ, SRAD). Quintal et al. also worry that the initial formulation of the privacy principle suggests that data be shared by default (the principle insists on being able know what becomes of personal data, without objecting to the data collection in the first place). "The Declaration should include improved safeguards for privacy of user data."

Lastly the SRAD notes that data anonymity techniques are not yet mature enough to be used without risk. the SRAD also observes the link between data protection issues and the risks

of algorithmic discrimination. But that does not mean that protected data (for example gender or race) should not be collected insofar as fighting discrimination usually implies having access to this kind of information.

SUGGESTED RECOMMENDATIONS

The privacy principle, which includes the concern for transparency, leads to more specific recommendations:

> People need to be informed of, authorized to and able to check use of their personal data at any time (MAIEM).

> We must introduce a culture of "data privacy by default" as is the case with neuroethics, meaning that by default, personal data should not be shared (Quintal et al.).

> The "burden of consent", meaning ensuring that consent is truly free and informed, should fall on companies/organizations and not citizens, just like correcting erroneous information (CAIQ).

> People need to be able to understand how their personal information is being used (MAIEM, CAIQ, Hernandez, Parent).

> People must be able to withdraw their consent regarding use of their personal information (MAIEM).

> Computer codes for interpreting results and algorithm training methods must be made public and open. (OIQ)

> We must make people aware of privacy protection issues (CAIQ).

> People should be able to know the monetary value of their personal information at all times (Hernandez).

Lastly, an original and detailed proposal from Wark answers, to a certain extent, a question put forward by Hernandez : How can we create a private digital space? Basically, it is a matter of using AI to protect against AI.

> Indeed, Wark suggests using "smart contracts" technology to protect personal information and make business transactions and social interactions easier. This can be achieved by developing a secure personal profile and a "loyal AI" that would serve as a personal data manager, thereby solving many of the challenges previously identified. "For example, a loyal AI-agent must not compromise its loyalty to its owner through overt or covert association with a business, such as an online store." To find out more, refer to Wark's document which gives a detailed presentation of what loyal AI might look like.

Many papers discuss how these recommendations can be implemented. From a public policy standpoint, there are at least three ways to respond to this concern for respect of privacy and transparency: through regulation, self-regulation or incentives.

Both the CAIQ and the MAIEM agree that self-regulation is insufficient. Updating existing regulations is more important. Both organizations (as well as Parent) also stress the importance of conducting business and organizational audits. This update could go in a different direction: OIQ supports "flexible regulation mechanisms", which aligns with the preventive approach adopted by the CAIQ.

Finally, we can consider financial incentives for companies that develop technologies to protect privacy, and promote those who make efforts, namely through labels or certifications—a sentiment that seems to be shared by the Association québécoise des technologies (AQT).

JUSTICE

PROPOSED PRINCIPLE

"AI development should promote justice and seek to eliminate discrimination, especially when it comes to gender, age, mental and physical capacity, sexual orientation, ethnic and social origins and religious beliefs."

GENERAL OBSERVATIONS

Like the privacy principle, justice was also present in many of the proposals: MAIEM, SRAD, OIQ, Hernandez, Parent, McNally, Ravet.

The SRAD suggests distinguishing between the various meanings of justice (according to Aristotle): commutative justice that oversees exchanges between people who are considered equal, and distributive justice which is linked to merit. Who deserves what in society? This second meaning is the one that appears mobilized in the submissions received, and it raises a number of questions.

Is it possible to identify a universal justice principle to regulate AI development? Should it not be limited to principles that apply only to a given community? This delicate issue lies at the core of many political philosophy debates.

The MAIEM leans towards a non-universal approach, or at least one which tries to take cultural and historical variations on the notion of justice into consideration:

"The development and utilization of AI-enabled solutions should promote justice and human agency as transparently defined by the target community's welfare-defining organization (e.g. democratically elected government), in concert with the target community. It should seek to eliminate inequality and discrimination within that community."

The counterpart to this reformulation exists in Ravet's more universal approach, which identifies a universal principle in Kant's idea of human dignity and life: "AI innovations must be based on the principle of non-instrumentalization of humans and be careful not to crush life." This approach is also favoured by the SRAD who, in addition to the notion of equal dignity of human beings, introduces the idea of social justice: "AI development should promote social justice and respect equal human dignity, particularly by seeking to eliminate all forms of discrimination especially with regard to gender, age, ethnic origins, social status, etc."

One way to articulate social justice and justice as non-discriminating would be to see the first as correcting (socioeconomic) inequalities, whereas the second seeks to prevent inequalities from appearing and guarantees equal chances. Social justice can also be considered in greater context, as the MAIEM does when it underscores the need to consider different perspectives on justice, especially those from marginalized communities.

SUGGESTED RECOMMENDATIONS

The question of biases (already discussed in the previous section) and the opacity of algorithms (the "black box" problem) also caught people's attention. This is unsurprising given that the issue has received a great deal of media coverage. For example, Parent notes that "assisted, or even automated, decision-making systems in medicine, finance, defence or justice, will give biased results if their input is biased". The OIQ also insists on the need to implement "control and protection mechanisms" to correct the bias.

Other recommendations are also worth mentioning:

> We must train students and AI practitioners in law and ethics. (Parent, OIQ)

> We must foster diversified and female hires in AI system development. (OIQ)

> We must ensure quick and transparent processing of claims by users/citizens who have been negatively impacted by an AI system (OIQ).

Many proposals call for the creation of an **independent oversight body** (Parent, McNally, Hernandez, OIQ, AQT). Its role would not be limited to applying the justice principle, but as it often appears in discrimination issues, this a good opportunity to discuss it.

The form this will take varies from one document to the next. The OIQ talks about an AI observatory, Hernandez describes, "a regulatory body whose task would be to ensure that citizens have a good understanding of the decisions made by AI"; as for the AQT, it advocates "the implementation of a multisectoral advisory committee whose purpose would be to reflect on the opportunities and challenges for Quebec's technology industry in the matter of ethics in artificial intelligence". One could also envision, as McNally does, an oversight organization that would work closely with the government and whose mandate would be to anticipate problems that AI might cause for the society of tomorrow.

RESPONSIBILITY

SUGGESTED PRINCIPLE

"The different stakeholders in AI development should assume their responsibilities by working to counter the risks of these technological innovations."

GENERAL OBSERVATIONS

The responsibility principle does not appear as often as the previous two in the submissions received, but it tends to overshadow the question of the relationship between humans and AI. Who will be responsible for AI, especially its adverse effects? As the SRAD remarked, AI development could extend all the way to using killer robots. This possibility raises, in turn, a commonly shared concern: that humans are handing their responsibilities over to AI. Here we find the theme of AI as a tool: it should be viewed as an extension of human intentionality, but not as an autonomous intentionality (MAIEM).

Among the people and groups responsible, we can include the researchers who, because they possess the knowledge, must start the debate (SRAD). To this we can add those who sponsor the researchers, such as universities, the military or the industries. Being responsible namely means implementing the knowledge and tools to "understand the functioning of AI and anticipate its reactions" (MAIEM).

In an essay that offers a broader outlook on the prevailing understanding of AI rather than expanding any specific principles of the Declaration, Jean-Claude Ravet, editor in chief of the magazine *Relations*, cautions against human instrumentalization in the age of AI and believes that AI development is our collective responsibility, and that we must maintain a global perspective that is both historical and ideological. Thereby, the motive itself of AI as a tool is worth questioning, since "the line between using the technique and the technique itself is blurrier than ever". Most importantly, notes Ravet, we should not kid ourselves about the ideology behind AI development that serves the interests of powerful multinational corporations. For Ravet, this ideology, which tries to pass itself off as scientific speech rather than a social project, is characterized by "an extremely reductive vision of humans and life". (Hernandez also questions the specificity of humans).

The transhumanism movement or the book *Homo Deus* by Yuval Noah Harari are good examples of this reductionist ideology that Ravet condemns: "Under the pretext of making humans more, we must not make them less and make it a means to an end. The sole criteria of making money isn't enough. Nor is the respect of individual choice. Because the issues affect life and humanity itself." We need to look critically look at what often appears obvious: that humanity progresses because of AI and that it is inevitable these machines will make their way into our lives. In other words, we are collectively responsible and that is why humans must always have the last word "as beings capable of speech, feelings, sensations, who are aware of human fragility and the ties that bind them to others, to life and to the Earth" (Ravet).

SUGGESTED RECOMMENDATIONS

- Human beings must ultimately be held accountable for AI-assisted legal decisions (SRAD, Parent, Ravet).
- In the case of engineers, we must ensure professional accountability (OIQ)
- From the standpoint of legal responsibility, we must anticipate eventual disputes over AI systems with non-Canadian jurisdictions (e.g. components designed or built elsewhere than where the system was used) (OIQ).
- To avoid attributing undue responsibility to AI, they should not have the misleading appearance of a moral patient (meaning an individual that can be wronged) that deserves our empathy (MAIEM).
- The formulation of the principle, the intention to "counter the risks" does not go far enough: the people responsible must assume the results of AI development (MAIEM).

WELL-BEING

SUGGESTED PRINCIPLE

"AI development should ultimately aim for the well-being of all sentient beings."

GENERAL OBSERVATIONS

Like responsibility, the well-being principle is often present implicitly, especially in health, safety or even the equal distribution of AI benefits. In fact, according to certain approaches in moral philosophy, this principle could even serve as general criteria for decision-making: when given the choice, we should act so as to create as much well-being as possible. Obviously, as the MAIEM observes, other values may conflict with well-being, especially autonomy. For example, we can find situations where paternalism seemed justified, such as when a moral patient's autonomy is constrained for their well-being. It is hardly surprising, then, that such conflicting values—often discussed by philosophers in moral dilemmas—are considered in the submissions received. However, it is nonetheless true that a principle on well-being must be simple, easy to understand and leave some room for future interpretation (MAIEM).

For its part, the OIQ states that the well-being principle is aligned with one of the main tenets of the engineer's code of ethics (article 2.02) which stipulates that the "engineer must respect their obligations towards mankind and take into consideration the consequences their work will have on the environment and on everyone's life, health and property". For this reason, promoting well-being implies evaluating, to the greatest extent possible, risks related to the deployment and operations of AI applications, keeping in mind that "there is no such thing as zero risk" (OIQ).

We should note that the very inclusive character of this principle, which not only targets the well-being of humans, but of sentient beings as a whole, was not questioned. It may be a sign of our changing mentalities and our relationships with non-human (sentient) animals. The MAIEM and the SRAD expand their notion of the domain of morality to sentient beings while Parent brings up AI interference with animal life. We also note that some papers (Ravet, MAIEM, Parent) seem interested in considering the criteria of life and extending the circle of morality to non-sentient entities (such as plants and ecosystems). These concerns, however, which could be qualified as biocentric, have not been adequately developed to be said to reflect a (fairly radical) moral position: it may be a concern for an anthropocentric environment.

Ideas also seem divided on whether the capacity for AI to be sentient (or sensitive) would be adequate criteria on which to grant it rights or, at the very least, moral consideration. If, for example, a robot could suffer, it would have a legitimate interest in being protected. This point remains highly speculative since AI systems are still very far from having feelings or emotions.

Lastly, in a somewhat speculative and programmatic text, Museau attempts to articulate the notion of moral minimalism developed by the philosopher Ruwen Ogien and the well-being principle recommended by the Declaration. What stands out is that the goal of AI development should be to not harm others nor to improve itself—self-improvement belonging, according to Museau, to both moral maximalism and transhumanism.

SUGGESTED RECOMMENDATIONS

> Formulate the principle to reduce suffering rather than promote well-being (which corresponds with what is sometimes called negative utilitarianism (MAIEM).

> Out of safety concerns, blocking/disengagement devices must be planned when designing AI systems to maintain control in case of failure (OIQ).

AUTONOMY

SUGGESTED PRINCIPLE

"AI development should promote the autonomy of humans and responsibly control that of computer systems."

GENERAL OBSERVATIONS

With regard to autonomy, consensus was reached on promoting human autonomy. This idea especially translates into the theme of AI at the service of humans, as mentioned earlier. The OIQ notes that "robots and AI systems must be seen as tools to assist or help with decision-making, not as a replacement for human judgment". For his part, Ravet insists that humans should not be reduced to machines nor become a means to an end, while Hernandez wonders if AI won't one day replace humans to the point that they become obsolete.

Nonetheless, the truth is that the idea of autonomy is subject to multiple interpretations. The SRAD proposes a detailed analysis grid of the types of autonomy ("condition of an entity which chooses itself the rules to which it submits") divided into moral, political and functional (non-dependence) autonomy. These three types of autonomy can be cross-referenced with three types of situations: the autonomy of a human assisted by AI (for example a person with a disability), the autonomy of a human in an environment populated by AI, and finally the autonomy of AI in a human environment. The SRAD suggests a reformulation, therefore, that further considers these diverse meanings: "AI systems must not harm the autonomy (moral, political and functional) of human beings, but rather seek to contribute to it. AI systems must not be made entirely independent of human beings, but remain under their control (moral, political and functional)". This being said, we should not jump to the conclusion that autonomy should systematically prevail over other values such as well-being, justice or knowledge. Each case must be examined in context. And as the MAIEM reminds us, people's consent remains a good way to guarantee their autonomy.

Although there was consensus on the value of human autonomy, the issue of "AI system autonomy" was more sensitive in that its guardianship could be contested. Therefore, citing an article on digital evolution and artificial life, the MAIEM reminds us that situations where the autonomy and creativity of AI systems could contribute to the general well-being are foreseeable. Nonetheless, the MAIEM states that the autonomy of an AI system should not be sought out in itself if it conflicts with the well-being of a sentient being. These remarks, though relevant, are rather isolated in the documents; they give the impression that we need to keep a close watch on AI systems or risk losing control. Reconciling these seemingly divergent considerations does, however, seem possible: we could maintain control of AI at a certain level on an AI system while authorizing—at a lower level and within a defined framework—AI to find certain solutions to problems in a free and creative fashion.

SUGGESTED RECOMMENDATIONS

The notion of autonomy triggered more philosophical reflections than concrete recommendations, even if some recommendations from other sections are not in the report (for example ones on consent in the privacy section).

KNOWLEDGE

SUGGESTED PRINCIPLE

"AI development should foster critical thinking and protect us from propaganda and manipulation."

GENERAL OBSERVATIONS

A number of links can be made between AI and knowledge. First, from the perspective of cognitive sciences, artificial intelligence can help us understand natural intelligence, each being defined by what guides their capacity for action (SRAD). We may, then, wonder why natural intelligence should prevail over artificial intelligence because at a certain analytical level, humans and animals, just like machines, are causal systems.

In many proposals, the knowledge principle provides us with an opportunity to discuss the issues of propaganda and fake news. Seen in this light, the issue is as much about democracy as it is knowledge. We can, however, question how AI or those who produce and market it are in a position to decide what propaganda or manipulation is. It seems illegal, even dangerous to entrust them with such a responsibility. That is why the MAIEM suggests reformulating the principle to place greater emphasis on transparency: "The development of AI should not hamper critical thinking. It must also proceed in a transparent and open manner, to enable public participation in its development, scrutiny, and education."

Among other themes related to knowledge are public access to AI study results, critical thinking (MAIEM warns against echo chambers), AI education and the opacity of the algorithms, previously mentioned in the justice section. On this last point, the SRAD calls for efforts to not only improve data and algorithm transparency, but also to publish the source codes behind AI.

SUGGESTED RECOMMENDATIONS

> Measures to promote public access to the results of academic studies should be implemented. (MAIEM)

> We must encourage competition and diversity in AI applications so that they benefit society as a whole. (MAIEM)

> We must rethink the business model for social media from other news sites (MAIEM).

> All AI students and practitioners should receive advanced ethics training. (Parent)

DEMOCRACY

SUGGESTED PRINCIPLE

"AI development should foster informed participation in public life, cooperation and democratic debate."

GENERAL OBSERVATIONS

With regard to democracy, many documents (Robert, Parent, OIQ, AQT, SRAD) welcome the Declaration initiative and the opportunity it gives them to have their voice heard. MAIEM sees it as an "important contribution" to international discussions on the subject.

Others are more critical. Quintal et al. contest the very process of how the Montréal Declaration was produced. Although they are in favour of public consultation efforts, they question whether it is a way to render an existing document legitimate. More specifically, they fear that the preliminary version of the Declaration (the seven proposed principles on which this summary is built) may have strongly influenced the citizen debates: "the public should have been meaningfully engaged in deliberating the contents of the Declaration from the very beginning". For Quintal et al., this risks compromising the ultimate legitimacy of the Declaration.

These concerns are, evidently, a call for greater democracy (and transparency and critical thinking) in AI development which, to a certain extent, supports the democracy principle. Furthermore, Quintal et al. specify that democratic good will remain an empty promise if does not come with industry regulations. We also run the risk of companies using algorithms to limit the debate to issues only they deem acceptable (what we, along with the SRAD, could qualify as an epistemic issue with adverse effects on democracy). A similar argument is made by the MAIEM, which notes external regulations appear to be the best solution since, in order to protect their intellectual property, it is highly unlikely that companies will share their algorithms.

As for the principle itself, the MAIEM finds its formulation somewhat vague and deplores that it focuses on democracy when all humans do not live under this kind of regime. The MAIEM therefore suggests replacing it with a "public participation principle" which would read as follows: "The development of AI should promote the dissemination of clear and accurate information to the public to enable open and educated debate on AI and its applications, and encourage open and transparent research collaboration."

Finally, the SRAD mentions that major technology companies (such as the GAFA) hold considerable power nowadays, both political and economic—particularly because they have direct access to a tremendous amount of personal data. This can present a serious threat to democracy, as evidenced in the wake of the Cambridge Analytica affair. Furthermore, insofar as democracy demands a certain socioeconomic equality—at the risk of spiralling into an oligarchy—we must remain watchful of the growing inequalities that will automatically result from AI development. Indeed, the SRAD states that automating a task by AI comes down to transferring wealth to capital (thereby concentrating it in the hands of shareholders rather than employees replaced by AI). Unless there is a framework or regulations, AI risks amplifying the growing economic inequalities that have been observed since the 1950s.

SUGGESTED RECOMMENDATIONS

> Leading researchers in the field in our public universities must remain independent of the private sector. (Parent)

> We must break up major monopolies in the technological industry (SRAD).

> We must seriously consider the possibility of a guaranteed basic income funded by a tax on automation or on capital (SRAD).

> We must encourage new company ownership structures such as cooperatives to fight the concentration of wealth (SRAD).

FINAL REPORT CREDITS

The Montréal Responsible AI Declaration was prepared under the direction of:

Marc-Antoine Dilhac, the project's founder and Chair of the Declaration Development Committee, Scientific Co-Director of the Co-Construction, Full Professor, Department of Philosophy, Université de Montréal, Canada Research Chair on Public Ethics and Political Theory, Chair of the Ethics and Politics Group, Centre de recherche en éthique (CRÉ)

Christophe Abrassart, Scientific Co-Director of the Co-Construction, Professor in the School of Design and Co-Director of Lab Ville Prospective in the Faculty of Planning of the Université de Montréal, member of Centre de recherche en éthique (CRÉ)

Nathalie Voarino, Scientific Coordinator of the Declaration team, PhD Candidate in Bioethics, Université de Montréal

Coordination

Anne-Marie Savoie, Advisor, Vice-Rectorate of Research, Discovery, Creation and Innovation, Université de Montréal

Content contribution

Camille Vézy, PhD Candidate in Communication Studies, Université de Montréal

Revising and editing

Chantal Berthiaume, Content Manager and Editor

Anne-Marie Savoie, Advisor, Vice-Rectorate of Research, Discovery, Creation and Innovation, Université de Montréal

Joliane Grandmont-Benoit, Project Coordinator, Vice-Rectorate of Student and Academic Affairs, Université de Montréal

Translation

Rachel Anne Normand and François Girard, Linguistic Services

Rebecca Sellers, Copywriter, Translator, ESL Teacher

Graphic design

Stéphanie Hauschild, Art Director

This report would not have been possible without the input of the citizens, professionals and experts who took part in the workshops.

OUR PARTNERS

Université de Montréal

CRE
CENTRE DE RECHERCHE EN ETHIQUE

LAB VILLE PROSPECTIVE

IVADO

SAT

Mila

CIFAR
AI &
Society
Program

Université de Montréal | designNsociété

polEtHics
CHAIRE DE RECHERCHE DU CANADA
ÉTHIQUE PUBLIQUE ET THÉORIE POLITIQUE

CIRANO
Allier savoir et décision
Centre interuniversitaire de recherche en analyse des organisations

Québec
Fonds de recherche – Nature et technologies
Fonds de recherche – Santé
Fonds de recherche – Société et culture

UNIVERSITÉ LAVAL
Institut d'éthique appliquée

MUSÉE DE LA CIVILISATION
Québec

CENTRE D'ÉTUDES ET DE RECHERCHES INTERNATIONALES
CÉRIUM
Université de Montréal

crdm_ul
CENTRE DE RECHERCHE EN DONNÉES MASSIVES DE L'UNIVERSITÉ LAVAL

Canada

Centre Culturel Canadien Paris
Canadian Cultural Centre Paris

MEC
Maison des étudiants canadiens

< >

Montréal Declaration Responsible AI_

</ >

montrealdeclaration-responsibleai.com