TOWARDS A PROPORTIONATE AND RISK-BASED APPROACH TO FEDERATED DATA ACCESS IN CANADA

TANIA BUBELA
REGIANE GARCIA
IVAN BESCHASTNIKH
ALINE TALHOUK

JULY 2023

CIFAR

AI Insights

ABOUT THE AUTHORS

TANIA BUBELA

BSc (Hons), PhD, JD, FCAHS, FRSC. Professor & Dean, Faculty of Health Sciences, Simon Fraser University.

IVAN BESCHASTNIKH

BSc, MSc, PhD. Associate Professor, Department of Computer Science, University of British Columbia.

REGIANE GARCIA

LLB, LLM, PhD. Research Associate, Faculty of Health Sciences, Simon Fraser University.

ALINE TALHOUK

BA, MSc, PhD. Assistant Professor and Michael Smith Health Research BC Scholar, Department of Obstetrics & Gynecology, Faculty of Medicine, University of British Columbia and British Columbia's Gynecological Cancer Research Program (OVCARE).

ACKNOWLEDGMENTS

The authors wish to thank the 47 experts in ethics, privacy, data governance and security who took time from their busy schedules to be interviewed for this work and those who attended the validation workshop and completed the online survey. We also wish to thank the research assistants who have assisted with interviews, literature reviews and/or preparation of this policy brief: Elise Abi Khalil [University of British Columbia (UBC)], Mishaal Kazmi (UBC), Kalli Leung (UBC), Matheus Stolet (UBC) and Howard Zhang (SFU).

LAND ACKNOWLEDGMENT

We wish to acknowledge this land on which CIFAR operates. For thousands of years it has been the traditional territory of many nations including the Mississaugas of the Credit, the Anishnabeg, the Chippewa, the Haudenosaunee and the Wendat peoples and is now home to many diverse First Nations, Inuit and Métis peoples. We are grateful to have the opportunity to work on this land.
We also acknowledge we are all responsible for reconciliation. CIFAR's AI & Society program seeks to advance our understanding of the societal implications of AI to design a future of responsible AI. A future of responsible AI includes one that centres the concerns of Indigenous communities. CIFAR is committed to prioritizing Indigenous perspectives in the development and design of responsible AI.

TABLE OF CONTENTS

Correspondence to
Aline Talhouk: a.talhouk@ubc.ca

EXECUTIVE SUMMARY

Artificial intelligence (AI) algorithms that employ machine learning (ML) are in development for a range of health applications but require massive amounts of data to learn hidden patterns. Such data are often siloed across multiple sites and jurisdictions and can be challenging to pool, curate, and access due to concerns over ethics, privacy, and security. Federated learning (FL) is an emerging type of ML that allows multiple parties to collaborate on model training without sharing their data. As such, FL can alleviate some of the privacy, security and ethical concerns typically associated with pooling data for learning.

In this policy brief, we explore how FL may be implemented. Our discussion of the twelve technical and ethical-socio-legal challenges and policy options for FL consortia derives from the synthesis of analyses of four data sources: a document and literature review, expert interviews, a validation workshop, and a survey of solutions to privacy, ethics and security challenges raised by FL. We provide policy options to address each challenge.

Our hope is first to build understanding for policy makers of the technical intricacies and societal impacts of FL, which may be enhanced through regular discussions with technologists, ethicists, legal experts, and other stakeholders. Engaging with the public and affected communities will be essential to build trust, gain valuable input, and ensure that the benefits of FL are distributed equitably. We contend that the implementation of the outlined policy options should be proportionate to the realized risks and potential benefits of FL applications. Policymakers should facilitate the development of an ecosystem that encourages innovation in FL while ensuring the protection of individuals' privacy and the promotion of social good. This involves adopting a risk-based approach to its regulation and governance, where only higher-risk applications are subject to more stringent scrutiny and oversight. Policy makers will need to consider mechanisms to foster collaborations between public and private sectors, promote technical and policy research, and provide training and dialogue opportunities for all stakeholders. Collaboration, both nationally and internationally, will be crucial in sharing expertise, mitigating risks, and formulating common standards. Finally, it is important to note that FL is a rapidly evolving field, and as such, policies and regulatory frameworks will need to be regularly reassessed and updated to keep pace with technological advancements and emerging challenges.

KEY ISSUES OF CONCERN FOR POLICY MAKERS

In this policy brief, we discuss 8 technical and ethical-socio-legal challenges and policy options for FL consortia. We first discuss ethics, privacy and data governance, followed by security challenges.

ETHICS, PRIVACY AND DATA GOVERNANCE

1 PRIVACY CONCERNS WITH FEDERATED LEARNING

Health data are protected by a tapestry of privacy legislation, and institutional or organizational policies and practices that reflect the fiduciary duties of data custodians to data subjects. FL could alleviate privacy concerns, because only models are shared, and data do not leave secure local environments. Policy options include: the development of template agreements for FL; the implementation of complaints, breach and audit processes; the evaluation of privacy risks that is proportionate potential harm based on principled privacy metrics; the update of the legislative environment; and the consideration of insurance or no-fault compensation models for breaches.

2 APPROVAL PROCESSES FOR ETHICS, PRIVACY, AND ACCESS

As a new and potentially disruptive technology, FL researchers may lack understanding of laws, policies and practices to represent the risks of their technologies under development. Similarly, decision-makers for ethics, privacy, and operational approvals may lack the technical competence and processes to proportionately evaluate risks and mitigate potential harms. Policy options include to: harmonize approval processes; implement training and education for decision-makers on FL and researchers on processes; facilitate access to technical expertise by decision-makers; and pre-vet certifications and technological solutions.

3 ENGAGE PATIENTS, EMPOWER COMMUNITIES AND ESTABLISH/ MAINTAIN TRUST

A key question for federated models is whether they can enable data sovereignty, community empowerment and community governance, while mitigating harms? Because data do not leave their local environments, opportunities arise to be more inclusive of patients and publics in the governance of data and their use. Potential policy options for public engagement are to encourage patients and communities to participate in governance of FL networks and consortia and engage with them about FL, its applications and its risks and benefits. FL may align with initiatives to enhance data sovereignty for Indigenous communities and organizations but would require engagement with and capacity building for Indigenous communities and organizations.

4 CONSENT FOR FEDERATED LEARNING

FL involves using various types of de-identified data with different levels of risk for re-identification. Consent is a preferred option to enable the use and sharing of data interprovincially or internationally because it provides opportunities to withdraw consent. Consent is an ongoing process, not a one-time transaction and can establish relationships and ongoing trust. We discuss various models of consent in order of their applicability: consent to governance, broad consent, dynamic consent, community consent, opt-out of consent, waiver of consent, specific consent, historic consent, and no consent.

5 GOVERNANCE OF DATA INFRASTRUCTURE FOR FL

FL introduces governance challenges specific to the distributed nature of the data, but it has the potential to improve others. Federated networks empower each participating site by giving them full control of their data and the ability to revoke access at any time. Federated networks can also facilitate international collaboration, overcoming some legislative and policy restrictions of data leaving Canada, but only if users from outside a jurisdiction are not restricted from accessing local secure data environments. Governance challenges discussed are roles and responsibilities for data controls, data privacy, data linkage, model ownership, and model validation. Potential policy options include to: be transparent about data handling with accountability processes in place; harmonize data control, access, and use agreements as well as policies and procedures across all sites; develop agreements for potential commercial use or partners; and pre-determine each site's contributions.

6 SPECTRUM OF DE-IDENTIFICATION/IDENTIFIABILITY

A challenge that is unique to FL is site-level re-identification. This refers to the possibility of determining whether a particular site has participated in the training of a model, as well as the possibility of identifying specific individuals whose data were included at a given site. Site-level re-identification can potentially compromise the privacy and confidentiality of individuals involved in the training process. Potential policy options are to: establish best practice guidance on de-identification of data and data classification and handling standards based on de-identification risk; develop and use terms in contractual/licensing/collaboration-consortium agreements that prohibit use for re-identification; harmonize de-identification practices across centers for proportionate risk; legislate prohibitions with penalties; define metrics for de-identification for data/models; and require minimum cohort size for reporting results.

SECURITY CHALLENGES FOR FL

1 DATA BLINDNESS

In FL, the central server and participating sites have limited visibility into the training data used, which can make it difficult to diagnose and fix issues related to data quality or bias. Collaborating and harmonizing unique datasets and data structures from different sites into a single model is also challenging, and there is an increased risk of bad data entering the system. Potential policy options are to: harmonize to a common data model/standard through governance standards and contractual terms; collaborate with trusted partners; implement local and global quality controls; audit models regularly; generate federated data summaries and visualizations; share meta-data; and share differentially private, synthetic data.

2 CYBERSECURITY

A security breach occurs when unauthorized individuals gain access to a system or data without permission, due to hacking, malware, phishing, or any other malicious activity. Security breaches may lead to privacy breaches. Cybersecurity threats are real and can have significant financial and reputational consequences for organizations. FL security breaches can be categorized into network security, access control, and model security. Policy options include to: enhance network security; restrict unauthorized access to the model and use technologies to control users; identify any weak links; create secure research environments; monitor and test security, regularly; and use encryption methods.

CONCLUSION

While security, privacy, and research ethics threats will inevitably require some technical developments, we argue that a proportioned, governance-based approach to federated learning systems should prevail.

1.0 INTRODUCTION

Artificial intelligence (AI) and machine learning (ML) algorithms are in development for a range of health applications. AI enables precision medicine, which uses data to predict the best treatment for individual patients based on their genomic or molecular profile. AI is also gaining currency in the recognition of images, video, audio, and text, both for real-time diagnosis and the prediction of the likelihood of disease onset within a given time-period. Diagnostic applications of ML could improve reproducibility for human evaluation of images, for example, as aids to a pathologist's characterization of medical imaging. Applications under development include detection of damage to the retina caused by diabetes from photographs[1]; and clinical-grade pathology support for various cancers from slides of tumours.[2,3,4,5,6,7]

The clinical potential of ML depends on access to the massive amounts of data collected by health systems. Such data are often siloed across multiple sites and jurisdictions and can be challenging to pool, curate, and access due to concerns over ethics, privacy, and security. Federated learning (FL) is a promising approach to governance for data. Its use to train models that may alleviate some of these concerns, because FL enables multiple parties to collaborate on model training without sharing their data.[8] Rather than pooling data in a central repository, an FL algorithm processes local data to train a global model across a federated network of research centres.

Conversations about data federation are central to Canadian aspirations to accelerate improvements in healthcare, health system performance and population health across the continuum of care. While Canada aims to modernize its health care systems with standardized health data and digital tools,[9] it is timely to discuss access to those data for research and innovation to improve the health of Canadians. Federating data sources is an important step towards maximizing the value of data, not only for clinical, but also for research purposes. This will require the development of global policy frameworks for the collection and sharing of data, data oversight and governance in federated systems.

In this policy brief, we explore how FL may be implemented. We discuss findings from our document review, expert interviews, a validation workshop, and a survey of solutions to privacy, ethics and security challenges raised by FL. In evaluating solutions to potential challenges, we focus on a proportionate response to realized risks, specifically the frequency and magnitude of harm caused by ethical, privacy, and security breaches of health data. We discuss the trade-offs between protections and the utility of data for FL and recommend enabling governance models.

1.1

GOVERNANCE OF POOLED VERSUS FEDERATED MODELS

Currently, most ML algorithms are developed using data pooled in a central location from distributed sites (Figure 1a). These data are used to train a global model, which may then be shared with collaborators for validation on other datasets.[10] The governance is centralized and hierarchical; the institution/organization that hosts the central pool sets the terms for input, access, and use of the data, because it bears the risks associated with safeguarding the data from privacy and security threats. Data pooling is necessary because it increases the power of ML. However, once shared, access to data is difficult to revoke. Data pooling can be challenging from a legal perspective.[11] It requires data sharing agreements and compliance with privacy laws and consents, and from a security perspective, it requires anonymization without loss of data fidelity, access control and transfer-safety. Commercial or commercialization interests may also limit willingness to pool data, especially if value has been added to data through aggregation or curation.

Alternatives to data pooling do not require data to leave their sites of origin. In their simplest form, peer-to-peer model sharing (Figure 1b), collaborators share their partially trained models developed using only local data. These models are then shared with peers for evaluation using local data at a different site. Peer-to-peer model sharing, therefore, has a distributed governance structure, with agreements for model sharing and use, including evaluation and improvement. This model is efficient as no copies of the data are made, and analyses can be run at sites asynchronously. Depending on scale, peer-to-peer model sharing may be simple to set-up, but the lack of a central coordinator may result in poor model accuracy and generalizability. Models can still leak private information and all sites require computational capacity.

In contrast, FL is reliant on one institution/organization to coordinate the generation of a global model (Figure 1c). FL models can perform as well as ones trained on either a pooled dataset or smaller-scale peer-to-peer collaborations.[12] The collaborative training process of the global model is iterative; it is developed from partially trained, locally-developed models, which are returned to the central FL coordinator intermittently for aggregation. Only model characteristics, such as parameters and gradients, are shared. Training of the global model continues until the model converges on data across the sites. The data never leave their firewall-protected institutions, and no copies are made. However, coordination requires a hierarchical governance structure for the consensus-based process of developing the global model among networked sites. FL requires computational capacity at all sites, models can leak private information and iterative training may take longer. Legal agreements may be difficult to negotiate while scope of research and development remain uncertain for this new mode of technology development.

In summary, FL has benefits over pooled data for ML.[13] The primary benefit is that data are not transferred or shared. FL retains more control for institutions in revoking access to data, which is difficult once data are shared. FL is more secure and protective of privacy, as it is harder for bad actors to attack a distributed system, but information leakage can never be fully prevented. Shared models may indirectly expose private health or other information through reverse-engineering processes by adversaries. Federated networks require consortium agreements among all participating sites that are more difficult to negotiate than peer-to-peer data sharing agreements.

2.0

OUR ANALYTICAL APPROACH

We combined three data collection approaches to analyze the technical and ethical-socio-legal challenges for FL consortia and the potential best practices for their governance:[i]

1. Forty-three key-informant interviews with experts in research ethics, privacy, network security and data governance across Canada and the Pacific North-West of the United States (US).[ii]
2. A document review of laws, policies, legal cases, and literature
3. A virtual validation workshop with 19 participants and an online survey open to all interviewees.

Our analysis relies on experts and literature, not public deliberations, but strengthening public engagement is a clear recommendation. Our engagement with Indigenous researchers or representatives of Indigenous health, research and data-related organizations was minimal. As settler researchers, we note the significant issues of Indigenous data sovereignty, but it is not our place to make recommendations in this domain.

FIGURE 1A: DATA POOLING (CURRENT STATE OF PRACTICE)

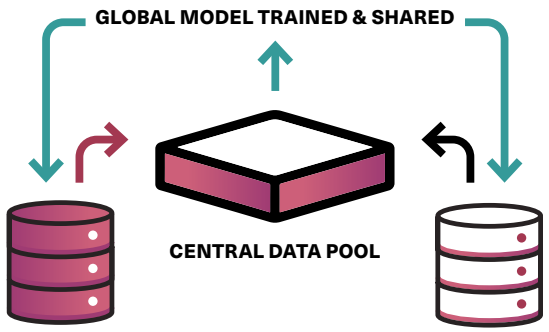

1. Pool data from sites.
2. Access all data from the pooled site.
3. Train a single model on pooled data.
4. Share model with collaborators.
   - Data leave sites of origin.
   - Hierarchical governance.

FIGURE 1B: PEER-TO-PEER MODEL SHARING

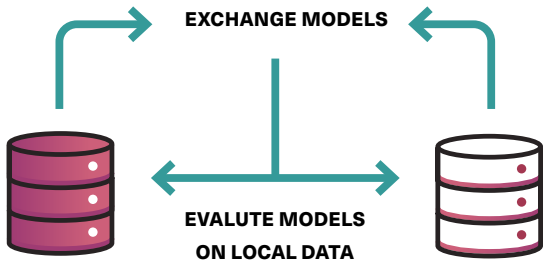

1. Train local models on local data.
2. Exchange trained models with collaborators.
3. Each site uses their local data to evaluate the trained models.
   - Data never leave local sites.
   - Distributed governance.

FIGURE 1C: FEDERATED LEARNING – AGGREGATION SERVER

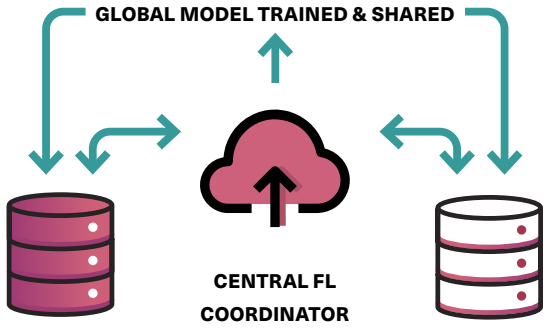

1. Select a central coordinator for FL (e.g., in the cloud) that builds a global model, which it distributes to training sites.
2. Federation of training sites submit their partially trained models back to central FL coordinator intermittently for aggregation.
3. Continue iterative training of the global model returned from the central server.
   - Data never leave local sites.
   - Network/consortium governance - the coordinating site bearing the greatest risk and should lead the agreements and their negotiations.

Note that there may be more than two data sites.

FIGURE 1

Three models for training of ML algorithms: (a) data pooling, where data sharing agreements are negotiated with the central data pool; (b) peer-to-peer model sharing, where model sharing agreements are negotiated between sites; and (c) Federated learning, which operates best under a network/consortium agreements between all sites, with the Central FL Coordinator taking the lead as the site which bears the greatest risk. Blue arrows represent model training. Black & red arrows represent data sharing. Silos represent local sites that hold data.

i See Supplemental Methods for a full description of our analytical approach.

ii We have edited quotes for length and grammar.

3.0

POLICY DISCUSSION

Numerous approaches to privacy-preserving ML/FL are in development, but all require a trade-off with utility (Figure 2). In addition to governance, technical solutions include use of synthetic data, training on encrypted data, differential privacy frameworks, or combinations of these.[14,15,16] Synthetic data are artificially model-generated data that imitate the statistical properties of real-world data.[17] They allow for the analysis and manipulation of data without compromising the privacy of individuals.

FIGURE 2

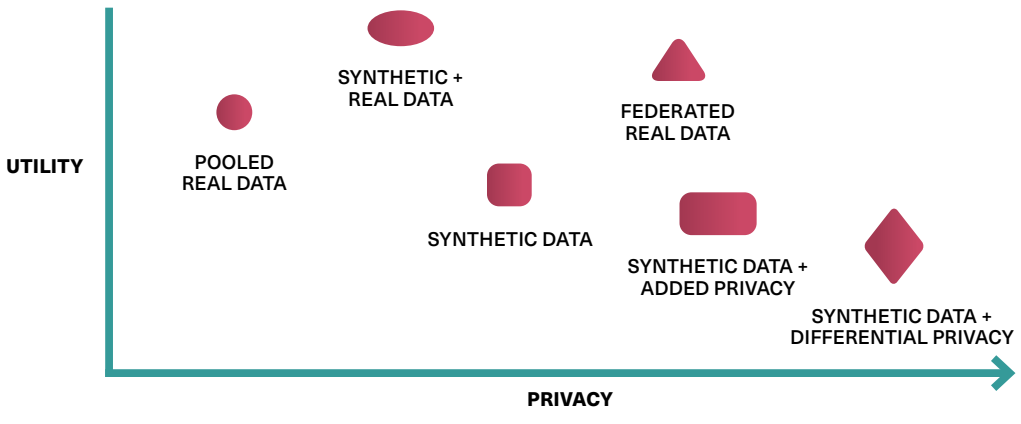

Privacy and utility trade-offs for different types of data, governance models and privacy-protecting approaches.

Differential privacy is a mathematical framework that enables the analysis of sensitive data while providing a calculable privacy guarantee (Figure 3). It works by adding noise to the results of a computation so that even if an attacker gains access to the results, it cannot determine which specific individuals were included in the original dataset. While the additional noise in differential privacy provides greater privacy, it may compromise the accuracy of the model. These privacy-protecting approaches may be difficult to scale in FL as they incur communication and computation costs at each iteration.

FIGURE 3

MORE NOISE = GREATER PRIVACY

Differential privacy operates by adding tactical noise to create plausible deniability, however, it may compromise the accuracy of the model.

Our analysis considers a proportionate response in accounting for these trade-offs between protection mechanisms and utility. We first discuss ethics, privacy, and data governance, followed by security challenges that were identified by our analysis.

3.1

ETHICS, PRIVACY, AND DATA GOVERNANCE

3.1.1 ETHICS, PRIVACY, AND DATA GOVERNANCE

Health data are protected by a tapestry of privacy legislation, and institutional or organizational policies and practices that reflect the fiduciary duties of data custodians to data subjects.[18] FL could alleviate privacy concerns, because data are not shared, and data do not leave secure local environments.[19] Only the models are shared.

Privacy legislation in its current form was not designed for FL or other big data health initiatives. Projects involving FL need to address the same issues as all projects that involve private information, and Research Ethics Boards (REBs)/Institutional Review Boards (IRBs) and other decision-makers need to apply the same legal standard for protecting health information. Those standards would suggest that, to protect individuals, data should be minimal, not comprehensive, but the big data approach for AI/ML/FL is collect all data possible. De-identifying data, discussed below, then becomes the gold standard.

While privacy and ethics reviews should be proportionate to the potential harms, some harms may be difficult to measure. Some harms are simple, for example identity theft leading to financial harm, but others are more intangible. Privacy breaches may lead to loss of trust of citizens in the system and researchers, which harm is difficult to measure, and some fields of research, such as genomics have paid attention to privacy risks related to genetic information.[20] These principles can equally apply to other forms of data sharing and analyses, such as imaging data.

Some privacy concerns are specific to FL. First, shared models may contain identifiable information, such as PHI, which may be considered a privacy breach. Second, while synthetic data - artificially model-generated data that imitates the statistical properties of real-world data - may be an option to enhance privacy, methods that work best probably have the least privacy controls in the generation of the applicable synthetic data. This raises the potential for leakage of real patient data into synthetic data sets, and the ability to infer missing data in real patient data from the synthetic data using the similarity of the synthetic data patients to real patients.

POLICY OPTIONS TO PROTECT PRIVACY WHILE ENABLING FL

Develop template data access agreements for FL. These templates could facilitate timely use of and access to local data. They should encompass non-disclosure, confidentiality, and prohibit re-identification. Templates could accelerate negotiations and conclusion of consortium/network agreements and will need to be international in scope.

Develop patient/public complaints processes to an identified official and a breach-reporting and adjudication mechanism as legislated. Complaints related to privacy breaches may be managed differently in different organizations/institutions, but there should be a clear process for reporting privacy breaches and following applicable privacy legislation for notifying affected individuals.

Adapt privacy impact assessment (PIA) to FL in a way that is proportionate to the risk. This would empower a shift away from a culture of protectionism to a culture of data stewardship, thus making data stewards accountable for the end-to-end data life cycle, including use and value.

Develop a principles-based approach for privacy metrics, especially for international programs of work with different privacy legislation. Privacy concerns may be less if data are de-identified. This may include differential privacy principles,[21,22] or synthetic data generation. Differential privacy depends on the granularity of the data and on the nature of the question that's being asked. However, REB/IRBs and privacy decision-makers are not yet familiar with these methods. The alternative is to build in the privacy protection through data security, and federated models may better protect privacy.

Update legislation/policies for ML/FL and artificial intelligence applications, possibly with omnibus legislation relevant to research use and legislative penalties for breaches. Privacy legislation in its current form was not designed for FL or other big data health initiatives. Some provincial legislation is reformed to consider the interplay between federal and provincial privacy legislation and the ability to share data across jurisdictional boundaries, but these reform processes are slow.

Perform regular privacy audits to conform with institutional/organizational privacy and security frameworks combined with an incidence management and reporting framework. Privacy audits are a best practice and need to trace where researchers have taken data when moving between institutions. These are augmented by regular reviews by provincial privacy commissioners and mandatory reporting to the privacy commissioner about non-trivial privacy breaches. Privacy Commissioners may provide guidance on how to disclose and mitigate any breach and on the development of Privacy Impact Assessments (PIA).

Deploy insurance or no-fault compensation models for breaches. High-profile privacy breaches are leading to class action lawsuits. The Ontario Court of Appeal ruled in 2015 that the *Personal Health Information Protection Act*[iii] was not an exhaustive code for remedial action on privacy breaches and that plaintiffs could bring a common law action for damages for the common law tort of intrusion, which includes invasion of privacy rights in relation to patient records.[iv] The tort does not require proof of harm to the plaintiff's economic interests, nor that the personal information has been published or disseminated by the defendant. It focuses on the act of intrusion through intentional or reckless conduct, without lawful justification and that a reasonable person would regard the invasion as highly offensive, causing distress, humiliation, or anguish.[v] Potential settlements are in the range of $10.00 to $15,000.00 per person. An alternative to litigation, however, is to establish a tiered no-fault compensation model for privacy breaches or for research institutions/companies developing FL models to insure against privacy breaches.

3.1.2 APPROVAL PROCESSES FOR ETHICS, PRIVACY, AND ACCESS

FL requires collaboration across sites and/or jurisdictions. As a new and potentially disruptive technology, FL researchers may lack understanding of laws, policies and practices to represent the risks of their technologies under development. Similarly, decision-makers for ethics, privacy, and operational approvals may lack the technical competence and processes to proportionately evaluate risks and mitigate potential harms.

POLICY OPTIONS TO ENHANCE APPROVAL PROCESSES

Harmonize Approval Processes.
Harmonization of approval processes is key to addressing delays in access and use of data that include health information.[23] The issue of harmonization is not unique to FL. However, federated structures amplify the need for harmonization of approval processes and consideration of legal-interoperability across institutions or jurisdictions. Approval processes include ethics approvals by

[iii] Personal Health Information Protection Act, 2004, SO 2004, c 3, Sch A
[iv] Hopkins v. Kay (2015) 124 O.R. (3d) 481.
[v] Jones v. Tsige (2012) 108 O.R. (3d) 241.

REBs/IRBs, PIAs by data stewards and any operational approvals for use of health authority resources. Advances have been made in harmonization of ethics approvals wherein one REB is the REB of record with responsibility for the review according to the Tri-Council Policy Statement: Ethical Conduct for Research Involving Humans - TCPS-2 (2018) guidelines. Local REBs acknowledge and may advise the REB of record to reconsider its decision in-light of local circumstances or unaddressed substantive issues.

While progress is being made towards ethics harmonization, review for privacy and security issues is more problematic. Privacy acts have been interpreted differently by different organizations and work in clinical contexts may additionally require operational approvals. Streamlining and harmonizing evaluation processes has the added benefit of enabling access by decision-makers to the same information and criteria for evaluation.

Training and Education. FL is a new and disruptive technology that is still unfamiliar to many. Both those applying the technology and those responsible for making decisions about it may not have the technical expertise needed to fully understand and manage the potential risks. There is a need for continuing education for ethics staff and researchers about potential risks and best mitigation practices. Such training needs to be tailored to the different contexts of review (ethics/privacy/operations). Common myths and misconceptions about the technology need to be addressed. Training of researchers can assist them in providing clear and concise information, written at a level that can be easily understood by decision-makers.

Access to Technical Expertise. Decision-makers could benefit from access to technical expert advisors who can offer guidance and support. Having subject matter experts who can speak the language of different fields can help to promote a more nuanced and comprehensive evaluation of the technology and bridge the gap between experts and researchers in the privacy and security space and the issues and concerns faced by data stewards and health authorities. How experts can be deployed without having to be formally appointed to REBs/IRBs needs thought.

Pre-Vetting Certification and Technological Solutions. A pre-vetting process for ethics boards could ensure that details and minutia have been thoroughly considered before the technology is presented to the REB/IRB. Certification can provide additional assurance that the technology has undergone rigorous testing and meets certain security standards. Approaches can be pre-vetted by a group of experts, including an analysis of the proposed architecture globally and how it will be governed at each of the training sites, because the ethics board does not need to understand that level of detail. Finally, formal methods exist to prove that systems have been designed and then built in a way that is secure. These methods provide a way to generate certifications for system implementation.

3.1.3 ENGAGE PATIENTS, EMPOWER COMMUNITIES AND ESTABLISH/ MAINTAIN TRUST

A key question for federated models is whether they can enable data sovereignty, community empowerment and community governance, while mitigating harms? Because data do not leave their local environments, opportunities arise to be more inclusive of patients and pub-

lics in the governance of data and their use.[24] There is a need to engage with and empower patients and communities and referred to the "social license" to make good use of their data and to establish and maintain trust.[25,26] Trust is difficult to earn and difficult to regain once lost. This is exacerbated by an environment, amplified by the pandemic and social media, of distrust with science, such that engagement processes should remain apolitical or non-partisan.

Big data collection and use to date has largely been driven by a small group of health leaders. Patients and publics are generally supportive of use of data for research purposes. For example, a December 2013 poll indicated: "a vast majority of British Columbians are willing to let health researchers study their medical information, as long as the records are anonymous."[27] It is time for inclusive platforms for public deliberation and inclusion of diverse voices in data governance structures.

POTENTIAL POLICY OPTIONS FOR PUBLIC ENGAGEMENT

Encourage patients and communities to participate in governance. We need to empower those represented in datasets, for example, rare diseases communities. Citizen or patient partners can be included on data and research committees or in the research itself and can be provided with sufficient understanding of the technologies to make meaningful contributions.[28] The pan-Canadian Health Data Strategy, in its consultation document, calls for a citizen's assembly and further public deliberations and related federal health agencies and networks have public advisory councils. Models exist for public/citizen engagement and deliberative fora, and across Canada the need for these has

become critical, but the approach has been highly disorganized and uncollaborative, or the questions asked have been too narrow or too focused on privacy. Engagement has the benefit of joining interests with patients. Participants can help develop appropriate frameworks and processes. Patients and communities can additionally be consulted in legislative developments or regulatory reforms about research and other uses for data.[29,30,31,32]

Engage with patients about FL, its applications, and its risks/benefits. Irrespective of how data are shared, including using federated systems, there is a need to be transparent about de-identification standards and to develop strong disincentives to re-identification to build and maintain trust. Similarly, there is a need to be transparent about potential commercial uses of data, conflicts of interest, and return of benefits, if any. Patients and members of the public can be more hesitant with use of their data by commercial and cross-border entities. Commercial uses raise concerns that the technology developed will be proprietary, expensive, or only available to certain segments of the population. These issues of distributive justice reduce trust.[33,34,35,36]

Transparency is needed in underlying model assumptions and performance that may not be validated and may entrench inequities or stereotypes.[37,38,39,40,41] Inherent biases at the point of data collection, the point of data sharing, the data shared, the granularity at which it is shared, the actual algorithms themselves, and the use of those algorithms may result in ethical risks at each stage along the process.

Consider whether FL could align with initiatives to enhance data sovereignty for Indigenous communities and organizations. Such consideration will require engagement and recognition of foundational concepts of collective rights that are inherent, constitutionally protected, and increasingly recognized in legislation and international treaties (United Nations High Commissioner for Refugees, undated), as well as principles of "Ownership, Control, Access, Possession" OCAP® principles which "assert that First Nations have control over data collection processes, and that they own and control how this information can be used."[42]

3.1.4 CONSENT FOR FEDERATED LEARNING

FL uses large datasets of disparate forms of de-identified data that have variable risks for re-identification.[43] Personal health information that may be accessed for FL may generally fall into two categories, some data are consented for collection, access, and use, and some are accessible for research purposes under legislation. Consent is one mechanism to enable inter-provincial or international movement and/or use of data. However, consent processes are problematic in the complexity of the language used and the ambiguities in the risks that are set out.[44] Despite the need to simplify consent documents, requiring consent is preferred because of the opportunity for individuals to withdraw their consent. Consent is an ongoing process, not a one-time transaction and can establish relationships and ongoing trust.

TABLE 1

FORMS OF CONSENT RANKED IN ORDER OF APPLICABILITY FOR FEDERATED LEARNING (FL)

1 CONSENT TO GOVERNANCE

Individuals consent to governance of their data by the institution/ organization that hosts data. The governance model and decision-framework for data access and use are described in the consent form, including collection and storage of the data and the structure and representation of the institutional committee that makes decisions about access to and use of the data. The committee may include public or patient representation. Withdrawal of data by individuals is possible prospectively, but individuals are not contacted for each new use.

2 BROAD CONSENT

Participants are consented to enable a broad spectrum of future uses of the data by various actors, including industry. TCPS-2 has recently undertaken consultation on the issue of broad consent, including how much information will be needed for the consent and any limits on potential future uses of the data.

3 DYNAMIC CONSENT

Individuals have an ongoing option to consent to specific research projects and withdraw consent. This option maximizes participant autonomy and creates an ongoing research relationship. The approach is the most transparent and builds trust. Dynamic consent may be operationalized using technologies such as block-chain, particularly in support of biobanks that house 'omics data.[45,46,47]

4 COMMUNITY CONSENT

In specific circumstances, it may be more appropriate to seek consent from communities or community organizations that represent the interests of individuals within those communities.[48]

5 OPT-OUT

The routinely collected health data of participants are included in the research unless they give their express decision to be excluded. The benefits are greater participation and representativeness at the cost of patient autonomy. In a publicly-funded health system, there is an expectation that patient data should be used for system improvement. Some hospitals have an opt-out program where people are informed when they register for hospital admission.

6 WAIVER OF CONSENT

Researchers may request a waiver consent for secondary use of health data, which may involve a transfer of authority from a principal investigator to an institutional/organizational data custodian/steward. Waiver requests are assessed by ethics boards and data custodians/stewards in line with legislation and policy. Often data may not leave the custodial site and analyses must be performed in secure settings. Only the analyses may leave the site, not the originating data. Data may be accessed across jurisdictions in some places but are centrally managed in a secure facility, and an internal team cleans the data.[49]

7 SPECIFIC CONSENT

Some studies are based on specific consent that do not extend beyond use in that study, or they may consent only to specific activities within a larger data-generating initiative. For example, participants in a clinical trial may grant permission for their data to be linked to health administrative data.

8 HISTORICAL CONSENT

Legacy data must be managed in conformity with the specific consents obtained at the time the data were collected, which may their utility and ability to link them to other data or use them in federated contexts. Historical constraints may be overcome in some instances if secondary use of data under a waiver of consent can be justified. Historical consents may constrain international or commercial use of data, which should be specifically consented to.

9 NO CONSENT

There is a need to clarify the distinction between use of data for health system improvement/program evaluation and research. Data may also be collected and used without consent for statutory purposes such as public health surveillance, however, lack of consent and statutory constraints limit the use of these data. Some data that are otherwise publicly available do not require consent for use.

3.1.5 GOVERNANCE OF DATA INFRASTRUCTURE FOR FL

FL introduces new governance challenges that are specific to the distributed nature of the data, but it has the potential to ameliorate others.[50] Federation may enable virtual data trusts without needing to pool data in siloed platforms that create honey pots for hackers (Figure 1). Federation could facilitate international collaboration, of legislative and policy restrictions on health data leaving Canada, but only if restrictions to local secure data environments are not restricted. Impediments to data access are due to layers of authorities and complicated mechanisms for approvals.

Many initiatives are currently underway to improve health digital infrastructure in Canada and to enable access to these data. We restrict our discussion to those governance challenges specific to FL, which include:

Roles and Responsibilities for Data Controls: In FL, data are distributed across multiple sites, each of which may have different governance structures and be subject to different legal regimens (Figure 1c).[51] As a result, roles and responsibilities may be poorly defined or conflicting with respect to control of the data, their use for training, as well as scope and accountability. The role of data stewards should not be to evaluate the value of the research, but to determine whether the data can be released safely, because other governance and quality-evaluating approval bodies have determined protections are in place.

Data privacy: Because FL involves training models on data that are distributed across multiple sites, it is important to ensure that the privacy of the data is protected, but that privacy is balanced against data use for public health or social benefit. Privacy may be protected through appropriate governance mechanisms but also through the technological means discussed further below. The governance challenges can derive from a culture of protectionism amidst privacy concerns. However, controls may exist at the level of data, which may be subject to different access criteria based on their sensitivity and consent terms.

Data linkage: Challenges for data linkage include multiple layers of enabling legislation, agreements and authority levels, inadequate infrastructure to share, move or link data between research and clinical environments, and lack of funding, compounded by cumbersome processes.

Model ownership: In FL, multiple parties contribute to the training of the model, which can make it difficult to determine who owns the resulting model, creating challenges around intellectual property rights and licensing agreements.

Model validation: Because FL involves training models on distributed data, it can be challenging to validate the accuracy and reliability of the resulting models, requiring new techniques for model validation and testing.

POTENTIAL POLICY OPTIONS

Be transparent and include patient/public representation. FL systems that involve sensitive data must be transparent about their data handling practices and have the appropriate accountability processes in place. It is important to communicate with participants and the public about the data that are being collected, how data are being used, and how privacy is being protected.[52]

Harmonize data control, access, and use agreements. Harmonize agreements to ensure that data are accessed and used in a consistent and secure manner across all sites. Data Use/Transfer Agreements should include differential privacy controls for varying types and sensitivities of data. A privacy review can divide data into levels of identifiability and associated access, with clear rules about who can access data and under what conditions.[53] Standardized data access management plans enable different groups to apply different ML approaches to the same set of data and enables cross-comparison of the models but may be difficult to negotiate across jurisdictions, however, it is possible to use hybrid language to address variations in and reference to privacy provisions. The agreement should be written in the interest of the party that bears the highest risk.

Adopt harmonized policies and procedures across all sites. Establish and enforce standardized policies and procedures for model development across all participating sites, including guidelines for data collection, data sharing, and model training. Constraints on standardization may derive from the need to account for the policies and procedures that are locally specific and legally compliant.

Develop agreements for potential commercial use or partners. These include guidance on conflicts of interest, development of intellectual property and commercialization strategies, and return of benefits, if not to individuals/community, then to the system.

Pre-determine each site's contributions to ensure that each site is appropriately compensated for their contributions, which fosters trust and cooperation.

3.1.6 SPECTRUM OF DE-IDENTIFICATION/IDENTIFIABILITY

A challenge that is unique to FL is site-level re-identification - whether a particular site has participated in the training of a model or whether specific individuals were included at a given site.[54] Site-level re-identification can potentially compromise the privacy and confidentiality of individuals involved in the training process. Questions that arise are: What does de-identification mean in a Federated system for data governance? Does federation mitigate or increase potential risks, for example due to inconsistencies in de-identification standards or increased opportunities to combine data in ways that can lead to re-identification? What is the appropriate level of de-identification that does not render the data unfit for proposed uses?

POTENTIAL POLICY OPTIONS

Best practice guidance on de-identification of data and data classification and handling standards based on de-identification risk. De-identifying data is a best practice, but not all centres do so consistently. Inconsistencies may create issues if data are sourced from different locations and combined. REBs can assist, because they apply standardized guidelines/rules to human subjects research based on the level of identifiability and determine if consent may be waived.

Most centres deploy considerable efforts to de-identify data, for example stripping header information from scans, removing patient names, scrambling the faces of MRI scans through elaborate image processing. Despite such measures, access to multiple variables may enable re-identification.

Metadata can pose additional de-identification challenges. For example, image data contains fielded data, that may contain sensitive data due to human error. Systems need to be in place to catch such errors.

Harmonize de-identification practices across centers for proportionate risk. Practices within federated networks need to balance the risk to utility ratio, which may require novel methods for de-identification, including:

- Removal of personally identifying information;
- Application of differential privacy techniques by adding noise to the data, making it difficult to link individual data points to specific individuals (Figure 2). It is unclear how IRBs/REBs and privacy reviews will account for these methods;
- Use of output controls, which may be most useful for descriptive statistics and statistical products like regression models;
- Methods to analyze the risk for re-identification that evaluate whether data are safe and document the basis for making that judgment prior to products leaving the environment;
- Remove sensitive information in the model that can be linked back to individuals;
- Require minimum cohort size for results reporting; and
- Define metrics to measure the probability of potential reidentification by an attacker.

Develop and use terms in contractual/licensing/collaboration-consortium agreements that prohibit use for re-identification to augment existing or enhanced legislative prohibitions on re-identification.

3.2

SECURITY CHALLENGES FOR FL

3.2.1 DATA BLINDNESS

The distributed nature of FL can make it more challenging to ensure data quality and integrity. In FL, the central server and participating sites have no visibility into all the data that are being used to train the model (data blindness).[55] Sites are therefore unable to directly access or analyze the training data at other sites, making it difficult to diagnose and fix issues with data quality or bias. Participating FL sites will have unique datasets and data structures, making collaboration challenging to harmonize the data into a single model. Additionally, distributed sites create more opportunities for bad data to enter the system.[56,57]

POTENTIAL POLICY OPTIONS

Harmonize to a common data model/ standard through governance standards and contractual terms.

Collaborate only with trusted partners to minimize the risk of data poisoning.

Implement local and global quality controls that can compute on data without directly observing it.

Audit models regularly to evaluate model quality.

Generate federated data summaries and visualizations, which can provide insights into the distribution of the data and help identify potential biases or data quality issues.[58]

Share meta-data, which can inform the design of the model and quality issues.

Share synthetic data with differential privacy. Synthetic data are artificial data that are generated based on statistical models of the original data. Any leakage of private information can be mitigated through differential privacy.[59,60] These techniques may reduce data utility, but should still enable data exploration, harmonization, and other high-level actions.

3.2.2 CYBERSECURITY

A security breach occurs when unauthorized individuals gain access to a system or data without permission, due to hacking, malware, phishing, or any other malicious activity. Security breaches may lead to privacy breaches.[61] Cybersecurity threats are real and can have significant financial and reputational consequences for organizations. Some forms, such as ransomware attacks on healthcare delivery organizations, have more than doubled between 2016 and 2021 in the US.[62,63] Nevertheless, most security breaches occur due to human error, such as improper access or storage, or improper system configurations. Security breaches in the context of FL can arise due to:

- Network security: FL performs training across a network. Each iteration of the training process sends data between participating sites and the central coordinator. The risk of breaches is not greater than in other variants of ML, and FL may, in fact, reduce the risk, especially with the use of secure communication protocols. However, risks may derive from:
    - Unauthorised access to the model: The final FL trained model may need to be moved over the network to another

site, which increases the risk of leaks or unauthorized access. During training, the model may be shared with participating sites, and these need to be trusted not to compromise the model or reverse engineer the private data through a process known as model inversion.[64]

- Cost associated with performing the training over the network with multiple sites due to insufficient bandwidth or computation power.[65]

- Weakest link: FL involves multiple parties, such as data owners, clients, and the central server, all of whom must cooperate to ensure the security of the system. The security of the entire FL system can be compromised by a single weak link in the system.[66]

POTENTIAL POLICY OPTIONS

Access control and authentication. In FL, collaborating sites do not directly access the data but submit a query which is distributed and run locally at each individual site with results released. An essential aspect of governance is to determine who can submit queries and access their results and how much risk do these collaborators pose. Levels of trust are continuously monitored and determined through an agreed upon governance model, which dictates access and enforcement. Access controls restrict access to the model to authorized users and sites and include multi-factor authentication. Controls may be adjusted to allow specific uses of the final model that decrease information gain about the underlying model.

Sandboxing and secure research environments provide input and output controls and do not allow data to leave, if vulnerabilities at entry and exit points for queries are well protected.[67] This is the model employed by most federal and provincial data centres in Canada such as the Canadian Institute for Health Information (CIHI), Ontario's the Institute for Clinical Evaluative Sciences Ontario (ICES), and Population Data British Columbia (PopData BC). [vi]

Regular monitoring is a standard security and reliability practice and is essential during FL model training, when processes need to run uninterrupted for long periods of time.[68] Many issues, such as network outages or performance problems, can be caught by monitoring infrastructure. However, distributed monitoring across multiple sites is challenging, because the sited belong to different administrative domains. Monitoring can catch unauthorized access and suspicious activity and should focus on points of entry into the network.

Security penetration testing such as competitions, and events where students and experts attempt to break the system, can help identify vulnerabilities, improve security, and protect sensitive data.[69]

Encryption methods. such as homomorphic encryption, allow computations to be performed on encrypted data without the need to decrypt it, thereby enhancing security.[70]

[vi] CIHI - https://www.cihi.ca/
ICES - https://www.ices.on.ca/
PopulationData BC - https://www.popdata.bc.ca/

4.0 CONCLUSION

Federated Learning holds great promise in advancing the development of health technologies for the benefit of Canadian patients, Canadian health systems, and Canada's health innovation ecosystem. It allows multiple parties to collaborate on model training without sharing data or pooling data into a central repository. To date, much of the focus of public discussion has been on protecting data. It is time for meaningful public deliberations about the appropriate balance between privacy risks and the harms of not using data for health research and innovation. Inclusive of deliberations should be the ethical limits of models that drive health decisions, considering potential for stigmatization and inequities in model development and access.

Options are available to govern FL consortia that are inclusive of multiple stakeholders in decision-making. Technological development to advance dynamic consent provides options for individuals to determine their level of participation, while others may provide broad consent to appropriately constituted governance bodies. FL simplifies governance with a lead institution that bears the risks of global model development.

However, development and implementation of FL will require training for all partners and decision-makers involved in approvals and governance. Collaboration with computer scientists will be critical to translate advances in security research to facilitate FL and its applications for health and existing research prototypes that have undergone small-scale evaluations. Real-world application will be necessary to demonstrate the effectiveness of these techniques and address their utility in addressing ethics, privacy, and security challenges for the use of health data.

5.0 REFERENCES

[1] Gulshan, V., Peng, L., Coram, M., Stumpe, M. C., Wu, D., Narayanaswamy, A., Venugopalan, S., Widner, K., Madams, T., Cuadros, J., Kim, R., Raman, R., Nelson, P. C., Mega, J. L., & Webster, D. A. (2016). Development and Validation of a Deep Learning Algorithm for Detection of Diabetic Retinopathy in Retinal Fundus Photographs. JAMA, 316(22), 2402. https://doi.org/10.1001/jama.2016.17216

[2] Campanella, G., Hanna, M. G., Geneslaw, L., Miraflor, A. P., Silva, V., Busam, K. J., Brogi, E., Reuter, V. E., Klimstra, D. S., & Fuchs, T. (2019). Clinical-grade computational pathology using weakly supervised deep learning on whole slide images. Nature Medicine, 25(8), 1301–1309. https://doi.org/10.1038/s41591-019-0508-1

[3] Coudray, N., Ocampo, P. S., Sakellaropoulos, T., Narula, N., Snuderl, M., Fenyö, D., Moreira, A. L., Razavian, N., & Tsirigos, A. (2018). Classification and mutation prediction from non–small cell lung cancer histopathology images using deep learning. Nature Medicine, 24(10), 1559–1567. https://doi.org/10.1038/s41591-018-0177-5

[4] Bejnordi, B. E., Veta, M., Van Diest, P. J., Van Ginneken, B., Karssemeijer, N., Litjens, G., Van Der Laak, J. a. W. M., Hermsen, M., Manson, Q. F., Balkenhol, M., Geessink, O., Stathonikos, N., Van Dijk, M. C., Bult, P., Beca, F., Beck, A. H., Wang, D., Khosla, A., Gargeya, R., . . . Venâncio, R. (2017). Diagnostic Assessment of Deep Learning Algorithms for Detection of Lymph Node Metastases in Women With Breast Cancer. JAMA, 318(22), 2199. https://doi.org/10.1001/jama.2017.14585

[5] Esteva, A., Kuprel, B., Novoa, R. A., Ko, J. S., Swetter, S. M., Blau, H. M., & Thrun, S. (2017). Dermatologist-level classification of skin cancer with deep neural networks. Nature, 542(7639), 115–118. https://doi.org/10.1038/nature21056

[6] Mobadersany, P., Yousefi, S., Amgad, M., Gutman, D. A., Barnholtz-Sloan, J. S., Vega, J. A., Brat, D. J., & Cooper, L. (2018). Predicting cancer outcomes from histology and genomics using convolutional networks. Proceedings of the National Academy of Sciences of the United States of America, 115(13). https://doi.org/10.1073/pnas.1717139115

[7] Halpern, A. C., Janda, M., Lallas, A., Longo, C., Malvehy, J., Paoli, J., Puig, S., Rosendahl, C., Soyer, H. P., Zalaudek, I., & Kittler, H. (2020). Human–computer collaboration for skin cancer recognition. Nature Medicine, 26(8), 1229–1234. https://doi.org/10.1038/s41591-020-0942-0

[8] Rieke, N., Hancox, J., Li, W., Milletari, F., Roth, H. R., Albarqouni, S., Bakas, S., Galtier, M., Landman, B. A., Maier-Hein, K. H., Ourselin, S., Sheller, M. J., Zhang, D., Trask, A., Xu, D., Baust, M., & Cardoso, M. J. (2020). The future of digital health with federated learning. Npj Digital Medicine, 3(1). https://doi.org/10.1038/s41746-020-00323-1

[9] Health Canada (2023, February) Working together to improve health care for Canadians. https://www.canada.ca/en/health-canada/news/2023/02/working-together-to-improve-health-care-for-canadians.html

[10] Rieke, N., Hancox, J., Li, W., Milletari, F., Roth, H. R., Albarqouni, S., Bakas, S., Galtier, M., Landman, B. A., Maier-Hein, K. H., Ourselin, S., Sheller, M. J., Zhang, D., Trask, A., Xu, D., Baust, M., & Cardoso, M. J.

(2020). The future of digital health with federated learning. Npj Digital Medicine, 3(1). https://doi.org/10.1038/s41746-020-00323-1

[11] Public Health Agency of Canada. (2021, November). Building Canada's health data foundation: Pan-Canadian Health Data Strategy Expert Advisory Group. Canada.ca. https://www.canada.ca/en/public-health/corporate/mandate/about-agency/external-advisory-bodies/list/pan-canadian-health-data-strategy-reports-summaries/expert-advisory-group-report-02-building-canada-health-data-foundation.html

[12] Rieke, N., Hancox, J., Li, W., Milletari, F., Roth, H. R., Albarqouni, S., Bakas, S., Galtier, M., Landman, B. A., Maier-Hein, K. H., Ourselin, S., Sheller, M. J., Zhang, D., Trask, A., Xu, D., Baust, M., & Cardoso, M. J. (2020). The future of digital health with federated learning. Npj Digital Medicine, 3(1). https://doi.org/10.1038/s41746-020-00323-1

[13] Li, T., Sahu, A. K., Talwalkar, A., & Smith, V. (2019). Federated Learning: Challenges, Methods, and Future Directions. IEEE Signal Processing Magazine, 37(3), 50–60. https://doi.org/10.1109/msp.2020.2975749

[14] Li, T., Sahu, A. K., Talwalkar, A., & Smith, V. (2019). Federated Learning: Challenges, Methods, and Future Directions. IEEE Signal Processing Magazine, 37(3), 50–60. https://doi.org/10.1109/msp.2020.2975749

[15] Wood, A., Altman, M., Bembenek, A., Bun, M., Gaboardi, M., Honaker, J., . . . Vadhan, S. (2018). Differential Privacy: A Primer for a Non-Technical Audience. Vanderbilt Journal of Entertainment & Technology Law, 21(1), 209. Retrieved from https://scholarship.law.vanderbilt.edu/cgi/viewcontent.cgi?article=1058&context=jetlaw

[16] Zhao, Y., & Chen, J. (2022). A Survey on Differential Privacy for Unstructured Data Content. ACM Computing Surveys, 54(10s), 1–28. https://doi.org/10.1145/3490237

[17] Kokosi, T., & Harron, K. (2022). Synthetic data in medical research. BMJ Medicine, 1(1), e000167. https://doi.org/10.1136/bmjmed-2022-000167

[18] Gstrein, O. J., & Beaulieu, A. (2022). How to protect privacy in a datafied society? A presentation of multiple legal and conceptual approaches. Philosophy & Technology, 35(1). https://doi.org/10.1007/s13347-022-00497-4

[19] World Economic Forum. (2020) Sharing Sensitive Health Data in a Federated Data Consortium Model: An Eight-Step Guide. https://www.weforum.org/reports/sharing-sensitive-health-data-in-a-federated-data-consortium-model-an-eight-step-guide/

[20] Wan, Z., Hazel, J. W. G., Clayton, E. W., Vorobeychik, Y., Kantarcioglu, M., & Malin, B. A. (2022). Sociotechnical safeguards for genomic data privacy. Nature Reviews Genetics, 23(7), 429–445. https://doi.org/10.1038/s41576-022-00455-y

[21] Information and Privacy Commissioner of Ontario. (2011) Privacy by Design - The 7 Foundational Principles - Implementation and Mapping of Fair Information Practices. https://www.ipc.on.ca/wp-content/uploads/resources/pbd-implement-7found-principles.pdf

[22] Wood, A., Altman, M., Bembenek, A., Bun, M., Gaboardi, M., Honaker, J., . . . Vadhan, S. (2018). Differential Privacy: A Primer for a Non-Technical Audience. Vanderbilt Journal of Entertainment & Technology Law, 21(1), 209. Retrieved from https://scholarship.law.vanderbilt.edu/jetlaw/vol21/iss1/4/

[23] Public Health Agency of Canada. (2021, December 13). Building Canada's health data foundation: Pan-Canadian Health Data Strategy Expert Advisory Group. Canada.ca. https://www.canada.ca/en/public-health/corporate/mandate/about-agency/external-advisory-bodies/list/pan-canadian-health-data-strategy-reports-summaries/expert-advisory-group-report-02-building-canada-health-data-foundation.html

[24] Rieke, N., Hancox, J., Li, W., Milletari, F., Roth, H. R., Albarqouni, S., Bakas, S., Galtier, M., Landman, B. A., Maier-Hein, K. H., Ourselin, S., Sheller, M. J., Zhang, D., Trask, A., Xu, D., Baust, M., & Cardoso, M. J. (2020). The future of digital health with federated learning. Npj Digital Medicine, 3(1). https://doi.org/10.1038/s41746-020-00323-1

[25] Kalkman, S., Van Delden, J. J. M., Banerjee, A., Tyl, B., Mostert, M., & Van Thiel, G. J. M. W. (2019). Patients' and public views and attitudes towards the sharing of health data for research: a narrative review of the empirical evidence. Journal of Medical Ethics, 48(1), 3–13. https://doi.org/10.1136/medethics-2019-105651

[26] Teng, J., Bentley, C., Burgess, M. M., O'Doherty, K. C., & McGrail, K. (2019). Sharing linked data sets for research: results from a deliberative public engagement event in British Columbia, Canada. International Journal for Population Data Science, 4(1). https://doi.org/10.23889/ijpds.v4i1.1103

[27] Culbert, L. Would you let researchers access your health care data?; B.C.'s ‚medical database cited as a treasure trove of information. Vancouver Sun (5 December 2013).

[28] Velarde, M. R., Tsantoulis, P., Burton-Jeangros, C., Aceti, M., Chappuis, P. O., & Hurst-Majno, S. (2021). Citizens' views on sharing their health data: the role of competence, reliability and pursuing the common good. BMC Medical Ethics, 22, 62. https://doi.org/10.1186/s12910-021-00633-3

[29] Aitken, M., de St Jorre, J., Pagliari, C., Jepson, R., & Cunningham-Burley, S. (2016). Public responses to the sharing and linkage of health data for research purposes: a systematic review and thematic synthesis of qualitative studies. BMC Medical Ethics, 17(1), 73. https://doi:10.1186/s12910-016-0153-x

[30] Harrison, J., Auerbach, A. D., Anderson, W. G., Fagan, M., Carnie, M. B., Hanson, C., Banta, J. E., Symczak, G., Robinson, E. J., Schnipper, J. L., Wong, C., & Weiss, R. B. (2019). Patient stakeholder engagement in research: A narrative review to describe foundational principles and best practice activities. Health Expectations, 22(3), 307–316. https://doi.org/10.1111/hex.12873

[31] Teng, J., Bentley, C., Burgess, M. M., O'Doherty, K. C., & McGrail, K. (2019). Sharing linked data sets for research: results from a deliberative public engagement event in British Columbia, Canada. International Journal for Population Data Science, 4(1). https://doi.org/10.23889/ijpds.v4i1.1103

[32] Tripp, L., Vanstone, M., Canfield, C., Leslie, M., Levasseur, M. A., Panday, J., Rowland, P., Wilson, G. A., You, J., & Abelson, J. (2022). The impact of COVID-19 on patient engagement in the health system: Results from a Pan-Canadian survey of patient, family and caregiver partners. Health Expectations, 25(2), 744–753. https://doi.org/10.1111/hex.13421

[33] Lysaght, T., Ballantyne, A., Xafis, V., Ong, S., Schaefer, G., Ling, J., Newson, A. J., Khor, I. W., & Tai, E. S. (2020). "Who is watching the watchdog?": ethical perspectives of sharing health-related data for precision medicine in Singapore. BMC Medical Ethics, 21, 118. https://doi.org/10.1186/s12910-020-00561-8

[34] McCradden, M. D., Sarker, T., & Paprica, P. A. (2020). Conditionally positive: a qualitative study of public perceptions about using health data for artificial intelligence research. BMJ Open, 10(10), e039798. https://doi.org/10.1136/bmjopen-2020-039798

[35] Paprica, P. A., De Melo, M. N., & Schull, M. J. (2019). Social licence and the general public's attitudes toward research based on linked administrative health data: a qualitative study. CMAJ Open, 7(1), E40–E46. https://doi.org/10.9778/cmajo.20180099

[36] Teng, J., Bentley, C., Burgess, M. M., O'Doherty, K. C., & McGrail, K. (2019). Sharing linked data sets for research: results from a deliberative public engagement event in British Columbia, Canada. International Journal for Population Data Science, 4(1). https://doi.org/10.23889/ijpds.v4i1.1103

[37] Carter, S. M., Rogers, W. A., Win, K. T., Frazer, H., Richards, B., & Houssami, N. (2020). The ethical, legal and social implications of using artificial intelligence systems in breast cancer care. The Breast, 49, 25–32. https://doi.org/10.1016/j.breast.2019.10.001

[38] Mantelero, A. (2022). Regulating AI. In A. Mantelero (Ed.), Beyond Data: Human Rights, Ethical and Social Impact Assessment in AI (pp. 139-183). Springer Nature.

[39] Naik, N., Hameed, B.M.Z., Shetty, D. K., Swain, D., Shah, M., Paul, R., Aggarwal, K., Ibrahim, S., Patil, V., Smriti, K., Shetty, S., Rai, B. P., Chlosta, P., & Somani, B. K. (2022). Legal and Ethical Consideration in Artificial Intelligence in Healthcare: Who Takes Responsibility? Frontiers in Surgery, 9. https://doi.org/10.3389/fsurg.2022.862322

[40] UNESCO. (2023) https://en.unesco.org/artificial-intelligence/ethics

[41] Cooley, O., Pestrue, J., Phillips, M., Konye, J., Penoza, C., Ghous, M., & Singh, K. (2021). External Validation of a Widely Implemented Proprietary Sepsis Prediction Model in Hospitalized Patients. JAMA Internal Medicine. https://doi.org/10.1001/jamainternmed.2021.2626

[42] First Nations Information Governance Centre. (2014). Ownership, Control, Access and Possession (OCAP™): The Path to First Nations Information Governance. https://achh.ca/wp-content/uploads/2018/07/OCAP_FNIGC.pdf

[43] World Health Organization. (2021) Ethics and governance of artificial intelligence for health: WHO guidance. https://www.who.int/publications/i/item/9789240029200

[44] Kassam, I., Ilkina, D., Kemp, J., Roble, H., Carter-Langford, A., & Shen, N. (2022). Patient Perspectives and Preferences for Consent in the Digital Health Context: State-of-the-art Literature Review. Journal of Medical Internet Research, 25, e42507. https://doi.org/10.2196/42507

[45] Alghazwi, M., Turkmen, F., Van Der Velde, K. J., & Karastoyanova, D. (2021). Blockchain for Genomics: A Systematic Literature Review. Distributed Ledger Technologies: Research and Practice, 1(2), 1–28. https://doi.org/10.1145/3563044

[46] Mamo, N., Martin, G. M., Desira, M., Ellul, B., & Ebejer, J. (2020). Dwarna: a blockchain solution for dynamic consent in biobanking. European Journal of Human Genetics, 28(5), 609–626. https://doi.org/10.1038/s41431-019-0560-9

[47] Teare, H., Prictor, M., & Kaye, J. (2021). Reflections on dynamic consent in biomedical research: the story so far. European Journal of Human Genetics, 29(4), 649–656. https://doi.org/10.1038/s41431-020-00771-z

[48] ten Have, H., & Patrão Neves, M. d. C. (2021). Community Consent. In ten Have, H., & Patrão Neves, M. D. C. (2021). Dictionary of Global Bioethics. Springer Nature.

[49] Cumyn, A., Barton, A., Dault, R., Cloutier, A., Jalbert, R., & Ethier, J. (2020). Informed consent within a learning health system: A scoping review. Learning Health Systems, 4(2). https://doi.org/10.1002/lrh2.10206

[50] Broes, S., Lacombe, D., Verlinden, M., & Huys, I. (2018). Toward a Tiered Model to Share Clinical Trial Data and Samples in Precision Oncology. Frontiers in Medicine, 5. https://doi.org/10.3389/fmed.2018.00006

[51] Wolf, L. E., Hammack, C. M., Brown, E. G., Brelsford, K. M., & Beskow, L. M. (2020). Protecting Participants in Genomic Research: Understanding the "Web of Protections" Afforded by Federal and State Law. Journal of Law Medicine & Ethics, 48(1), 126–141. https://doi.org/10.1177/1073110520917000

[52] World Economic Forum. (2020) Sharing Sensitive Health Data in a Federated Data Consortium Model: An Eight-Step Guide. https://www.weforum.org/reports/sharing-sensitive-health-data-in-a-federated-data-consortium-model-an-eight-step-guide/

[53] Broes, S., Lacombe, D., Verlinden, M., & Huys, I. (2018). Toward a Tiered Model to Share Clinical Trial Data and Samples in Precision Oncology. Frontiers in Medicine, 5. https://doi.org/10.3389/fmed.2018.00006

[54] Mothukuri, V., Parizi, R. M., Pouriyeh, S., Huang, Y., Dehghantanha, A., & Srivastava, G. (2021). A survey on security and privacy of federated learning. Future Generation Computer Systems, 115, 619–640. https://doi.org/10.1016/j.future.2020.10.007

[55] Mothukuri, V., Parizi, R. M., Pouriyeh, S., Huang, Y., Dehghantanha, A., & Srivastava, G. (2021). A survey on security and privacy of federated learning. Future Generation Computer Systems, 115, 619–640. https://doi.org/10.1016/j.future.2020.10.007

[56] Nguyen, T. V., Dakka, M. A., Diakiw, S., VerMilyea, M. D., Perugini, M., Hall, J. M. M., & Perugini, D. (2022). A novel decentralized federated learning approach to train on globally distributed, poor quality, and protected private medical data. Scientific Reports, 12, 8888. https://doi.org/10.1038/s41598-022-12833-x

[57] Scheibner, J., Sleigh, Ienca, M J., & Vayena, E. (2021). Benefits, challenges, and contributors to success for national eHealth systems implementation: a scoping review. Journal of the American Medical Informatics Association, 28(9), 2039–2049. https://doi.org/10.1093/jamia/ocab096

[58] Rieke, N., Hancox, J., Li, W., Milletari, F., Roth, H. R., Albarqouni, S., Bakas, S., Galtier, M., Landman, B. A., Maier-Hein, K. H., Ourselin, S., Sheller, M. J., Zhang, D., Trask, A., Xu, D., Baust, M., & Cardoso, M. J. (2020). The future of digital health with federated learning. Npj Digital Medicine, 3(1). https://doi.org/10.1038/s41746-020-00323-1

[59] Behera, M. R., Upadhyay, S. K., Shetty, S. K. B., Priyadarshini, S., Patel, P., & Lee, K. F. (2022). FedSyn: Synthetic Data Generation using Federated Learning. arXiv (Cornell University). https://doi.org/10.48550/arxiv.2203.05931

[60] Chen, H., Tu, C., Li, Z., Shen, H., & Chao, W. (2022). On the Importance and Applicability of Pre-Training for Federated Learning. arXiv (Cornell University). https://doi.org/10.48550/arxiv.2206.11488

[61] Ghimire, B. K., & Rawat, D. B. (2022). Recent Advances on Federated Learning for Cybersecurity and Cybersecurity for Federated Learning for Internet of Things. IEEE Internet of Things Journal, 9(11), 8229–8249. https://doi.org/10.1109/jiot.2022.3150363

[62] Al Kinoon, M., Omar, M., Mohaisen, M., & Mohaisen, A. (2021). Security Breaches in the Healthcare Domain: A Spatiotemporal Analysis. In Lecture Notes in Computer Science (pp. 171–183). Springer Science+Business Media. https://doi.org/10.1007/978-3-030-91434-9_16

[63] Seh, A. H., Zarour, M., Alenezi, M., Sarkar, A., Agrawal, A., Kumar, R., & Khan, R. A. (2020). Healthcare Data Breaches: Insights and Implications. Healthcare, 8(2), 133. https://doi.org/10.3390/healthcare8020133

[64] Fredrikson, M., Jha, S., & Ristenpart, T. (2015). Model Inversion Attacks that Exploit Confidence Information and Basic Countermeasures. CCS '15: Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, 1322-1333. https://doi.org/10.1145/2810103.2813677

[65] Kairouz, P., et al. (2021) Advances and Open Problems in Federated Learning. Now Foundations and Trends Books | IEEE Xplore. https://ieeexplore.ieee.org/document/9464278

[66] Rieke, N., Hancox, J., Li, W., Milletari, F., Roth, H. R., Albarqouni, S., Bakas, S., Galtier, M., Landman, B. A., Maier-Hein, K. H., Ourselin, S., Sheller, M. J., Zhang, D., Trask, A., Xu, D., Baust, M., & Cardoso, M. J. (2020). The future of digital health with federated learning. Npj Digital Medicine, 3(1). https://doi.org/10.1038/s41746-020-00323-1

[67] Banner, N. (2022). NHS data breaches: a further erosion of trust. BMJ, o1187. https://doi.org/10.1136/bmj.o1187

[68] Mothukuri, V., Parizi, R. M., Pouriyeh, S., Huang, Y., Dehghantanha, A., & Srivastava, G. (2021). A survey on security and privacy of federated learning. Future Generation Computer Systems, 115, 619–640. https://doi.org/10.1016/j.future.2020.10.007

[69] Rieke, N., Hancox, J., Li, W., Milletari, F., Roth, H. R., Albarqouni, S., Bakas, S., Galtier, M., Landman, B. A., Maier-Hein, K. H., Ourselin, S., Sheller, M. J., Summers, R.M., Trask, A., Xu, D., Baust, M., & Cardoso, M. J. (2020). The future of digital health with federated learning. npj Digital Medicine, 3, 119. https://doi.org/10.1038/s41746-020-00323-1

[70] Phong, L. T., Aono, Y., Hayashi, T., Wang, L. et Moriai, S. (2017). Privacy-preserving deep learning: Revisited and enhanced. In Batten, L., Kim, D. S., Zhang, X. et Li, G (éd.). Applications and Techniques in Information Security. ATIS 2017. Communications in Computer and Information Science, vol. 719, Springer, Singapore. https://doi.org/10.1007/978-981-10-5421-1_9