SCHEDULE 1
PUBLIC FEEDBACK FROM MAIEI WORKSHOPS ON PRIVACY LEGISLATIONS AMENDMENTS RELATIVE TO ARTIFICIAL INTELLIGENCE

Dates: 24[42] and 27[43] February, 2020
Editor: Mirka Snyder Caron
Authors: Montrealers
Total # registrants: 73.

Topic 1.: AI in privacy law: Should AI (artificial intelligence) be governed by the same rules as other forms of processing, (ie. no need of defining AI and maintaining a technologically neutral law), or should certain rules be limited to AI (artificial intelligence) due to its specific risks to privacy and, consequently, to other human rights?

Short Reply:

Defining AI at law is difficult and may fall short of preventing actors from circumventing the definition to avoid regulation. However, a solely technologically neutral law is not the answer either.

As such, specific sections focusing on technical and organizational processing of data, as well as increased transparency, and greater control of the data to be shifted to the user, would be technology-specific sections to legislate about for appropriate protection and guidance to actors and stakeholders concerned.

Comments:

There are unique characteristics of AI, which go beyond automation, and by which processing of data is "scaled up" and may have effects when applied to decision-making processes. In terms of determining whether it is possible to define AI, the group

[42] https://www.eventbrite.ca/e/ai-ethics-quebec-and-canada-ai-privacy-legislations-part-1-tickets-95677119841
[43] https://www.eventbrite.ca/e/ai-ethics-quebec-and-canada-ai-privacy-legislations-part-2-tickets-95689171889

identified that attempts to circumvent AI-specific regulations were made by introducing a “human in the loop”, and this should be prevented.

It was concluded that an AI definition would be very vague, as it is not easy to define automation in a way which would restrict privacy laws to processing AI algorithms.

Instead, the group explored what would be necessary to define or regulate. They identified bias and discrimination as a priority. They identified it was mostly the dataset that is biased and not the algorithm, although the algorithm can significantly reinforce “bias” as processing is “scaled up” without human intervention, which problems go against merely having a technologically neutral law.

Questions arose in terms of the method to prevent bias in datasets by introducing specific privacy laws for AI, and whether the target should move away from AI specifically and more towards the methods or processes of collection and structure the data goes through.

Additionally, there are a lot of actors and stakeholders in the private sector using AI technologies, and the group concluded that actors such as data brokers and processors should be subject to higher security standards, which goes against the precepts of technologically neutral law.

Users should be made aware as to why and how their data is used by AI technologies. Different important factors or conditions were identified, mainly:

- The readability is important
- Education to raise awareness on these disclosures
- Focus should be on protection of all personal data and not just as used by AI, which required more transparency and control, in a manner that goes against the precepts of technologically neutral law.

Concerns were raised as to whether the logic of an AI system could be adequately explained to a layperson, and in conclusion, it was found that due to increased risk in standardizing discrimination, users should have more control through transparency, and higher security standards to be implemented throughout the data lifecycle.

Some specific use cases of AI were identified as most important:

- Facial recognition
- social credit scores

- law enforcement

The conclusion was that it is difficult to define AI at law with a sufficiently broad scope, but that a technologically neutral law would be risky, considering the risk present technologies have in reinforcing bias and illegal discrimination, and affecting human rights and freedoms. As such, specific sections imposing greater transparency and control to the user over the data were identified as primary areas of concern.

Topic 2.: Right to object: Should PIPEDA (Personal Information Protection and Electronic Documents Act) include a right to object as framed in this proposal? If so, what should be the relevant parameters and conditions for its application?

Short reply:

Yes, there should be a right to object, and/or another concept such as a right to negotiate an automated decision with a human, and/or a right to a reasonable inference made the automated agent.

Parameters and conditions for this right are as follows:

- Such right must be made explicit by law
- Transparent disclosure to individual that the decision was made by a human or by an automated agent
- To be weighed against the hierarchy in types of decisions made, in other words:
    - Whether the decision has, or has the potential to, have a direct impact or effect on individuals
    - Whether the decision has, or has the potential to, have an indirect impact of effect on individuals
    - Whether the decision has no, or has no potential to, have any impact or effect on individuals
- The intensity of the impact on the individual, community, or collectivity
- The scale of the impact on a same individual, community, or collectivity
- The type and sensitivity of data
- Disclosure of objective and subjective criteria programmed in the automated agent, and the right to contest the input of subjective data which may negatively affect the result of the decision for an individual
- To be weighed against the types of impact, the following being non-exhaustive examples:

    - Whether the decision has, or has the potential to, have an impact on the rights and freedoms of individuals and/or communities
    - Whether the decision has, or has the potential to, have an impact on the health and well-being of individuals and/or communities
    - Whether the decision has, or has the potential to, have an impact on the economic interests of individuals, communities and/or entities
    - Whether the decision has, or has the potential to, have an impact on the sustainability of an ecosystem, social, environmental or otherwise.
- The efficiency of anti-discrimination laws

Comments:

The general answer to the above mentioned question was **yes,** the individual should be able to object to a specific business decision made about him by an automated agent and which may have an effect on him. Other concepts of such rights were discussed, such as the “right to negotiate” an automated decision with a human, as well as the right to a reasonable inference by such automated agent, as in how much would be considered “reasonable” or not, and also how much arbitrariness or “freedom” in the business criteria ought be provided to businesses to build their models.

Greater transparency was requested for identifying whether a decision was made by a human or an automated agent.

It was also proposed that due to the lack of individual’s power against businesses, a legislative framework was needed to level the playing field and balance powers between businesses and individuals.

Considerations were given into determining the hierarchy in the types of decisions that could be made. For instance, some decisions have a direct impact or effect on individuals, while others do not or may impact the individuals indirectly. Such impact may also vary in intensity for the same individual, and in scale for the same community.

Concerns were raised as to determine whether or not the criteria embedded within the automated agent was objective or subjective, and how to mitigate or prevent risks associated with detrimental or discriminatory subjectivity.

The hierarchy of the criticality to such right to object would have to be generally considered in conjunction with the following parameters to be weighed:

- the rights of individuals and/or communities (individual good AND collective good)
- the health and well-being of individuals and/or communities
- the economic interests of individuals, entities, and/or communities
- the ongoing sustainability of the ecosystem.

Furthermore, the right to object would be hierarchized based on the type of data that is being analyzed, or in other words, the sensitivity of the data. However, it was considered that alone such criteria was not good, since even non traditional sensitive data, one can reconstruct very precise profiles.

Another consideration was that of the importance of effective anti-discrimination laws, as part of the protection of human rights and freedoms of individuals, including that of the right to privacy. Arguments were raised in saying that to avoid gaps or competencies overlaps between different legislative frameworks and regulatory mandates, the potential opportunity of overhauling the siloed privacy approach and create an effective ecosystem of data protection with a broader legislative framework than that focused solely on the right to privacy, would be more efficient, but would require major restructuring of existing regulatory entities.

The right to object was linked with the right to explainability, by saying that to accept a decision, there needs to be sufficient trust in the system. It did not seem at the time of the discussion that there was much trust in the manner, method and governance businesses were going to design and deploy automated agents without more stringent legal requirements.

As such, the right to object was linked to protecting and preserving the ability to think, to make personal decisions (ie. right to self-determination), and to be creative which, due to the potential loss of flexibility in the automation of decision-making systems, this could be negatively impacted. For instance, AI based decision-making tools can introduce at least 2 kinds of bias, data bias and cognitive bias, cognitive bias being defined as absolute reliance on decisions made by the automated agent, and in such case you eventually lose critical cognitive reflexes (this is also linked to an inertia and over reliance biases).

Topic 3.: Right to explainability: If incorporated, what should it entail? And would enhanced transparency measures significantly improve privacy protection, or would more traditional measures suffice, such as audits and other enforcement actions by regulators?

Short reply:
Yes, it should be incorporated.

3 different categories of factors to be met:
-Factors pertaining to data
-Factors pertaining to the result generated by the automated agent
-Factors pertaining to cybersecurity risks

1) Factors to be disclosed and questions to be answered to ensure explainability were identified as follows:

- What data is being used?
- What data specifically is necessary for the identified purpose?
- Why is the data needed?
- Re-disclosure when purpose is changed/expanded -> new consent required
- How is the data used?
- How is the data accessed?
- From where is the data being accessed?
- Who has access to the data?

2) Factors explaining how the output of algorithmic data processing were generated:

- Logic of decision-making
- Decision criteria
- How is data displayed (aggregated or individualized)?
- Explanation of access rights and consent protections
- How can users have access to the data collected on them?
- How can users retract or limit their consent?
- Other consequences on rights and interests
- Explanation of safeguards against bias in algorithmic decision-making -> would necessarily be a continuous improvement, even disclosure of safeguards would always need to be updated
- How was the data set collected?

- What variables are used in the algorithm?
- What variables may serve as proxies for identity-markers that may be a ground for discrimination?
- What efforts are being made to track improvements in rooting out bias?

3) **Factors to respond to cybersecurity concerns**:

- What is the company doing to keep personal data safe?
- What is the level of “trackability”, based on provenance, lineage and vulnerability?

Yes, enhanced transparency measures will improve protection, **and no,** traditional measures will not suffice and present audit and enforcement measures need to be **more deterrent and credible.**

Comments:

Enhanced transparency will ensure more trust between individuals and businesses and, an incentive for more responsible and efficient privacy governance by businesses concerned.

Considering cybersecurity issues and implementation of automated agents, present developments require enhanced transparency and enforcement mechanisms by governmental actors.

It appears that pertaining to cybersecurity and human rights risk of automated agents, businesses would be held to obligation of means rather than obligation of results. It is expected that the government would commit itself in making the best efforts to safeguard against bias, and protect against data breaches, etc., and a similar commitment is expected from businesses as well. Realistically, an absolute commitment to precluding bias or data breaches from occurring in the first place is not possible, since the technology evolves too fast to be able to commit to an obligation of results.

Measures of transparency were found necessary, as findings were that the current traditional measures do not suffice. It was proposed to put in place an enhanced OPC (Office of Privacy Commissioner).

A possible recommendation was to put in place a specific governmental agency responsible for evaluating new AI technologies and applications in areas of public interest, in other words, having an AI-specific version of Pest Management Regulatory Agency **(PMRA)** and other regulatory agencies.

Responsibilities of such AI regulatory agency would be to oversee the enforcement and compliance with transparency measures. There would be a need to have 2 different implementation models, one for the private sector, and another for the public sector. Other quality control and expert inspection would be available to validate accuracy of models and to efficiently track efforts in improving such accuracy.

To address and oversee safeguards implemented against bias, it was proposed that there should be tracking of such obligation of means, and a disclosure of the algorithmic design when the private company has a public sector mandate.

Concerns were raised pertaining to the use of personal data without consent. It was proposed that there should be a generalization of AI regulations beyond just AI applications: as such, the regulation would cover any use of personal data. It was suggested that public institutions which are not held to a duty to disclose the use made of personal data were typically the ones hoarding the most personal data.

It was also proposed that for adequacy reasons and other reasons, there should be comparative brainstorming sessions made with the *EU Guidelines on Artificial Intelligence and Data Protection,* in particular pertaining to the right to obtain information on the reasoning underlying AI data processing operations applied to them, and the consequences of such reasoning.

Topic 4.- Privacy by Design & Mandatory testing: Should Privacy by Design be a legal requirement under PIPEDA? Would it be feasible or desirable to create an obligation for manufacturers to test AI (artificial intelligence) products and procedures for privacy and human rights impacts as a precondition of access to the market?

Short reply:
Yes, privacy by design should be made a mandatory legal requirement under PIPEDA.
Yes, it is both feasible and desirable to impose such an obligation.

Proposition:

- Prior to testing, establishing:
    - Ethical expert within the company (having both tech & ethics expertise), modelled on resident medical ethicist
    - Establish a self-regulatory professional board of experts (diversity in knowledge), modelled on existing professional bodies (medical, lawyers, banks, etc.): they would oversee the testing of AI systems.
    - Apply for a license for each project
        - License differs according to the sensitivity of the data & vulnerability of the audience

Implementing a point system like the driver's license (incentives, etc.)

During testing:

- proceed in an iterative instead of adversarial way(questions; incentives; etc.)
- Building in a feedback system (for explainable AI)
- Pushing a disclaimer at the release

General framework to ensure efficiency:

- Give more power to the Privacy Commissioner:
    - (General recommendation) Selected by multi-party committee in Parliament instead of the party in power
- Ongoing consent of the user/audience
- Compliance with ethics as a requirement for obtaining a tax break or fiscal incentive

- Community service if not compliant
- Each project: ongoing ethics training for tech experts. Like getting an ethics approval. (Makes sure that we are in sync with technology)

Comments:

Privacy by Design:
Privacy by Design was emphasized as being the obligation to include privacy protections into the technologies themselves. It was found that the difficulty lied not in implementing, but ensuring efficient auditing and monitoring of such. In terms of building in privacy, the example of the feedback system like XAI in the States was identified, to explain decisions.

It was brought up that there are often public statements made by businesses assuring that “Your privacy is important to us”. However, despite terms and conditions and privacy policies, at times it is still not possible to know how the data is used. Additional details ought to be provided to tell the consumer or individual concerned the way to find out how it is used, and if there is a breach, a way to know how the breach happened.

Mandatory obligation for testing AI prior to access to market:

It was concluded that manufacturers should have a mandatory obligation to test their AI systems against privacy and human rights impacts before gaining access to the market or marketing their products or services. There was awareness at how difficult this could be, and what would be the best way to ensure this be done.

Another difficulty was identified, namely, that in practice, tech designed wait until the problem happens because the technology is not designed to foresee problems, but to deal with them as they come. The proposal to remedy this issue was to put in place a self-regulating board, with sanctions and licences, such as a professional order that can be found for doctors, engineers and lawyers. In such a way, a certain degree of professional investment is required to maintain a certain level of ethics and professional conduct.

Augmenting deontology and ethics obligations of engineers and programmers:

It was perceived that engineers’ deontology code, as compared to doctors or lawyers, seemed to be the weakest. Furthermore, many programmers to date have learnt to

code and program without any standard academic curriculum or licensed training, which would ensure that adequate training and testing be made pertaining to ethical and responsible AI design and development, including that about preventing and mitigating impacts on privacy and other human rights.

Various examples were provided from the practice and field of medicine. It was found that there may be a lack of teaching the people who deal with the data about 1) the tech literacy issues and 2) the implication of using such technology surrounding automated and machine-learning models and ensuring privacy protection when using such sensitive health, biometric and genetic data.

It was proposed that within hospitals and other research labs and manufacturing labs, there ought to be a medical ethics and/or bioethics expert, which was defined as someone who would be an expert with training in applied ethics & in computer science, in other words, someone who know how to ask the right questions, as well as someone who understands how people code, so that such expert may process questions to coders. It was deemed mandatory that such an expert was to be a coding expert.

General framework for testing AI:

The proposal stated above was then generalized to recommend there to be an ethical expert within a company to assess the impact and oversee the testing of the AI, one that would have the expertise in AI & ethics, tech & philosophy. To create an analogy, the same way a data protection officer needs to be designated for a company, a similar title and position would be designated for AI systems.

Another proposal was to implement a board of experts to oversee such testing, which would be provided with the explanation of how the AI was built and how it worked, how it was structured, and to disclose this to the public, without necessarily revealing the intellectual property protected aspects of it.

Additional organizational measures could be taken by training the board members to augment their technical literacy, and to regularly test such knowledge and maintain it up to date.

Particular to the testing, it was proposed that it ought to follow a step-by-step process as the board members or other experts go through the approval of projects, favouring the iterative process over and adversarial design in technologies. Furthermore, the process should take into consideration whether the project would take in a particular sector of

the market or economy, identified as the 3Ps: Private sector, Public sector, and Plural (combined partnership).

It was also suggested that the regulatory body in charge should be diverse in terms of knowledge, and have an explicit obligation to ensure such diversity. More power should be delegated to the Privacy Commissioner in terms of a reward and penalty system, akin to that of a driver's license points and score, and to treat responsible AI ethics, and privacy ethics as a baseline to such scoring.

General conditions for such system were identified as follows:

- penalties to be proportional to the repercussions implied (principle of proportionality)
- punishment should not be only financial but should expand to such bans against ongoing or future commercial projects and/or governmental relations.
- punishment and compensation measures should also be able to impose community services: if they are non-compliant, they must do something for the community.
- Putting supervision mechanisms in place to ensure the OPC is not, and is not becoming a partisan from lobbying groups, political parties or businesses: considerations as to designating the representative by a multi-party committee and not only the party in power.

General conditions for testing process were identified as follows:

- Provide coaching and measures within the company to ensure that there is meaningful awareness of potential social impacts in the design and scaling of an AI product, as this is the responsibility of the tech company.
- Including & pushing a disclaimer informing the user that an AI state is being used in the product, and not wait for the user to request such information.
- Leveraging Canadian diversity as a national strength, to ensure diverse feedback at the time of testing, from a vast array of the population.
- Creating a mandatory interactive training of the company with the user, to provide for a mandatory formation of the user by the company, when using product or service. (as an analogy, banks already a mandatory obligation to augment financial literacy; this can be done to and for tech businesses and users)
- Mandatory training of tech experts in ethics and vice versa.

Topic 5.-Can the legal principles of purpose specification and data minimization work in AI (artificial intelligence) context and be designed for at the outset? If yes, would doing so limit potential societal benefit to be gained from use of AI (artificial intelligence)? If not, what are the alternatives or safeguards to consider?

Short Reply:

Purpose specification and data minimization are insufficient in terms of mechanisms for dealing with AI systems, and in some cases are impractical.

The group recommended:

- Imposing respect of human rights explicitly in the legal framework for governance
- Identifying clear Go/No Go Zones, or Go/No Go Rules, for practical guidance.

Comments:

The group proposed that human rights be used as a “ceiling” for limiting specific AI applications within the general data protection and privacy framework. Such limitations are to include limitations on the use of data. The proposed test to protect individuals was identified as a rule: Data cannot be used to detriment anyone’s life.

The group was also aware that there were limits as to the practicality of the consent model which became near infeasible in the context of AI systems evolving daily, and for which it is potentially impossible to know for which purpose it is achieving its tasks.

It was proposed that a list of things, activities or purposes which would be clearly and explicitly not permitted or unauthorized would be more effective, similarly to that of the Go/No Go zones. Despite limits of such identification, some purposes and some applications appear reasonably clear that there are permissible and not, and these should be highlighted for practical guidance. An example of such No Go Zone - or No Go Rule- is that data cannot be used to harm society, and should therefore not exacerbate differences or political uses.

Topic 6.-Is it fair to consumers to create a system where, through the consent model, they would share the burden of authorizing AI (artificial intelligence) versus one where the law would accept that consent is often not practical and other forms of protection must be found?

Short Reply:

Consent model should remain but should be augmented, while being supported by other forms of protection as well. A combined approach would be optimal.

The group recommended:

- Using plain language instead of “legalese” for terms and conditions for consent.
- Imposing mandatory opt-outs for “optional” data versus necessary data (data minimization), as well as imposing more flexible settings in terms of different options for use of data, instead of an “All-in” or “All-out” approach.

Comments:

Challenges linked with the consent model were identified as unclear, and very long and complex terms and conditions written in legalese. It was proposed that instead of catering to protect the liability of the business or corporation, the privacy policy and consent models used should be catered in a manner understandable and useful for the user’s interests as well, since he has to provide consent in a meaningful manner.

It was proposed that the use of general language should be made for conditions applicable to the consent model or mechanism, as understandable by any user -and not as readable by lawyers only-, in order to provide for actual consent.

It was additionally proposed that in cases where consent may be absolutely impractical, access to data audits could be seen as an alternative method for consent mechanism to be withdrawn.

Finally, it was proposed that for further options to the user or individual concerned for the use of data should be available, or an opt-out opportunity, and that general guidelines should let a user know how data can or should be used even in the case of AI.

Topic 7.-What could be the role of de-identification or other comparable state of the art techniques (synthetic data, differential privacy, etc.) in achieving both legitimate commercial interests and protection of privacy?

Short Reply:

De-identification is necessary and should be held at the highest level of priority.

Proposal:

- Define tangible criteria to use to de-identify information
- Government should work with technical regulator to monitor non-compliance and judge upon new cases and non-compliant cases

Comments:

The group brought up the unclear scope of the concept of legitimate commercial interests and identified potential dilemmas between what needs to be used and what the business would want to use to better exploit its business. The question that arose was: what type of data would be considered legitimate, and which one would not?

The general agreement was that there should not be a stringent prohibition preventing the collection of data since it is not possible to know how the data could be use -for social and individual good- in the future.

Concerns were raised about AI systems and de-identification techniques being like “moving targets”, rendering the imposition of a specific technique in regulation as difficult, but then which inevitably pushed businesses into using risk management models. It was said that de-identification may at times even lose all meaning since technologies are evolving and it is not possible to know in the long run what will be possible in terms of re-identification.

Despite the above, de-identification techniques and organizational measures, were seen as a high priority, and the group converged towards proposing a fixed security measure combined with a certain flexibility and judgement-based assessment or

mechanism to monitor the practices. It was suggested that IT auditors would become a necessity to verify compliance with such security criteria.

One proposal was to have tiers of data, or different categories of data, which should be held with different security standards depending on the sensitivity of the data.

Topic 8.-Is data traceability necessary, in an AI (artificial intelligence) context, to ensure compliance with principles of data accuracy, transparency, access and correction and accountability, or are there other effective ways to achieve meaningful compliance with these principles?

Short Reply:

Yes, data traceability is necessary, and should be imposed as a mandatory business practice, but technical and organizational measures should be put in place, to ensure that such traceability records would or could not be used for re-identification of personal data.

However, it should be combined with other ways to effectively achieve compliance.
A few examples that were explored were:

- Creating a presumption, or explicitly shifting the burden of proof at law relative to compliance, non-negligence, respect of human rights and of privacy upon the business, and away from the individual.
- Mandatory internal and external audit mechanisms

Comments:

The question in the manner it was drafted was found to be ambiguous in meaning and in scope, and it was suggested it would be good for it to be more specific. For instance, the group had difficulty if the traceability related to the manner in which the data was input, processes, weighed and recorded within the AI systems, or if it referred to the manner and method of tracing back the data sources to their origins, as well as to whom received such data.

It was proposed that every step should be documented, including the metadata in an unaltered fashion, and it should be made easily ready for audit from a legitimate body. This would provide for additional protection pertaining to explainability, accountability

and auditability. The need for guidelines and standards for documentation arose from the discussion.

Concerns were raised about how the concept of audit, from traditional accounting, was increasingly shifting, and the need to consult specific technical agencies was identified. The solution of auditing compliance was justified on grounds that audit from an independent and neutral third-party would treat everybody equally.

In conclusion, data traceability should be insured to the extent that data was de-identified before cannot become identifiable with tracing.