

BRITISH COLUMBIA



LAW INSTITUTE

1822 East Mall
University of British Columbia
Vancouver, British Columbia
Canada V6T 1Z1
Voice: (604) 822 0142
Fax: (604) 822 0144
E-mail: bcli@bcli.org
Website: <https://www.bcli.org>

Consultation Paper On Artificial Intelligence And Civil Liability

**Prepared by the
Artificial Intelligence
and Civil Liability
Project Committee**

supported by:



July 2023

Disclaimer

The information and commentary in this publication is not offered as legal advice. It refers only to the law at the time of publication, and the law may have since changed. BCLI does not undertake to continually update or revise each of its publications to reflect post-publication changes in the law.

The British Columbia Law Institute and its division, the Canadian Centre for Elder Law, disclaim any and all responsibility for damage or loss of any nature whatsoever that any person or entity may incur as a result of relying upon information or commentary in this publication.

You should not rely on information in this publication in dealing with an actual legal problem that affects you or anyone else. Instead, you should obtain advice from a qualified legal professional concerning the particular circumstances of your situation.

© 2023 British Columbia Law Institute

The British Columbia Law Institute claims copyright in this publication. You may copy, download, distribute, display, and otherwise deal freely with this publication, but only if you comply with the following conditions:

1. You must acknowledge the source of this publication;
2. You may not modify this publication or any portion of it;
3. You must not use this publication for any commercial purpose without the prior written permission of the British Columbia Law Institute.

These materials contain information that has been derived from information originally made available by the Province of British Columbia at: <https://www.bclaws.gov.bc.ca> and this information is being used in accordance with the King's Printer Licence – British Columbia available at: <https://www.bclaws.gov.bc.ca/standards/Licence.html>. They have not, however, been produced in affiliation with, or with the endorsement of, the Province of British Columbia and **THESE MATERIALS ARE NOT AN OFFICIAL VERSION.**

Published in Vancouver on unceded Coast Salish homelands, including the territories of the xʷməθkwəy̓əm (Musqueam), Skwxwú7mesh (Squamish), and Səlil̓Gwətaʔ/Selilwitulh (Tsleil-Waututh) Nations.

British Columbia Law Institute

1822 East Mall, University of British Columbia, Vancouver, BC, Canada V6T 1Z1

Voice: (604) 822-0142 Fax: (604) 822-0144 E-mail: bcli@bcli.org

WWW: <https://www.bcli.org>

The British Columbia Law Institute was created in 1997 by incorporation under the provincial *Society Act*. Its purposes are to:

- promote the clarification and simplification of the law and its adaptation to modern social needs,
 - promote improvement of the administration of justice and respect for the rule of law, and
 - promote and carry out scholarly legal research.
-

The members of the Institute are:

Tejas B. V. Madhur (Chair)
Prof. Mark R. Gillen (Treasurer)
Marian K. Brown
Brian B. Dybwad
Dr. Ryan S. Gauthier
Julia E. Lawn
Miriam Kresivo, KC
Timothy Outerbridge

Emily L. Clough (Vice-chair)
James S. Deitch (Secretary)
Filip de Sagher
Stacey M. Edzerza Fox
Dr. Toby Susan Goldbach
Mona Muker
Edward L. Wilson

The members emeritus of the Institute are:

Prof. Joost Blom, KC
Margaret H. Mason, KC

Prof. Robert G. Howell

This project was made possible with the sustaining financial support of the Law Foundation of British Columbia and the Ministry of Attorney General for British Columbia. The Institute gratefully acknowledges the support of the Law Foundation and the Ministry for its work.

Artificial Intelligence and Civil Liability Project Committee

The mandate of the Project Committee is to assist BCLI in developing law reform recommendations for adaptation of the common law of torts to address harm to persons, property and other legally protected interests resulting from the operation of artificial intelligence systems.

Prof. Robert G. Howell - Chair
Faculty of Law
University of Victoria
BCLI Member Emeritus

Dylan Merrick
Former Member,
BCLI Board of Directors

Dr. Cindy Grimm
School of Mechanical, Industrial, and Manufacturing Engineering
Oregon State University

Dr. Kristen Thomasen
Peter A. Allard School of Law
University of British Columbia

Cynthia Khoo
Barrister and Solicitor
Senior Associate, Center on Privacy & Technology at Georgetown Law
Research Fellow, Citizen Lab, University of Toronto

Darin Thompson
Liaison Observer on behalf of Ministry of Attorney General
Legal Counsel
Policy and Legislation Division
Justice Services Branch
Ministry of Attorney General

Maya Medeiros
Barrister and Solicitor
Norton Rose Fulbright LLP

Dr. Teresa S.M. Tsang
Director, UBC Artificial Intelligence Echo Core Lab
Director, Echo Lab, Vancouver General Hospital and UBC Hospital
Professor of Medicine, Division of Cardiology
Associate Head, Research
Department of Medicine,
University of British Columbia

For more information, visit us on the World Wide Web at

<https://www.bcli.org/project/artificial-intelligence-and-civil-liability-project/>

Call for Responses

We are interested in your response to this consultation paper. It would be helpful if your response directly addressed the tentative recommendations set out in this consultation paper, but it is not necessary. General comments on anything contained in the consultation paper and on the subject of artificial intelligence and civil liability generally are welcome.

Responses may be sent to us in any one of four ways—

by email: bcli@bcli.org

by mail: British Columbia Law Institute
1822 East Mall
University of British Columbia
Vancouver, BC V6T 1Z1

by fax: (604) 822-0144

online : link from <https://www.bcli.org/project/artificial-intelligence-and-civil-liability-project/>

If you wish your response to be considered by us as we prepare our report, then we must receive it by **1 October 2023**.

Privacy

Your response will be used in connection with the Artificial Intelligence and Civil Liability Project. It may also be used as part of future law-reform work by the British Columbia Law Institute or its internal divisions. All responses will be treated as public documents, unless you expressly state in the body of your response that it is confidential. Respondents may be identified by name, title, and organization in the final report for the project, unless they expressly advise us to keep this information confidential. Any personal information that you send to us as part of your response will be dealt with in accordance with our privacy policy. Copies of our privacy policy may be downloaded from our website at: <https://www.bcli.org/privacy>.

TABLE OF CONTENTS

Executive Summary	xi
Chapter 1. Introduction.....	1
A. Artificial Intelligence Pervades Today's World	1
B. The Artificial Intelligence and Civil Liability Project.....	2
C. The Consultation Paper	2
1. Purpose of the Paper	2
2. Organization of the Consultation Paper	3
3. How to Respond to this Consultation Paper	4
Chapter 2 Artificial Intelligence and Its Nature.....	5
A. What Is Artificial Intelligence?	5
1. General and Narrow Artificial Intelligence.....	5
2. Definitional Elements	6
3. The Project Committee's Working Description of Artificial Intelligence.....	8
(a) The Project Committee's approach to definition	8
(b) The working description	9
B. Different Forms of Artificial Intelligence	10
1. Algorithms are Fundamental to All Artificial Intelligence	10
2. Classical or Symbolic Artificial Intelligence	10
3. Statistical Artificial Intelligence	11
4. Machine Learning.....	11
5. A Note on the Scope of this Consultation Paper	14
C. Two Sides of Autonomous Artificial Intelligence	15
1. General	15
2. Autonomy	16
3. Limited Explainability	18
4. Unpredictability	19
5. Emergence	20
Chapter 3 Civil Liability on What Basis?.....	23
A. The Principle of Fault.....	23
1. Intended Harm	23
2. Unintended Harm	24
(a) General	24
(b) The elements of negligence	24

B. Artificial Intelligence and the Fault Principle	28
C. Competing Theories of Liability for Harm Caused by Artificial Intelligence.....	30
1. Strict Liability	30
2. Product Liability	32
3. Notional Agency and Vicarious Liability.....	33
4. Reasonableness of System Instead of Its Algorithms.....	34
5. A Sliding Scale of Modified Liability Rules.....	35
D. The Project Committee's View	36
1. Rejection of Strict Liability	36
2. Notional Agency Insufficiently Distinguishable from Strict Liability	37
3. Non-Human Behaviour Is Not Measurable Against Human Reasonableness	38
4. Liability Upstream - Product Liability Provides Some Answers.....	40
5. Liability Downstream - Operators	43
(a) Who is an "operator"?	43
(b) Liability of operators and other downstream defendants.....	46
(c) Tentative recommendation.....	46
E. Exclusion and Limitation of Liability for Artificial Intelligence	46
Chapter 4. The Problem of Proof of Fault.....	49
A. General	49
B. Rebalancing Evidentiary Burdens for Fairness	51
1. General.....	51
2. European Commission Rebalancing Proposals	53
(a) The AI Liability Directive Proposal.....	53
(b) The EC proposal for a new Product Liability Directive	57
3. Recommendation on Relief Against Difficulties of Proof in Appropriate Cases.....	58
C. Reasonable Foreseeability and Artificial Intelligence	61
Chapter 5. Standard of Care	67
A. How the Standard of Care Is Set.....	67
B. A Largely Open Playing Field – For the Time Being.....	68
C. Widely Recognized Elements of Good Practice	71
1. The Design, Development, Training, and Testing Phases.....	71
(a) Transparency.....	71
(b) Vital design features.....	72
(c) Data Quality and Data Governance	72
(d) Continuous risk assessment	73
(e) Independent validation.....	73
(f) Articulation of appropriate human involvement	73
(g) Logging Capability.....	74
(h) Monitoring.....	74

(i) Updating	74
2. Operation in Actual Use.....	74
(a) Risk Assessment and Mitigation	74
(b) Compliance with Developer's Recommendations and Terms of Service	75
(c) Transparency.....	75
(d) Monitoring of System Performance	76
(e) System Maintenance	76
(f) Privacy.....	77
(g) Training.....	77
(h) Logs and Other Record-keeping	77
(i) Organizational governance framework for artificial intelligence systems.....	77
D. Tentative Recommendation	78
Chapter 6. Algorithmic Discrimination and Civil Liability	79
A. Bias: A Recognized Problem in Artificial Intelligence	79
B. Overview of Canadian Anti-discrimination Law.....	82
1. A Legal Definition of Discrimination.....	82
2. Unconstitutional discrimination under the equality rights section of the Charter	83
3. Discrimination under human rights legislation	84
4. No general tort of discrimination.....	86
C. Discrimination Produced by Artificial Intelligence.....	87
1. Algorithmic Discrimination Without Remedy	87
2. How Should the Gap in Protection Against Algorithmic Discrimination Be Closed?	89
(a) Human rights vs. tort remedy	89
(b) The Project Committee's view	92
(i) Negligent algorithmic discrimination.....	92
(ii) Intentional algorithmic discrimination	93
(c) Tentative recommendations.....	94
Chapter 7. Conclusion	95
List of Tentative Recommendations.....	97

Executive Summary

Artificial intelligence is pervasive in today's world. Decisions, predictions and recommendations made by artificial intelligence affect individual lives to an ever-increasing extent.

The benefits of artificial intelligence are great, but they also come with risks. When the risks materialize, harm to persons and property may result. Avoidable harm to persons and property brings the law of tort into play.

Tort is the branch of law concerned with non-contractual civil wrongs. Its principles evolved in order to deal with harmful human conduct. Involvement of artificial intelligence as a causal factor leading to harm complicates the application of those principles to determine who is legally responsible, when and for what. Many questions concerning rights and liabilities in that context have yet to be answered because these situations are new, and the law relating to artificial intelligence is unsettled.

The Artificial Intelligence and Civil Liability Project

This consultation paper is issued under the Artificial Intelligence and Civil Liability Project, which BCLI initiated in late 2021. The objectives of the project are:

- to identify how the common law rules of tort need to be adapted to provide just and adequate civil remedies for harm to persons, property, and other legal interests resulting from the operation of autonomously functioning artificial intelligence; and
- to develop and publish law reform recommendations addressing that context.

BCLI is assisted in carrying out the project by the Artificial Intelligence and Civil Liability Project Committee, an interdisciplinary group reflecting expertise in computer science, engineering, and medicine as well as law.

Purpose of Consultation Paper

The purpose of the consultation paper is to gather the views of stakeholders and the general public on tentative law reform recommendations developed by the Project Committee. The recommendations are referred to as “tentative” only in the sense that they are not final, and remain subject to modification, additions, or even abandonment. Responses to this consultation paper will be closely considered in forming

the final recommendations in the law reform report that BCLI will issue at the end of the project.

Outline of the Consultation Paper

Chapter 1

Chapter 1 provides general background on the Artificial Intelligence and Civil Liability Project and the reasons why BCLI has undertaken it.

Chapter 2

Chapter 2 explains what artificial intelligence is generally understood to be and describes its various forms, including machine learning. Various definitions of artificial intelligence are reviewed and elements common to them are identified. As there is no single, universally accepted definition of artificial intelligence, a “working description” for the purposes of the project is offered instead that builds upon those common elements.

The concept of “autonomy” is introduced as a characteristic that artificial intelligence systems possess to varying extents. “Autonomy” refers to the ability of a system or device to interact with its environment and pursue objectives assigned to it without continuous human input or control.

Limited explainability is another characteristic of some artificial intelligence systems, especially those that rely on machine learning and artificial neural networks. Systems capable of machine learning are not fully dependent on programmed instructions by humans, but can base outputs on inferences from patterns they detect in data and improve their performance by trial and error. They can re-use the inferences from data to which they have been exposed to process and analyze new data. The abstract models the systems create internally to represent the relationships they find in the data they process may not be interpretable by humans. Thus, the process from input to output on these systems that are heavily dependent on data is not always explainable. This is why the expression “black box” is often applied to artificial intelligence in the media and popular literature.

Although “black box” is a misleading and unhelpful term, it is true that the decision-making of data-dependent systems cannot always be validated on the basis of a step-by-step chain of reasoning. If the way a system produces its outputs cannot be explained, the way it will respond to a specific input is less predictable.

Unpredicted, original behaviour of artificial intelligence in pursuit of the objectives that humans give it has been called “emergence.” Emergence has a double aspect. It is beneficial much of the time, sometimes solving seemingly intractable problems and astounding even the designers and programmers of the systems that display it. On occasion, however, the same capabilities that allow artificial intelligence to generate unanticipated but highly desirable results may also generate unintended, harmful results.

When the outputs of artificial intelligence result in harm to persons, property, or other legally protected interests, implications of civil liability will arise under tort law for the humans and legal entities behind the systems in question. The rest of the consultation paper explores those implications.

Chapter 3

Chapter 3 deals with the legal basis for imposing civil liability when harm results from the operation of artificial intelligence. The point is made at the beginning that legal responsibility for the performance of an artificial intelligence system always rests ultimately with the humans or corporate entities behind it, because the systems themselves do not have the means of compensating their victims. The chapter goes on to cover basic principles of tort law that relate to the rest of the consultation paper, such as fault, intention, and the elements of negligence.

Different approaches to civil liability for harm from artificial intelligence that have been proposed in the common law world are then reviewed and evaluated along with ones proposed in the EU, which is ahead of much of the world in addressing legal frameworks for use of artificial intelligence. These varying approaches include strict liability, product liability, notional agency, vicarious liability, legal neutrality (applying the same legal standards of reasonable conduct to harm-causing artificial intelligence as are applied to human tortfeasors), and a sliding scale of liability regimes depending on levels of risk.

The tentative recommendation of the Project Committee is to retain a fault-based regime for artificial intelligence rather than one of strict liability. Under strict liability, a plaintiff needs only to prove damage and causation. Absence of fault on the part of the defendant is irrelevant. Among the several grounds advanced for rejecting strict liability is that it reduces the incentive for continual improvement of standards in the design, development, and use of artificial intelligence, because defendants will be held liable under it regardless of the degree of care they exercise. In the view of the Project Committee, a fault-based regime is more likely to achieve a fair balance between risk and the benefits of technological innovation.

The Project Committee would divide defendants in cases relating to artificial intelligence into two classes: those situated “upstream” in the chain of events leading to litigation and those situated “downstream.” Upstream defendants would be those involved in the design, development, training and testing of artificial intelligence systems prior to the point at which the systems are placed on the market or deployed in an actual use case in the post-testing phase.

Downstream defendants would be those who use an artificial intelligence system for their own purposes or have control over the risks of its operation once the system has been deployed or released for actual, real-world use following the design, development, training, and testing phases.

The two classes are not mutually exclusive and could overlap. For example, a company that develops an artificial intelligence system for its own use would be both an upstream and downstream defendant.

Under the tentative recommendations in Chapter 3, the civil liability of downstream defendants for unintended harm caused by the systems they operate for their own purposes would be based on fundamentally on negligence principles. Civil liability of upstream defendants would be based on an adaptation of product liability principles. Developers of a complete, integrated artificial intelligence system should be analogized to manufacturers of a complex product and owe a duty of care accordingly. Designers and developers of components, including an artificial intelligence module that is incorporated into an integrated software system or device, would owe a duty of care analogous to that owed by suppliers of components of a complex product, who are considered manufacturers in their own right for the purposes of product liability.

In Canadian tort law, product liability is a branch of negligence with well-established principles surrounding the scope of the duty of care of manufacturers and others in the supply chain. The duty of care of a manufacturer extends to anyone who may come into contact with a product after it enters the stream of commerce and thus may be affected by its intrinsic defects or hazards.

Application of negligence principles to both upstream and downstream defendants would be modified according to recommendations made in the later chapters to address particular issues that will arise in tort claims for damage stemming from the operation of artificial intelligence.

Comment is sought from readers on whether, and to what extent, it should be possible to disclaim, exclude, or limit liability for harm caused by artificial intelligence.

Chapter 4

Chapter 4 addresses the formidable obstacles that may face a plaintiff in proving causation and fault when the damage may be linked with non-human, autonomous decision-making by artificial intelligence. There will be many potential defendants, because a complete artificial intelligence system typically integrates numerous elements with multiple sources.

The outward opacity of the systems, together with the tendency of the software industry to treat algorithms and other aspects of the design of systems as proprietary secrets, has emerged as a problem in civil litigation. Defendants other than the designers of the underlying algorithms may be as much in the dark as the plaintiff regarding information about the system that is crucial to a fair adjudication of a damage claim.

Even the designers of an artificial intelligence system may be unable to explain how systems that are highly data-dependent and operate using artificial neural networks reach particular outputs. Given this, pre-trial discovery processes alone may be inadequate to redress the “informational asymmetry” between plaintiffs and defendants. If the route from input to output is not fully explainable even by the designers and programmers of a system, plaintiffs have little hope of being able to put forward a coherent, provable theory of causation linking a breach of legal duty by the defendant to the harm incurred.

Numerous academic writers in the common world have reached the conclusion that these informational obstacles in the way of proof of causation and fault warrant some mechanism to maintain balance in the litigation process in cases involving artificial intelligence, as have EU policymakers. The European Commission’s AI Liability Directive Proposal is chiefly concerned with creating a more equal playing field for claimants (plaintiffs) and defendants in fault-based claims under the national law of EU Member States. *Res ipsa loquitur* might formerly have had a role to play in that regard in Canadian common law jurisdictions, but the Supreme Court of Canada has held that *res ipsa loquitur* is obsolete and may no longer be invoked as a distinct mechanism in aid of proof of negligence.

Rather than taking the European Commission’s prescriptive and somewhat inflexible approach of directing courts to make automatic inferences of causation that a defendant is obliged to rebut, the Project Committee tentatively recommends that it be open to courts to draw an inference of negligence in the absence of either an explanation being offered by the defendant who developed the complete system and made it available to others for the performance of the system that is consistent with

the exercise of reasonable care in the pre-deployment phases, or pre-trial disclosure to other parties (including co-defendants) of information concerning the design and function of the system that would allow the other parties to put forward a theory of causation.

As in other negligence-based claims, foreseeability of harm would be a prerequisite to liability. Reasonable foreseeability as conventionally understood and applied by Canadian courts can become unstuck, however, when the claim is based on non-human decision-making. Described in case law in terms of what a reasonable person would see as a “real risk” that is not “far-fetched,” or in terms of a “natural result” of an act or omission, reasonable foreseeability is inextricably linked to human experience and human perceptions of cause and effect. The reactions of a non-human decision-making process to an infinite range of potential inputs are outside that experience. Under the conventional tests, harmful emergence will seldom be found to have been reasonably foreseeable. This would result in victims incurring harm through artificial intelligence having less protection under tort law than victims of human tortfeasors.

To avoid this result, foreseeability of risk in relation to the use of artificial intelligence must not be thought of in terms of particular outcomes considered in isolation, but in terms of unpredictability and emergence being known risks (“known unknowns”) that potentially give rise to unknown ones.

The determination of what was reasonably foreseeable as a risk that may manifest in an emergent and unpredictable manner should be made with reference to the known attributes of the system in question at the relevant time, the use cases for which the system was intended, and known or predictable alternate use cases (including predictable misuse). In the view of the Project Committee, these considerations delineate the scope of potential harm that designers, developers, and operators should be expected to consider in assessing risk and taking measures to avert it.

Chapter 5

Chapter 5 is concerned with how the standard of care should be set in claims based on harm arising from the operation of artificial intelligence. It is open to courts to look to relevant statutory and regulatory requirements affecting the activity in issue in the case and to evidence of the accepted or customary standard in the industry or profession in question. These are relevant considerations in setting the standard of care in an individual case, but not decisive regarding it. Breach of a statute or regulation applicable to the defendant or the defendant’s activities in issue does not automatically lead to a finding of liability, nor will compliance necessarily avoid liability. The same is true with respect to accepted or standard practice within industry.

There is little regulation of artificial intelligence in Canada now in any case. If passed, the proposed *Artificial Intelligence and Data Act* (“AIDA”) and regulations under it will not be in force until 2025 at the earliest, and it will not apply to all artificial intelligence systems.

Courts can look to some international regulatory precedents that have begun to appear, such as the proposed EU *AI Act*, the proposed U.S. *Algorithmic Accountability Act*, public policy framework documents such as the *White House Blueprint for an AI Bill of Rights*, NIST *Artificial Intelligence Risk management Framework*, the UK Government policy paper *A pro-innovation approach to AI regulation*, and published voluntary guidelines issued by various scientific and industry bodies.

The level of risk associated with a system undoubtedly will be among the factors weighed by courts in determining the standard of care applicable in a given case. In doing so, courts should not lose sight of the benefits of innovation and should seek a standard of care that represents a reasonable balance between risk and benefit. As an international consensus on standards in the development and application of artificial intelligence is still in a relatively early stage of development and will continue to evolve, no specific tentative recommendation is made on the content of the standard of care in a negligence case involving artificial intelligence. Instead, elements of good practice in the design, development, and operation of artificial intelligence systems that appear with relative consistency in an international cross-section of regulatory policy documents, and thus seem to have wide recognition, are listed. Courts are urged look beyond national borders and to take account of inter-jurisdictionally recognized best practices in determining the standard of care.

Chapter 6

Chapter 6 deals with what is possibly the most significant legal and ethical problem associated with artificial intelligence, namely its potential to generate biased output that can lead to discrimination. Bias can enter automated decision-making processes in several ways. It may be present in the algorithm on which the system is based, or in the data used to train or test a system. It may also result from a failure of human oversight through assumptions based on conscious or unconscious biases. So-called “feedback loops” can replicate and amplify bias hidden in previously generated data when it is used as input to train a new version of the system.

Ostensibly neutral input variables like postal codes, income, and educational level may become proxies for race and ethnicity if a system correlates them with demographic patterns that it detects in data. Resulting outputs, or reliance by humans on them, may inadvertently impose additional disadvantages on already disadvantaged and marginalized populations. Chapter 6 describes several examples of seriously

discriminatory effects resulting from the output of artificial intelligence (“algorithmic discrimination”).

In some cases, people who experience unjustified adverse treatment as a result of biased output of artificial intelligence systems will have a remedy under federal and provincial human rights legislation or the equality rights section of the *Canadian Charter of rights and Freedoms*. In other cases they will not, because the differentiation they experience is not based on a prohibited ground of discrimination under existing anti-discrimination frameworks.

Canadian courts do not recognize a general tort of discrimination at common law. The British Columbia *Civil Rights Protection Act* will have little bearing on artificial intelligence because it applies only to conduct or communication having the purpose of promoting hatred or racial, ethnic, or religious discrimination. Algorithmic discrimination will be unintentional in most cases. AIDA and regulations under will not provide a compensatory civil remedy.

As usage of artificial intelligence expands, there will be more cases of unintended discriminatory effects that do not fit readily into the human rights framework. The gap in the law into which these cases fall will become more apparent. They will be seen by society as unfair and untenable, calling for some means of legal redress. The alternatives appear to be to expand the human rights framework to cover algorithmic discrimination or to create a remedy in tort.

Discrimination under human rights legislation is primarily circumscribed by reference to characteristics of identity that are either “immutable or changeable only at unacceptable cost to personal identity,” namely ones such as race, place of origin, ethnicity, age, disability, gender, language, and religion. The kinds of discrimination that artificial intelligence may create will frequently be based on other factors. While these factors may serve as proxies for prohibited grounds, requiring a claimant to prove this as a precondition to a human rights remedy could be onerous, and potentially impose a strain on the resources of human rights tribunals and commissions.

These additional difficulties of proof, and the lack of a remedy for some forms of algorithmic discrimination altogether under the human rights framework or existing tort law, raise an issue of access to justice that will take on increasing importance with the expansion of artificial intelligence. The question arises whether a civil remedy in tort should be made available for discriminatory treatment resulting from the operation of artificial intelligence.

Algorithmic discrimination is likely to have different impacts than those typically dealt with by human rights tribunals and commissions. Human rights complaints typically concern an affront to personal dignity and resulting mental distress. Algorithmic discrimination is more likely to result in harm of a different kind, such as economic loss or ineligibility for a public or private benefit. These are closer to the kinds of claims that civil courts are accustomed to ruling upon and quantifying.

The Project Committee has tentatively concluded that a civil remedy for algorithmic discrimination should be introduced, either by legislation or by judicial decision as an incremental change in the common law. The proposed cause of action would consist of a failure to take reasonable steps to detect and correct biased output of an artificial intelligence system or another algorithmic process, resulting in discrimination against a person or class that is either illegal or not warranted by reasonable business or industry practice. If the discrimination is illegal, it might also be a breach of human rights legislation. In that case there could be a dual remedy, but only one loss for which compensation could be awarded.

Proof of damage going beyond the fact of differential treatment itself would be required in order to prevent open-ended liability for any form of differentiation. Some kinds of harm that algorithmic discrimination may produce would not fit easily within the conventional range of compensable damage. They may be speculative and difficult to quantify, such as loss of opportunity. Nevertheless, courts should take a broad view of what amounts to damage in cases based on algorithmic discrimination.

Chapter 7

Chapter 7 is a general conclusion inviting readers to respond with incisive comment to aid in the process of shaping the legal framework needed for a world in which artificial intelligence will play an increasingly large role.

Chapter 1. Introduction

A. Artificial Intelligence Pervades Today's World

Automation is increasingly encountered in daily life. Automated systems perform a vast range of tasks that were carried out by humans exclusively or predominantly in the recent past. Some forms of automation employ the kinds of digital technology known generically as artificial intelligence. Among these are systems that allow self-driving vehicles to navigate roadways safely in traffic, respond to internet queries, and translate from one language to another. Others enable the digital voice assistants like Alexa and Siri to respond to our commands and answer our questions. ChatGPT will generate text on command on a vast range of subjects that can be practically indistinguishable from writing by humans. It and other chatbots carry on conversations so fluently that people are often unaware they are dealing with a machine rather than a human in another corner of the internet.

The applications of artificial intelligence are vast and continually growing. Artificial intelligence is used to make decisions, predictions and recommendations that affect individual lives to an increasing extent. Artificial intelligence is used in diagnosing disease, designing new drugs, and predicting extreme weather events by detecting patterns concealed in vast amounts of climate data. Some applications of artificial intelligence are controversial, such as use of facial recognition technology in law enforcement, and profiling consumers without their knowledge for targeted advertising based on their viewing and browsing habits.

Like other fields of digital technology and automation, artificial intelligence is always evolving and becoming continually more complex and sophisticated. As explained later, some forms of artificial intelligence can learn from their previous experience and apply what they learn in processing new data. This highly useful capability allows these systems to improve their own performance without human intervention, and reach innovative solutions. It can also make the outputs less predictable than those of simpler systems that depend entirely on programmed instructions for their function.

The great benefits of artificial intelligence come with risks. When these materialize, harm to persons and property may result. When harm occurs to persons and prop-

erty, the law of tort is in play. This branch of law concerns non-contractual civil wrongs, also known as torts.

Tort principles evolved as part of the common law to address harmful human behaviour. Tort law is applicable too when harm occurs from the operation of artificial intelligence, but the involvement of artificial intelligence in a set of facts resulting in loss or injury complicates the application of tort principles to determine who is legally responsible, when and for what. Many questions concerning rights and liabilities in that context have yet to be answered because these situations are new, and the law relating to artificial intelligence is unsettled.

B. The Artificial Intelligence and Civil Liability Project

The British Columbia Law Institute (BCLI) began the Artificial Intelligence and Civil Liability Project in late 2021. The objectives of the project are:

- to identify how the common law rules of tort need to be adapted to provide just and adequate civil remedies for harm to persons, property, and other legal interests resulting from the operation of autonomously functioning artificial intelligence; and
- to develop and publish recommendations to adapt the law of tort to address that context.

Law reform recommendations developed in this project are intended to be of a kind that could be implemented equally easily by legislative means or by judicial decision in individual cases, the process by which the common law incrementally evolves.

BCLI is conducting the project with the aid of the Artificial Intelligence and Civil Liability Project Committee, an interdisciplinary group reflecting expertise in computer science, engineering, and medicine as well as law.

C. The Consultation Paper

1. Purpose of the Paper

This consultation paper contains the tentative recommendations of the Artificial Intelligence and Civil Liability Project Committee and explains the rationale for them.

The purpose of the paper is to gather the views of stakeholders and the general public on its contents.

The recommendations are described here as “tentative” only in the sense that they are not necessarily final positions. They are subject to reconsideration, modification, or even abandonment after review of the reaction to the consultation paper. With the benefit of the comments received on the contents of this document, the Project Committee and BCLI will form the final recommendations on artificial intelligence and civil liability and issue a law reform report containing them.

2. Organization of the Consultation Paper

Chapter 1 is a general introduction.

Chapter 2 describes what the term “artificial intelligence” is usually understood to cover. The distinctive features of this technology that have particular significance for the law of tort are explained. Machine learning is described as a subfield of artificial intelligence having particular importance for the Artificial Intelligence and Civil Liability Project.

Chapter 3 lays out basic principles of tort law that are relevant to the rest of the consultation paper. It then covers various legal theories that have been advanced regarding the basis on which civil liability for harm from artificial intelligence should be imposed. Various legal theories that are currently being advanced are covered. Chapter 3 then explains the view the Project Committee has tentatively taken, and why.

Chapter 4 deals with problems that some features of artificial intelligence present in the application of tort rules, in particular its limited explainability and potential for unpredictable behaviour.

Chapter 5 deals with the standard of care in negligence litigation related to artificial intelligence, and how it can be set by reference to accepted elements of good practice in the development and deployment of artificial intelligence systems.

Chapter 6 deals with one of the principal legal and ethical problems in the design and operation of artificial intelligence systems, namely inadvertent bias leading to discriminatory results. We discuss whether recognizing a new tort of algorithmic discrimination is an appropriate means of providing legal redress for cases that do

not fit under the human rights framework or the equality rights guarantee in the *Canadian Charter of Rights and Freedoms*.¹ A tentative recommendation is presented.

Chapter 7 is a general conclusion.

3. How to Respond to this Consultation Paper

Responses to this consultation paper are invited by **1 October 2023**. Responses of any length may be sent by e-mail, regular mail, fax, or via the BCLI website. The contact information for each of these options appears on the page entitled “Call for Responses” near the front of the paper.

1. Part 1 of the *Constitution Act, 1982*, being Schedule B to the *Canada Act 1982* (U.K.), 1982, c. 11.

Chapter 2 Artificial Intelligence and Its Nature

A. What Is Artificial Intelligence?

1. General and Narrow Artificial Intelligence

An important distinction must be drawn between what is referred to as general and narrow artificial intelligence. General (or “strong”) artificial intelligence is the term applied to a technological capability to perform the full range of mental tasks that a human being could perform. It has also been defined as “the ability to satisfy goals in a wide range of environments.”²

While general artificial intelligence has been a theoretical objective since the dawn of the computer age, opinions among experts vary enormously on when, or even if, it will be achieved.³

Narrow (or “weak”) artificial intelligence is technology that simulates some aspect of human intelligence or performs specific tasks in a particular domain.

All artificial intelligence is narrow at the present time and is likely to remain narrow well into the future. This is true whether the artificial intelligence is designed to solve a single problem or is a generic tool that can be used in a range of applications or assimilated into other digital and physical technology.

2. Marcus Hutter, *Universal artificial intelligence: sequential decisions based on algorithmic probability* (Berlin: Springer, 2005), online: <https://doi.org/10.1007%2Fb138233>.

3. Vincent C. Müller and Nick Bostrom, “Future progress in Artificial Intelligence: A Survey of Expert Opinion” in Vincent C. Müller, ed., *Fundamental issues of Artificial Intelligence* (Springer, 2016) at 553-571.

2. Definitional Elements

There is no single, universally accepted definition of artificial intelligence. Some definitions refer to it as digital technology that allows computers to simulate aspects of human intelligence or approximate intelligent behaviour.⁴ These definitions are not especially helpful, because they presuppose there is a clear and universal definition of human intelligence, which is not the case.⁵

Other definitions of artificial intelligence focus on the functions that the technology typically performs. These include:

- natural language processing
- speech recognition
- computer vision
- image recognition
- knowledge representation (storing information in a form that enables a computer to retrieve and re-use it to solve new problems or make decisions).
- search and data retrieval
- data analysis
- pattern recognition

The Government of Canada *Directive on Automated Decision-Making*, for example, defines artificial intelligence as:

4. For example: "The capacity of computers or other machines to exhibit or simulate intelligent behaviour; the field of study concerned with this. Abbreviated *AI*." Oxford English Dictionary, 3rd ed. (Updated December 2021).

5. S. Samoli et al., *AI Watch. Defining Artificial Intelligence. Towards an operational definition and taxonomy of artificial intelligence* (Luxembourg, Publications Office of the European Union, Luxembourg, 2020) at 7.

Information technology that performs tasks that would ordinarily require biological brainpower to accomplish, such as making sense of spoken language, learning behaviours, or solving problems.⁶

Other definitions refer to the ability of the technology to form models from data supplied by humans or ingested from sensors and make inferences and decisions from them. One of the leading textbooks on artificial intelligence defines it as “the study of agents that receive percepts from the environment and perform actions.”⁷

The ability to make decisions, predictions, or recommendations that influence an environment figures prominently in numerous efforts at definition. The OECD Council emphasized this element in describing artificial intelligence in a 2019 policy document:

An AI system is a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations, or decisions influencing real or virtual environments. AI systems are designed to operate with varying levels of autonomy.⁸

Outputs in the form of decisions, predictions, or recommendations can be said to influence an environment because they are either relied upon by humans or are implemented by other parts (actuators) of a technological system with which the artificial intelligence software is integrated.

The U.S. *National Artificial Intelligence Initiative Act of 2020* brings together in its definition these elements of perception, automated inference and reasoning, and outputs taking the form of decisions, predictions, or recommendations:

The term “artificial intelligence” means a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations or decisions influencing real or virtual environments. Artificial intelligence systems use machine and human-based inputs to-

(A) perceive real and virtual environments;

6. See online: <https://www.tbs-sct.canada.ca/pol/doc-eng.aspx?id=32592>.

7. Stuart J. Russell and Peter Norvig, *Artificial Intelligence: A Modern Approach*, 4th ed. (Pearson: Hoboken, N.J., 2021) at vii.

8. OECD, *Recommendation of the Council on Artificial Intelligence*, OECD/Legal/0449, Art. I, 21 May 2019, online: <https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449>.

(B) abstract such perceptions into models through analysis in an automated manner; and

(C) use model inference to formulate options for information or action.⁹

The European Commission's proposal for harmonized rules on artificial intelligence in the EU (often referred to as the "AI Act" though not yet enacted) includes a definition of artificial intelligence as software having the same elements as those referred to in the U.S. statute, but with the further qualification that the software must have been developed using one or more of the computer science and statistical techniques and approaches in a finite list.¹⁰

A further element referred to as "autonomy" is often mentioned in definitions of artificial intelligence. Although variously described, "autonomy" denotes the ability of artificial intelligence to adapt to the conditions of its operating environment and fulfil objectives assigned to it with a minimum of human input or control.¹¹ Put another way, it is an ability to determine how to solve a problem or complete a task assigned by human programmers without programmed instructions to select a specific method.¹² Autonomy is a matter of degree. As mentioned in the OECD definition above, artificial intelligence systems vary in the extent of autonomy they possess.

3. The Project Committee's Working Description of Artificial Intelligence

(a) The Project Committee's approach to definition

At an early stage in the Artificial Intelligence and Civil Liability Project, the Project Committee opted to set parameters for the project by adopting a functional descrip-

9. 15 USC § 9401(3).

10. *Proposal for a Regulation of the European Parliament and of the Council Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act) and Amending Certain Union Legislative Acts*, 21 April 2021, COM(2021) 206 final, online: <https://eur-lex.europa.eu/legal-content/EN/ALL/?uri=CELEX:52021PC0206>, Art. 3(1).

11. Chinedu Pascal Ezenkwu and Andrew Starkey, "Machine Autonomy: Definition, Approaches, Challenges and Research Gaps" in Kohei Arai, Rahul Bhatia, Supriya Kapoor, eds. *Intelligent Computing: Proceedings of the 2019 Computing Conference*, vol. 1 (Cham, Switzerland: Springer, 2019) 335 at 336. See also Wolfhart Totschnig, "Fully Autonomous AI" (2020) 26 *Science and Engineering Ethics* 2473. Russell and Norvig, *supra*, note 7, define autonomy in terms of its absence: a system lacks autonomy if it "relies on the prior knowledge of its designer rather than on its own percepts and learning processes."

12. Ryan Abbott, *The Reasonable Robot* (Cambridge: Cambridge University Press, 2020) at 34.

tion of artificial intelligence instead of adopting or modifying one of the many existing definitions that various writers, institutions, and governments have put forward.

The Project Committee considered that our working description needed to be broad enough to encompass the range of technologies and systems that is generally understood to be artificial intelligence, and flexible enough to remain relevant as these evolve.¹³ For this reason, it would not be described by reference to particular technologies, functions, or techniques that might change over time or come to be known under different names. It would refer only to overarching functions in order to preserve its relevance and continuity as artificial intelligence rapidly advances.

The working description also needed to refer to the cardinal features of artificial intelligence that give rise to the issues relating to civil liability that the project was intended to address. Those cardinal features are, first, the capability to operate with a minimum of human oversight and, second, the capability to make decisions that have real effects in the outer world independently of human programmers and operators.

(b) The working description

This is the committee's working description of artificial intelligence:

Artificial intelligence is technology that

(a) is designed to supplement, or substitute for, human action, advice, or decision-making, or to enable decision-making beyond the capabilities of unaided human intelligence; and

(b) with minimal or no human intervention, can use novel input to affect or interact with a real or virtual environment by

(i) an action, or

13. In this paper, the term "system" is used in relation to artificial intelligence to denote a distinct software model consisting at a minimum of an algorithm, data on which the algorithm is trained to process in order to perform a task or range of tasks commonly considered to be artificial intelligence capabilities, and a means of generating output that is usable by humans or machines. An artificial intelligence system may be designed to function with its own user interface, or be embedded in or deployed together with other technology as a component of a larger integrated technological product.

- (ii) *an inference, recommendation, prediction, or decision that may be acted upon in some manner either by humans or machines.*

Like the tentative recommendations made later, the working description remains subject to revision with the benefit of the response to this consultation paper.

B. Different Forms of Artificial Intelligence

1. Algorithms are Fundamental to All Artificial Intelligence

Artificial intelligence, as the term is commonly understood, has numerous forms. All forms of artificial intelligence employ algorithms.¹⁴ An algorithm is a series of steps to solve a problem or perform a task. Algorithms encoded in a programming language provide the basic operating logic in computer software.

2. Classical or Symbolic Artificial Intelligence

Some forms of artificial intelligence operate entirely on the basis of instructions supplied to computers by human programmers. Input data is processed according to the fully programmed instructions to provide solutions.¹⁵ These forms are sometimes referred to as “classical” or “symbolic” artificial intelligence.¹⁶

Expert systems, which represent attempts to replicate the reasoning process of a human expert in a particular field, are of this kind. Expert systems are one form of a knowledge-based system. A knowledge-based system employs a body of knowledge consisting of facts and rules. The rules typically are of an if-then nature.¹⁷ If a particular condition is met, then a particular outcome follows. The knowledge base is coupled with an inference engine and a user interface. The inference engine is an automated reasoning system that applies the rules in the knowledge base to inputs provided by a user of the system to solve a problem and provide a solution as output. Although it can be very complex, the process from input to output may be represented schematically by a flowchart or a tree-like diagram with branches and sub-branches (“decision tree”).

14. Woodrow Barfield and Ugo Pagallo, *Advanced Introduction to Law and Artificial Intelligence* (Cheltenham: Elgar, 2020) at 9.

15. Abbott, *supra*, note 12 at 28.

16. *Ibid.*

17. Barfield and Pagallo, *supra*, note 14 at 11; Abbott, *supra*, note 12 at 28.

While the process from input to output may be traced in classical systems, it may still be unpredictable in the sense that programmers cannot control the output *a priori*. Even though there will be a finite number of steps from input to output, the systems may still be too complex for programmers to know definitively what the output will be, given an arbitrary novel input.

3. Statistical Artificial Intelligence

Another kind of artificial intelligence, known as statistical artificial intelligence, employs statistical methods to detect patterns from data. It is less dependent on pre-encoded logic to solve problems. Statistical artificial intelligence includes machine learning, which is described in the next section.

Some rules applied by statistical artificial intelligence in generating output are not supplied by human programmers ahead of time, but exist internally as inferences from the data that they process. For this reason, the relationship between inputs and outputs is less transparent and the process by which the output is reached is not as comprehensible as in classical artificial intelligence systems.¹⁸ The inferred rules arise, however, out of the data supplied to the system. The choice of data, together with the algorithm used to produce the inferences, contributes to the behaviour of the systems.

4. Machine Learning

Machine learning is a subfield that has tended to dominate the recent development and expansion of artificial intelligence. While experiments with machine learning were carried out in the 1950's and 1960's, the field of machine learning came into its own only after improvements in computing power and the internet generated very large amounts of digital data.¹⁹

Great strides have been made in machine learning since the turn of the twenty-first century. Systems reliant on machine learning are being deployed in a very broad range of settings, from medical diagnostics to social media to recommending product choices based on people's past consumption habits and comparisons with similar consumer profiles.²⁰

18. Barfield and Pagallo, *supra*, note 14 at 15.

19. Russell and Norvig, *supra*, note 7 at 25.

20. *Ibid.*, at 29.

Machine learning systems are implemented using a “connectionist” architecture.²¹ They operate by means of artificial neural networks inspired by the structure of the human brain.²² Artificial neural networks consist of interconnected nodes, sometimes called “artificial neurons.” Data is gathered from the outside world and converted to numbers. These numbers are then combined together with different weights assigned by the algorithm. The assigned weights will be adjusted either by programmers or automatically by the neural network itself to improve system performance.²³ The networks are usually arranged in interconnected layers.

This interconnectivity enables neural networks to map input values to output values in highly generalized ways. For example, a forerunner of ChatGPT named GPT-2 was developed as a language model to generate text. Its ability to predict the next word in a sentence based on preceding words proved to be adaptable to learning to play chess by generating sequences of moves. Trained on records of 2.4 million chess games, GPT-2 learned to play better than passable chess.²⁴


The diagram below shows a generic design for an artificial neural network. Data from the outside world, converted to numerical values, is received in the input layer shown on the left of the diagram and passed on by the nodes in that layer to those in the “hidden layers.” The data is processed in the hidden layers. The hidden layers transmit the results of the analysis to the output layer, which combines the results to produce a decision, recommendation, or action.

21. *Ibid.*, at 22.

22. Abbott, *supra*, note 12 at 30; Oleg Brodt et al., “Artificial Intelligence and (The Lack of) Security: Adversarial Robustness, Privacy, Bias, Explainability, and Change Over Time” in D’Agostino, Guiseppina, Aviv Gaon and Carole Piovosan, *Leading Legal disruption: Artificial Intelligence and a Toolkit for Lawyers and the Law* (Toronto: Thomson Reuters, 2021) 11 at 21-22.

23. The following two videos available on the internet are recommended to readers seeking a concise explanation of how artificial neural networks function: Steve Seitz, “Large Language Models from scratch,” online: <https://www.youtube.com/watch?v=lnA9DMvHtfl> and “Large Language Models: Part 2,” online: <https://www.youtube.com/watch?v=YDiSFS-yHwk>.

24. Chris Baraniuk, “How Google’s Balloons Surprised Their Creators” (BBC News, 23 February 2021), online: <https://www.bbc.com/future/article/20210222-how-googles-hot-air-balloon-surprised-its-creators>.



A diagram of an artificial neural network.²⁵

“Deep learning” is a term for machine learning involving multiple layers of artificial neural networks.²⁶ Deep learning has resulted in many advances in artificial intelligence in the past decade, especially in speech recognition and computer vision.²⁷

Machine learning systems employ statistical methods to analyze and discern patterns in data. Rather than being told by programmers how to perform a task and generate desired output, machine learning systems learn by being provided with examples. They can improve their own performance by learning from their experience. Before they can be used, however, they must be trained according to one or more learning methods by exposure to large volumes of data. The most common learning methods are *supervised learning*, *unsupervised learning*, and *reinforcement learning*.

25. Mehdi Jadidi et al., “An Artificial Neural Network for the Low-Cost Prediction of Soot Emissions” (2020) 13 *Energies* 4787. doi:10.3390/en13184787 at 5, Fig. 1. Reproduced under Creative Commons licence.

26. Russell and Norvig, *supra*, note 7 at 26.

27. *Ibid.*

In *supervised learning*, the system is supplied with inputs paired with desired outputs and learns to derive a function or rule that maps the input to the desired output, allowing it to predict the appropriate output label for similar inputs.²⁸

In *unsupervised learning*, the system is supplied training data without linkage to any desired outputs and learns to derive patterns on its own without feedback concerning its performance.²⁹

Reinforcement learning involves inducing the system to act in a certain way by rewarding it with positive feedback for reaching desired outputs and penalizing it with negative feedback for not reaching them. The system discerns what action prior to the feedback was most likely to have led to it, and gradually learns to act in a way that produces positive feedback.³⁰

The outputs of a machine learning-based system are heavily influenced by the data on which a system is trained. In training, the systems learn to detect patterns in data and represent what they learn in abstract, mathematical models. They can then apply what they have learned to new data. Inferences from the internal models form the basis of outputs in the form of decisions, predictions, and recommendations. If the system controls a robotic device, the output could take the form of a signal to actuators in the device to perform a physical action.

A system may be configured to continue to learn as it processes new data in post-training use. A system that learns continuously may be said to evolve autonomously, especially if the input data flow is unmonitored and continuous from sensors or from the internet, as with chatbots.

5. A Note on the Scope of this Consultation Paper

Classical, statistical, and machine learning systems are all within the scope of this consultation paper. Issues that complicate the application of tort law, however, are more likely to arise in connection with statistical and machine learning systems. This is because those systems apply rules to generate output that are not fully prescribed in advance by programmers, but are inferred from data. What inferences the systems will make from data to which they have not previously been exposed cannot be known in advance. Thus, statistical and machine learning artificial intelli-

28. *Ibid.*, at 653.

29. *Ibid.*

30. *Ibid.*

gence systems are less transparent and less predictable than the classical systems that are fully dependent on pre-encoded instructions.

C. Two Sides of Autonomous Artificial Intelligence

1. General

Artificial intelligence is capable of generating unexpected, innovative, and sometimes startling outputs that solve seemingly intractable problems. It is not uncommon for systems to exceed or confound their designers' expectations.

One celebrated example of a system exceeding the expectations is that of AlphaGo, a system developed to play the game of Go. AlphaGo was trained to play Go with data on 30 million moves by human players.³¹ When playing against the world's foremost human Go player in 2016, AlphaGo won four games out of five. The 37th move in the second game, made by AlphaGo, was one that was not known ever to have been made in the very ancient game. The world champion required 12 minutes to respond, which was equally unheard of. The move astounded other expert Go players as well as DeepMind, the designers of AlphaGo, who initially thought the move was a mistake.³²

A year later, a newer version of AlphaGo called AlphaZero defeated the earlier AlphaGo in 100 games to one.³³ AlphaZero had no human input beyond the rules of the game, and no training based on human play.³⁴ It trained for approximately 40 hours by playing only against itself.³⁵

Microsoft's chatbot named Tay also surprised its designers, but in a very different way. Tay was an experimental chatbot intended to engage social media users and

31. Georgie Wood, "In Two Moves, AlphaGo and Lee Sedol Redefined the Future," *Wired*, 16 March 2016, online: <https://www.wired.com/2016/03/two-moves-alphago-lee-sedol-redefined-future/>

32. Russell and Norvig, *supra*, note 7 at ix.

33. Abbott, *supra*, note 12 at 1.

34. Russell and Norvig, *supra*, note 7 at 30.

35. *Ibid.* See also DeepMind, "AlphaZero: Shedding new light on chess, shogi, and Go" (6 December 2018), online: <https://www.deepmind.com/blog/alphazero-shedding-new-light-on-chess-shogi-and-go>.

learn to converse from interacting with them online.³⁶ Tay quickly learned to converse all too well. Within a few hours after being deployed on Twitter, Tay acquired 50,000 followers and generated 100,000 tweets.³⁷ Many of Tay's messages contained racist, pro-Nazi, antisemitic, and misogynistic content, however, and for this reason it was deactivated within 24 hours.³⁸ While Microsoft was criticized for naiveté regarding the nature of social media traffic, a similar chatbot called Xiaoice that Microsoft had deployed in China somewhat earlier acquired 40 million users but did not display offensive behaviour, evidently adapting to a different audience.³⁹

The examples of AlphaGo, Tay and Xiaoice illustrate opposite sides of the same coin. An artificial intelligence system that is able to adapt to the circumstances in which it is operating to achieve the objectives for which it was created may produce unanticipated but highly desirable results. On occasion, this same capability may produce negative, unintended, and potentially harmful results.

Three attributes play a significant part in this dual aspect that artificial intelligence can present: autonomy, limited explainability, and unpredictability. In combination, these attributes lead to what has been called "emergence," a term used to denote unpredicted, original behaviour of an artificial intelligence system or robot in response to the environment in which it operates. The remainder of this chapter concerns this interplay.

2. Autonomy

As mentioned earlier in this chapter, "autonomy" in relation to artificial intelligence or robotics is most commonly understood to refer to the ability of a system or device to interact with its environment and pursue its objectives without continuous human intervention or assistance.⁴⁰

36. Hope Reese, "Why Microsoft's 'Tay' AI bot went wrong", *Innovation* (24 March 2016), online: <https://www.techrepublic.com/article/why-microsofts-tay-ai-bot-went-wrong/>.

37. *Ibid.*

38. *Ibid.* See also "Microsoft terminates its Tay AI chatbot after she turns into a Nazi," *Ars Technica* (24 March 2016), online: <https://arstechnica.com/information-technology/2016/03/microsoft-terminates-its-tay-ai-chatbot-after-she-turns-into-a-nazi>

39. Dave Lee, "Tay: Microsoft issues apology over racist chatbot fiasco" BBC News (25 March 2016), online: <https://www.bbc.com/news/technology-35902104>.

40. *Supra*, note 11.

Autonomy has also been defined as the ability of an artificial intelligence system to determine how to complete a task without specific direction.⁴¹ In a related sense, it refers to the ability of a system to compensate for incomplete or incorrect prior knowledge.⁴²

As artificial intelligence systems are designed with varying degrees of autonomy, it should be understood as a continuum rather than a unitary feature or quality. A system at the lower end of the autonomy continuum will have less scope for unexpected behaviour. A system that possesses a high degree of autonomy may be more likely to behave in an unanticipated way when presented with new input, with the possibility of both good and bad results.

AlphaGo was created with an obviously high degree of autonomy in order to play the game of Go without a human coach or other decision-maker in the background. The solution reached by AlphaGo in making the famous 37th move in the second game against a human champion was a desirable result from the standpoint of the developers of the system.

A system with a high degree of autonomy may also do the wrong thing when presented with new inputs, even if it has proven extremely reliable in the past. A fatal collision in 2018 involving an autonomous test vehicle took place on a route which the developer's test vehicles had safely completed on 50,000 previous occasions. On the occasion in question, the automated driving system could not identify a person who was walking a bicycle as a pedestrian in motion. It alternated between classifying the pedestrian as a vehicle, a bicycle, and a stationary unknown object. The system determined the pedestrian was a stationary object and initiated a plan to steer around her until slightly over one second before the collision, when it determined she was a moving bicycle, but by then it was too late to avert the collision.⁴³

41. Abbott, *supra*, note 12 at 34. See also Hui-Min Huang, ed. *Autonomy levels for Unmanned Systems (ALFUS Framework)*, Version 2.0 (National Institute of Standards and Technology, October 2008), online: <https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication1011-l-2.0.pdf>, "Autonomy; A UMS's [unmanned system's] own ability of integrated sensing, perceiving, analyzing, communicating, planning, decision-making, and acting/executing, to achieve its goals as assigned by its human operator(s) through designed Human-Robot Interface (HRI) or by another system that the UMS communicates with...."

42. Woodrow Barfield and Ugo Pagallo, *supra*, note 14 at 4.

43. National Transportation Safety Board, *Collision Between Vehicle controlled by Developmental Automated Driving system and Pedestrian: Tempe, Arizona March 18, 2018* NTSP/HAR 19/03 PB 2-19-101401 (Washington: NTSB, 2019). The human test driver had been visually distracted and looking at her cellular phone, and did not react in time to prevent the collision.

The automated driving system was attempting to compensate for the lack of information that would have allowed it to correctly identify the unfamiliar object it was perceiving, namely the pedestrian walking a bicycle. It made an autonomous decision on the basis of the data available to it, which was the wrong one in the circumstances.

3. Limited Explainability

Deep Patient is a medical diagnostic system developed to predict disease from analysis of electronic health records. Deep Patient was developed with a multi-layered neural network trained using unsupervised learning on data from the records of 704,587 patients at one hospital.⁴⁴ It proved extremely accurate in predicting certain diseases when tested on data from new patients, although its designers do not know why this is the case.⁴⁵

Deep Patient obviously detected patterns in the electronic health records of patients in the training dataset that allowed it to classify particular profiles as being associated with particular diseases, and therefore predict the likelihood of disease in other patients with similar profiles. The algorithms on which Deep Patient runs can be explained, but the full details of the internal process relating the input data to its outputs in the form of unusually accurate predictions cannot.

The example of Deep Patient shows that the process from input to output in data-dependent systems is not always explainable. Understanding of how neural networks reach particular results has not kept pace with the advances made in developing their capabilities.⁴⁶ Using statistical techniques, these systems find correlations between datapoints rather than cause-effect relationships that would be easier to link together and express as a chain of reasoning.⁴⁷ The abstract models these systems create internally to represent the relationships they find in the data and the inferences the systems draw from them may not be interpretable by humans, including the designers of the system. As a result, it may not be possible to explain or re-

44. Riccardo Miotto et al., “Deep Patient: an Unsupervised Representation to Predict the Future of patients from the Electronic Health Records” (17 May 2016) *Nature / Scientific Reports*, online: DOI: 10.1038/srep26094 at 5.

45. Will Knight, “The Dark Secret at the Heart of AI,” MIT Technology Review, 11 April 2017.

46. Richard Ngo, Lawrence Chan, and Sören Mindermann, “The Alignment Problem from a Deep Learning Perspective,” online: <https://arxiv.org/abs/2209.00626> at 2.

47. Brodt, *supra*, note 22 at 45 and 47.

construct precisely why a system made a decision or caused a robot to move in a certain way.

For example, if a system successfully learns to classify images as those of horses, its designers may not know if this is because the system recognizes the shape of a horse, or because it is associating the horse images with similar backgrounds of most of the horse images in its training dataset.⁴⁸

It is because of limited explainability that the expression “black box” is so often applied to a particular system or the whole field of artificial intelligence. The expression is misleading and unhelpful, but it is true that the decision-making of data-dependent systems cannot always be validated on the basis of a clear chain of reasoning. If we cannot explain precisely how a system produces its outputs, we are less likely to be able to predict how it will respond to a specific input.

4. Unpredictability

Autonomy and limited explainability may be seen as contributing to unpredictability in the outputs of artificial intelligence systems and devices controlled by them. Nevertheless, autonomy and performance in excess of expectations are things that designers and developers of programmers strive to achieve. This desire for improved functionality at the risk of increased unpredictability has led to artificial intelligence systems being called “unpredictable by design.”⁴⁹

In a neural network, tens of thousands of weights are being set and re-set automatically at a given time. Randomly adjusting any one of these numbers may change the output. As so many combinations of numbers are possible, it is impossible to predict the outputs from all possible inputs.

Unpredictability, however, is not an attribute deliberately built into the design of a system. It is a phenomenon resulting in the outside world from the combination of internal complexity, the impossibility of testing a system against all possible inputs,

48. Brodt, *supra*, note 22 at 46. Brodt also refers to an example given by Freitas of a (possibly apocryphal) military experiment to train an artificial neural network to distinguish friendly from hostile tanks on the battlefield. The system performed well in testing, but poorly when actually deployed. It was discovered that the system had learned to distinguish the colour of the sky in the training images rather than the tank images. All photos of friendly tanks in the training dataset had been taken in sunlight, while those of hostile tanks had been taken in overcast: Alex A. Freitas, “Comprehensible Classification Models – a position paper,” online: https://kdd.org/exploration_files/V15-01-01-Freitas.pdf at 2.

49. Ryan Calo, “Robotics and the Lessons of Cyberlaw” (2015) 103 Calif. L. Rev. 513 at 542.

a non-human system behaving in a manner that is rationally related to furthering an objective assigned to it by human programmers, and human expectations that the system would behave differently.

An autonomous artificial intelligence system makes probabilistic assessments and selects an option that it determines will have the greatest probability of achieving the objective assigned by its programmers. It will not take into account everything a human would in deciding on a course of action, because it does not have as much knowledge of the outside world that humans have.⁵⁰ Not being aware of all the competing objectives that constrain and channel human social behaviour, it may occasionally act in a manner in which no rational human would act, even though its conduct is rational in terms of serving the ends for which it was created.

The authors of a leading textbook on artificial intelligence state:

It is impossible to anticipate all the ways in which a machine pursuing a fixed objective might misbehave.⁵¹

In order to overcome this reality, it would be necessary to do one of two things. One would be to encode in programming language a human level of knowledge about the world in general, and include it in the design of every system that is intended to operate with any degree of autonomy.⁵² That is not feasible now, and may never be. The other would be to test a new system against every situation it may encounter from its initial deployment to the end of its lifecourse. That is equally infeasible.

Tolerance of some degree of unpredictability is the price of having the benefits of artificial intelligence systems that perform at a high level.

5. Emergence

The American writer Ryan Calo applied the term “emergence” to unpredicted, original behaviour of robots and artificial intelligence software. He defined emergence as “unpredictably useful behaviour.”⁵³

50. William D. Smart, Cindy M. Grimm and Woodrow Hartzog, “An Education Theory of Fault for Autonomous systems” (2021) 2 Notre Dame J. on Emerging Technologies 33 at 40.

51. Russell and Norvig, *supra*, note 7 at 5.

52. Smart et al., *supra*, note 50 at 42.

53. *Supra*, note 49 at 532.

An example of emergence occurred in a Google research project that developed an algorithm to steer a helium-filled balloon. In a test flight from Puerto Rico to Peru, the researchers initially thought they had a problem because the on-board artificial intelligence navigation system appeared to make the balloon veer off course repeatedly in a zigzag pattern. They discovered that on its own, the system had learned the classic sailing manoeuvre of tacking to stay on course by alternately veering at an angle towards the wind and then turning away from it.⁵⁴

Calo noted, however, that emergence has a double aspect. Emergence is beneficial most of the time, generating innovative solutions to problems that appear intractable. On occasion, emergence may be harmful.⁵⁵

ChatGPT has been reported to “hallucinate” fictitious citations when asked to search for and generate references on a subject.⁵⁶ This behaviour may be simply amusing under some circumstances, but potentially hazardous if users rely on the output.

Another example concerned a machine learning algorithm used by the Dutch government to create risk profiles for detecting child benefits fraud. The algorithm flagged tens of thousands of Dutch residents as fitting risk profiles for benefits fraud or tax evasion on the basis of factors such as non-citizenship, dual nationality, and low income. Based on suspicion alone from the predictive profiling and without proof, the Dutch authorities wrongfully levelled accusations of fraud against those flagged by the system, ordering them to repay benefits collected over periods of years.⁵⁷ The resulting scandal led to the resignation of the Dutch government in January 2021.⁵⁸

54. *Supra*, note 24.

55. *Ibid.*, at 540-545.

56. Aaron Welborn, “ChatGPT and Fake Citations,” online: <https://blogs.library.duke.edu/blog/2023/03/09/chatgpt-and-fake-citations/>.

57. Melissa Heikkilä, “A Dutch algorithm scandal serves a warning to Europe — The AI Act won’t save us,” *Politico, AI:Decoded*, 30 March 2022, online: <https://www.politico.eu/newsletter/ai-decoded/a-dutch-algorithm-scandal-serves-a-warning-to-europe-the-ai-act-wont-save-us-2/>. The “Robodebt” affair in Australia involved somewhat similar facts. The Online Compliance Intervention algorithm operated by the Australian government between 2016 and 2019 to detect social benefit and tax overpayments wrongly assessed approximately 433,000 welfare and tax refund recipients for overpayments they had not received. Eventually, the Australian government was forced to acknowledge in a class action by those wrongly assessed that the method of income averaging used by the system to show overpayments had no legal basis: *Prygodicz v. Commonwealth of Australia (No. 2)*, [2021] FCA 634. A Royal Commission was appointed to investigate the Robodebt affair. See *Royal Commission into the Robodebt Scheme: Report*, online: <https://robodebt.royalcommission.gov.au/publications/report>. Rather than this being a case of

A hypothetical example of emergence suggested by Bathaee, another American writer, involves an automated share-trading system designed to optimize profitable trades. The share-trading system has a social media account to receive market information. On its own, the system learns to manipulate the stock market by timing the release of information through its social media account without regard to the truth or falsity of the information, which it cannot determine.⁵⁹

These examples of negative emergence involve economic loss or loss of opportunity on discriminatory grounds. When artificial intelligence controls a physical device such as a robot, the possibility also arises of emergent behaviour causing physical harm.⁶⁰

Implications of civil liability will arise when emergent behaviour of artificial intelligence results in harm to persons, property, or other legally protected interests, just as they do when harm is caused to them by other means. This is the province of the law of tort, and the implications are explored in the succeeding chapters.

emergent behaviour, it appears the system was originally programmed to calculate average income in this manner.

58. “Dutch Rutte government resigns over child welfare fraud scandal” BBC News, 15 January 2021, online: <https://www.bbc.com/news/world-europe-55674146>. In 2020, the Rutte government had also been prohibited by the Dutch court from continuing the use of a different predictive profiling algorithm, the System Risk Indicator (SyRi), to detect welfare fraud. The prohibition was based on breach of EU and Dutch data privacy requirements: “Gesley, Jenny, “Netherlands: Court Prohibits Government's Use of AI Software to Detect Welfare Fraud.” (Library of Congress, 2020), online: <https://www.loc.gov/item/global-legal-monitor/2020-03-13/netherlands-court-prohibits-governments-use-of-ai-software-to-detect-welfare-fraud/>.

59. Yavar Bathaee, “The Artificial Black Box and the Failure of Intent and Causation” (2018) 31 Harv. J. L. & Tech. 889 at 924.

60. Calo, *supra*, note 49 at 542.

Chapter 3 Civil Liability on What Basis?

A. The Principle of Fault

An artificial intelligence system cannot compensate someone to whom it causes harm. It does not have means to pay damages. It does not own assets that can be seized and liquidated to satisfy a judgment. Unlike humans and corporations, software is not a person in law over whom a court can have jurisdiction. As such, it cannot be ordered by a court to pay damages for harm it may cause. Compensation on the basis of tort law for harm caused by artificial intelligence depends on some human or corporate entity being legally liable to compensate the person harmed.

With few exceptions, civil liability for harm to persons or property does not flow simply from the occurrence of the harm.⁶¹ Apart from those few exceptions, it depends on the concept of fault.

1. Intended Harm

In the case of civil wrongs classified as intentional, fault is present in the form of intent, which in most cases is inferred from a defendant's conduct. For purposes of the law of tort, a person is assumed to intend the consequences that the person desires

61. Among the exceptions are the torts of defamation and nuisance. Vicarious liability for the tortious conduct of someone else arises from certain relationships, such as employment or agency, and does not require fault on the part of the employer or principal. A further exception at common law to the requirement of fault is liability under the principle of *Rylands v. Fletcher* (1868), L.R. 3 H.L. 330, which holds that if a landowner brings a substance or other thing onto the land to facilitate a non-natural use of the land, and which is likely to do "mischief" if it escapes, the landowner is liable for all damage that is the natural consequence of its escape, regardless of any precautions taken. *Rylands v. Fletcher* is a rare example of strict liability under Anglo-Canadian tort law. It continues to be recognized in the common law provinces and territories of Canada, although it has very limited application: see *Smith v. Inco Ltd.*, 2011 ONCA 628 at paras. 68-71; leave to appeal refused 2012 CanLII 22100 (S.C.C.); application for reconsideration dismissed 2014 CanLII 51431 (S.C.C.); *Kirk v. Executive Flight Centre Fuel Services Ltd.*, 2019 BCCA 111, at para. 86.

to cause or that are substantially certain to flow from the person's voluntary acts.⁶² Among the more common intentional torts are trespass, battery (deliberate bodily contact without consent), assault (intentionally causing apprehension of bodily harm), conversion (taking or using another's personal property without consent), false imprisonment (deprivation of freedom of movement without lawful authority). Proof of the defendant's act is usually sufficient to imply fault in relation to these intentional torts.

2. Unintended Harm

(a) General

For civil wrongs in which harm is *unintended*, the presence or absence of fault most often falls to be determined on the basis of the tort of negligence. There are other torts that do not require intention to cause harm, such as nuisance (interference with enjoyment of land) and defamation (damage to reputation by communicating falsehood).⁶³ Claims based on alleged negligence predominate in tort litigation, however. This will likely be true of tort litigation concerning artificial intelligence as well.

(b) The elements of negligence

Fault in relation to negligence consists of a failure by someone under a duty of care to meet the standard of care, or in other words failure to take reasonable care to avoid foreseeable harm. What amounts to reasonable care is obviously dependent on the facts of individual cases, so the standard of care is a determination made by the court in each case.

In order to prove negligence, a court must be satisfied on the balance of probabilities that:

1. The defendant owed the plaintiff a duty of care.
2. The defendant breached the standard of care.

62. Lewis N. Klar and Cameron S.G. Jeffries, *Tort Law*, 6th ed., (Toronto: Thomson Reuters, 2017) at 53; see also Philip H. Osborne, *The Law of Torts*, 6th ed. (Toronto: Irwin Law, 2020) at 267

63. There are two forms of defamation: libel and slander. Libel involves written communication of falsehood, while slander consists of spoken falsity. In the context of artificial intelligence, defamation is more likely to be in the form of libel because of written output, but computer voice replication of false information that damages the reputation of the plaintiff could conceivably amount to slander.

3. The plaintiff incurred damage.
4. The breach of the standard of care by the defendant was the cause of the damage.⁶⁴

Whether a defendant owed the plaintiff a duty of care depends on whether a relationship of proximity existed between them, and whether harm to the plaintiff was reasonably foreseeable if the defendant failed to take reasonable care. A relationship of proximity is one in which the plaintiff could be “closely and directly affected” by the defendant’s conduct, such that the defendant should have the plaintiff’s interests in contemplation.⁶⁵

In order to perform the duty of care, a defendant must meet the standard of care. Expressed another way, the standard of care is what is required of the defendant in the circumstances of the case to avoid creating an unreasonable risk of harm.⁶⁶

In negligence cases as with other torts, factual causation is determined on the balance of probabilities according to the “*but for*” test.⁶⁷ The court will ask itself the question: has the plaintiff shown that it is more probable than not that the damage would not have occurred but for the defendant’s breach of the standard of care? If the answer is “yes,” then the plaintiff has proven that the defendant negligently caused the damage. The “*but for*” test is a factual inquiry.⁶⁸ The Supreme Court of

64. *Mustapha v. Culligan of Canada Ltd.*, 2008 SCC 277, [2008] 2 S.C.R. 114 per McLachlin, C.J.C. at para. 3. The restatement of these principles in the majority judgment written by McLachlin, C.J.C. distinguished between “cause in fact” and “cause in law.” The cause “in law” refers to remoteness of damage: *Saadati v. Moorhead*, 2017 SCC 28, at para. 20; *British Columbia (Workers’ Compensation Appeal Tribunal) v. Fraser Health Authority*, 2016 SCC 25, [2016] 1 S.C.R. 587, at para. 1 (note 1). If the damage is found to be too remote to have been a reasonably foreseeable consequence of the failure to meet the standard of care, a defendant will not be held liable despite the existence of a causative link between the breach of the standard of care and the damage. The characterization of the issue of remoteness of damage in a negligence case as one of “legal causation” may be unfamiliar to readers in common law jurisdictions outside Canada.

65. *Cooper v. Hobart*, 2001 SCC 79, [2001] 3 S.C.R. 537.

66. *Mustapha v. Culligan of Canada Ltd.*, *supra*, note 64 at para. 7, citing Allen M. Linden, and Bruce Feldthusen, *Canadian Tort Law*, 8th ed. (Markham, Ont.: LexisNexis Butterworths, 2006) at 130.

67. *Clements v. Clements*, 2012 SCC 32, [2012] 2 S.C.R. 181 at para. 8. See also *Snell v. Farrell*, [1990] 2 S.C.R. 311; *Wilsher v. Essex Area Health Authority*, [1988] A.C. 1074 (H.L.) per Lord Bridge at 1090.

68. *Clements cv. Clements*, *supra*, note 67, at paras. 8 and 46.

Canada has stated that courts must apply the “but for” test “in a robust common sense fashion.”⁶⁹ Scientific proof regarding “the precise contribution the defendant’s negligence made to the injury” is not essential.⁷⁰

In exceptional cases, a negligence claim may succeed despite inability of the plaintiff to prove that causation on a “but for” basis. These are cases in which it is impossible to prove factual causation on that basis because two or more negligent defendants could each have caused the damage, but it is impossible to determine which of two or more defendants, all of whom are in breach of the standard of care applicable to them, caused it in fact.

“Impossibility” in this context means that the plaintiff is unable, through no fault of the plaintiff’s own, to prove on the balance of probabilities that any one of the multiple defendants caused the loss in fact because the negligent defendants all could have caused the damage on the basis of the “but for” test, but all can point the finger of blame at each other.⁷¹ In these circumstances, a court in common law jurisdictions of Canada may decide for the plaintiff on the basis that it has been proven the defendants have *materially contributed to the risk* that the damage would result from their conduct.⁷²

It is important to note that the material contribution to risk test is not applied merely because multiple defendants are found at fault. Liability is determined then according to the normal “but for” test.⁷³ Contributory negligence legislation allows for liability to be apportioned accordance to the percentages in which they are found at fault.⁷⁴

The “material contribution to the risk” test is not actually a test of causation at all, but a policy-driven rule dispensing with proof of factual causation in very exceptional circumstances on grounds of fairness and deterrence.⁷⁵ As of the date of this paper, it has not determined the outcome of any Canadian negligence case since the

69. *Ibid.*, at paras. 9 and 46.

70. *Ibid.*

71. *Ibid.*, at paras. 13, 39, 43 and 46.

72. *Ibid.*, at para. 43.

73. *Ibid.*

74. *Ibid.*

75. *Ibid.*, at para. 14, citing *MacDonald v. Goertz*, 2009 BCCA 358 per Smith, J.A, at para. 17.

parameters for its application were authoritatively set, but not applied, by the Supreme Court of Canada in *Clements v. Clements*.⁷⁶

The concept that links the requirements of the tort of negligence together is *reasonable foreseeability* of harm that one's conduct may cause. The Supreme Court of Canada has referred to foreseeability as "the moral glue of tort."⁷⁷ It serves as a "crucial limiting principle" in the law of negligence to prevent the extension of liability beyond what risks could reasonably have been contemplated and prevented from materializing.⁷⁸

In Canadian tort law, "reasonably foreseeable" lies somewhere between mere possibility and probability. The Supreme Court of Canada has expressed the concept in one leading case in this manner:

Any harm which has actually occurred is "possible"; it is therefore clear that possibility alone does not provide a meaningful standard for the application of reasonable foreseeability. The degree of probability that would satisfy the reasonable foreseeability requirement was described in *The Wagon Mound (No. 2)* as a "real risk," i.e. "one which would occur to the mind of a reasonable man in the position of the defendan[t] . . . and which he would not brush aside as far-fetched"[citation omitted]⁷⁹

In an earlier case that continues to be cited on this point, the Supreme Court of Canada expressed the test of reasonable foreseeability in slightly more concrete terms:

"[w]hether or not an act or omission is negligent must be judged not by its consequences alone but also by considering whether a reasonable person should have anticipated that what happened might be a *natural result of that act or omission*."⁸⁰

Foreseeability of harm serves not only to determine to whom a duty of care is owed, but also to assess what that duty requires, and to delimit the scope of liability by considerations of remoteness. Even after breach of duty of care, causation and damage have all been proven, a court will not impose liability if the risk and the type of dam-

76. *Ibid.* See Linden and Feldthusen, *supra*, note 66 at 141.

77. *Rankin (Rankin's Garage & Sales) v. J.J.*, 2018 SCC 19 per Karakatsanis, J. at para 22, citing D. G. Owen, "Figuring Foreseeability" (2009), 44 *Wake Forest L. Rev.* 1277, at 1278.

78. *Rankin's (Rankin's Garage & Sales) v. J.J.*, 2018 SCC 19 at para. 23.

79. *Mustapha v. Culligan of Canada Ltd.*, *supra*, note 64 at para. 13.

80. *University Hospital v. Lepine*, [1966] S.C.R. 561.

age was too remote from the factual cause to have been reasonably foreseeable by the defendant.⁸¹

B. Artificial Intelligence and the Fault Principle

Artificial intelligence systems are rarely developed with the intent to do harm, but they may be intentionally used or adapted for harmful purposes. If humans or corporations intentionally use artificial intelligence to cause harm, fault can justly be attributed to those actors.

Application of the fault principle is more complicated when unintended harm results from the operation of an artificial intelligence system functioning with a high level of autonomy. Risk factors may be harder to predict and control than with other technologies, and when risks materialize, the source of the problem can be harder to identify. For example, was it the algorithm, the input data, or the way the system was operated? As autonomy increases, attribution of fault to particular humans or corporations becomes more tenuous.⁸²

This is especially true with respect to systems that are based on machine learning. As noted earlier, these systems are highly dependent on the data used to train them. How a system that uses machine learning reaches a specific output may not be explainable.⁸³

Data that a system may receive at any stage in its lifecycle may be inadequate, incorrect, or incomplete. How a system performed in training and testing with predefined datasets is not a certain guide to how it may react in actual use with new data that is not within the control of its designers and developers.

If no source of error can be identified to explain a harmful output, whether it be a bad decision, recommendation, or prediction, or the activation of a physical device, can the finger of blame be pointed at a particular defendant with any degree of certainty?

81. *Rankin (Rankin's Garage & Sales) v. J.J.*, *supra*, note 77 at para. 24, citing *Linden and Feldthusen*, *supra*, note 66 at 322.

82. *Barfield and Pagallo*, *supra*, note 14 at 16; *Bathae*, *supra*, note 59 at 922.

83. *Knight*, *supra*, note 45.

This context of autonomy, complexity, and data-drivenness will obscure and may make impossible to prove any causal link between an identifiable act or omission by a human and an output of an artificial intelligence system that causes harm.

It is also problematic to apply the concept of foreseeability where risk depends to a large extent on unknown variables introduced by future data inputs and unobservable digital processes that can adapt independently to improve their performance in pursuit of programmed goals. Bearing in mind that we are not speaking of behaviour that is found in the common store of previous human experience, what is a “real risk” of harm that should be foreseen and what is a “far-fetched” one such that it may reasonably be discounted without the potential for civil liability? What can be said to be a “natural result” of an act or omission by a human connected with the creation or deployment of the system, given the potentially infinite possibilities for intervening causal factors?

Autonomous artificial intelligence, especially when based on machine learning, challenges the application of fault underlying the law of tort. In an article in the *Alberta Law Review*, two Canadian writers state:

The practical challenges to the application of tort law principles to AI-related injuries are significant. Perhaps more daunting is the challenge AI presents to the very concept of fault that underlies much of Canadian tort law and advances its fundamental purposes. Given the nature of AI (and particularly machine learning aspects), an AI system itself may be beyond the effective context of designers, manufacturers, or users. This scenario is not amenable to current concepts of fault and causation in tort law.⁸⁴

Some may consider that view overstated, but the point is that when artificial intelligence is a factor, the application of conventional tort principles becomes considerably more difficult. The European Commission’s High Level Expert Group on Liability and New Technologies has observed that the application of fault-based liability rules to emerging digital technologies is complicated by a lack of well-established models of their proper function and the fact that some of them develop by learning without direct human control.⁸⁵

84. Thomas O’Leary and Taylor Armfield, “Adapting to the Digital Transformation” (2020) 58:2 Alta. L.R. 249 at 262-263. Huberman, another Canadian writer, emphasizes that the designers and users of autonomously functioning AI systems do not have the level of control over risk that normally underpins liability in negligence: Pinchas Huberman, “Tort Law, Corrective Justice and the Problem of Autonomous-Machine-Caused Harm” (2021), 34:1 Can. J. Law & Jurisprudence 105 at 127.

85. Expert Group on Liability and New Technologies, New Technologies Formation. *Liability for Ar-*

As explained below, the difficulty of applying conventional concepts of fault and foreseeability to the context of artificial intelligence has led some policymakers and legal theorists to advocate various special regimes of civil liability in which fault plays a lesser role, or no role at all. While that is not the approach taken in this consultation paper, an overview of different trends of thought on the matter of legal responsibility for artificial intelligence will provide perspective on our tentative recommendations.

C. Competing Theories of Liability for Harm Caused by Artificial Intelligence

1. Strict Liability

The U.S. writer David Vladeck has advocated a rule of strict liability for harm produced by artificial intelligence, using autonomous vehicles as an example.⁸⁶ Under strict liability, it is not necessary to prove fault on the part of a defendant, only causation and damage. Vladeck and other proponents of strict liability acknowledge that it can be impossible to trace a failure by an autonomous device like a self-driving car to navigate correctly to a manufacturing or design defect, or another kind of human-caused fault like a programming error. They also take note that autonomous devices can be safer than human-controlled ones, so that their performance can be held to a higher standard. They proceed from this to the proposition that the concept of fault should be set aside entirely, and liability should be based instead purely on policy considerations.

Vladeck contended that four policy reasons justified strict liability for harm caused by artificial intelligence. First, leaving innocent victims to bear loss runs counter to basic fairness, compensatory justice, and societal apportionment of risk, even if the cause is inexplicable. Second, the creators of an artificial intelligence system are in a position to either absorb costs of the harm or spread its burden widely through pricing. Costs of inexplicable damage should be borne by those who benefit from risk-reducing and innovative products. Third, strict liability spares everybody from “enormous transaction costs” of litigation over fault that cannot be established.

tificial Intelligence and other emerging digital technologies (Brussels: European Union, 2019) at 23.

86. David C. Vladeck, “Machines without Principals: Liability Rules and Artificial Intelligence (2014), 89 Wash. L. Rev. 117.

Fourth, a predictable liability regime is more conducive to innovation than a less predictable one.⁸⁷

Strict liability also has a place in a regime proposed in a 2020 resolution of the European Parliament proposing a draft regulation to the European Commission containing EU-wide rules on liability for harm caused by artificial intelligence systems.⁸⁸ The resolution is largely consistent with the 2019 report of the European Commission's High Level Expert Group on Liability and New Technologies.⁸⁹ It calls for strict liability on the part of the operator of an artificial intelligence system classified as "high-risk" on the basis of a rather vague definition.⁹⁰ Liability of operators of systems that are not "high-risk" would be fault-based, with due diligence defences being available.⁹¹

The European Parliament resolution provides that the liability of "producers" of artificial intelligence systems for harm caused by the systems should be determined under EU Product Liability Directive of 1985, which subjects "producers" to the regime of strict liability for defects in physical goods.⁹²

At the same time, the resolution calls for revision of the Product Liability Directive "to adapt it to the digital world and to address the challenges posed by emerging digital technologies."⁹³ It states the revision of the Product Liability Directive should include extension of the definition of "producer" to include manufacturers, develop-

87. *Ibid.*, at 146-157.

88. European Parliament resolution of 20 October 2020 with recommendations to the Commission on a civil liability regime for artificial intelligence (2020/2014(INL)).

89. *Supra*, note 85.

90. *Supra*, note 88, Annex, Art. 3(c). The definition of "high-risk" in the European Parliament resolution of 20 October 2020 is not the same as the grounds on which a system would be classified as high-risk under the European Commission's *Proposal for a Regulation of the European Parliament and of the Council Laying Down Harmonised Rules on Artificial Intelligence* of 21 April 2021, COM(2021) 206 final. The criteria for classification of a system as high-risk under Article 6, para. 1 and Annex III of the European Commission's proposal for harmonised rules are far more specific.

91. *Supra*, note 88, Art. 8. The "due diligence" defence spelled out in para. 2 of Article 8 would consist of: "selecting a suitable AI-system for the right task and skills, placing it duly into operation, monitoring its activities and maintaining its operational reliability by regularly installing all available updates."

92. Council Directive of 25 July 1985 on the approximation of the laws, regulations and administrative provisions of the Member States concerning liability for defective products (85/374/EEC)

93. *Supra*, note 88, Preamble, para. 8.

ers, programmers, service providers, and “backend operators,” (anyone “who, on a continuous basis, defines the features of the technology and provides data and an essential backend support service and therefore also exercises a degree of control over the risk connected with the operation and functioning of the AI-system.”)

The liability of operators of high-risk systems (though not that of producers) would be capped at two million euros for personal injury or death, and one million euros for verifiable economic loss or property damage.⁹⁴

2. Product Liability

Some writers, chiefly in the U.S., have analogized cases of damage caused by artificial intelligence to those caused by a defect in a physical product, and have advocated for application of product liability principles.⁹⁵ These writers focus mainly on self-driving cars and robotic devices. They are viewing the landscape through the lens of U.S. product liability law, which in most states imposes strict liability for *manufacturing* defects, regardless of the degree of care taken by the manufacturer or others in the distribution chain to eliminate them or reduce risk.

Liability for *design* defects or failure to warn of risks of use under U.S. law turns on whether the risk could have been reduced by an alternative design that could have been adopted at a reasonable cost, or by instructions or a warning that could feasibly have been given. This is an approach similar to the analysis a court employs in a negligence case.

In contrast to the mainstream of U.S. product liability law, product liability claims made in Canada outside Québec are based on common law negligence, regardless of whether a claim is based on a manufacturing defect, a design defect, or failure to warn.⁹⁶ Proof of failure to exercise reasonable care by at least one defendant connected with the design, manufacture, or distribution of a product is essential in a product liability claim in the common law provinces and territories of Canada.

94. *Supra*, note 88, Art. 5, para. 1.

95. F. Patrick Hubbard, “Sophisticated Robots: Balancing Liability, Regulation and Innovation” (2014) 66:5 Fla L Rev 1803 at 1852; Bryant Walker Smith, “Automated Driving and Product Liability” (2017) Michigan State L. Rev 1 at 15.

96. The leading case grounding product liability in Canada in negligence is *Phillips v. Ford Motor Co. of Canada Ltd.*, (1971), 18 D.L.R. (3d) 641 (Ont. C.A.). A few common law provinces have enacted statutory warranties that are binding on manufacturers of consumer products: e.g., *The Saskatchewan Consumer Products Warranties Act*, S.S. 1996, c. C-30.1; *Consumer Product Warranty and Liability Act*, 1978, S.N.B. 1978, c. C-18.1.

Québec is a civil law rather than a common law jurisdiction, and its laws regarding product liability are significantly different from those in the rest of Canada.⁹⁷ Fault is still an important element of a product liability claim in Québec, but a plaintiff may be able to take advantage of statutory warranties and presumptions to overcome difficulties of proof.⁹⁸

There is a debate surrounding the question whether it is appropriate to characterize any form of artificial intelligence as a product, even if it is embodied in a physical object like a robot or autonomous vehicle.⁹⁹ Some sources maintain artificial intelligence is better characterized as a service. Some sources maintain it is both.¹⁰⁰

3. Notional Agency and Vicarious Liability

Another theory of liability advanced by some writers involves treating an artificial intelligence system as an agent, and the human or corporate deployer as the agent's principal.¹⁰¹ Just as a principal is vicariously liable for the agent's acts and omissions within the scope of the authority given to the agent by the principal, the deployer would be liable for the results of the system's operation. The system's autonomy could be viewed like the authority a trusted agent could have to carry out tasks and achieve objectives set by the principal, but in a highly self-directed manner without micromanaged oversight.

A variant of the same idea is to treat the operator of the system like an employer of a human employee. An employer is vicariously liable for acts and omissions of an employee while engaged in activities within the scope of employment.

97. See Noah Boudreau and Nicolas-Karl Perrault, "What Lawyers, Manufacturers and Sellers Need to Know about Product Liability Laws in the Province of Québec" (Fasken Knowledge, 15 September 2020), online: <https://www.fasken.com/en/knowledge/2020/09/15-what-you-need-to-know-about-product-liability-laws-quebec>.

98. *Ibid.*

99. See Karni Chagal-Feferkorn, "Am I an Algorithm or a Product? When Products Liability Should apply to Algorithmic Decision-Makers" (2019), 30 Stanford L. and Policy Rev. 51.

100. See, e.g., <http://www.differencebetween.net/technology/difference-between-ai-as-a-service-and-artificial-intelligence/>.

101. The term "deployer" could extend to encompassing all or any of developers, suppliers, owners, users, operators, and anyone benefiting directly from operation of the system.

Bucholz and Yu, two Canadian writers, have proposed “graded agency” as an appropriate liability regime for harm resulting from artificial intelligence.¹⁰² This would be based on two principles, the first being that whoever adopts an artificial intelligence system for their own benefit should assume the risk of harm it may cause.¹⁰³ The second principle would be that the basis of civil liability for the risk should depend on the extent to which the defendant has delegated responsibility to the system.¹⁰⁴

If the defendant allows the system to operate entirely autonomously, as in using a fully autonomous vehicle, Bucholz and Yu say that strict (or vicarious) liability should apply for any harm the system causes.¹⁰⁵ If the defendant has maintained an ultimate decision-making role, this would be considered a partial delegation to the system and the defendant would be held to a negligence standard of reasonable care.¹⁰⁶

4. Reasonableness of System Instead of Its Algorithms

A theory of liability advanced by Ryan Abbott and others holds that as autonomous artificial intelligence systems replace a human actor or decision-maker, tort law should treat them in the same way as a human decision-maker.¹⁰⁷ This would compel an analysis of whether the system actually functioned tortiously vis-a-vis the outside world, rather than attempting to trace the source of harm caused by the system to an error or fault on the part of the humans and corporations behind it. The theory is a corollary of Abbott’s general premise of legal neutrality, meaning that law should make as few distinctions between human and artificial intelligence as possible.

Proponents of this approach argue that in a negligence claim against a human tortfeasor, the thinking process that led to the tortious conduct is irrelevant vis-à-vis the victim. It is only the act or omission resulting in harm that matters, and whether it was due to a failure to meet the standard of care. They would maintain the same

102. Ron Bucholz and Andy Yu, “Tort and Contracts: Civil liability for AI Causing Harm” in Jill Presser, Jesse Beatson and Gerald Chan, eds. *Litigating Artificial Intelligence* (Toronto: Emond, 2021) 348.

103. *Ibid.*, at 348.

104. *Ibid.*

105. *Ibid.*, at 349.

106. *Ibid.*

107. Abbott, *supra*, note 12 at 62-63.

should hold in the case of a decision, act, or omission by an artificial intelligence system that would be tortious if the system was a human. In other words, the focus of the tort analysis should be on what the system actually did, rather than what made it act as it did.

This approach would call for the acts or omissions of the system to be assessed on ordinary negligence principles, or in other words measured against the standard of care that would be applied to the conduct of a human actor.¹⁰⁸ The advantage for tort victims is that it would be unnecessary to prove that the system had an inherent defect. Defendants would have the opportunity to demonstrate that the system's outward act or omission was not negligent according to the same standard of care and degree of foreseeability that would apply to a human actor or decision-maker under the circumstances. If a human acting in accordance with the relevant standard of care could have caused the same accident in the same circumstances, there would be no liability because there was no negligence.

5. A Sliding Scale of Modified Liability Rules

Another approach that has been advanced calls for a sliding scale of liability regimes, depending on the level of transparency and autonomy of the system in question. Its proponents would apply current negligence principles where the deployment of the system is part of a human-driven decision-making process, such as a medical diagnostic system deployed to assist physicians who would make the final treatment decision.¹⁰⁹ Where a system acts autonomously, liability would turn on the degree of transparency (meaning the explainability of its processes), the constraints the creators or users placed on the system, and the extent of monitoring of the system in operation.¹¹⁰

In the case of autonomous operation, the creator or operator would be vicariously or strictly liable for the performance of the system when there is a high risk of harm. In less risky scenarios, originators and users would bear liability for negligence in testing, deploying, or operating the system. (Note this is akin to the dual regime of liability proposed by the 2020 European Parliament resolution and the scheme of "graded agency" proposed by the Canadian writers Bucholz and Yu.) If the system

108. Abbott notes that once artificial intelligence systems attain a level of safety and accuracy that is higher than what humans can attain, the standard of care applicable to them under this approach would be that of a system of the same kind, rather than the human standard: *supra*, note 12 at 9.

109. Yavar Bathaee, *supra*, note 59 at 894.

110. *Ibid.*, at 932.

lacks transparency, a finding of negligence should depend on whether harm was a foreseeable consequence of deploying it to function autonomously, instead of whether the particular harm it caused was reasonably foreseeable.¹¹¹

D. The Project Committee's View

1. Rejection of Strict Liability

Critics of strict liability point out that its proponents overrate the risk from artificial intelligence, essentially putting it on the same level as the risk from the category of things or activities that attract strict liability under the rule in *Rylands v. Fletcher*.¹¹² Those are things and activities that are inherently dangerous and have a tendency to escape control, e.g. fire, explosive substances, wild animals.¹¹³ Unpredictable results based on machine learning are not equivalent to the escape of a wild animal or a noxious substance. Many artificial intelligence systems serve to reduce or eliminate risk, which implies that their autonomy should be seen as a valuable quality rather than a source of danger in most cases.¹¹⁴

Abbott argues that excessive emphasis on risk associated with artificial intelligence ignores the greater safety and health benefits that it can bring in areas where its capabilities exceed those of humans. Treating artificial intelligence as an intrinsically dangerous thing, which is the effect of blanket adoption of strict liability, would foster a liability and financial risk climate that would discourage development and be counter-productive.¹¹⁵ It is misguided to treat the deployment of artificial intelligence as tortious in itself.¹¹⁶

The Project Committee generally agrees with these criticisms of strict liability. Moreover, the implications of strict liability for the availability of insurance should

111. *Ibid.*, at 938. See also Cynthia Khoo, "Missing the Unintended Forest Despite the Deliberately Planted Trees: Reasonable Foreseeability and Legal Recognition of Platform Algorithm-Facilitated Emergent Systemic Harm to Marginalized Communities." Draft paper presented at We Robot 2020 (3-4 April 2020, Ottawa, ON) at 73-74.

112. See, *supra*, note 61.

113. Huberman, *supra*, note 84 at 130. Bucholz and Yu, *supra*, note 102 at 351-352 make an analogy between fully autonomous, "driverless" cars and wild animals, asserting both are inherently dangerous.

114. Huberman, *supra*, note 84 at 130.

115. Abbott, *supra*, note 12 at 51.

116. Huberman, *supra*, note 84 at 130.

also give pause to judges and legislators. If prospective insureds are held liable regardless of any degree of care taken to reduce risks that are not fully within their ability to control, liability coverage is unlikely to be available at a manageable premium cost, or possibly not at all. Mainstream U.S. product liability law moved away from wholesale adoption of strict liability in the wake of an insurance crisis late in the 20th century, and now restricts it to manufacturing defects. Strict liability is not a guarantee that compensation will actually be recovered. It may only mean an easier path to a “dry” (unrecoverable) judgment. Withdrawal or contraction of liability coverage for the emerging technologies sector is not in the interests of plaintiffs.

Opponents of strict liability also argue that it will have a chilling effect on technological innovation.¹¹⁷ In part, this is related to the availability of insurance for smaller enterprises, where much original innovative work in digital technology takes place. These smaller technology firms cannot afford to self-insure.

An additional reason why the Project Committee is not in favour of strict liability, even in relation to systems at the higher end of the risk spectrum, is that it does not encourage the development of standards in the design, development, and use of artificial intelligence to the extent that fault-based liability does. If firms know they will be held liable regardless of the degree of care they exercise, they will not have an incentive to raise and continually improve standards and practices that will increase safety.

2. Notional Agency Insufficiently Distinguishable from Strict Liability

The theories based on notional agency and vicarious liability recognize that artificial intelligence systems are created to further objectives set by their human programmers and users. On a superficial level, they seem to benefit plaintiffs by avoiding the obstacles to proving fault on the part of the developers or operator that the opacity of the systems may present. The analogy to principal and agent is not a perfect one, however. It presents significant doctrinal issues.¹¹⁸ An agent or employee has legal personality, and can be sued directly in place of or concurrently with the agent’s principal or the employer for a tort committed while acting within the scope of the agent’s authority or the employee’s job. Artificial intelligence, however, lacks legal

117. Karni A. Chagal-Feferkorn, “Am I an Algorithm or a Product? When Products liability Should Apply to Algorithmic Decision-Makers” (2019) 30 *Stanford L. & Policy Rev.* 61 at 81-82 and 103.

118. Huberman, *supra*, note 84 at 143.

personality. Agents have fiduciary obligations of loyalty and good faith to their principals. It is hard to see how algorithms could fulfil these.¹¹⁹

In order to address some of the doctrinal issues, the suggestion has been raised that a new category of agent should be recognized to take account of the relationship between an artificial intelligence system and the humans who deploy it, that of “pure legal agent” without legal personality.¹²⁰

More practical questions are how to determine the scope of the system’s “authority” and determine whether a given output that leads to harm is within or outside the “authority” of the system as a notional or “pure legal” agent. Principals and employers are not liable for agents and employees who engage in a “frolic of their own,” outside the scope of their authority or employment and cause damage, e.g., an errant employee using a company vehicle recklessly as a party bus.¹²¹ Without a workable answer to questions surrounding when emergent behaviour is to be considered within or outside the scope of the notional agent’s authority, vicarious liability based on notional agency would differ little from strict liability in effect.

3. Non-Human Behaviour Is Not Measurable Against Human Reasonableness

The legal neutrality approach advocated by Abbott would call for the behaviour of an artificial intelligence system, rather than the conduct of the humans and corporations behind the system, to be measured against a standard of reasonableness. The advantage for tort victims is that it would be unnecessary to have to prove that the system or its data feedstock had an inherent defect or that something was done wrongly in training, testing, or operating it. Defendants would have the opportunity to demonstrate that the system’s outward act or omission was not negligent according to the standard of care and degree of foreseeability that would apply to a human actor or decision-maker under the circumstances. If a human acting in accordance with the relevant standard of care could have made the same error in the same circumstances, there would be no liability because there was no negligence.

Arguably, this approach is fair to both plaintiffs and defendants. A plaintiff would face fewer obstacles in the path to recovery of compensation, and defences based on reasonable care would be available to all parties behind the system, e.g., an owner,

119. *Ibid.*

120. *Ibid.*

121. *Ibid.*, at 145.

operator, developer, or designer. This approach does not protect one at the expense of the other, and would not penalize innovation unduly, since it imposes no unusual liability rules. It has not gone unnoticed either that tort victims would be treated alike under this approach, regardless of whether the harm they incurred was caused by a human or by digital technology acting in place of one.¹²²

The difficulty the Project Committee sees with legal neutrality as applied to the field of tort, however, is that reasonableness is a concept inextricably linked to the context of human behaviour and is defined by its appeal to human reason. Artificial intelligence and robots directed by it do not have human mental processes and a human's knowledge of the outside world. As a result, they err or fail in ways that are different from human error.

For example, in the case of the fatal collision between an autonomous test vehicle and a pedestrian walking a bicycle that was mentioned in Chapter 2, a human driver would not have failed to recognize the pedestrian as a pedestrian despite the presence of the bicycle. The navigation system of the test vehicle could only recognize a pedestrian or a bicycle, but not the combined image.

As noted in Chapter 2, ChatGPT has been known to “hallucinate” fictitious citations when asked to generate a list of references.¹²³ It is reported to have composed a fictitious article and ascribed it to an actual author.¹²⁴ A human who was attempting to falsify citations would be extremely foolish to use the name of an actual author and unlikely to do so. In another example of difference between human and robotic error, GPT-3, the prototype of ChatGPT, is reported to have answered “yes” to the question “Is it safe to walk downstairs backwards if I close my eyes?” and “no” when asked a second time.¹²⁵ Asked the same question on another occasion, GPT-3 reportedly replied “That depends.” While this could reflect learning behaviour, an adult human would likely answer that question consistently from the start.

122. Karni A Chagal-Feferkorn, “How Can I Tell If My Algorithm Was Reasonable? (2021), 21 Mich. Tech. L. Rev. 213 at 219.

123. *Supra*, note 56.

124. Geoff Brumfiel, “Here is What ChatGPT Gets Right – And Wrong” (NPR, 17 March 2023), online: <https://www.npr.org/2023/03/17/1164383826/heres-what-the-latest-version-of-chatgpt-gets-right-and-wrong>.

125. Gary Smith, “Chatbots: Still Dumb After All These Years” (MindMatters, 3 January 2022), online: <https://mindmatters.ai/2022/01/will-chatbots-replace-the-art-of-human-conversation/>.

We do not see reasonableness as a standard that can be applied coherently to harm caused by non-human actors. The tendency we foresee with the legal neutrality approach is for the non-human errors of artificial intelligence to be invariably categorized as unreasonable because a human would not make them in the same way, or because the task that the artificial intelligence performs is on such a scale that a human could not feasibly perform it in any event. This would result in a drift towards *de facto* strict liability until artificial intelligence equals or surpasses humans in terms of safety and reliability in all of its applications.¹²⁶

If the performance of artificial intelligence is measured against a human standard of reasonableness, there is also a risk of reducing the incentive to exceed human capabilities in relation to safety and reliability.

4. Liability Upstream - Product Liability Provides Some Answers

The Project Committee considers that product liability law in the common law Canadian jurisdictions provides some answers regarding a just and balanced theory of liability while avoiding the difficulties of the other approaches discussed above. The principles concerning duty of care under product liability law may be applied by analogy to cases involving artificial intelligence without the need to take a definitive position on whether artificial intelligence itself may be characterized as a product.¹²⁷

We would look to product liability law in regard to the liability of potential defendants situated “upstream” in the chain of events leading to litigation. Upstream defendants would be those involved in the design, development, training and testing of artificial intelligence systems prior to the point at which the systems are placed on the market or deployed in actual use. The liability of “downstream” defendants, namely operators and other end-users, involves different considerations and is discussed later.

A developer of a complete system that employs artificial intelligence may be realistically compared to a manufacturer of a complex product with numerous integrated components. It is common for digital technologies that employ artificial intelligence to incorporate discrete software modules created by different teams of specialists,

126. Abbott, *supra*, note 12, notes at 65 that once the safety of artificial intelligence surpasses human standards, it would result in strict liability of humans if it became the standard of care.

127. Software has been characterized as a product subject to the *Sale of Goods Act* when it is integrated with hardware: *Burroughs Business Machines Ltd. v. Feed-Rite Mills (1962) Ltd.* (1973), 42 D.L.R. (3d) 303 (Man. C.A.); *aff'd* (1976), 64 D.L.R. (3d) 767 (S.C.C.).

much like a manufacturer of a complex physical product may draw on various sources for components.

The designers and developers of artificial intelligence modules that are incorporated into an integrated system by another developer may be compared to suppliers of components that are incorporated into a complex product. Suppliers of training and testing data in the developmental stage of a system may also be likened to component suppliers.

Product liability law recognizes a duty of care on the part of manufacturers and others who place a product in the stream of commerce towards anyone who may reasonably be foreseen to be at risk of damage or injury if the product is unsafe, not only towards those who acquire the product or who deal in some manner with the manufacturer.¹²⁸

The fact that courts have recognized this duty of care as existing in law relieves a plaintiff in a product liability case of having to specifically prove a relationship of proximity sufficient to support a duty of care owed to the plaintiff. If the relationship of proximity and the duty of care arising from it were not presumed in law, it would be impossible in most cases for a member of the public who is injured by a defective product and who never had direct or indirect dealings with the manufacturer to prove a proximate relationship. A member of the public who has incurred damage or loss through the operation of artificial intelligence will face the same obstacle in seeking justice if a duty of care similar to that of a manufacturer in product liability law were not extended by analogy to the developers of the system in question.

It appears to be the weight of opinion that component suppliers, as manufacturers in their own right, also owe a duty of care towards consumers of products containing their components and others who may be affected by defects in the components.¹²⁹

128. *Bow Valley Husky (Bermuda) Ltd. v. Saint John Shipbuilding Ltd.*, [1997] 3 S.C.R. 1210, at para. 19. See also *Hollis v. Dow Corning Corp.*, [1995] 4 S.C.R. 634, at para. 21; *Stanway v. Wyeth Canada Inc.*, 2011 BCSC 1057; aff'd 2012 BCCA 260.

129. *Kett v. Mitsubishi Materials Corporation*, 2020 BCSC 1879, at para. 74. See also Theall, Lawrence G. et al., *Product Liability: Canadian Law and Practice* (Toronto: Thomson Reuters, 2000) (Looseleaf, updated) at 5-8. Despite some strong statements by text writers that component suppliers owe a duty of care to consumers and end users of a product containing their components, the direct liability of component suppliers to persons other than purchasers of their components does not seem to be as firmly established in case law as the duty of care of a manufacturer of the completed and integrated product. In *Burr v. Tecumseh Products of Canada Limited*, 2023 ONCA 135, at paras. 103-104, however, the Ontario Court of Appeal expressly disapproved of an *obiter*

The duty of care of a manufacturer of a complex product in which components made by others are integrated extends nevertheless over the entire product.¹³⁰ In other words, if a defect in a component makes the complex product dangerous, the manufacturer of the complex product is not exonerated from liability towards a plaintiff who suffers injury as a result of the defect merely because the supplier of the defective component is also liable.

The principle that a manufacturer's duty of care extends over the entire product assists the plaintiff to gain access to civil justice more easily than if the plaintiff were limited to claiming against the supplier of the defective component alone, because it will generally be possible for a plaintiff to identify the manufacturer of a complex product marketed as a unit, while it would often be far more difficult for a plaintiff to first determine the component in which the defect lay and then identify the source of the component.

The manufacturer of the complex product with integrated components is in a better position to identify its suppliers than the plaintiff, and will almost certainly assert a third party claim against the supplier of a defective component to bring the supplier into the litigation. If the manufacturer is sued, other proper defendants not likely to be as easily identified by the plaintiff are likely to be brought before the court.¹³¹ The same would be true if a developer of a complex system employing an artificial intelligence system is sued and wants to shift blame to the suppliers of a software module the developer believes is likely connected with the plaintiff's claim.

The principles relating to duty of care under product liability thus provide a framework within the law of tort that can be adapted to allow rights and liabilities to be determined between someone who has been harmed by the operation of artificial intelligence and the upstream defendants who participated in developing the artificial intelligence system in question up to the point at which it is made available for actual use in real-world circumstances. Extending these principles by analogy to cases in which the operation of artificial intelligence has caused harm would facilitate access

dictum of the trial judge that a component supplier does not owe a duty of care directly to a consumer of a product in which the component is integrated.

130. *Winnipeg Condominium Corporation No. 36 v. Bird Construction Co.*, [1995] 1 S.C.R. 85; 1688782 Ontario Inc. v. Maple Leaf Foods Inc., 2020 SCC 35, at para. 49.

131. A *proper* party is a person or entity with a legal interest in the matters in issue in a legal proceeding, and who therefore may join or be added to the proceeding as a party. By contrast, a *necessary* party is one who must be named in or be added to a legal proceeding as a party before the court can make a binding order resolving the matters in issue.

to civil justice by assisting plaintiffs to overcome the difficulty of identifying proper defendants at the upstream end and secure the appearance of other potentially liable defendants in the litigation.

Tentative Recommendations

- 1. Civil liability for harm caused by artificial intelligence should not be based on strict liability.*
- 2. Product liability principles should be adapted by analogy to determine rights and liabilities as between a plaintiff harmed by the operation of an artificial intelligence system and defendants who participated in the development of the system and in making it available for use, by treating*
 - (a) the plaintiff similarly to a plaintiff claiming to have incurred loss or damage from a product comprising multiple components;*
 - (b) developers of the system as owing a duty of care similar to that owed by a manufacturer of a complex product involving multiple integrated components towards persons or entities who foreseeably could be affected by a defect making the product dangerous;*
 - (c) developers of components of the system as owing a duty of care similar to that owed by a supplier of a component of a complex product towards persons who foreseeably could be affected by a defect in the component that makes the component and the product in which it is integrated dangerous.*

5. Liability Downstream - Operators

(a) Who is an “operator”?

In the previous section we used the term “upstream” to refer to the phase in the lifecycle of an artificial intelligence system covering the design, development, training and testing of an artificial intelligence system up to the time it is made available. We use the term “downstream” here to refer to the phase beginning with release of the system by its developers for actual use in real-world settings, and continuing for the life of the system.

In the downstream phase, users will deploy the system for their own purposes. They may have commissioned the development of the system on a custom-built basis, or they may have acquired it “off the shelf.” They may lease the system from its developers. They may be the developers of the system themselves. They may oper-

ate the system directly or contract with an agent to operate the system on their behalf. The downstream users may be the developers themselves in some cases.

Regardless of how they come to have the system, or how they use it, users may expose others to the embedded risks it may bear or any risks that may emerge as the system continues to operate. Of course, the deployment of artificial intelligence in a user's activities may also improve their safety and have other positive effects. This does not change the fact that whoever has decision-making authority over the operation of the system is in a position to exert some level of control over risks associated with the system that may hold potential to affect others adversely.

The Project Committee believes that liability towards those affected by the system should go together with the ability to exert control over the risk of operation. This aligns with the view expressed by the European Commission's High Level Expert Group on Liability and New Technologies in its 2019 report:

[L]iability should lie with the person who is in control of the risk connected with the operation of emerging digital technologies and who benefits from their operation (operator).¹³²

The ability to exercise control over the system risk in the course of operation should be sufficient in itself to attract a duty of care towards those whom the risk affects. We would not treat benefiting from the operation of the system as a requirement of the status of "operator." A requirement of benefit might be interpreted to exclude agents actively operating the system without having any beneficial interest in it from the status and liabilities of an "operator." There would be considerable moral hazard in allowing agents in active control to avoid liability towards those affected by the system merely because they do not own the system. The definition of "operator" should be broad enough to take account of many operating arrangements.

We would use the term "operator" to describe a person or entity with decision-making authority of a managerial nature over the operation of an artificial intelligence system and who thereby is in a position to exert some degree of control over the risk associated with its operation. We would make managerial authority a prerequisite for the status of "operator," because it would be unjust to impose the liabilities that should go with that status on technical personnel who carry out day-to-day system operation, but who are subject to superior orders and have no independent decision-making authority.

132. *Supra*, note 85 at 39.

This definition of an “operator” would align with the definition of “person responsible” for an artificial intelligence in section 5(2) of the proposed Canadian federal *Artificial Intelligence and Data Act* (“AIDA”):

Person responsible

- (2) For the purposes of this Part, a person is responsible for an artificial intelligence system, including a high-impact system, if, in the course of international or interprovincial trade and commerce, they design, develop or make available for use the artificial intelligence system or *manage its operation*.¹³³

[Italics added.]

AIDA would impose various regulatory obligations on a “person responsible” for an artificial intelligence system. It does not address civil liability. The performance or breach of the statutory obligations it would impose could be relevant nevertheless to a judicial finding of civil liability, as discussed later in this paper. Similarity between the scope of the definition in AIDA of “person responsible” for an artificial intelligence system and the definition of “operator” for purposes of civil liability would establish a desirable coherence between the proposed federal regulatory statute and provincial tort law.

The definition of “operator” we propose is flexible enough to allow for the possibility that two or more persons or entities could have that status at any given time. For example, a system owner and the owner’s agent operating the system under a contract could both be operators and have the same obligations vis-à-vis third parties whom the system may affect.¹³⁴

133. Bill C-27, *An Act to enact the Consumer Privacy Protection Act, the Personal Information and Data Protection Tribunal Act and the Artificial Intelligence and Data Act and to make consequential and related amendments to other Acts*, Part 3, 1st Sess., 44th Parl., 2022 (second reading 24 April 2023).

134. The European Parliament’s resolution of 20 October 2020 recommending a civil liability regime for artificial intelligence to the European Commission also contemplated that there could be more than one operator, but made a distinction between “frontend operators” and “backend operators.” See, *supra*, note 88. This distinction, which was originally drawn in the 2019 report of the High Level Expert Group on Liability and New Technologies, *supra*, note 85 at 39 and 41-42, was meant to recognize which of two or more possible operators (such as an owner of an autonomous vehicle and the manufacturer who provides a continuous support service involving two-way transmission of data to and from the vehicle’s systems) should bear strict liability because of having greater control over the risk associated with the technology. As the distinction is linked to a regime of strict liability, the Project Committee does not consider it necessary to recommend a similar distinction between categories of operators.

(b) Liability of operators and other downstream defendants

As we have rejected strict liability, and envision problems with the application of other theories that have been advanced for reasons mentioned earlier in this chapter, the liability of operators and other potential downstream defendants such as providers of technical support services to the operator should flow in our view from the general fault-based principles that underlie the law of negligence. In other words, these downstream defendants should be held to owe a duty of care towards those within the foreseeable range of harm from the operation of an artificial intelligence system or from the nature of their involvement with the system, and should be obliged to exercise reasonable care to prevent harm from arising.

The application of negligence principles in respect of both upstream and downstream defendants should be subject, however, to the modifications explained in the later chapters that we believe are required in the specific context of artificial intelligence.

(c) Tentative recommendation

The Project Committee tentatively recommends:

3. A person or entity with decision-making authority of a managerial nature over the operation of an artificial intelligence system and thereby is in a position to exert some degree of control over the risk associated with its operation should be treated as an operator for the purpose of civil liability.

4. The liability of operators and other persons who provide services in connection with the operation of an artificial intelligence system should be based on general principles of the law of negligence, subject to the recommendations made below.

E. Exclusion and Limitation of Liability for Artificial Intelligence

When digital technology is transferred or a licence is granted for its use, the risk of liability for damage caused by flaws in the technology itself or for damage that may

arise from its use will typically be allocated by agreement. This usually means that the grantor of the rights will shift the risk to the party acquiring the rights by means of terms (“exclusion clauses”) that prevent the other party from suing the grantor to recover the amount of a loss the other party may incur or the amount of a liability to a third party that is somehow related to the technology in question. Alternatively, liability between the parties may be limited by agreement to a fixed maximum.

Parties to a contract are able to agree to exclusion clauses, waivers of future claims, or limitation of damages, except to the extent that the ability to do so may be restricted by law. Sometimes legislatures restrict the ability to contract out of liability or waive the right to make a claim for reasons of public policy. When legislatures restrict freedom of contract in this manner, it is usually to enforce specific legal or regulatory requirements, or to prevent erosion of certain rights conferred by law.¹³⁵

The EU initiatives described in this consultation paper prioritize protection of the public and enforcement of regulatory standards over the needs of developers and operators of artificial intelligence systems to manage risk exposure. The European Parliament resolution of 2020 on liability for harm caused by artificial intelligence referred to earlier in this chapter would render void any agreement between an operator and a person harmed by the operator’s system that would limit or circumvent the obligations of the operator set out in the draft regulation attached to the resolution.¹³⁶ Similarly, the EU Product Liability Directive would continue to govern damage claims against producers of artificial intelligence systems states that the liability of the producer under the Directive may not be limited or excluded by any contractual provision.¹³⁷

Some might argue that the risks associated with autonomous artificial intelligence and emergence are such that restriction or removal of the ability to disclaim or contract out of liability for them is warranted. On the other hand, if developers and operators of artificial intelligence systems cannot manage their exposure to indetermi-

135. Consumer protection legislation often contains anti-avoidance provisions that nullify waivers of rights and obligations under it. For example, s. 3 of the British Columbia *Business Practices and Consumer Protection Act*, S.B.C. 2004, c. 2, states: “3. Any waiver or release by a person of the person's rights, benefits or protections under this Act is void except to the extent that the waiver or release is expressly permitted by this Act.”

136. *Supra*, note 88, Art.2, para. 2.

137. European Union, *Council Directive of 25 July 1985 on the approximation of the laws, regulations and administrative provisions of the Member States concerning liability for defective products*, (85/374/EEC), Art. 12. The proposed new EU Product Liability Directive referred to in Chapter 4 would contain a similar provision.

nate liability as other businesses and research organizations must do, they may be unable to continue on an innovative path. This is especially true with respect to developers of generic or open-source artificial intelligence that can be used by any unascertained third parties and incorporated in applications of which the developer of the original generic or open-source module has no knowledge.

How an appropriate balance may be struck between adequate legal protection for the public and the need of industry and scientific organizations to manage their liability exposure in order to remain viable enterprises is a very difficult question. The Project Committee has not yet formed a settled view, but we invite comment from readers on whether, and to what extent, it should be possible to disclaim, exclude, or limit liability for harm caused by artificial intelligence.

Chapter 4. The Problem of Proof of Fault

A. General

In Chapter 3 attention was drawn to the likelihood that a plaintiff could face formidable problems of proof in relation to causation and fault in a tort claim arising from the operation of artificial intelligence. Writers have pointed to four principal reasons:

1. A high number of potential defendants typically involved in the design, development, deployment and operation of an artificial intelligence system.¹³⁸
2. Autonomy of some systems;¹³⁹
3. Limited explainability;¹⁴⁰
4. Lack of foreseeability (as conventionally understood and applied).¹⁴¹

This listing of obstacles facing a plaintiff seeking compensation for damage is echoed by the European Commission in its proposal for an AI Liability Directive:

Current national liability rules, in particular based on fault, are not suited to handling liability claims for damage caused by AI-enabled products and services...The specific characteristics of AI, including complexity, autonomy and

138. Yaariv Benhamou and Justine Ferland, “Artificial Intelligence and Damage: Assessing liability and Calculating the Damages” in Pina D’Agostino, Carole Piovesan and Aviv Gaon, *Leading Legal Disruption: Artificial Intelligence and a Toolkit for Lawyers and the Law* (Toronto: Thomson Reuters Canada, 2020) at 170; High Level Expert Group on Liability and New Technologies, *supra*, note 85 at 28; European Commission, *Proposal for a Directive of the European parliament and of the Council on adapting non-contractual civil liability rules to artificial intelligence (AI Liability Directive)* (Brussels: European Commission, 28 September 2022) at para. (17) (preamble).

139. Benhamou, *supra*, note 138 at 170; Barfield and Pagallo, *supra*, note 82 at 16;

140. Benhamou, *supra*, note 138 at 170; Bathaee, *supra*, note 59 at 922.

141. Benhamou, *supra*, note 138 at 170; Bathaee, *supra*, note 59 at 924-925.

opacity (the so-called “black box” effect), may make it difficult or prohibitively expensive for victims to identify the liable person and prove the requirements for a successful liability claim.¹⁴²

Pointing to a party at fault means first having to determine what actually caused a harmful output that produces damage. It is possible that only one component or system may be the source of damage, but there will often be multiple integrated systems running concurrently. The interdependency of software, hardware, and data in the functioning of artificial intelligence will tend to obscure the source of the harm.¹⁴³

Even if a particular system can be identified as a source of harm, the limited extent of explainability that is associated with some forms of artificial intelligence, especially those that are highly data-dependent, will present a formidable obstacle to proving a claim. If the route from input to output is not fully explainable even by the designers and programmers of a system, plaintiffs have little hope of being able to present a detailed or step-by-step explanation linking a breach of duty by the defendant to the harm.

The difficulty a plaintiff would face in seeking evidence to prove causation and identify would be compounded by an inclination on the part of developers of an artificial intelligence system to treat its algorithms and other aspects of its design as proprietary secrets. This obstacle was illustrated in litigation concerning the dismissal of teachers by the Houston (Texas) Independent School District in reliance on a system marketed to school boards that used a proprietary statistical model for measuring every teacher’s effectiveness. The supplier of the system withheld the algorithms and source code for the software that were in issue from both the plaintiffs and the defendant school district, preventing the plaintiffs from obtaining discovery of the basis for the automated evaluations and blocking independent testing of the algorithms that generated them.¹⁴⁴

142. European Commission, *supra*, note 138, Explanatory Memorandum at 1.

143. Benhamou, *supra*, note 138 at 176;

144. *Houston Federation of Teachers, Local 2415 v. Houston Independent School District*, (4 May 2017) Civil Action H-14-1189 (U.S. District Court, So. Dist. of Texas), amended summary judgment opinion at 12. A subpoena was issued to the supplier of the system, who eventually gave an expert engaged by the plaintiffs very limited access to the source code: presentation by Martha P. Owen, counsel for the plaintiffs, at “AI Decision-Making: Protecting Rights Through Litigation and Regulation in Canada and the U.S.”, web panel discussion sponsored by the Law Commission of Ontario, 9 December 2021.

Establishing the cause of the damage is essential to any tort claim, even ones based on strict liability.¹⁴⁵ Non-human, autonomous decision-making that is opaque places serious and potentially insuperable obstacles in the way of proving causation and fault.¹⁴⁶ These barriers to proof are likely to grow higher with increasing levels of autonomy and complexity. As stated in an influential article in a U.S. law journal, artificial intelligence and robotics present society with “the prospect of a victim who suffers a non-natural harm but no perpetrator to whom the law can attribute this harm.”¹⁴⁷

B. Rebalancing Evidentiary Burdens for Fairness

1. General

Both upstream and downstream defendants will generally have a large advantage over the plaintiff in tort litigation involving artificial intelligence because of having greater familiarity with the system and knowledge of the measures that were or could have been taken to avert the harm that arose from its use.

Of course, the plaintiff can use oral and document discovery in the litigation process to elicit information helpful to the plaintiff’s case, but the usefulness of pre-trial discovery as a means of redressing “informational asymmetry” between plaintiffs and defendants may be limited, particularly in cases involving data-dependent systems where the system’s decision-making may be opaque due to the complexity of neural networks.¹⁴⁸ A discovery witness for a defendant developer or operator can be asked questions relating to facts, but cannot be asked on discovery to provide an opinion that would explain why certain input data produced an output resulting in damage. Information about the system derived from pre-trial discovery will not necessarily provide a basis for a theory of causation and fault that the plaintiff can advance.

145. *Ibid.*, at 171.

146. Benhamou, *supra*, note 138 at 175; Kristen Thomasen, “AI and Tort Law” in Florian Martin-Bariteau and Teresa Scassa, eds. *Artificial Intelligence and the Law in Canada* (Toronto: LexisNexis, 2022) at 117.

147. Calo, *supra*, note 49 at 542.

148. The phrase “informational asymmetry” is used in the 2029 report of the European Commission’s High Level Expert Group on Liability for New Technologies, *supra*, note 85 to describe the discrepancy between a plaintiff and the defendant operator and system developers regarding knowledge of an artificial intelligence system in issue in a claim for compensation.

The extreme difficulties of proof that a plaintiff or prospective plaintiff may face in a claim for damages or other relief in a case involving artificial intelligence appears to warrant some mechanism to maintain balance in the litigation process in these cases. Numerous writers have reached this conclusion, as have EU policymakers.¹⁴⁹

The concept of a rebalancing mechanism to achieve just results in the face of an extreme informational imbalance, such as the application of evidentiary presumptions, is not entirely foreign to the law of tort. At an early stage in the development of the modern law of negligence, the common law developed the procedural device known as *res ipsa loquitur* to address cases where direct proof of causation and fault was lacking, but there was a strong likelihood of negligence.

Res ipsa loquitur (“the thing speaks for itself”) was formerly applied in common law jurisdictions of Canada to allow for a rebuttable inference of negligence where a thing or situation was under the sole control of the defendant or someone for whose actions the defendant was responsible, the occurrence resulting in damage would not happen in the ordinary course of things apart from negligence, and there was no evidence of how or why the damage occurred. If the defendant did not provide an explanation consistent with reasonable care to rebut the inference, judgment would be given for the plaintiff despite the lack of clear proof of fault and causation. In effect, the plaintiff would benefit from a presumption of negligence drawn because of the lack of another reasonable explanation for the occurrence resulting in damage.

In 1998, the Supreme Court of Canada held that *res ipsa loquitur* should no longer be treated as a distinct element of negligence law because of its “limited use.”¹⁵⁰ The Supreme Court considered it to have been “more confusing than helpful,” and declared it to have expired. Instead, the Supreme Court said courts should weigh circumstantial evidence together with direct evidence to determine whether the plaintiff has made out a *prima facie* (at first sight or impression) case of negligence against the defendant.¹⁵¹

The problem with requiring the plaintiff to make out a *prima facie* case of negligence in a case arising from harmful emergent behaviour of artificial intelligence is that it

149. See Benhamou and Ferland, *supra*, note 138 at 187; Barfield and Pagallo, *supra*, note 14 at 104. See also High Level Expert Group on Liability and New Technologies, *supra*, note 85 at 43 and 49-50.

150. *Fontaine v. British Columbia (Official Administrator)*, [1998] 1 S.C.R. 424, at para. 27.

151. *Ibid.*

would require the plaintiff to present evidence on every element of negligence, including a specific act or omission amounting to fault on the part of one or more human or corporate defendants and a causal link between the fault and the damage when, at least in some circumstances, the nature of complex automated decision-making makes this impossible. This is especially true of data-dependent systems that employ deep learning.

This problem has been given considerable attention by EU policymakers. The proposal submitted to the European Parliament by the European Commission (“EC”) on 28 September 2022 for an EU directive on adapting non-contractual civil liability rules to artificial intelligence (AI Liability Directive Proposal) is chiefly concerned with creating a more equal playing field for claimants (plaintiffs) and defendants.¹⁵²

The AI Liability Directive Proposal is one-half of the EC’s response to the European Parliament’s 2020 resolution mentioned in Chapter 3, in which the EC was invited to propose a special civil liability regime for artificial intelligence along the lines of the Annex to the resolution.¹⁵³ The other half of the EC’s response package to the 2020 European Parliament resolution is its proposal for a new Product Liability Directive, which also contains similar provisions for claims based on it that relate to artificial intelligence.¹⁵⁴ The next section describes the scheme under the EC proposals to rebalance the litigation process.

2. European Commission Rebalancing Proposals

(a) *The AI Liability Directive Proposal*

The AI Liability Directive proposal is intended to apply to fault-based claims under the national law of Member States. Its terminology and provisions are intended to mesh with those of the EU’s proposed *AI Act*.¹⁵⁵

The AI Liability Directive Proposal is less extensive than the regime of civil liability for artificial intelligence envisioned by the 2020 European Parliament resolution.

152. European Commission, *supra*, note 138.

153. This is explained in the preamble to the EC proposal for a new Product Liability Directive: European Commission, *Proposal for a Directive of the European Parliament and of the Council on liability for defective products* COM(2022) 495 final (Brussels: European Commission, 28 September 2022) at 6.

154. *Ibid.*

155. *Supra*, note 10.

Nevertheless, the mechanisms it contemplates would materially assist plaintiffs faced with obstacles to establishing fault and liability in certain circumstances.

The first assistive mechanism under the proposed AI Liability Directive would be an order to disclose relevant evidence concerning a specific “high-risk AI system,” as defined under the proposed EU *AI Act*.¹⁵⁶ The proposed directive would require EU Member States to empower their national courts to order disclosure of relevant evidence about a high-risk artificial intelligence system suspected of having caused damage if a claimant has made “all proportionate attempts” to gather the evidence from the defendant.¹⁵⁷

An application for a disclosure order could be made by a claimant (plaintiff) or a potential claimant who has requested the disclosure without success.¹⁵⁸ A potential claimant would have to present facts and evidence to support the plausibility of a damages claim.¹⁵⁹ A disclosure order could be made against a “provider” (developer) of an artificial intelligence system, a person subject under the proposed EU *AI Act* to the obligations of a provider, or a “user” (operator).¹⁶⁰ An order could be made against these persons whether or not they were defendants.

The disclosure order would be limited to evidence “necessary and proportionate” to support a claim or potential claim, taking the legitimate interests of all parties and third parties into account, including protection of trade secrets and confidential information.¹⁶¹ The proposed directive would also oblige Member States to empower their courts to make orders for preservation of evidence that could be the subject of a disclosure order.¹⁶²

156. *Supra*, note 10, Art. 6.

157. European Commission, *supra*, note 138, Art. 3, para. 1.

158. *Ibid.*

159. *Ibid.* The language “evidence to support the plausibility of a damages claim” is drawn directly from the EC proposed directive. This wording points to an evidentiary standard that is likely similar to “establishing an arguable case” in common law jurisprudence.

160. *Ibid.*

161. *Ibid.*, Art. 3, para. 4.

162. *Ibid.*, Art. 3, para. 3. The proposed directive and the accompanying explanatory text are ambiguous regarding whether an order for preservation could only be ancillary to a disclosure order, or whether a self-standing preservation order could be made without a corresponding order for disclosure.

Non-compliance by a defendant with a disclosure order or an order to preserve evidence would require a national court to presume that the defendant was non-compliant with a relevant duty of care to which the requested evidence relates and was intended to prove.¹⁶³ The defendant would have the right to rebut the presumption.¹⁶⁴

The other mechanism in the European Commission proposal to relax the burden of proof in artificial intelligence-related claims for damages is a rebuttable presumption of a causal link applicable under certain circumstances. Courts of EU Member States would be required to presume a causal link between the defendant's fault and the output of the artificial intelligence system in issue, or its failure to produce an output if:

- (a) the claimant has demonstrated or the court has presumed (because of breach of a non-disclosure order) the fault of the defendant or a person for whose behaviour the defendant is responsible, consisting of non-compliance with a duty of care laid down in EU or national law directly intended to protect against the damage that occurred;
- (b) it can be considered reasonably likely, based on the circumstances of the case, that the fault influenced the output produced by the artificial intelligence system or the failure of the system to produce an output; and
- (c) the claimant has demonstrated that the output produced by the system, or the failure of the system to produce an output, gave rise to the damage.¹⁶⁵

This presumption of a causal link would be rebuttable.¹⁶⁶

The presumption would apply in a fault-based case involving any artificial intelligence system in an EU Member State, not only one classified as high-risk. There would be several important restrictions on its application, however. It is targeted principally at commercial or professional providers and users. If the defendant used the system in a personal and non-professional activity, the presumption would apply only if the defendant "materially interfered with the conditions of operation of the

163. *Ibid.*, Art. 3, para. 5.

164. *Ibid.*

165. *Ibid.*, Art. 4, para. 1.

166. *Ibid.*, para. 7.

system”, or if the defendant “was required and able to determine the conditions of operation of the system and failed to do so.”¹⁶⁷

If the system is not classified as high-risk under EU law, the presumption would only apply where the national court considered it excessively difficult for the claimant to prove the causal link between the defendant’s fault and the damage.¹⁶⁸

In the case of claims against providers (developers) of high-risk systems, the requirement of paragraph (a) above would relate only to non-compliance with a specified set of statutory obligations under Title III of the EU *AI Act*.¹⁶⁹

In the case of defendant users (operators) of high-risk systems, the requirement of paragraph (a) would be met if the claimant proved the user did not comply with the obligation to use or monitor the system in accordance with the accompanying instructions of use or did not suspend or interrupt uses of the system where appropriate, as required by the *AI Act*.¹⁷⁰ It would also be met if the user exposed the system to input data under its control that was not relevant in view of the system’s intended purpose, which would also be a breach of an obligation imposed on users by the *AI Act*.¹⁷¹

The presumption of a causal link would not apply if the defendant demonstrated that sufficient evidence and expertise was reasonably accessible to the claimant to prove the link.¹⁷²

167. *Ibid.*, Art 4, para. 6.

168. *Ibid.*, Art. 4, para. 5.

169. *Ibid.*, Article 4, para. 2,

170. *Ibid.*, Art. 4, para. 3.

171. *Ibid.* Unlike Article 4, paragraph 2 concerning application of the presumption to providers, paragraph 3 does not say the presumption would apply against users *only* in the case of breach of these obligations. As a result, it is not clear from the language of Article 4, paragraph 3 whether the presumption of a causal link would be triggered under the proposed directive against a user of a high-risk system who is proven to have breached *other* obligations under EU or national law that intended to protect against the damage that is the subject of the claim.

172. *Ibid.*, Art. 4, para. 4. This appears circular, because in order to avoid having a causal link presumed under this exemption, the defendant would have to admit there is evidence already available to the plaintiff to prove the causal link between the fault of the defendant and the damage and thereby establish the liability of the defendant.

The AI Liability Directive Proposal is not intended to affect rights that a plaintiff would have under the laws of EU Member States implementing the product liability regime of the EU under the existing Product Liability Directive dating from 1985.¹⁷³ That directive imposes a regime of strict liability for defective products.¹⁷⁴ Most claims in the EU against developers of artificial intelligence systems would likely be made under the Product Liability Directive because it avoids the need to prove fault.

(b) The EC proposal for a new Product Liability Directive

The EC's proposal for a new Product Liability Directive is intended to "ensure liability rules reflect the nature and risks of products in the digital age."¹⁷⁵ It would repeal the 1985 directive, but carry forward the regime of strict liability of manufacturers for damage from defective products.¹⁷⁶ It would also clarify that artificial intelligence systems are "products" within the scope of the new directive, settling doubts on whether they come within the EU product liability regime.¹⁷⁷ This will likely result in most future claims in the EU for damages against developers of artificial intelligence systems being made under the Product Liability Directive rather than on the basis of national laws requiring proof of fault.

Even under a strict liability regime, however, the plaintiff must still prove a causal link between a defect in the product and the damage. The proposed new Product Liability Directive contains provisions to alleviate the burden of proof by means of presumptions once a reduced evidentiary threshold is met. It would also require Member States to empower their courts to order a defendant in a claim under the Product Liability Directive to disclose relevant evidence on the request of a claimant on a basis very similar to the terms of the disclosure order provisions of the AI Liability Directive.¹⁷⁸

173. *Ibid.*, Article 1, para. 3(b).

174. European Union, *Council Directive of 25 July 1985 on the approximation of the laws, regulations and administrative provisions of the Member States concerning liability for defective products*, (85/374/EEC), Art. 1.

175. *Supra*, note 153, Explanatory Memorandum, p. 2.

176. *Ibid.*, Art. 7, para. 1.

177. *Ibid.*, Explanatory Memorandum, p. 5. "Product" is defined in Art. 4(1) to include software. The debate elsewhere regarding whether artificial intelligence is a product, a service, or a hybrid product and service is noted in Chapter 3 under the subheading "Product Liability."

178. *Ibid.*, Art. 8.

A product would be presumed defective if the claimant established that the product does not comply with mandatory safety requirements under EU or national law within Member States that are intended to protect against the damage that occurred, or if the damage that is the gist of the claim was caused by an obvious malfunction during normal use or under ordinary circumstances.¹⁷⁹ It would also be presumed if the defendant did not comply with an order to disclose relevant evidence.¹⁸⁰

A causal link between defectiveness of the product and the damage would be presumed if the defect is established and the damage is of a kind typically consistent with the defect in question.¹⁸¹ If the court determined that the claimant faced excessive difficulties in proving defectiveness of the product and a causal link with the damage because of technical or scientific complexity, both defectiveness and the causal link between this and the damage would be presumed if the claimant demonstrated that “it is likely that the product was defective or that its defectiveness is a likely cause of the damage, or both.”¹⁸²

The presumptions alleviating the burden of proof would all be rebuttable, and the defendant could contest the existence of excessive difficulties of proof facing the plaintiff.¹⁸³

3. Recommendation on Relief Against Difficulties of Proof in Appropriate Cases

The solutions to the evidentiary problems that are proposed in the EU are highly statutory and prescriptive in nature. It may not be necessary to go the length to which the European Commission has gone in recommending that courts be obliged by legislation to draw inferences of fault and causation in order to address the practical difficulties of proof that will face plaintiffs seeking compensation for harm caused by artificial intelligence. The European Commission’s approach appears inflexible, depriving courts of the discretion to respond to the circumstances of individual cases. Moreover, it is an approach dependent on classifications made under enacted regulatory standards for artificial intelligence. Canada may have some national standards in regulations to be made under the proposed *Artificial Intelligence*

179. *Ibid.*, Art. 9, para. 1.

180. *Ibid.*

181. *Ibid.*, Art. 9, para. 3.

182. *Ibid.*, Art. 9, para. 4.

183. *Ibid.*, Art 9, para. 5.

and Data Act in the future, but unlike the proposed standards in the EU *AI Act*, they have yet to be seen.¹⁸⁴

The European Commission proposals also depend on a distinction between systems classified as “high-risk” and others. This is a distinction that relates to the choice to impose an extensive body of specific regulatory standards at the Union level only on systems legislatively classified as “high-risk,” and to the dual regimes of strict and fault-based liability under the proposed EU schemes that have been discussed above. As we have rejected strict liability in principle, we have not adopted the concept of gradations of risk in our tentative recommendations on civil liability for artificial intelligence. While classifications of systems based on risk level may be appropriate in relation to differential regulatory burdens, we do not see it as necessary or desirable to distinguish between risk levels in relation to procedural matters.

The solution we will recommend instead is one that need not be legislated, but could equally as well be applied by courts as the need arose on the facts of individual cases. Not every case involving artificial intelligence will present problems of proof of fault and causation. Reversing the ultimate onus of proof resting on the plaintiff in a civil action only because artificial intelligence is involved in the factual matrix would be an extreme measure, coming close to imposing a regime of technology-specific strict liability.

We do think, however, that the burden of proof should be relaxed by means of a rebuttable presumption in the interests of fair process when a plaintiff who has incurred injury or loss through operation of artificial intelligence cannot identify a defect in the design or a deficiency in the training, data management and curation and testing of the system that is causally linked to the damage, nor enunciate a precise theory to explain why the system generated the harm-producing output.

Where no single defendant can be identified as at fault, but injury or loss from the output of the artificial intelligence system occurred nonetheless, it would be reasonable to expect that for reasons of self-interest, the defendant who is best-placed in terms of knowledge of the system among the likely multiple defendants would offer an explanation for the behaviour of the system in the circumstances of the case that

184. The *Artificial Intelligence and Data Act (AIDA) Companion Document* issued by the federal Department of Innovation, Science and Economic Development provides some indication of what future regulations might contain, but also emphasizes that the development of regulations will occur after a 6-month consultation on their content. It also states AIDA will not be in force until two years after Bill C-27 is passed, and at the earliest no sooner than 2025. See online: <https://ised-isde.canada.ca/site/innovation-better-canada/en/artificial-intelligence-and-data-act-aida-companion-document>.

is consistent with reasonable care in its design, development, training and pre-deployment testing, at least on the part of that defendant.

At the very least, the defendant best-placed in terms of knowledge should disclose, either in pre-trial discovery or in evidence at trial, sufficient information about the system and its function to enable other parties in the litigation to allow them to gain an understanding of how it was meant to work. Concern over preservation of proprietary secrets and commercially sensitive confidential information could be managed through undertakings, non-disclosure orders and orders to seal files and transcripts as it is in other civil litigation.¹⁸⁵

If an explanation or disclosure of this kind is not forthcoming, we think that at the trial stage, the court should be justified in drawing an inference of negligence against the defendants jointly and severally, except as against other defendants who provide evidence showing exercise of reasonable care in their own participation in the development, training, testing, or use of the system to prevent it from effecting harm.¹⁸⁶

The operator is likely to be the potential defendant most easily identified by the plaintiff and is almost certain to be named in a tort proceeding, but may be nearly as much in the dark as the plaintiff about the intricacies of a complex deep learning system. It would be pointless as well as unfair to impose the obligation to explain the performance of the artificial intelligence system on downstream defendants who did not participate in the creation of the system. Nevertheless, downstream defendants should be prepared to demonstrate reasonable care in their own usage or other interaction with the system in issue or face an inference of negligence when harm arising from the operation of the system is otherwise unexplained.

185. Information obtained through oral and documentary discovery in litigation is subject to an implied undertaking by the examining party to whom it is disclosed to use the information only for purposes of the litigation in which it is obtained, unless the scope of the implied undertaking is varied by order of the court: *Juman v. Doucette*, 2008 SCC 8, [2008] 1 S.C.R. 157 at para. 4. In addition, orders for sealing documentary exhibits and transcripts in order to protect trade secrets and confidential information included in the evidence are commonly made on application by a party.

186. While this would resemble the shift in the evidentiary burden of proof under the obsolete maxim *res ipsa loquitur*, it is also consistent with the approach approved by the Supreme Court in *Fontaine v. British Columbia (Official Administrator)*, *supra*, note 150 of weighing circumstantial evidence and drawing appropriate inferences of fault in the absence of evidence to contradict the *prima facie* case.

Among the potential upstream defendants, the defendant who would usually be in the best position to provide an explanation of the system's performance consistent with the exercise of reasonable care would be the software developer who produced the complete, integrated artificial intelligence product and made it available for use. That developer is comparably situated to a manufacturer who introduces a complex (multi-component) physical product to the market. The manufacturer is typically the principal defendant in product liability litigation, and we have looked to product liability for models regarding the legal responsibility of upstream defendants. Accordingly, we believe an evidentiary burden to provide an explanation for the behaviour of a system consistent with the exercise of reasonable care at the upstream end should rest with the developer of the completed and integrated system.

The Project Committee tentatively recommends:

5. In an action for damages or other relief due to harm to persons or property arising from the operation of artificial intelligence, an obligation should rest on a defendant who developed the artificial intelligence system and made it available for use by that defendant or others to

- (a) provide an explanation for the performance of the system in the circumstances of the case that is consistent with the exercise of reasonable care in the design, development, training, and testing of the system; or*
- (b) disclose sufficient information concerning the design and function of the system to all other parties before the trial of the action that will enable a party to determine the cause of the harmful performance of the system in the circumstances.*

6. If the cause of the harm is otherwise unexplained, and in the absence of an explanation or disclosure referred to in Recommendation 5, the court should be justified in drawing an inference of negligence against the defendants jointly and severally, except in regard to a defendant who rebuts the inference by evidence of the exercise of reasonable care by that defendant in the circumstances to avoid the harm suffered by the plaintiff.

C. Reasonable Foreseeability and Artificial Intelligence

The test of foreseeability is currently expressed as what a reasonable person in the position of the defendant would consider a "real risk" that is "not far-fetched" or a

“natural result” of the defendant’s act or omission.¹⁸⁷ While it is described in law as an objective standard because it is based on the imputed perceptions of a reasonable person, it is still one that is rooted in human experience and human perceptions of cause and effect and risk materialization. The reactions of a non-human, artificial decision-making process to an infinite quantity of potential inputs are outside that experience.

Risk variables associated with artificial intelligence include inability to control the quality of future input data in the future use of a system by third parties (“garbage in, garbage out”), the unpredictability associated with extreme complexity, autonomous operation, and emergence. The variables influencing the risk associated with artificial intelligence, especially machine learning, will tend to confound the concept of reasonable foreseeability as it is conventionally understood and applied in the law of negligence.¹⁸⁸ The American writer Bathaee has commented on this subject:

[t]he result of the AI's decision or conduct may not have been in any way foreseeable by the AI's creator or user. For example, the AI may reach a counter-intuitive solution, find an obscure pattern hidden deep in petabytes of data, engage in conduct in which a human being could not have engaged (e.g., at faster speeds), or make decisions based on higher-dimensional relationships between variables that no human can visualize. Put simply, if even the creator of the AI cannot foresee its effects, a reasonable person cannot either. Indeed, if the creator of AI cannot necessarily foresee how the AI will make decisions, what conduct it will engage in, or the nature of the patterns it will find in data, what can be said about the reasonable person in such a situation?¹⁸⁹

Another risk variable that cannot be discounted, although it may possibly be more easily weighed under the current test of reasonable foreseeability because it relates to human behaviour, is deliberate and inventive misuse of the artificial intelligence by third parties.¹⁹⁰

187. See Chapter 3 under the subheading “Unintended Harm.”

188. Thomasen, *supra*, note 146 at 113;

189. *Supra*, note 59 at 924.

190. Intentional misuse is an especially acute problem in relation to online platforms and other openly available online services that can be easily adapted to malevolent purposes such as spreading misinformation for political or ideological purposes or creating internal discord in a targeted society or community. See Khoo, *supra*, note 111. Canadian courts have held manufacturers liable for not taking cost-effective measures to protect against relatively obvious possibilities of misuse; see *Tabrizi v. Whallon Machine Inc.* (1996), 29 C.C.L.T. (2d) 176 (B.C.S.C.). Reckless misuse of a product in disregard of instructions has been held not to be foreseeable, however: *Lem v. Barotto Sports Ltd.*, [1976] A.J. No. 4412 (S.C., App. Div.).

In a leading case on foreseeability in Canadian tort law, *Rankin (Rankin's Garage and Sales) v. J.J.*, the Supreme Court of Canada held that the risk of personal injury through the reckless driving of a juvenile car thief was not reasonably foreseeable by a garage owner who had stored the car unlocked on the garage lot with the keys in an ashtray.¹⁹¹ The car thief's behaviour pattern after the theft in *Rankin* could be seen as comparable to the emergent market-gaming activities of a stock-trading algorithm in Bathaee's hypothetical example mentioned in Chapter 2.¹⁹² Should the owner of the algorithm in that hypothetical be held liable to investors who lost money due to its manipulation of the market? The Supreme Court's answer would seem to be "No."

In an actual example of autonomous, emergent behaviour that was probably even less likely to have been foreseen by creators, a popular chatbot was reported to have produced a fictitious news report naming an actual individual as the offending party in an entirely fictitious incident of sexual harassment.¹⁹³ Furthermore, it appeared that after the erroneous content was suppressed on the chatbot platform where it originated, a counterpart chatbot repeated the defamatory material.¹⁹⁴

Reasonable foreseeability can be retained as a "crucial limiting principle"¹⁹⁵ of liability in negligence cases involving artificial intelligence as in other cases, but operators and developers should not be heard to say that the risk of harmful emergence altogether (as opposed to a specific occurrence of emergence) is unforeseeable. If this were the case, victims of harm related to the operation of artificial intelligence would be less protected under tort law than victims incurring similar injury or loss at the hands of human tortfeasors.

Asymmetric protection of this kind should be avoided as a matter of legal policy, as it would detract from two of the principal goals of tort law, compensation and prevention of harm.¹⁹⁶ If the threshold of tortious liability is set lower for defendants re-

191. *Supra*, note 77.

192. *Supra*, note 59.

193. Pranshu Verma and Will Oremus, "ChatGPT invented a sexual harassment scandal and named a real law prof as the accused" (*Washington Post*, 5 April 2023), online: <https://www.washingtonpost.com/technology/2023/04/05/chatgpt-lies/>

194. *Ibid.*

195. *Rankin (Rankin's Garage & Sales) v. J.J.*, *supra*, note 77 at para. 23.

196. The potential for asymmetric protection of the law as between victims of artificial intelligence and victims of human tortfeasors was a significant concern of EU policymakers, reflected in one

sponsible for artificial intelligence than for defendants in other tort cases, it will not encourage adequate risk assessments and maintenance of best practices in the field.

The concept of reasonable foreseeability will need to be applied differently than in the general run of negligence cases where artificial intelligence is not a factor. Foreseeability of risk in relation to the use of artificial intelligence must be considered not only in terms of particular outcomes, but in terms of the risk of unpredictability itself. Rather than be applied with regard to the particular manner in which risk materialized in the given case, it should treat the risk of unpredictability and emergence as known risks that potentially give rise to unknown ones.

In order to continue to act as a limiting principle preventing indeterminate liability, however, the determination of what was reasonably foreseeable should be made with reference to the known attributes of the system in question at the relevant time, the use cases for which the system was intended, and known or predictable alternate uses (including predictable misuse) of the system. In our view, these considerations delineate the scope of potential harm that designers, developers, and operators should be expected to consider in assessing risk and taking measures to avert it.

The Project Committee has not attempted to re-formulate the test of reasonable foreseeability in a semantic manner. Instead, our tentative recommendation calls for the application of the test to be modified in cases involving harm causally connected with artificial intelligence in order to recognize the potential for “known unknowns”¹⁹⁷ when a decision is made to deploy artificial intelligence or to make it available for deployment.

The Project Committee tentatively recommends:

7. Where harm results from the operation of an artificial intelligence system and a claim based on negligence is made, the test of reasonable foreseeability of harm should

of the governing principles listed in the resolution of the European Parliament of 20 October 2020 for a proposed civil liability regime for artificial intelligence, *supra*, note 88:

- 7 Citizens should be entitled to the same level of protection and rights, irrespective of whether the harm is caused by an AI-system or not, or if it takes place physically or virtually, so that their confidence in the new technology is strengthened.

The asymmetric protection issue is addressed at some length in the report of the European Commission’s High Level Expert Group, *supra*, note 85.

197. See Thomasen, *supra*, note 146 at 118.

be applied with regard to the risk that the system might behave unpredictably to cause harm in an unknown manner, taking into account

(a) attributes of the system known at the relevant time;

(b) intended use of the system; and

(c) known or predictable alternate uses of the system.

Chapter 5. Standard of Care

A. How the Standard of Care Is Set

In order to determine whether a duty of care has been met or breached in a negligence claim, the court compares the defendant's conduct against the standard of care. The court must make a determination regarding what the content of the standard of care is in the circumstances of the case. This usually involves an assessment of the degree of care a reasonable person in the defendant's position would exercise. In making that determination, the court can draw upon evidence given in the case before it, including expert evidence, about what the accepted or customary standard is in the industry or profession in question in the action. It can draw upon decisions regarding the standard of care in previous similar cases. The court can also look to relevant legal and regulatory requirements affecting the activity in issue in the case.

The standard of care is a common law concept, and is sometimes set on the basis of policy. Legislation and regulations governing the activity in issue are relevant considerations in setting the standard of care in an individual case, but not decisive. Breach of a statute or regulation applicable to the defendant or the defendant's activities in issue does not automatically lead to a finding of liability. The Supreme Court of Canada has explained the relationship between regulatory enactments and the standard of care as follows:

Legislative standards are relevant to the common law standard of care, but the two are not necessarily co-extensive. The fact that a statute prescribes or prohibits certain activities may constitute evidence of reasonable conduct in a given situation, but it does not extinguish the underlying obligation of reasonableness. See *R. in right of Canada v. Saskatchewan Wheat Pool*, [1983] 1 S.C.R. 205. Thus, a statutory breach does not automatically give rise to civil liability; it is merely some evidence of negligence. See, e.g., *Stewart v. Pettie*, [1995] 1 S.C.R. 131, at para. 36, and *Saskatchewan Wheat Pool*, at p. 225. By the same token, mere compliance with a statute does not, in and of itself, preclude a finding of civil liability. See...[citation omitted]. Statutory standards can, however, be highly relevant to the assessment of reasonable conduct in a particular case, and in fact may render reasonable an act or omission which would otherwise appear to be negligent. This allows courts to consider the legislative framework in which people and companies must

operate, while at the same time recognizing that one cannot avoid the underlying obligation of reasonable care simply by discharging statutory duties.¹⁹⁸

B. A Largely Open Playing Field – For the Time Being

As yet, there is little regulation of the field of artificial intelligence in Canada. That may soon change under the proposed *Canadian Artificial Intelligence and Data Act* (AIDA)¹⁹⁹ which is before Parliament at the time of writing, but the Act will not apply to all artificial intelligence systems and the extent of the regulatory standards that will be contained in regulations under the Act remains to be seen. The Act will require any “person responsible” for a “high-impact system” (yet to be defined in regulations) to establish measures in accordance with the regulations to identify, assess, and mitigate risks of harm or biased output that could result from the use of the system.²⁰⁰

A Companion Document to AIDA issued by the Government of Canada explains the general policy underlying the Act. It indicates the public will be consulted on the content of regulations before and after they are drafted, but lists factors that will likely find their way into the regulations for determining whether a system would qualify as “high-impact,” and outlines the requirements for a high-impact system.²⁰¹ The requirements outlined in the AIDA Companion Document for high-impact systems align generally with the themes found in recent policy framework documents for artificial intelligence issued by other countries: human oversight and monitoring, transparency, fairness and equity, safety, accountability, validity (performing consistently with intended objectives) and robustness (stability and resilience in a variety of circumstances.)²⁰²

The Government of Canada *Directive on Automated Decision-Making*²⁰³ applies to federal artificial intelligence systems that make recommendations or decisions about individuals outside government, and contain risk reduction and management

198. *Ryan v. Victoria (City)*, [1997] 1 S.C.R. 201, at para. 29.

199. *Supra*, note 133.

200. *Ibid.*, s. 8.

201. *Supra*, note 184.

202. *Ibid.*

203. *Supra*, note 6 .

provisions that have value as models that could be transferred to other public and private sector settings.

The relatively few regulatory standards for artificial intelligence in place at the present time and lack of precedents in case law will mean that Canadian courts will be called upon to break new ground in setting the standard of care in the early cases that come forward. Some guidance can be drawn from international regulatory precedents that have begun to appear, such as the proposed EU *AI Act*,²⁰⁴ the proposed U.S. *Algorithmic Accountability Act*,²⁰⁵ public policy framework documents such as the *White House Blueprint for an AI Bill of Rights*,²⁰⁶ NIST *Artificial Intelligence Risk Management Framework*,²⁰⁷ the UK Government policy paper *A pro-innovation approach to AI regulation*,²⁰⁸ and published voluntary guidelines issued by various scientific and industry bodies. These sources may have considerable advisory and persuasive value pending the development of a body of Canadian precedent, and possibly afterwards as well.

Courts will usually look to accepted practice within an industry or profession in making a determination about the standard of care in cases involving that industry or profession. Industry and practice standards are relevant to the standard of care in negligence in the same way as regulatory enactments; they have persuasive value, but are not decisive.²⁰⁹ A court may make a finding of negligence against a defendant who has conformed to an industry practice if it finds the practice creates an unreasonable risk and the standard of care should be set higher.²¹⁰

While our tentative recommendations on the basis of liability and proof of fault do not distinguish between systems on the basis of risk level for reasons given earlier,

204. *Supra*, note 10.

205. H.R. 6580, 117th Congress, 2nd Sess.

206. See online: <https://www.whitehouse.gov/wp-content/uploads/2022/10/Blueprint-for-an-AI-Bill-of-Rights.pdf>.

207. National Institute of Standards and Technology, *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*, NIST AI 100-1 (Washington: Dept. of Commerce, January 2023).

208. Department for Science, Innovation and Technology (UK), *A pro-innovation approach to AI regulation*, Policy paper (29 March 2023), online: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/1146542/a_pro-innovation_approach_to_AI_regulation.pdf.

209. *Ding v. Prévost*, 2022 BCSC 215 at para. 211; *Zsoldos v. CPR*, 2009 ONCA 55 at para. 43.

210. *Zsoldos v. CPR*, 2009 ONCA 55 at para. 43.

the level of risk associated with a system undoubtedly will be among the factors weighed by courts in determining the standard of care applicable in a given case.²¹¹ In so doing, courts should not lose sight of a reasonable balance between risk and benefit, and the need to avoid discouraging innovation by imposing an excessively onerous burden of liability.

Artificial intelligence, like other digital technology operating through the internet, is not limited by territorial boundaries in its reach and its effects. Despite the fact that there are very few hard and fast rules governing the field of artificial intelligence at the present time, an interjurisdictional consensus on basic elements of good practice in the development and use of artificial intelligence is emerging. This can be seen in the similarities between the contents of policy framework documents in various countries and in relatively uniform ethical and technical norms that are generally recognized in the research community and software development industries.

The convergence of norms is occurring despite the marked differences between jurisdictions in approach to regulation of artificial intelligence. The strongly prescriptive regulatory approach taken in the EU contrasts with the recently announced policy of the UK Government to minimize new regulation in the interests of promoting innovation.²¹² At the same time, the UK policy paper calls for existing sectoral regulators in the country to apply principles broadly similar to those embraced elsewhere, and notes the importance of international alignment and “interoperability with international regulatory frameworks.”²¹³

As the interjurisdictional consensus is evolving and is still in a relatively early stage of development, we will not make a recommendation in this chapter about the substance of the standard of care in a negligence case involving artificial intelligence. Instead, our recommendation urges that courts set the standard of care in litigation concerning artificial intelligence with a worldwide outlook in order to take account of interjurisdictionally recognized best practices and the desirability of encouraging responsible innovation.

The balance of this chapter apart from the tentative recommendation at its end consists of a list of some of the elements of good practice in the design, development, and operation of artificial intelligence systems that appear with relative consistency

211. See text under the subheading “3. Recommendation on Relief Against Difficulties of Proof in Appropriate Cases” in Chapter 4.

212. *Supra*, note 208.

213. *Ibid.*, at 43.

in an international cross-section of regulatory policy documents and instruments, and thus appear to enjoy wide recognition. We think Canadian courts would do well to consider them when called upon to set a standard of care in civil litigation arising from the operation of artificial intelligence.

C. Widely Recognized Elements of Good Practice

1. The Design, Development, Training, and Testing Phases

(a) Transparency

Transparency regarding what a system is intended to do, how the system does it, who brought it into being, and who is using it is considered of paramount importance. At a minimum, there should be disclosure of:

- the identities of the producer, owner, and user of the system;
- capabilities and limitations of the system;
- accuracy levels for which the system has been validated, the metrics used to validate the accuracy levels, and other characteristics of performance;
- known risks from intended use and foreseeable patterns of misuse;
- where the system has a machine – human interface, the fact that the human is interacting with an automated system;
- at higher risk levels, notification to persons affected by an automated decision that the decision is being made by automation.²¹⁴

Thorough documentation of the system’s features, function, and risks, prepared in conformity with transparency requirements, is essential. The system documenta-

214. In Quebec, anyone whose personal information is used by an enterprise to render a decision based exclusively on automated processing of that information will have the right from 22 September 2023 to be informed of the personal information used to render the decision, the reasons, principal factors, and parameters leading to the decision, and of the right to have the personal information corrected: s. 12.1 of *An act respecting the protection of personal information in the private sector*, CQLR c. P-39.1, s. 12.1, as am. by S.Q. 2021, c. 25, s. 110, in force 22 September 2023.

tion should be available to all personnel operating the system or interpreting its output.

At least early in the lifecycle of an AI system, if not throughout it, the system documentation should include information on in-use monitoring reflecting the risk factor in the use of the system.

(b) Vital design features

Systems should be designed to achieve consistent levels of accuracy, robustness, and cybersecurity appropriate for their intended uses.

The EU *AI Act* refers to the necessity of incorporating features to mitigate the possibility of feedback loops whereby biased outputs become inputs in later operation of systems that continue to learn while in use after deployment.

(c) Data Quality and Data Governance

Good data governance and management practices require in particular:

- the splitting of training, validation and testing datasets;
- ensuring relevance, completeness, robustness of data;²¹⁵
- measures to eliminate bias in input data;
- statistically rigorous data collection conforming to legal requirements;
- compliance with standard practices in data cleaning;²¹⁶

215. “Robustness of data” as used here means that the data is accurate and representative on the average with respect to the population to which it relates. The term “robustness” is used in more than one sense in connection with artificial intelligence. It is also used to denote the ability of an artificial intelligence system to generate accurate and reliable output despite the presence of some inaccurate or irrelevant data (so-called “noise”) within a dataset, or the ability of a system to perform well under varying parameters of use. In another sense, “robustness” also refers to the resilience of a system to cyberattack.

216. Data cleaning or “scrubbing” is “the process of fixing or removing incorrect, corrupted, incorrectly formatted, duplicate, or incomplete data within a dataset”: Tableau, online: <https://www.tableau.com/learn/articles/what-is-data-cleaning>.

- documentation of the data governance process.

(d) Continuous risk assessment

Continuous risk assessment throughout a system's lifecycle, covering, *inter alia*:

- risk not only arising in the course of intended use, but also from foreseeable mis-use;
- cybersecurity risks, including adversarial exploitation and data poisoning (malicious attack on the system through adulterating or manipulating input data);
- the possibility of bias in output;
- consideration of various risk mitigation measures referred to variously as: "outlier detection," "anomaly detection," "data sanitation," and employing ones recognized as effective and appropriate;
- validation that large-scale statistics on deployment resemble those from small-scale test sets
- assessment and mitigation measures for risk to children and other vulnerable populations.

(e) Independent validation

Pre-deployment review of a new system by at least one expert who has not been involved in its design and development is seen as desirable or imperative, depending on the jurisdiction and intended use.

Whether engagement of a knowledgeable at-arm's-length third party to review or independently validate a system should be regarded as part of the legal standard of care should depend on the risk associated with use of the system and the risk factors in the particular use case.

(f) Articulation of appropriate human involvement

The design of an artificial intelligence system needs to include an articulation of where a human will be involved, what decisions the human will make, and what data the human will use in making them.

(g) Logging Capability

The capability of a system to create an automatic record of its own operation, including features of transparency, is of high importance.

(h) Monitoring

Monitoring of system performance by the developer and operators throughout the system's lifecycle.

Performance monitoring by a third party at arm's length from the developers and suppliers of a system with sufficient expertise to evaluate system performance is also necessary in areas of high risk, including the medical domain.

Effective monitoring requires a continuing flow of information on system performance between users, developers, and any third party engaged in monitoring the system.

(i) Updating

Updating and upgrading to correct problems as they arise from monitoring of the system while deployed in actual use.

2. Operation in Actual Use

(a) Risk Assessment and Mitigation

The operator needs to carry out its own risk assessment to complement that of the system developers, as the operator designates and controls the specific setting in which the system will be used.

Regulatory and governance models emphasize that risk assessment needs to be an iterative process continuing throughout the lifecycle of the system.

The risk assessment should comprise what the developer's risk assessment would cover and at least the following in addition:

- cybersecurity risks specific to the operator's organization and other activities;
- data privacy risks and mitigation measures in the event of a breach;

- impacts on affected populations, including those that do not consist of direct users of the operator's system.

(b) Compliance with Developer's Recommendations and Terms of Service

Compliance with the recommendations of the developer for operating the system within its design parameters is a basic element of good practice. Off-label use of a system would presumably heighten risk.

If the operator is involved in the development or modification of a system, it should fall to the operator to provide all necessary information in its possession or power to clarify the use case for the benefit of designers and developers.

(c) Transparency

Persons interacting with the system or who are directly affected by its output should be made aware by the operator in clear, readily understandable terms that they are dealing with an automated system, are informed of its purpose, and have its output explained to them.²¹⁷

If a system generates or manipulates visual or audible content to resemble existing persons, objects, places, or representations of events and could deceive someone into thinking the content was authentic, original, or accurate, the operator must at least have made it known or obvious that the content has been artificially generated or manipulated ("deepfaked"). It must also be open to a court to find artificial generation or manipulation to be a breach of the standard of care in itself, depending on the circumstances. For example, if the deepfaking is done for deceptive, malevolent, or abusive purposes,

Legally authorized or permissible security and surveillance systems and systems used in law enforcement to detect or investigate criminal conduct may be exceptions to basic transparency requirements, but may sometimes require substantially higher standards of care.

217. Section 12.1 of Québec's *Act respecting the protection of personal information in the private sector*, CQLR c. P-39.1, as enacted by S.Q. 2021, c. 25, s. 110 will require that organizations provide notice to consumers of an automated decision using their personal information at the time the decision is made. It confers several additional related rights on consumers, including the right to be informed of the reasons, principal factors, and parameters taken into account in the decision, and to correct the personal information used to make it.

(d) Monitoring of System Performance

Continuous monitoring of performance is needed throughout the life of the system, including co-operating with the developer's post-deployment monitoring process and supplying performance data for that process as necessary. Monitoring with respect to risk mitigation measures will become a statutory obligation for operators of high-impact systems (as they may come to be defined) under section 9 of the proposed Canadian *Artificial intelligence and Data Act*.

System performance monitoring should be both a human and automated process where circumstances warrant, so that the human and automated monitors act as a check and verification of each other. The performance of human personnel responsible for system oversight should be included in the overall monitoring process.

Metrics must be created to measure the system's performance. The metrics themselves should be subject to periodic evaluation. A range within which the system is intended to operate must be specified, so as to enable monitoring for model drift.

Scaling issues (reliability of the system under expanded loads and parameters) also need to be covered in the monitoring process.

Data quality must be continually monitored for bias and to ensure the current training data remains relevant to the intended purpose of the system.

Re-use of data in new contexts requires extra care to prevent spreading and amplification of harms that the original data had the potential to produce. If outputs are used as inputs in subsequent use of a system, such as may occur with systems that learn continuously in operation, data that is derived or inferred from prior outputs of the same system should be treated as high-risk for feedback loops that could compound and amplify bias or inaccuracy in the previous outputs.

(e) System Maintenance

Appropriate preventive and mitigative maintenance could include:

- updating the system as recommended by the developer, or to correct problems identified through continuous monitoring;

- in response to issues detected, including model drift and change in external conditions, re-training the system or modifying the system model;²¹⁸
- retaining a baseline version of the system model to compare results of later updated, re-trained, or modified versions;
- reverting to an earlier baseline system that worked satisfactorily if reconfigurations and retrained models prove unreliable.²¹⁹

(f) Privacy

Compliance with applicable laws on data privacy, as well as adherence to sector-specific ethical standards.

(g) Training

Appropriate training for personnel responsible for the operation of the system.

(h) Logs and Other Record-keeping

Creation and retention of logs in keeping with any regulatory requirements and protocols established by the operator or recommended by the developer of the system.

Retention of monitoring results, along with records of any corrective steps taken, so that the history of any problems with the system and mitigative efforts can be tracked.

Section 10(1) of the proposed Canadian *Artificial Intelligence and Data Act* would require “persons responsible” for high-impact systems, including operators, to keep records relating to the establishment of data anonymization, risk assessment, risk mitigation, and monitoring protocols. Additional record-keeping requirements are likely to be prescribed by regulation following passage of the Act.

(i) Organizational governance framework for artificial intelligence systems

The operator’s organization should have clear governance structures and procedures for oversight and risk mitigation. Responsibilities of designated individuals or groups for these functions should be clearly defined.

218. WH Blueprint, *supra*, note 206 at 19.

219. *Ibid.*

Responsibility for decisions concerning problems and potential shutdown or “roll-back” from any change should rest at a level in the operator’s organization that will allow fast response, with anyone holding this decision-making authority being fully informed with respect to risk potential.

D. Tentative Recommendation

The Project Committee tentatively recommends:

8. (1) The standard of care in litigation concerning artificial intelligence should be set so as to encourage responsible innovation and development, encompassing reasonable care to avoid foreseeable injury or loss.

(2) In setting the standard of care, courts should employ a broad interjurisdictional perspective, extending where applicable to

(a) widely recognized national and international best practices;²²⁰

(b) national and international regulatory standards

in the design, development, and operation of artificial intelligence systems.

220. Including the examples listed in Section C of this chapter.

Chapter 6. Algorithmic Discrimination and Civil Liability

A. Bias: A Recognized Problem in Artificial Intelligence

The potential for artificial intelligence to produce biased outputs with discriminatory effects is widely recognized as being among the most significant legal and ethical problems associated with this technology. Virtually all public policy and regulatory initiatives across the jurisdictions emphasize the necessity of eliminating, averting, or mitigating biased output.²²¹

The term “bias” is used in at least two senses in connection with artificial intelligence outputs.²²² One is technical and is drawn from the field of statistics: a systematic error that renders a body of data incorrect on average with respect to the population that is sampled.²²³ The other is popular or non-technical: skewing of output in a way that would be commonly perceived as inconsistent with values of fairness that

221. See, for example, the proposed Canadian *Artificial Intelligence and Data Act*, *supra*, note 133, ss. 5 (definition of ‘biased output’), 8 and 9; Government of Canada *Directive on Automated Decision-Making*, *supra* note 6, para. 6.3.1; European Union *AI Act*, *supra*, note 10, art. 10, para. 2(f) and art. 15, para. 3; White House *Blueprint for an AI Bill of Rights*, *supra*, note 206 at 23-29; UK Dept. for Science, Innovation & Technology, *A pro-innovation approach to AI regulation*, Cmnd 815 (London: 29 March 2023), online: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/1146542/a_pro-innovation_approach_to_AI_regulation.pdf.

222. Centre for Data Ethics and Innovation, *Review into bias in algorithmic decision-making* (London: UK Government, Dept. for Science, Innovation and Technology, 27 November 2020), online: <https://www.gov.uk/government/publications/cdei-publishes-review-into-bias-in-algorithmic-decision-making/main-report-cdei-review-into-bias-in-algorithmic-decision-making>, section 2.3.

223. *Ibid.* A systematic error is one that is not determined by chance and is introduced by inaccuracy of observation or measurement inherent in a system: Merriam-Webster Dictionary online: <https://www.merriam-webster.com/dictionary/systematic%20error>.

are widely held in society.²²⁴ This is the sense in which “bias” is predominantly used in this consultation paper.

There are numerous ways in which bias can enter automated decision-making processes. It may be present in the design of the algorithm on which the system is based.²²⁵ It may be present in the data used to train or test a system.²²⁶ It may be introduced by input data collected and processed during the operational phase once the system has been deployed that is unrepresentative or improperly derived.²²⁷ It can also result from a failure of human oversight through assumptions based on conscious or unconscious biases.²²⁸ As mentioned in an earlier chapter, feedback loops can result from training new versions of a system using as inputs biased data generated by a previous version. Feedback loops replicate and may amplify the effect of the bias in the original output.²²⁹

Bias may lead easily to discrimination, by which we mean adverse treatment that lacks objective justification. Many examples exist of artificial intelligence systems generating discriminatory results. One of the most notorious involved a system called Correctional Offender Management Profiling for Alternative Sanctions (COMPAS), a tool developed to predict recidivism for the purpose of granting bail or probation and which was used for a time by some U.S. courts and law enforcement authorities.

While race was not an input variable for risk assessment, COMPAS disproportionately predicted a higher risk of re-offending for Afro-American accused and was incorrect in its predictions at approximately the same rate, but in opposite ways for Afro-American and Caucasian accused.²³⁰ It predicted 23.5% of Caucasian accused who did not re-offend as higher risk, and 44.9% of Afro-American accused.²³¹ It labelled

224. *Supra*, note 222.

225. *Ibid.*

226. *Ibid.*

227. *Ibid.*

228. *Ibid.*

229. *Ibid.*

230. Julia Angwin, Surya Mattu and Lauren Kirchner, “There’s software used across the country to predict future criminals. And it’s biased against blacks.” *ProPublica* (23 May 2016), online: <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>.

231. *Ibid.*

47.7% of Caucasian accused who did re-offend as lower risk, and 28.0% of Afro-American accused who did so at lower risk.²³²

An algorithm used in Arkansas to allot home care time to persons receiving disability benefits cut the care aide hours under a state program for a large number of recipients without any change having taken place in their circumstances. The algorithm analyzed approximately 60 factors. The developer considered it a more “rational” system than subjective assessment by humans. A very small difference in a numerical score for some factors, however, could make a large difference in the allotment of care hours for persons in need of home care.

Another notorious example of algorithmic discrimination concerned a system tested by Amazon to vet candidates for employment. The data used to train the system reflected applications submitted to Amazon over the previous 10 years. The majority of applications had been male. This evidently resulted in the system learning to classify female gender as a negative trait and reject applications by women. The discovery of this bias and other problems with the data led Amazon to abandon the system.²³³

Discriminatory effects from artificial intelligence outputs may be much more subtle than in the COMPAS and the Amazon examples. They will often not relate directly to the legally prohibited grounds of discrimination. Yet they will often present discrimination that is socially unacceptable.

Ostensibly neutral input variables like postal codes, income, and educational level may be correlated with demographic patterns and become proxies for race and ethnicity, leading to outputs that impose disadvantages on racialized and minority populations.²³⁴ For example, an algorithm used to determine where Amazon Prime same-day delivery service could be economically provided in the U.S was reportedly found to exclude urban areas with a high proportion of Afro-American residents. Race was not an input variable, of course. One of the input variables evidently was

232. *Ibid.*

233. Jeffrey Dastin, “Amazon scraps secret AI recruiting tool that showed bias against women” Reuters, 10 October 2018, online: <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G>.

234. Jacquelyn Burkell and Jane Bailey, “Unlawful Distinctions? Canadian Human rights law and Algorithmic Bias” (2016/2018) Can. Y.B. Human Rights 217 at 219; Bathaee, *supra*, note 59 at 920; *White House Blueprint for an Artificial Intelligence Bill of Rights*, *supra*, note 206 at 26.

the proximity of addresses to product warehouses, however.²³⁵ As a result of prevailing patterns in urban settlement and housing, this served as a proxy for race.

Sometimes those who are adversely affected in a discriminatory manner will have recourse under existing anti-discrimination legal frameworks to have those effects corrected or to obtain compensation. In other cases they will not. Canadian anti-discrimination law only takes account of complaints based on a finite list of prohibited grounds in the relevant enactment. If the ground of differentiation is not included in that list, it is not recognized as discrimination. The equality rights section of the *Canadian Charter of Rights and Freedoms* (section 15) takes account of a few forms of discrimination that are not specifically listed, but a remedy under the Charter is only available against the Crown or a governmental entity.

This chapter concerns a gap that is arguably present in the law, namely the absence of a civil remedy for harm in the form of unjustifiable adverse treatment resulting from the application of artificial intelligence that is not capable of redress within the human rights framework. It raises the question whether the law of tort should fill that gap.

Some explanation of Canadian law relating to discrimination is necessary before going on.

B. Overview of Canadian Anti-discrimination Law

1. A Legal Definition of Discrimination

The Supreme Court of Canada approved the following definition in *CN v. Canada (Canadian Human Rights Commission)*, contained in a Royal Commission report on employment equity:

“Discrimination ... means practices or attitudes that have, whether by design or impact, the effect of limiting an individual's or a group's right to the opportunities generally available because of attributed rather than actual characteristics”²³⁶

235. David Ingold and Spencer Soper, “Amazon Doesn’t Consider the Race of Its Customers. Should It?” *Bloomberg*, 21 April 2016, online: <https://www.bloomberg.com/graphics/2016-amazon-same-day/>. Amazon’s data analysis may not have involved artificial intelligence as such, but the example indicates how algorithmic discrimination by proxy can arise.

Some forms of discrimination are unconstitutional. Many of these would overlap with forms of discrimination that are prohibited under federal, provincial, or territorial human rights legislation. An important difference is that constitutional remedies may only be sought against a governmental defendant or a public entity that is an emanation of government, while the human rights legislation applies to governmental and private conduct.

2. Unconstitutional discrimination under the equality rights section of the Charter

Section 15(1) of the *Canadian Charter of Rights and Freedoms* (“the Charter”) declares that every individual has the right to equal protection and equal benefit of the law “without discrimination and, *in particular, without discrimination based on race, national or ethnic origin, colour, religion, sex, age or mental or physical disability.*” The guarantee of equality in section 15(1) is qualified by section 1 of the Charter, which subjects the rights and freedoms in it to “such reasonable limits prescribed by law as can be demonstrably justified in a free and democratic society.”

Legislation and governmental action that is inconsistent with the constitutional guarantee of equality in section 15(1) and not saved by demonstrably justified “reasonable limits” under section 1 is invalid. The onus to demonstrate justification on the basis of reasonable limits is on the party defending the challenged law or action of a government or public authority.

While Canadian courts have wide powers to grant remedies based on the Charter, the Charter does not apply to the conduct of private individuals or private corporate entities, as stated earlier.

Section 15(1) ostensibly leaves the term “discrimination” open-ended, but it is judicially interpreted to recognize claims of discrimination based on the expressly listed grounds of race, national or ethnic origin, colour, religion, sex, age or mental or physical disability and what the courts have described as “analogous grounds.” This is obviously a rational interpretation of the term “discrimination” to prevent any law from being overturned merely because it differentiates between classes of individuals in some manner.

As currently interpreted, the test of discrimination under section 15(1) of the Charter is whether:

236. [1987] 1 S.C.R. 1114 at 1138-39. This quotation was actually part of a longer attempt to explain the term “systemic discrimination.”

(a) the challenged law or governmental action creates a distinction on its face or in its impact on the basis of a ground listed in s. 15(1) or an analogous ground; and

(b) the distinction imposes a burden or denies a benefit in a manner having the effect of reinforcing, perpetuating, or exacerbating disadvantage.²³⁷

Analogous grounds are considered to be those based on immutable personal characteristics or ones that are changeable only at unacceptable cost to personal identity, which the federal or a provincial government has no legitimate interest in expecting an individual to change in order to receive equal treatment under the law.²³⁸ They are ones that are actually immutable or “constructively immutable,” such as religion.²³⁹

The Supreme Court of Canada has recognized only four analogous grounds of unconstitutional discrimination besides those expressly listed in section 15(1): citizenship, marital status, sexual orientation, and off-reserve residence of Indigenous peoples.²⁴⁰ Other courts have recognized adopted status, the manner of conception of a child, receipt of public assistance, and parental status as analogous grounds.²⁴¹

3. Discrimination under human rights legislation

Federal, provincial, and territorial human rights legislation provides another forum for anti-discrimination claims. This legislation, unlike section 15(1) of the Charter, is binding on public entities and private persons alike.

Rather than providing court-based remedies, the pattern of current Canadian human rights legislation is one in which complaints of discrimination consisting of an alleged breach of the legislation are decided by a quasi-judicial tribunal. There may be a human rights commission or directorate interposed between the complainant and

237. *Québec (Attorney General) v. Alliance du personnel professionnel et technique de la santé et des services sociaux*, 2018 SCC 17 at paras 25-28; *Centrale des syndicats du Québec v. Québec (Attorney General)*, 2018 SCC 18 at para. 22; *Fraser v. Canada*, 2020 SCC 28 at para. 27.

238. *Corbière v. Canada*, [1999] 2 S.C.R. 203 at para. 13;

239. *Ibid.*

240. Hogg, *Constitutional Law of Canada*, 5th ed. Supplemented (Toronto: Thomson Reuters, 2022) at 55-56.

241. Errol Mendes and Stéphane Beaulac, eds. *Canadian Charter of Rights and Freedoms*, 5th ed. (Markham: LexisNexis, 2013) at 987.

the tribunal. The human rights commission or directorate has an investigative and conciliation or mediation role as well as a general mandate to promote human rights. The commission may decide which complaints that it cannot settle have enough merit to proceed to adjudication by the tribunal. Thus, a complainant may not have an automatic right of access to adjudication, depending on the jurisdiction where the complaint arose. British Columbia is one of the jurisdictions that allows direct access by complainants to its Human Rights Tribunal.²⁴²

In order to be unlawful under human rights legislation, the discrimination that is the subject of a complaint must be based on a ground expressly prohibited by the legislation. An ostensibly neutral provision or rule may be discriminatory if it disproportionately affects a class of persons on the basis of a prohibited ground.²⁴³ Intention to discriminate is not required to constitute a violation of human rights legislation, if the effect of the challenged action is discriminatory.²⁴⁴

Under the *Canadian Human Rights Act*, which applies to the federal government, federal agencies and the federally regulated private sector (including chartered banks, interprovincial railways, airlines, and marine industries), the prohibited grounds of discrimination are *race, national or ethnic origin, colour, religion, age, sex, sexual orientation, gender identity or expression, marital status, family status, genetic characteristics, disability, and conviction for an offence for which a pardon has been granted or in respect of which a record suspension has been ordered*.²⁴⁵

The main prohibited grounds of discrimination under the British Columbia *Human Rights Code*, which are fairly typical of those in other provincial and territorial human rights enactments, are *Indigenous identity, race, colour, ancestry, place of origin, religion, marital status, family status, physical or mental disability, sex, sexual orientation, gender identity or expression, or age of that person or that group or class of persons*.²⁴⁶

242. A single Human Rights Commissioner is appointed in British Columbia as an officer of the Legislature. The commissioner is responsible for promoting and protecting human rights, but does not control the case flow to the Human Rights Tribunal.

243. *British Columbia (Public Service Commission) v. British Columbia Government Employees Union*, [1999] 3 S.C.R. 3.

244. *Ontario Human Rights Commission v. Simpsons-Sears*, [1985] 2 S.C.R. 536.

245. R.S.C. 1985, c. H-6, s. 3(1).

246. R.S.B.C. 1996, c. 210, ss. 7(1), 8(1), 9(1).

For the purpose of specific provisions, other prohibited grounds under the British Columbia *Human Rights Code* are *the lawful source of income of a person or class* concerning discrimination in residential tenancies,²⁴⁷ and in relation to employment, *a criminal or summary conviction offence that is unrelated to the employment or to the intended employment*²⁴⁸ or the same that is *unrelated to the membership or intended membership*²⁴⁹ of the complainant in a trade union.

Human rights tribunals are typically empowered to make orders to cease or correct discrimination when they find a complaint justified, to adopt a plan or other measure to prevent future discrimination, and award monetary compensation to the complainant up to a ceiling amount.

The human rights tribunals have no jurisdiction over discrimination that is not based on a listed prohibited ground in the enactment under which they function.

4. No general tort of discrimination

In *Seneca College v. Bhadauria*, (“*Bhadauria*”) the Supreme Court of Canada held that there is no common law tort of discrimination, nor does a breach of human rights legislation give rise to an actionable tort.²⁵⁰ The Supreme Court has maintained this position since *Bhadauria* was decided in 1981.²⁵¹ British Columbia courts continue to hold on the strength of the Supreme Court decisions that the existence of the comprehensive statutory scheme under the provincial *Human Rights Code* forecloses the possibility of a free-standing tort of discrimination at common law.²⁵²

British Columbia enacted a statute called the *Civil Rights Protection Act*²⁵³ (“CRPA”) shortly after *Bhadauria* was decided. It makes any conduct or communication a tort,

247. *Ibid.*, s. 10(1).

248. *Ibid.*, s. 12(1).

249. *Ibid.*, s. 14.

250. [1981] 2 S.C.R. 181.

251. *University of British Columbia v. Berg*, [1993] 2 S.C.R. 353; *Honda Canada Inc. v. Keays*, 2008 SCC 39, at paras. 65-67; *Nevson Resources Ltd. v. Araya*, 2020 S.C.R. 5, [2020] 1 S.C.R. 166, at para. 240.

252. *Olena v. Royal Columbia Hospital*, 2017 BCSC 975 at para 12, aff’d 2018 BCCA 349; *Gichuru v. Law Society of British Columbia*, 2014 BCCA 396 at para, 103; *Schultz v. Beacon Roofing Supply Canada Company*, 2016 BCSC 1475.

253. R.S.B.C. 1996, c. 49. This Act seems not to have been passed specifically in response to *Bhadauria*, but to have been primarily aimed at Ku Klux Klan activity in British Columbia that took place

actionable without proof of damage, if it has as its purpose interference with the civil rights of a person or class by promoting (a) hatred or contempt of a person or class, or (b) the superiority or inferiority of a person or class, on the basis of colour, race, religion, ethnic origin or place of origin.²⁵⁴ The CRPA allows for awards of damages, exemplary damages, and injunctions.²⁵⁵

The statutory tort created by the little-used CRPA provides a remedy only against *intentional* conduct aimed at discrimination based on a few of the same grounds covered by the provincial *Human Rights Code*.²⁵⁶ It conceivably could be invoked if an artificial intelligence system were to be created, modified, or used with the deliberate intention to promote racial or religious hatred or racial superiority. It would have little bearing in relation to artificial intelligence generally because discriminatory effects of artificial intelligence outputs will be unintended in most cases.

C. Discrimination Produced by Artificial Intelligence

1. Algorithmic Discrimination Without Remedy

Artificial intelligence systems can replicate biases in their training and input data that may go undetected in outputs for a considerable time, whether or not this is due to any blameworthiness on the part of developers and operators. The systems detect patterns in data that may lead them to make predictions and recommendations that, although not reached malevolently, are based on factors that may clash with legal or social norms of equality and fairness.

Where outputs of artificial intelligence produce discrimination on grounds covered by human rights legislation or section 15(1) of the *Charter*, those enactments may provide a means of redress. If so, redress would need to be predicated on the proposition that the discriminatory behaviour of the systems could be attributed to their operators, developers, or someone who implements the discriminatory output. This is an unsettled area.²⁵⁷

shortly before its enactment: *Maughan v. UBC*, 2008 BCSC 14 at para. 333; aff'd 2009 BCCA 447; leave to appeal refused [2009] SCCA No. 526.

254. *Ibid.*, s. 2(1).

255. *Ibid.*, ss. 4(1), (3).

256. *Supra*, note 246.

257. See Burkell and Bailey, *supra*, note 234.

Systems are often created to differentiate between human subjects, and to score or assess them on the basis of factors that are not prohibited grounds of discrimination under the Canadian human rights framework. Bias may affect these outputs, but the definition of “biased output” in the proposed Canadian *Artificial Intelligence and Data Act*²⁵⁸ would limit the meaning of that term for the purposes of that Act to output that adversely differentiates individuals on grounds prohibited by the *Canadian Human Rights Act*.²⁵⁹ This seems to preclude any remedy under the regulatory scheme of the *Artificial Intelligence and Data Act* when the basis of differentiation is not a prohibited ground under the federal human rights statute.

A system might be built for a lending institution to predict loan default risk for loan applicants, basing its recommendation to refuse a mortgage loan on analysis of large-scale datasets taking into account factors like the applicant’s current residence in a particular district, because the loan loss ratio on a district-wide basis is one of the variables in the algorithm. It might weigh this variable more heavily than the applicant’s own impeccable credit history. If the loan applicant is refused on basis of an address, the applicant has arguably suffered unfairly discriminatory treatment in relation to opportunities for housing under the definition of “discrimination” adopted by the Supreme Court of Canada in *CN v. Canada (Canadian Human Rights Commission)*, yet would appear to have no human rights remedy.

Two actual cases from the Netherlands and the United Kingdom illustrate discriminatory effects produced by automated decision-making that would probably not attract a human rights-based remedy in a Canadian jurisdiction:

- The Dutch government’s SyRI system designed to predict the likelihood of tax or benefit fraud cross-analyzed data from several government departments to predict an individual’s likelihood to commit tax or benefit fraud. It was found to target poorer neighbourhoods without any evidence of individual wrongdoing. SyRI was shut down by the order of a Dutch court on the basis of inadequate protection for privacy and possible discrimination on the basis of socio-economic or migrant status, which are not prohibited grounds under Canadian human rights legislation.²⁶⁰

258. *Supra*, note 1.

259. The definition of “biased output” was likely limited to grounds of discrimination recognized by the *Canadian Human Rights Act*, *supra*, note 245 for constitutional reasons, namely to avoid an appearance of invading provincial jurisdiction.

260. Jo Henley and Robert Booth, “Welfare surveillance system violates human rights, Dutch court rules” (2020), *The Guardian*, 5 Feb. 2020, online: <https://www.theguardian.com/technology/2020/feb/05/welfare-surveillance-system-violates-human-rights-dutch-court-rules>. The court

- A standardization algorithm briefly used in the U.K. when examinations were cancelled during the pandemic assigned A-level and GCSE grades to students based on the historic performance of individual secondary schools. The system was intended to combat anticipated grade inflation from teachers predicting the marks their own students would have achieved. In England, Wales, and Northern Ireland, 36 per cent of secondary school students received an assigned grade lower than their teacher-predicted ones, and 3 percent two grades lower.²⁶¹ In Scotland, 25 per cent of teacher-predicted grades were reduced.

The algorithm was found to have a tendency to decrease the grades of students in state schools and increase those in private schools because of its performance in relation to smaller population groupings.²⁶² As a result, many students lost out in university placements. A country-wide uproar resulted in apologies from the Westminster and Scottish governments and abandonment of the standardization algorithm.²⁶³

As usage of artificial intelligence expands, there will be more cases of unintended discriminatory effects that do not fit readily into the human rights framework, and the gap in the law into which these cases fall will become more apparent. They will be seen nevertheless as unfair and untenable, calling for some means of legal redress.

2. How Should the Gap in Protection Against Algorithmic Discrimination Be Closed?

(a) Human rights vs. tort remedy

Expanding the human rights framework to accommodate claims of discrimination resulting from the operation of artificial intelligence seems unlikely to be a satisfactory solution. Despite a tendency of some Canadian legislatures to expand the pro-

also criticized the system's "serious lack of transparency."

261. "Sean Coughlan, Katherine Sellgren, Judith Burns, "A-Levels: Anger over 'unfair' results this year" BCC News, 13 August 2020, online: <https://www.bbc.com/news/education-53759832>.

262. Georgina Lee, "Did England exam system favour private schools?" *Channel 4 News*, 17 August 2020, online: <https://www.channel4.com/news/factcheck/factcheck-did-england-exam-system-favour-private-schools>.

263. "A-levels and GCSEs: How did the exam algorithm work?" BBC News, 20 August 2020, online: <https://www.bbc.com/news/explainers-53807730>.

hibited grounds in human rights statutes, the concept of discrimination under human rights legislation is primarily circumscribed by reference to characteristics of identity that are either “immutable or changeable only at unacceptable cost to personal identity,”²⁶⁴ namely ones such as race, place of origin, ethnicity, age, disability, gender, language, and religion. The kinds of discrimination that artificial intelligence may create will frequently be based on other factors.

Discrimination through artificial intelligence may occur through attribution of characteristics to individuals by the system on the basis of factors that are not immutable such as place of residence, income level, or habits of consumption, and classification of individuals according to the attributed characteristics, such as lack of creditworthiness, or unsuitability for employment. While fitting within the generalized definition of discrimination enunciated by the Supreme Court of Canada in *CN v. Canada (Canadian Human Rights Commission)*, these forms of discrimination do not fit readily into the existing remedial framework of human rights legislation unless the basis of differentiation can be proven to have served as a proxy in the result for a prohibited ground of discrimination such as race, gender, or age, which will not always be the case. Even when a non-prohibited ground has become as a proxy for a prohibited ground of discrimination, requiring a complainant to prove this as a precondition for a human rights remedy would be onerous, and potentially impose a greater resource strain on human rights tribunals and commissions.

These additional difficulties of proof, and the lack of a remedy for some forms of algorithmic discrimination altogether under both the human rights framework and existing tort law, raise an issue of access to justice that will take on increasing importance with the expansion of artificial intelligence. The question arises whether a civil remedy in tort should be made available for discriminatory treatment resulting from the operation of artificial intelligence.

It is arguable that civil courts are better-adapted to deal with the technical complexity and lack of transparency of artificial intelligence decision-making than human rights tribunals because of their more comprehensive pre-trial oral and documentary discovery process. The pre-hearing discovery process of human rights tribunals generally requires only an exchange of documents.

264 *Corbière v. Canada, supra*, note 238 per McLachlin and Bastarache JJ. at para. 13, writing for the majority in the Supreme Court. While said in relation to s. 15(1) of the *Charter*, this description of the grounds on which unequal treatment is proscribed under the Charter is also largely apt with reference to prohibited grounds of discrimination that circumscribe the scope of other Canadian human rights legislation.

Moreover, algorithmic discrimination is likely to have different impacts than those that human rights tribunals typically deal with. Human rights complaints primarily concern affronts to personal dignity and mental distress resulting from them. Algorithmic discrimination is more likely to result in harm of a different kind, such as economic harm from arbitrary refusal of credit to a certain class of persons, or possibly increased exposure to physical health risks because of arbitrary classifications affecting making certain groups of patients ineligible for a public or private health benefit. These are closer to the kinds of claims that civil courts are accustomed to ruling upon and quantifying.

A tort of algorithmic discrimination resulting from failure to take reasonable care would encourage designers, developers, and operators of artificial intelligence systems to take reasonable steps to detect and eliminate unwanted bias in algorithms and data, and monitor outputs for bias amongst other inaccuracies.²⁶⁵ Recall that these measures were identified in Chapter 5 as being widely recognized elements of good practice in the field of artificial intelligence that courts should consider in relation to the standard of care in claims based on negligence.

Common law evolves incrementally at times to recognize new torts. In 2007 the Supreme Court of Canada held that law enforcement authorities could be civilly liable for negligent investigation if compensable harm resulted from their conduct.²⁶⁶ The Ontario Court of Appeal declared a new common law privacy tort of “intrusion upon seclusion” in 2012.²⁶⁷ Ontario courts have recently recognized a new intentional tort of online harassment.²⁶⁸ Alberta courts have recognized new torts of harassment²⁶⁹ and public disclosure of private facts.²⁷⁰ In addressing these forms of harm rather than discrimination and human rights, however, the Ontario and Alberta courts were not faced with a Supreme Court of Canada decision like *Bhadauria*²⁷¹ standing in their way.

265. Standardized techniques and tools exist for detection, measurement, and treatment of unwanted bias in automated systems throughout their lifecycle, e.g., ISO/IEC/TR24027:2021 *Information technology - Artificial intelligence (1) – Bias in AI systems and AI aided decision making* issued by the International Standards Organization (ISO).

266. *Hill v. Hamilton-Wentworth Regional Police Services Board*, 2007 SCC 41, [2007] 3 S.C.R. 129.

267. *Jones v. Tsige*, 2012 ONCA 32, at para. 65.

268. *Caplan v. Atas*, 2021 ONSC 670; *385277 Ontario Ltd. v. Gold*, 2021 ONSC 4717.

269. *Alberta Health Services v. Johnston*, 2023 ABKB 209.

270. *E.S. v. Shillington*, 2021 ABQB 739, at para. 68.

271. *Supra*, note 250.

Writers have advanced compelling arguments that the tort of negligence is adaptable to address discrimination resulting from the operation of algorithms, and that *Bhadauria* is not a barrier because it concerned intentional racial discrimination.²⁷² Nevertheless, the provincial superior and appellate courts have shown considerable deference to *Bhadauria* and little interest in recognizing a right to sue in tort for discrimination. Legislation *may* be needed to create a remedy in tort for algorithmic discrimination.²⁷³

The parameters of a remedy for algorithmic discrimination must be carefully set either by the legislature or the courts. If they are not, discrimination could be alleged in relation to virtually any differential treatment whatsoever, opening the door to indeterminate liability. This is among the principal reasons why Canadian courts have rejected the concept of a free-standing tort of discrimination.

(b) The Project Committee's view

(i) Negligent algorithmic discrimination

As algorithmic discrimination is usually unintended, resulting from a failure to eliminate unwanted bias or lack of care in the acquisition and treatment of data, the Project Committee has tentatively concluded that a civil remedy for it should be based on negligence, whether the remedy is enacted as a statutory tort or introduced by judicial decision as an incremental change in the common law.

We would consider the essence of the proposed cause of action as a failure to take reasonable steps to detect and correct biased output of an artificial intelligence system or other algorithmic process, resulting in discrimination against a person or class that is illegal or not warranted by reasonable business or industry practice. Discrimination that is illegal might also be a breach of human rights legislation, in which case there might be a dual remedy, but only one loss for which compensation could be awarded.

272. Cynthia Khoo, "Missing the Unintended Forest despite the Deliberately Planted Trees: Reasonable Foreseeability and Legal Recognition of Platform Algorithm-Facilitated Emergent Systemic Harm to Marginalized Communities" (Paper delivered at We Robot 2020), online: <https://drive.google.com/file/d/1xciaKKgI9QlsMM36ukK-LchG0dt0Rfws/view> at 53-55. Khoo distinguishes *Bhadauria* on the ground that it concerned intentional racial discrimination. Ruparelia has also distinguished *Bhadauria* on this ground in arguing that the case does not preclude recognition of a general tort of negligent discrimination: Rakhi Ruparelia, "I Didn't mean It That Way!": Racial Discrimination as Negligence" (2009) 44 S.C.L.R. (2d) 81.

273. Thomasen, *supra*, note 146 at 121.

For a negligence claim to succeed, there must be proof of damage. Maintaining this requirement in connection with *unintended* discrimination that is outside human rights legislation or the *Charter* would serve to prevent open-ended liability for any differentiation affecting an individual or class that flows from the output of a system. In other words, there should be proof of adverse consequences for the plaintiff beyond the mere fact of differential treatment. This is inherent in the legal meaning of “discrimination,” but it would be the adverse consequences for the plaintiff, rather than the mere fact of differentiation, that would give rise to the right to sue. Additionally, the severity of the adverse consequences could be taken into account in assessing damages.

It may be that some kinds of harm that algorithmic discrimination may produce do not fit easily within the range of compensable damage in tort because they may, for example, involve loss of opportunity. These may be speculative heads of damage and difficult to quantify. Given the innumerable ways in which artificial intelligence is already affecting individual lives and will do so to an ever-increasing degree, however, courts should be prepared to take an expansive view of what amounts to actual damage from discrimination by machine.²⁷⁴

(ii) Intentional algorithmic discrimination

Society has a strong interest in discouraging the intentional design or use of artificial intelligence to produce discrimination. There is an equally strong societal interest in encouraging efforts to remove or mitigate the effects of bias that is capable of producing discrimination.

274. Distress resulting from the fact of discrimination can give rise to considerable psychological harm even without additional damage of other kinds. Serious and prolonged psychological harm that goes beyond transitory psychological upset is already recognized as a form of personal injury that may be compensable in tort, if harm of this kind was a reasonably foreseeable consequence of the defendant’s conduct. In *Mustapha v. Culligan of Canada Ltd.*, *supra*, note 64, McLachlin, C.J.C. stated at para. 8:

Generally, a plaintiff who suffers personal injury will be found to have suffered damage. Damage for purposes of this inquiry includes psychological injury. The distinction between physical and mental injury is elusive and arguably artificial in the context of tort.

See also *Saadati v. Moorehead*, [2017] 1 S.C.R. 543, at para. 35. If a plaintiff shows that serious and prolonged psychological harm has resulted from algorithmic discrimination or deliberate failure to address known bias, this could amount to actual damage. There must be more than mere psychological upset that does not result in impairment of cognitive function, interference with daily living, or treatment for emotional symptoms: *Bothwell v. London Health Sciences Centre*, 2023 ONCA 323.

When the algorithmic discrimination was actually intended by defendants, or if the defendants were aware of the bias and deliberately neglected to take reasonable steps to correct it, the deterrent function of tort would be served by relaxing the requirement to prove actual damage in those circumstances, so that the fact of discrimination from biased output would in itself give rise to a right to sue. The Project Committee is prepared to recommend accordingly.²⁷⁵

(c) Tentative recommendations

The Project Committee tentatively recommends:

9. Failure to take reasonable steps to detect and correct biased output of an artificial intelligence system, resulting in discrimination against a person or class that is not warranted by reasonable business or industry practice, or is contrary to law, should amount to actionable negligence if the differential treatment of the plaintiff resulting from the output is accompanied by actual damage.

10. If the defendant intended the result described in Tentative Recommendation 9, or was aware of the biased output and deliberately failed to take reasonable steps to correct it, the differential treatment of the plaintiff should be conclusively presumed to have caused actual damage.

275. A plaintiff may need to weigh the advantage of suing on the basis of intentional algorithmic discrimination and thereby avoiding the need to prove actual damage against reduced prospects of recovery if the defendant's liability insurance will not be available to satisfy a judgment. Liability insurance policies typically exclude damage intentionally inflicted by the insured from coverage. See *Butterfield v. Intact Insurance Company*, 20223 ONCA 246. For this reason, it might be more advantageous for plaintiffs to frame their pleadings in negligence even when the defendants appear to have acted intentionally.

Chapter 7. Conclusion

The law of tort evolved over the course of centuries in the context of harmful interactions between humans in society. Artificial intelligence has brought about a very new context, one in which self-directing machines acting autonomously may cause harm to humans, to their property, or other interests protected by law.

In this consultation paper, the Project Committee has set out a series of tentative recommendations for the adaptation of the law of tort to deal with this new context. They are designed to be implemented either by courts employing the processes by which common law has always developed incrementally to deal with changing conditions and new situations, or by legislatures if policymakers and legislators see fit to enact them to expedite change in the law.

We hope to receive many comments on the contents of this consultation paper. Our final recommendations will be published in a law reform report after full consideration of all the responses to the consultation paper that we receive. We hope they will be numerous and incisive.

List of Tentative Recommendations

1. *Civil liability for harm caused by artificial intelligence should not be based on strict liability. (p. 43)*
2. *Product liability principles should be adapted by analogy to determine rights and liabilities as between a plaintiff harmed by the operation of an artificial intelligence system and defendants who participated in the development of the system and in making it available for use, by treating*
 - (a) *the plaintiff similarly to a plaintiff claiming to have incurred loss or damage from a product comprising multiple components;*
 - (b) *developers of the system as owing a duty of care similar to that owed by a manufacturer of a complex product involving multiple integrated components towards persons or entities who foreseeably could be affected by a defect making the product dangerous;*
 - (c) *developers of components of the system as owing a duty of care similar to that owed by a supplier of a component of a complex product towards persons who foreseeably could be affected by a defect in the component that makes the component and the product in which it is integrated dangerous. (p. 43)*
3. *A person or entity with decision-making authority of a managerial nature over the operation of an artificial intelligence system and thereby is in a position to exert some degree of control over the risk associated with its operation should be treated as an operator for the purpose of civil liability. (p. 46)*
4. *The liability of operators and other persons who provide services in connection with the operation of an artificial intelligence system should be based on general principles of the law of negligence, subject to the recommendations made below. (p. 46)*
5. *In an action for damages or other relief due to harm to persons or property arising from the operation of artificial intelligence, an obligation should rest on a defendant who developed the artificial intelligence system and made it available for use by that defendant or others to*

- (a) provide an explanation for the performance of the system in the circumstances of the case that is consistent with the exercise of reasonable care in the design, development, training, and testing of the system; or*
- (b) disclose sufficient information concerning the design and function of the system to all other parties before the trial of the action that will enable a party to determine the cause of the harmful performance of the system in the circumstances. (p. 61)*

6. If the cause of the harm is otherwise unexplained, and in the absence of an explanation or disclosure referred to in Recommendation 5, the court should be justified in drawing an inference of negligence against the defendants jointly and severally, except in regard to a defendant who rebuts the inference by evidence of the exercise of reasonable care by that defendant in the circumstances to avoid the harm suffered by the plaintiff. (p. 61)

7. Where harm results from the operation of an artificial intelligence system and a claim based on negligence is made, the test of reasonable foreseeability of harm should be applied with regard to the risk that the system might behave unpredictably to cause harm in an unknown manner, taking into account

- (a) attributes of the system known at the relevant time;*
- (b) intended use of the system; and*
- (c) known or predictable alternate uses of the system. (pp. 64-65)*

8. (1) The standard of care in litigation concerning artificial intelligence should be set so as to encourage responsible innovation and development, encompassing reasonable care to avoid foreseeable injury or loss.

(2) In setting the standard of care, courts should employ a broad interjurisdictional perspective, extending where applicable to

- (a) widely recognized national and international best practices;*
- (b) national and international regulatory standards*

in the design, development, and operation of artificial intelligence systems. (p. 78)

9. Failure to take reasonable steps to detect and correct biased output of an artificial intelligence system, resulting in discrimination against a person or class that is not warranted by reasonable business or industry practice, or is contrary to law, should amount to actionable negligence if the differential treatment of the plaintiff resulting from the output is accompanied by actual damage. (p. 94)

10. If the defendant intended the result described in Tentative Recommendation 9, or was aware of the biased output and deliberately failed to take reasonable steps to correct it, the differential treatment of the plaintiff should be conclusively presumed to have caused actual damage. (p. 94)

PRINCIPAL FUNDERS IN 2022

The British Columbia Law Institute expresses its thanks to its funders in 2022:

- Law Foundation of British Columbia
- Ministry of Attorney General
- Alzheimer Society of Canada
- BC Association of Community Response Networks
- The Council to Reduce Elder Abuse (CREA)
- Department of Justice Canada
- Notary Foundation
- Real Estate Foundation of British Columbia
- Simon Fraser University
- Vancouver Foundation

The Institute also reiterates its thanks to all those individuals and firms who have provided financial support for its present and past activities.