TECH INFORMED POLICY

AUGUST 18, 2020

FACIAL RECOGNITION MORATORIUM BRIEFING #2

Conditions for Lifting a Moratorium on Public Use of Facial Recognition Technology in Canada

Produced by

Taylor Owen, Policy Lead, Beaverbrook Chair in Ethics, Media and Communications, Director of the Centre for Media, Technology and Democracy, and Associate Professor in the Max Bell School of Public Policy, McGill University

Derek Ruths, Tech Lead, Director of the Network Dynamics Lab and Associate Professor of Computer Science, McGill University

Stephanie Cairns, Research Assistant

Sara Parker, Research Assistant

Charlotte Reboul, Research Assistant

Ellen Rowe, Research Assistant

Sonja Solomun, Research Director, Centre for Media, Technology and Democracy, McGill University

Kate Gilbert, Graphic Designer

Centre *for* MEDIA, TECHNOLOGY *and* DEMOCRACY

network dynamics @mcgill
measuring and predicting large-scale human behavior

ABOUT TIP

Tech Informed Policy (TIP) is an initiative spearheaded by two leading McGill researchers—Dr. Derek Ruths, Director of the Network Dynamics Lab and Associate Professor of Computer Science, and Dr. Taylor Owen, Beaverbrook Chair in Ethics, Media and Communications, Director of the Centre for Media, Technology and Democracy, and Associate Professor in the Max Bell School of Public Policy. TIP aims to demystify the technology underlying critical policy issues and to provide valuable, tech-based recommendations to Canadian policymakers.

For enquiries, please contact Derek Ruths.

Glossary of Terms

Artificial Intelligence (AI): Artificial Intelligence is a system designed to accomplish a task that normally requires human intelligence. AI systems "learn" how to do things by processing large amounts of information, finding patterns, and translating that knowledge to tasks.

Algorithm: A set of rules and procedures that a computer can follow to complete a certain task.

Database: A database consists of one or many datasets that have been organized and retained in a system/software program.

Dataset: A collection of data.

Distribution Curve: A graph of the frequency of different values.

Match Score: A score between 0 and 1, indicating the likelihood that a pair of images depict the same person.

Match Score Threshold: A value between 0 and 1 - pairs with match scores above this value will be flagged as matches.

EXECUTIVE SUMMARY:

This briefing is part two of two on Facial Recognition (FR) Technology. This briefing explores the conditions for lifting a federal moratorium. The first briefing addressed how FR works and is used, as well as the implications of a federal moratorium.

- This briefing outlines the technological, social, policy, and legal conditions required to lift a Canadian moratorium on FR systems.
- Amidst growing calls for Canada to impose a national moratorium on facial recognition (FR) technology, a holistic approach to both technical and policy conditions is needed.[1] [2]
- Some private companies, such as Microsoft and Amazon, have enacted moratoriums on selling FR technology to law enforcement, although they do not include limitations on current uses of their services.[3] [4] Clearview AI recently ceased all Canadian operations due to an investigation by the Office of the Privy Commissioner (OPC) regarding the use of its services by both the RCMP and the Toronto Police.[5]
- Private sector moratoriums are not a solution to addressing the policy implications of FR systems, nor are they adequate fixes for structural technological problems. Instead, industry-led moratoriums are an opportunity for the Canadian government to enact a national moratorium on the technology.
- A national moratorium would, however, afford governments time to evaluate and develop the necessary conditions FR tech companies and public sector actors should follow.
- These conditions should include data governance frameworks, accountability measures, privacy protections, and social impact assessments, among others.

    i. Technological conditions for lifting a moratorium are inseparable from the social and policy considerations detailed below and must always be implemented in tandem for any decision regarding FR use in the public sector (for instance, if bias and accuracy conditions are met, but data protection conditions are not, the moratorium should not be lifted).

    ii. For this to occur, strengthening existing laws (e.g. PIPEDA) may be required alongside new policies for biometric or FR-specific systems.

STRUCTURE OF BRIEFING NOTE #2

CONDITIONS FOR LIFTING A MORATORIUM

We give a number of conditions for lifting a moratorium, grouped into the following sections:

- Purpose conditions
- Data usage conditions
- Accuracy & bias conditions
- Review & oversight conditions
- Social conditions
- Legal conditions

As previously outlined, harms caused by FR systems[6]—including increased surveillance, identity discrimination, data abuse, and privacy infringement—cut across technological, social, policy, and legal spheres. As such, conditions for its safe use must likewise exist at the intersection of these areas.

Considerations/Questions

Every condition outlined in this document can only be developed, precisely formulated, and ultimately met through significant research by or consultation with experts in the fields of technology, social science, policy, and law. Some conditions are clear-cut but difficult to satisfy—FR technology must be proven to be unbiased before its safe use can be ensured, but knowing that FR systems should exhibit no bias doesn't directly enable us to reduce it or even to accurately assess its presence. Other conditions must be developed and defined by teams of experts—FR technology must adhere to a strong data governance framework, but the precise details of that framework remain unclear.

These sections outline some of the main questions and considerations that need to be addressed by policymakers and by teams of enlisted experts before a moratorium can be lifted.

Next Steps: Research and Evaluations During a Moratorium

Finally, we lay out the specific research, consultation, and assessment efforts that should be undertaken during a moratorium to solidify the outlined conditions and address their accompanying questions and considerations.

PURPOSE CONDITIONS

SPECIFYING USE OF SERVICE

A public institution wishing to use an FR service should specify their intended use cases. These requests for FR use should be assessed by a government regulatory body and specify:

- The desired output (e.g. a mugshot matching the profile of a suspect on CCTV footage)
- Actors who will have access to the service (e.g. high-level law enforcement officers with sufficient FR training, working on a office device)
- Situation(s) in which the service will be used (e.g. the RCMP's use of FR technology previously provided by Clearview AI for the NCECC)
- Potential future uses of the service that are not currently planned (e.g. the RCMP's testing of FR technology for potential use in other divisions)

Defining the use of FR will inform the specification of the service, the necessary data, and the required accuracy threshold.

PURPOSE LIMITATION

The institution should also track all actual uses of FR to ensure that the technology is being used for its intended purpose(s) and verified through periodic auditing.

Considerations/Questions:

How will use cases be assessed and by whom?

Policymakers must decide which organization(s), existing or yet to be created, will assess the uses of FR services and address and punish misuse. Policymakers must also determine who will assess and enforce prohibited use cases—an independent regulatory body or the institution itself.

How will potential future uses be assessed?

Any FR regulation must include provisions for potential future uses, while acknowledging that it is impossible to know what a future technological, political, and social environment will look like. This task, although difficult, is vital to ensure that any post-moratorium policy does not lag behind technological innovation.

What laws/policies need to be in place to ensure safe and proper use?

It is not sufficient to require that internal uses of FR follow Canadian law, as current law is insufficient at ensuring safe and proper use. Audits that seek only to verify that uses satisfy existing Canadian law are not a feasible condition to lifting a moratorium.

How should FR technology be used in law enforcement? Should it be used in its current form?

Because of its high risk of false arrests and racial bias, FR technology, in its current state, should not be used unless deemed essential to protecting the safety of Canadians.[7] [8] However, future iterations of the technology may have much lower risk levels. A use case framework for law enforcement should be equipped to assess risk and public good and to determine which uses of FR technology can be permitted. Setting up such a framework would require further research, but its adoption would be necessary to guarantee safe police use to any reasonable level of certainty.

DATA USAGE CONDITIONS

DATASET CONTENTS

An FR system requires two collections of facial images: a training dataset held by the FR provider and a database to search through to identify matches. The latter, which may also contain names and other personal information, is held either by the provider (e.g. Clearview AI's web scraped database) or by the institution (e.g. a police department's collection of mugshots).

If the search database is held by the provider, they should disclose its contents, uses, and origins to the institution, as well as whether they collect data generated by their users (i.e. the institutions using the service), and for what purpose that data is collected.

DATA HANDLING

Both the institution and the provider should be prudent with FR data, considering its sensitive nature, by following existing conventions about proper data handling.[9] Data should only be collected and stored according to necessity, as dictated by data minimization principles.[10] Data generated by an FR service should only be used for its intended purpose, kept as long as necessary, and individuals whose faces are in the search database must have the ability to request the removal of their likeness.[11] Furthermore, data should only be shared when necessary—this includes with the provider themselves.

Consideration/Questions:

What data should be allowed to be collected? How much data? For how long should it be retained?

While existing data usage conventions are insufficient to adequately regulate the use of FR, they should not be neglected; rather, these conventions should be researched, discussed, and expanded. Future regulation must balance data privacy protections, innovation within FR, and the specific environment in which the service is used. For example, data generated by an FR system may be valuable or even necessary for training the technology further. By restricting provider access to data generated by the user of an FR service, the institution may also be impeding the continued improvement of the FR service.

Furthermore, it may not be feasible for law enforcement to follow data conventions when using FR technology during criminal investigations, an especially prominent issue when determining how long data may be retained. These questions will likely only be answerable by each institution wanting to use FR technology; however, the answers will need to be assessed by an independent governing body.

Should individuals have the right to be informed when FR technology is used on them?

The large variety of potential uses of FR technology complicate the application of conventions of consent.[12] FR technology may infringe on Canadians' personal privacy including through unconsenting data collection or third party use. A healthy privacy environment is predicated on transparency, allowing users to understand where their data is being used and/

or collected. Depending on the specific use of FR, questions about who the service is used on must be answered. Will individuals be notified if an FR system is being used on their face—for example, will shoppers be made aware that security cameras in a mall are FR-enabled?[13] Will subjects of criminal investigations be notified that the use of FR technology led to their arrest? Will the data be considered admissible evidence if the service was used without a warrant, considering the subject likely did not consent?

Should individuals have the right to consent to their inclusion in a database?

Appearing in a training dataset presents little risk to individuals—a much greater risk emerges if an individual's photo is included in a search database. Companies like Clearview AI have enabled clients to search for subjects among billions of web-scraped photos, giving clients access to anyone with a visible online presence.[14] Any proposed data governance framework must therefore address how photos are added to search databases and whether individuals are able to consent to their inclusion. It is likely impossible for individuals to be notified if their images are included in a massive web-scraped database such as the one used by Clearview AI. Requiring that FR providers find the contact information of each person in such a database would not only be infeasible but would likely result in a greater breach of privacy than the original inclusion of the photo.

Further deliberation is required to decide which principle will be prioritized: the obligation of the provider and institution to obtain prior consent, or the right of the individual to request that their biometric data be deleted from databases. These conversations should include government, FR providers, and Canadian citizens.

Should Canadians have the right to be forgotten?[15]

Government regulation must dictate if and/or how individuals will be able to remove their likenesses from databases, and what method would protect their privacy. As a recent example from Clearview AI has shown, the removal of one's face from a database is not sufficient protection of biometric data. The only way to remove an individual's facial information from Clearview AI's database is to provide them with an additional photo of the individual to ensure that their face never re-enters the database.[16] This problem is not unique to Clearview AI—if an individual wished to be excluded from any search database, they would need to provide a photo, or if the database is labeled, some combination of uniquely identifiable information, like a name. This information could not be deleted; fully anonymized alternatives, such as retaining only the photo's numerical representation generated by the FR algorithm, would be infeasible, as the algorithm will change over time.

Because the process of removing an individual from the database requires a further privacy violation, allowing individuals to request the removal of their likeness does not guarantee privacy protection.

DATA GOVERNANCE FRAMEWORK

A data governance framework regulating public use of the technology should be developed, preferably by the government. This framework, centered around transparency and accountability would outline the terms and conditions of FR use in Canada and best practices of institutions and providers.[17]

Below is an example of an FR governance framework developed by the World Economic

Forum.[18] All four steps require extensive research and deliberation before implementation.

The EU's General Data Protection Regulation (GDPR) provides European institutions and corporations with a responsible data governance framework built on seven principles for the lawful processing of personal data.* Generally regarded as a global benchmark in responsible and accountable data protections, the GDPR was successfully applied for the first time by a French court to rule against the use of FR technology in the public sector. The court concluded that the use of FR is schools violated the GDPR's principles of proportionality (whether a less intrusive method could achieve the same objective) and data minimization (whether irrelevant or excessive data would be collected).[19]

ACCOUNTABILITY

In addition to FR-specific data governance for high-stakes uses in the public sector, (possibly new) accountability measures must be developed prior to lifting a moratorium. While data protection has been the global benchmark for AI governance in the last decade, policy around FR systems and AI more broadly is shifting toward decision-making with AI. To that effect, several new frameworks have been suggested, such as risk assessments and algorithmic accountability to account for certain limitations of data governance including the troubled distinction between personal data and non-personal data. While privacy protections are successful measures for data confidentiality and data security, harms of FR go beyond privacy.[20] Companies must be held accountable for not only ensuring data and personal information is protected and responsibly used but also for the decisions being made through their systems, especially if they cause undue harm, discrimination, human rights infringements, or risk to citizens. Risks may include exclusion from access to resources or care and are of particular importance to already vulnerable and racialized groups.

The safe use of FR systems in the public sector requires conditions beyond the data going in and out of the technical systems but also those that can account for the actors using it and their impacts on citizens.[21] For example, the Algorithmic Accountability Act in the United States shifts a focus in governance from data and privacy considerations to decision making using AI and would require "companies to evaluate the privacy and security of consumer data as well as the social impact of their technology, and includes specific requirements to assess discrimination, bias, fairness, and safety".[22] [23] This represents a shift in thresholds for law and policy applicability centered around how algorithmic systems impact peoples' rights and legal freedoms.

* These include fairness and transparency, retention, minimization, security, storage and purpose limitation, and accountability as broadly outlined in Europe's GDPR, Article 5.

Considerations/Questions:

Who will develop this data governance framework? What will it entail?

This guiding framework should be more robust than current Canadian examples, like PIPEDA, due to the highly sensitive nature of FR data and potential for harm. It should be informed by policy, law, and technology, and be jointly developed by experts in these fields. Policymakers must determine who will be responsible for creating such a framework and who the framework will regulate. There must also be enforcement mechanisms to handle violations.

SECURITY

FR data must be strictly protected from malicious actors and possible data breaches. Both the FR provider and institution must take adequate security measures, such as encryption, central storage, protected networks, and access restrictions.

Considerations/Questions:

How will adequate security be ensured? What are potential risks?

Institutions should be subject to audits to assess the reliability of their security measures, especially as reports of data breaches in Clearview AI are already emerging.[24] Security risks also extend beyond external hacking—actors within institutions (particularly law enforcement agencies) must also be prevented from accessing FR services and data for unauthorized purposes.

ACCURACY & BIAS CONDITIONS

Studies have shown that FR has a tendency towards identity-based discrimination, particularly against people of colour.[25] FR providers must therefore demonstrate that their facial recognition systems' overall accuracies and false positive rates for each demographic (gender, age, and skin colour) meet universal government-determined baselines.

Considerations/Questions:

What testing dataset should be used?

Auditors should measure accuracy and false positive rates using standardized, diverse datasets. Auditors must also consider that, if testing is conducted using a small number of publicly-available testing datasets, providers can easily train their algorithms to perfectly match them. Alternatives include classified datasets, or the frequent construction of new datasets that are released to companies and the public only after audits are performed.

How should demographic accuracy and bias be measured?

Measuring accuracy and bias in FR systems is a technically difficult task. A precise, systematic method for calculating accuracy and bias remains elusive; further computational research is required to determine, not only the best test to administer, but also how to ensure that algorithms pass.

At its core, FR technology compares facial images and detects potential matches between pairs of photos by generating a match score for each pair. The higher the score, the more likely that both images depict the same person. A potential match is flagged if a pair's match score exceeds a given match score threshold.

There are two primary approaches to estimating the accuracy of an FR system.

Method 1

Method 1 considers all pairs of facial images whose match score exceeds a particular match score threshold, counting how many of these pairs truly do represent the same face (true positive pairs) and how many do not (false positive pairs). The false positive rate (the percentage of pairs mistakenly flagged as matches) is a crucial metric; in policing, for instance, where false matches may lead to false arrests, a very low false positive rate is needed.

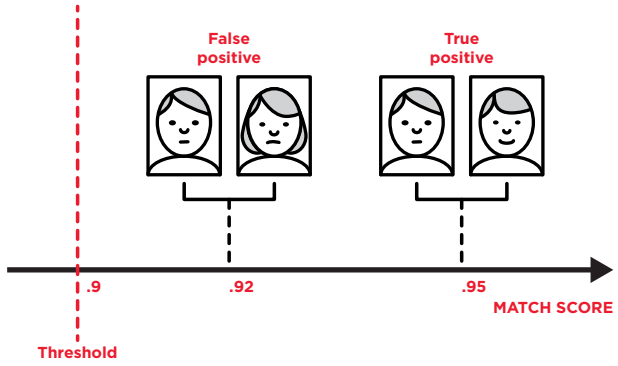

Method 1's failings

Method 1 is complicated by the fact that there is an important trade-off between an FR algorithm's false positive rate and its match score threshold. Increasing the threshold decreases the false positive rate, as doing so decreases the overall number of pairs whose match score exceeds that threshold (and by extension decreases the number of false matches). The threshold can therefore be adjusted to obtain a desired false positive rate.

To test an algorithm, there are thus two options:

1. Test each demographic against a fixed threshold.
2. Vary the threshold by demographic so that each group meets a fixed false positive rate.

Current FR systems are unlikely to satisfy the first test: a recent NSIST-affiliated paper found that a fixed threshold in Method 2 yielded a higher false positive rate for East Asians than Caucasians, mirroring similar results from other studies.[26] However, the second approach is equally flawed: instead of addressing the underlying causes of bias, it moves the goalposts so that each demographic appears equal. Moreover, this approach is not tenable for public use. If used to search through a database of unlabelled images captured by a CCTV camera, the algorithm would need to determine the appropriate demographic threshold to use for each individual photo. This would require assigning demographic information to each image, either manually or by using a potentially biased gender and race classification algorithm.[27]

Method 2 and its failings

A second method avoids this difficult trade-off by ignoring the match score threshold altogether and instead measuring the system's "general" accuracy, analyzing how all match scores generated by the system differ from each other. However, this method has been shown to fail to detect bias that the previous method identifies with ease.[28]

Both methods' failings:

Both methods are hindered by their reliance on arbitrary demographic categorizations. Race, age, and even gender cannot be neatly split into discrete classes. To measure demographic bias, however, a discrete split in the testing dataset is needed (i.e. male or female; young, middle-aged, elderly) so that the accuracy of each distinct group can be recorded. Photos of individuals who fit neatly into a particular category are more likely to be included in the testing dataset, while mixed-race or gender non-conforming individuals may be excluded. Consequently, testers may not learn how accurately the FR system performs when faced with these groups.

Moreover, bias can creep in when demographically labeling test datasets, regardless of whether labels are assigned manually or via an algorithm.

How can better methods for measuring demographic accuracy and bias be developed? How can bias in FR systems be reduced?

Developing suitable methods for measuring accuracy and bias in FR technology will require substantial research conducted by teams of computer scientists working alongside policy experts and demography researchers. But, for a moratorium to be lifted, it is not enough to possess the means by which to test bias—FR systems must be unbiased enough to pass said test. As summarized in briefing #1, the causes of—and by extension the solutions to—bias remain poorly understood. Extensive research into bias reduction must be undertaken by computer scientists and social scientists.

REVIEW & OVERSIGHT CONDITIONS

Prior to lifting the moratorium, an oversight and review body should be established. This body should be based on the EU's National Data Protection Authorities (DPAs), acting as an independent public authority tasked with investigating and correcting data protection law violations.[29] In addition to supervising the application of data protection law, this Canadian authority body could manage how privacy and data protection complaints are handled.

AUDITING

Many of the conditions outlined above require auditing to ensure compliance. Public institutions using FR technology should be audited to verify that their use cases adhere to a regulatory framework, and that their data collected via or for FR technology is correctly stored, handled, and follows adequate security measures. Companies selling FR systems to public institutions must also subject their products to audits evaluating their accuracy and bias and ensuring they are meeting established standards of data and privacy protections.

Considerations/Questions:

How will auditing be performed?

Policymakers must determine who will conduct audits of public institutions, with what frequency, and whether each condition—use cases, data usage, and security—should be audited separately. Potential auditors include a third-party, a government regulatory body, the OPC, or a regulatory body inside each institution. Similar deliberations must be held regarding the auditing of private FR providers. It must also be determined what power (e.g. fines or a moratorium on government contracts for providers, or restrictions on FR use for institutions) auditors have to hold providers and institutions accountable for malpractice.

Will the results of the audit be made available to the general public?

If the results of audits are not made available, public trust would need to be cultivated by other means.

SOCIAL CONDITIONS

Emerging global patterns show disproportionate harms of FR to (groups of) individuals detailed in briefing #1. The social impacts of FR technology should also be evaluated, including how they may affect society at large. Systems of ubiquitous public surveillance for instance, have come under increased public scrutiny amidst global protests for racial justice and civil liberties.[30] Prior to lifting a moratorium, the Canadian government should develop a tool or framework by which to evaluate the social validity of a given FR system. This might include asking whether other, less invasive and equally reliable systems exist for the intended purpose being replaced by a FR or biometric tracking system. One possible frame that may be appropriate under a moratorium could be a GAP analysis—a comparison of actual performance of existing FR systems with the desired future state. If the delta between them is too broad, a moratorium should not be lifted.

Considerations/Questions:

How could the social validity of an FR system be measured?

A social validity framework must be developed via research and collaborative consultations with those who stand to be most impacted by FR. It should directly address high-stakes public sector uses, including law enforcement, and should answer the following questions:

- ***How does the public impact of increased technological surveillance measure against the FR system's foreseen benefits?***
- ***Are there extant uses of the FR system in question, or similar systems, which have shown to cause undue harm, bias, and discrimination, especially along intersecting identity categories such as race, class, sexuality, and gender?***
- ***Can the intended purpose of the FR system be accomplished through alternative means?***

LEGAL CONDITIONS

SPECIFYING USE OF SERVICE

The Privacy Act details how federal government institutions must handle personal data, which can only be collected, used, and disclosed with the individual's consent and for limited and legitimate purposes, unless in very specific circumstances.[31]

The unique power of law enforcement organizations is subject to additional safeguarding mechanisms. For example, the collection of biometric information such as fingerprints on arrest is governed by the Identification of Criminals Act[32] and the collection of bodily substances for DNA analysis is regulated under the Criminal Code[33] and requires a warrant (additional laws such as those found in the Canadian Security and Intelligence Service Act and the Customs Act also apply depending on the context).

These provisions do not apply to the police's use of biometric digital imagery through FR, meaning that it can be collected without consent or court oversight.

CRIMINAL CODE AMENDMENT

The Government of Canada should consider amending the Criminal Code to make the use of FR technology subject to the same regulatory protections as other biometric characteristics such as fingerprints or bodily substances for DNA analysis. This would entail making the evidence acquired through FR technology inadmissible in court, unless particular standard procedures are followed for the collection and retention of such information (such as a warrant requirement).

PRIVACY ACT AMENDMENTS

The rights maintained by Canadians under the law should not be any different for data acquired using FR technology. The covert aspect of observation in the case of FR technology was found to remove the individual's ability to maintain control over how and when they are observed, as well as how the information about them is being used.[34] The ambiguity around current legal provisions' applicability to FR technology should be addressed as a priority.

NEXT STEPS: RESEARCH AND EVALUATIONS DURING A MORATORIUM

The previous sections outlined technological, social, policy, and legal conditions that must be in place prior to lifting a moratorium on FR technology. They also explored numerous considerations and questions arising from these conditions, all of which require significant research or deliberation. The following section lays out the steps the Canadian government should take during a moratorium to flesh out these conditions and address their considerations and questions. Lifting a moratorium without following these or substantially similar steps would be ill-advised, as many important conditions (such as ensuring a low degree of bias) would remain unmet, while other conditions (such as creating a data governance framework) would be difficult to define and develop without consultations and research efforts.

MULTISECTORAL PANEL

The Government of Canada should create a high-level panel made up of policy, technology, social science, and legal experts. This panel would be tasked with developing the optimal regulatory requirements for lifting a moratorium, building upon the conditions, questions, and considerations put forward in this briefing. Their mandate would also include studying the current use of FR technology in Canada and reviewing current data and privacy legislation to identify gaps.[35] This expert-led panel could spearhead broad consultation, research, and assessment projects, as described below.

PAN-CANADIAN CONSULTATION

The federal government should conduct large-scale consultations to assess the perspectives of Canadians—particularly those in marginalized communities—on FR technology in public use such as law enforcement. Stakeholders involved in these consultations should include: advocacy groups, national and international civil liberty associations, federal and provincial privacy Commissioners, relevant federal agencies and committees' representatives, as well as municipal, provincial and federal police forces.[36] This approach, which is in line with the federal government's policy process on technology and privacy, would enable policymakers to prioritize human rights protections.[37]

Coordination of the consultation by the federal government ensures consistency in the information provided to different levels of government. The primary output of these consultations should be an assessment of the desirability of FR technology use by law enforcement, as opposed to proposals to amend privacy and data protection laws.

COORDINATED INTERDISCIPLINARY RESEARCH COMMITTEE

Following consultations, the federal government should leverage the collected insight to coordinate a national research effort on the use of FR technology by the public sector, focusing on the impact of FR technology use on racialized communities. This research effort will serve to increase transparency around the current use of FR technology by law enforcement, addressing the grave lack of data which currently remains a substantial impediment to policymaking and regulation on the matter.

A series of reports, each framing its investigations through a human rights-based approach, could be jointly commissioned by the OPC, the provincial Privacy Commissioners, and the National Research Council. The most pressing of these would be a comprehensive report detailing which Canadian agencies, at the federal, provincial/territorial, and municipal levels, have used or tested FR systems and in what capacity, including an analysis of the potential societal risks and benefits of this use. This comprehensive report should indicate the outcome and lessons learned from the OPC's investigation in the RCMP's use of FR technology. A second report should examine all FR systems sold in Canada to test whether the increased awareness of FR's risks and biases has led to substantial improvements in the technology.

Getting access to detailed and disaggregated data on current FR technology shortcomings is imperative in getting buy-in from key stakeholders including federal and municipal law enforcement agencies on the need for increased regulations. Indeed, without sufficient information around the current use of FR technology, it is difficult to convince law enforcement agencies of the need for reform. These reports could also help position Canada as a leader in filling the knowledge gap around the technology's ethical and human rights implications in law enforcement.

Federal funding should also be provided to support universities and other research institutes to conduct more specific research, particularly on the topics of bias, its underlying causes, and the proposed solutions for combating it. Upon receiving any funds, explicit requirements for sharing insights, code, and techniques with an advisory committee should be established. As outlined in Public Safety's 2019-20 departmental plan, new requirements including but not limited to "hold[ing] discussions with provinces and territories to identify facilities that require immediate rehabilitation"[38] were developed upon receiving funds. The federal government should also consider providing financial support to provinces in their endeavor to assess their police force's use of the technology.

PRIVACY AND DATA PROTECTION IMPACT ASSESSMENTS

Privacy Impact Assessments (PIA) by the OPC review current uses of FR technology to ensure they respect federal privacy laws as outlined in the Privacy Act. Since 2004, the OPC has conducted PIAs for Passport Canada's Facial Recognition Project, detailing recommendations to mitigate the privacy risks of federal programs.[39] Prior to lifting a moratorium, PIAs on each relevant government institution

should be carried out, focusing on assessing justifiable need and consistent use (see Purpose conditions), access and retention and security (see Data usage conditions), and accuracy and bias. Similar to PIAs, Data Protection Impact Assessments (DPIAs) assess data-related risks and verify that data protection laws are adhered to[40]—DPIAs should likewise be conducted.

CONCLUSION

Success in navigating a moratorium is contingent on public trust—transparency about governments' decision(s) to enact a moratorium, to evaluate FR systems, and to implement conditions for its release must therefore be prioritized at every stage. All research, consultations, and assessments conducted during the moratorium should be made public, and their results should be communicated to citizens in a broad-reaching and accessible manner. The Government of Canada should likewise support efforts to improve digital literacy, equipping Canadians with knowledge of their own digital rights and responsibilities. We urge Canada to model global leadership in developing and enforcing specific technological, social, policy, and legal conditions which must all be met for any future moratoria on FR technology to be lifted.

REFERENCES

1 The Canadian Press. "NDP Calls for Moratorium on Clearview AI Facial Recognition Software." National Post, March 9, 2020, https://nationalpost.com/pmn/news-pmn/canada-news-pmn/ndp-calls-for-moratorium-on-clearview-ai-facial-recognition-software.

2 Taylor Owen and Nasma Ahmed. "Opinion: Let's Face the Facts: To Ensure Our Digital Rights, We Must Hit Pause on Facial-Recognition Technology." The Globe and Mail, February 14, 2020, https://www.theglobeandmail.com/opinion/article-lets-face-the-facts-to-ensure-our-digital-rights-we-must-hit-pause/.

3 Jay Greene. "Microsoft won't sell police its facial-recognition technology, following similar moves by Amazon and IBM." The Washington Post, June 11, 2020, https://www.washingtonpost.com/technology/2020/06/11/microsoft-facial-recognition/.

4 Mark Montgomery. "National Police to limit, but not stop use of facial recognition technology." RadioCanada International. March 10, 2020, https://www.rcinet.ca/en/2020/03/10/national-police-to-limit-but-not-stop-use-of-facial-recognition-technology/.

5 Office of the Privacy Commissioner of Canada. "Clearview AI ceases offering its facial recognition technology in Canada." July 6, 2020, https://www.priv.gc.ca/en/opc-news/news-and-announcements/2020/nr-c_200706/.

6 Kashmir Hill. "Wrongfully Accused by an Algorithm." The New York Times, June 24, 2020, https://www.nytimes.com/2020/06/24/technology/facial-recognition-arrest.html.

7 Bobby Allyn. "'The Computer Got It Wrong': How Facial Recognition Led To False Arrest Of Black Man." NPR, June 24, 2020, https://www.npr.org/2020/06/24/882683463/the-computer-got-it-wrong-how-facial-recognition-led-to-a-false-arrest-in-michig.

8 National Institute of Standards and Technology, Face Recognition Vendor Test (FRVT) Part 3: Demographic Effects, by Patrick Grother, Mei Ngan, and Kayee Hanaoka, Rep. 8280, US Department of Commerce, 2019, https://nvlpubs.nist.gov/nistpubs/ir/2019/NIST.IR.8280.pdf (accessed July 8, 2020).

9 Office of the Privacy Commissioner of Canada. "PIPEDA Fair information principles." May 2019, https://www.priv.gc.ca/en/privacy-topics/privacy-laws-in-canada/the-personal-information-protection-and-electronic-documents-act-pipeda/p_principle/.

10 Amba Kak and Rashia Richardson. "The Office of the Privacy Commissioner of Canada Consultation: Proposals for ensuring appropriate regulation of artificial intelligence." p. 11. AINow. March 12, 2020, https://ainowinstitute.org/ainow-comments-to-canadian-office-of-the-privacy-commissioner.pdf.

11 Ibid

12 Office of the Privacy Commissioner of Canada. "Consent." September 10, 2019, https://www.priv.gc.ca/en/privacy-topics/collecting-personal-information/consent/.

13 Heide Pearson. "Calgary mall defends use of facial-recognition technology after customer discovers they're being watched" Global News, July 28, 2018, https://globalnews.ca/news/4355444/chinook-mall-calgary-facial-recognition-technology/.

14 Kate O'Flaherty. "Clearview AI's Database Has Amassed 3 Billion Photos. This Is How If You Want Yours Deleted, You Have To Opt Out." Forbes, January 26, 2020, https://www.forbes.com/sites/kateoflahertyuk/2020/01/26/clearview-ais-database-has-amassed-3-billion-photos-this-is-how-if-you-want-yours-deleted-you-have-to-opt-out/#5665f59660aa.

15 Ben Wolford. "Everything you need to know about the 'Right to be forgotten'" GDPR.eu, accessed July 16, 2020, https://gdpr.eu/right-to-be-forgotten/.

16 Thomas Daigle. "Canadians can now opt out of Clearview AI facial recognition, with a catch." CBC News, July 10, 2020. https://www.cbc.ca/news/technology/clearview-ai-canadians-can-opt-out-1.5645089.

17 SAS. "The SAS Data Governance Framework: A Blueprint for Success." p. 5 2018. https://www.sas.com/content/dam/SAS/en_us/doc/whitepaper1/sas-data-governance-framework-107325.pdf.

18 World Economic Forum. "A Framework for Responsible Limits on Facial Recognition" February 2020. http://www3.weforum.org/docs/WEF_Framework_for_action_Facial_recognition_2020.pdf.

19 Theodore Christakis. "First Ever Decision of a French Court Applying GDPR to Facial Recognition." AI-Regulation, February 27, 2020, https://ai-regulation.com/first-decision-ever-of-a-french-court-applying-gdpr-to-facial-recognition/#:~:text=First%20Ever%20Decision%20of%20a%20French%20Court%20Applying%20GDPR%20to%20Facial%20Recognition&text=A%20French%20court%20canceled%20today,that%20this%20would%20be%20illegal.

20 AI Now. "The Office of the Privacy Commissioner of Canada Consultation: Proposals for ensuring appropriate regulation of artificial intelligence" March 2020, https://ainowinstitute.org/ainow-comments-to-canadian-office-of-the-privacy-commissioner.pdf.

21 Ibid.

22 Ibid.

23 US Congress, House, Algorithmic Accountability Act of 2019, HR 2231, 116th Cong., introduced in House April 10, 2019, https://www.congress.gov/bill/116th-congress/house-bill/2231/all-info.

24 Alfred Ng. "Clearview AI's entire client list stolen in data breach." CNET, February 26, 2020, https://www.cnet.com/news/clearview-ai-had-entire-client-list-stolen-in-data-breach/.

25 National Institute of Standards and Technology, Face Recognition Vendor Test (FRVT) Part 3: Demographic Effects, by Patrick Grother, Mei Ngan, and Kayee Hanaoka, Rep. 8280, US Department of Commerce, 2019, https://nvlpubs.nist.gov/nistpubs/ir/2019/NIST.IR.8280.pdf (accessed July 8, 2020).

26 Jacqueline Cavazos, P. Jonathon Phillips, Carlos D. Castillo, and Alice J. O'Toole. "Accuracy comparison across face recognition algorithms: Where are we on measuring race bias?" NSIST, June 4, 2020, https://arxiv.org/pdf/1912.07398.pdf.

27 Joy Buolamwini and Timnit Gebru. "Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification," (paper presented at 1st Conference on Fairness, Accountability and Transparency, Proceedings of Machine Learning Research), 81:77-91, 2018. https://dam-prod.media.mit.edu/x/2018/02/06/Gender%20Shades%20Intersectional%20Accuracy%20Disparities.pdf.

28 Ibid.

29 European Commission. "What are Data Protection Authorities (DPAs)?" https://ec.europa.eu/info/law/law-topic/data-protection/reform/what-are-data-protection-authorities-dpas_en.

30 Tawana Petty. "Defending Black Lives Means Banning Facial Recognition." Wired, July 10, 2020, https://www.wired.com/story/defending-black-lives-means-banning-facial-recognition/.

31 Privacy Act, Revised Statutes of Canada 1985, c. P-21. https://www.canlii.org/en/ca/laws/stat/rsc-1985-c-p-21/161168/rsc-1985-c-p-21.html#sec2.

32 Identification of Criminals Act, Revised Statutes of Canada 1985, c. I-1. https://www.canlii.org/en/ca/laws/stat/rsc-1985-c-i-1/161284/rsc-1985-c-i-1.html#sec2.

33 Criminal Code, Revised Statutes of Canada 1985, c. 46. https://www.canlii.org/en/ca/laws/stat/rsc-1985-c-c-46/161288/rsc-1985-c-c-46.html#sec487.05.

34 Office of the Information & Privacy Commissioner for British Columbia, Investigation into the Use of Facial Recognition Technology by the Insurance Corporation of British Columbia. Elizabeth Denham. February 16, 2012. https://www.oipc.bc.ca/investigation-reports/1245.

35 Taylor Owen and Nasma Ahmed. "Let's face the facts: To ensure our digital rights, we must hit pause on facial-recognition technology." The Globe and Mail, February 14, 2020, https://www.theglobeandmail.com/opinion/article-lets-face-the-facts-to-ensure-our-digital-rights-we-must-hit-pause/.

36 Nani Jansen Reventlow. "How Amazon's Moratorium on Facial Recognition Tech Is Different From IBM's and Microsoft's". Slate, June 11, 2020, https://slate.com/technology/2020/06/ibm-microsoft-amazon-facial-recognition-technology.html.

37 "Canada's Digital Charter: Trust in a digital world," Government of Canada, June 8, 2020, https://www.ic.gc.ca/eic/.site/062.nsf/eng/h_00108.html.

38 Public Safety. Public Safety Canada Departmental Plan 2019–20. Canada: Her Majesty the Queen in Right of Canada, 2019. Accessed August 10, 2020. https://www.publicsafety.gc.ca/cnt/rsrcs/pblctns/dprtmntl-pln-2019-20/index-en.aspx.

39 Office of the Privacy Commissioner of Canada, Automated Facial Recognition in the Public and Private Sectors, Gatineau, QC, 2014, https://www.priv.gc.ca/media/1765/fr_201303_e.pdf.

40 Amba Kak and Rashida Richardson. "Artificial Intelligence Policies Must Focus on Impact and Accountability." Centre for International Governance Innovation, May 1, 2020, https://www.cigionline.org/articles/artificial-intelligence-policies-must-focus-impact-and-accountability.