BRITISH COLUMBIA LAW INSTITUTE

ARTIFICIAL INTELLIGENCE AND CIVIL LIABILITY PROJECT

Backgrounder

Overview of Project

The purpose of the BCLI Artificial Intelligence and Civil Liability Project is to determine how the rules of tort law need to be adapted to deal with harm to persons and property caused by the autonomous operation of artificial intelligence. Tort is the branch of law concerned with liability to compensate for harm caused to persons and property.

The essential questions the project seeks to answer are: "Who is, or should be, liable for choices made by intelligent machines operating autonomously?" and "How does the law of tort need to be adapted to address the reality that intelligent machines cannot be right 100 per cent of the time?"

What is artificial intelligence?

There is no single, universally accepted definition of artificial intelligence ("AI"), and the term is commonly used in at least three different senses. AI can refer to:

(a) software systems (also referred to as "machines" or "agents") that simulate certain capabilities associated with human perception and cognition;

(b) the capabilities possessed by those systems;

(c) the branch of computer science concerned with developing and studying software systems that display those capabilities.

This project concerns AI mainly in the senses described in (a) and (b) above.

AI systems simulate functions associated with human intelligence like speech recognition, natural language processing, machine vision and image recognition. Some of them can learn from their past experience to improve their own performance to deal with tasks and situations they have not previously encountered, a process known as machine learning. They power digital voice assistants like Alexa and Siri, control the movement of self-driving vehicles, translate from one language into another, diagnose disease and recommend its treatment, and serve a vast range of other purposes.

Why is the project needed?

The use of artificial intelligence (AI) to substitute for or assist human action and decision-making is expanding in virtually every field of activity and endeavour. Digital technology based on AI is employed in medical diagnostics, critical infrastructure control, autonomous vehicles, new drug development, climate research, financial services, marketing, and of course robotics, to name only a few of its applications. More controversially, AI is being used in law enforcement and the administration of justice. It is also being used also to screen applicants for employment, immigration, and eligibility for public and private benefits. In these and many other "use cases," automated decision-making has potential for serious effects on the lives and well-being of individuals.

AI systems are designed to operate with varying degrees of autonomy from human programmers and controllers, but they are still machines. They make decisions in ways that differ from human decision-making. When they make mistakes, they do so in ways that differ from human error as well. The harm that can result to people and things from errors and failures of AI systems, however, can be equally damaging as the acts and omissions of human tortfeasors.

AI capabilities exceed human ones in many areas, especially in their ability to detect patterns in enormous volumes of data that humans cannot process on a feasible timescale. They rely, however, on probabilistic inferences that will not be correct 100% of the time. In addition, they are not necessarily dependent on rules pre-encoded by programmers for every step in the decision path leading to an output in the form of a decision, prediction, recommendation, or the activation of a machine. They may apply rules they have learned autonomously by inference from data to which they have previously been exposed.

AI systems have been described as "unpredictable by design." This feature allows them to assess new situations and react to them without constant human oversight or control, but also holds potential for significant harm. There may be no identifiable defect in the basic operational logic of an AI-based computer application or AI-guided machine that makes decisions or provides other output with damaging or even disastrous consequences. Errors by AI systems can have limited explainability. This raises difficult legal and policy issues concerning where liability should rest. It is not possible to test AI against every scenario in which it will be deployed, and many different players are typically involved at various stages in the development and deployment of a device or program that uses AI, not to mention the end user.

This BCLI project addresses the difficult issues that arise in applying tort principles of fault, causation, standards of care, foreseeability, and remoteness of damage when AI systems cause harm. It is being carried out with the aid of an interdisciplinary volunteer committee comprising expertise in computer science, engineering, medicine, and law. The project will generate recommendations for reform where existing tort law is found to fall short in providing practical and just solutions.

Project Committee

The members of the Artificial Intelligence and Civil Liability Project are:

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

Dylan Merrick
Software developer
Former Member,
BCLI Board of Directors

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

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

Cynthia Khoo
Barrister and Solicitor
Associate, Georgetown Center on Privacy
and Technology

Darin Thompson – Liaison observer
Barrister and Solicitor
Ministry of Attorney General

Maya Medeiros
Barrister and Solicitor
Norton Rose Fulbright LLP

Dr. Teresa S.M. Tsang
Director, Echo Lab
Faculty of Medicine,
University of British Columbia

Project Goals

The goals of the project are to:

(a) identify where the general rules of tort law require adaptation to provide adequate and just civil remedies for harm resulting from the operation of autonomously functioning AI systems to persons, property, and other interests protected by the common law of tort;

(b) develop law reform recommendations to address the needs for adaptation that are identified; and

(c) complete and publish a law reform report containing the recommendations.

Project Timeline

The project formally commenced in November 2021 and is anticipated to conclude in 2023. It will involve three phases. The first involves the development of tentative recommendations for

law reform through extensive deliberation and analysis by the Project Committee and informal consultation with a wider, multijurisdictional cadre of additional experts. The BCLI project staff will provide research, drafting, and logistical support services to the Project Committee.   The first phase will culminate with the completion of a detailed consultation paper explaining the tentative recommendations and the basis for them.

In the second phase, BCLI will publish the consultation paper and seek responses to its contents from other experts, stakeholders in industry and research, and the general public. This phase of formal consultation will last at least three to four months.

In the third phase, the Project Committee will consider the responses to the consultation paper, formulate final recommendations with the benefit of the responses, and prepare a final report. After the BCLI Board of Directors gives approval for publication, the project will conclude with issuance of the report.

Project Funding

Financial support for this project is generously provided by the Ministry of Attorney General of British Columbia.

About BCLI

BCLI was incorporated in 1997 as an independent, not-for-profit society dedicated to modernization and improvement of the law. Information about BCLI and its process can be found on our website at https://www.bcli.org.

Contact

Greg Blue, Q.C.
Senior Staff Lawyer
British Columbia Law Institute
1822 East Mall, University of British Columbia
Tel.: (604) 827-5337
E-mail: gblue@bcli.org