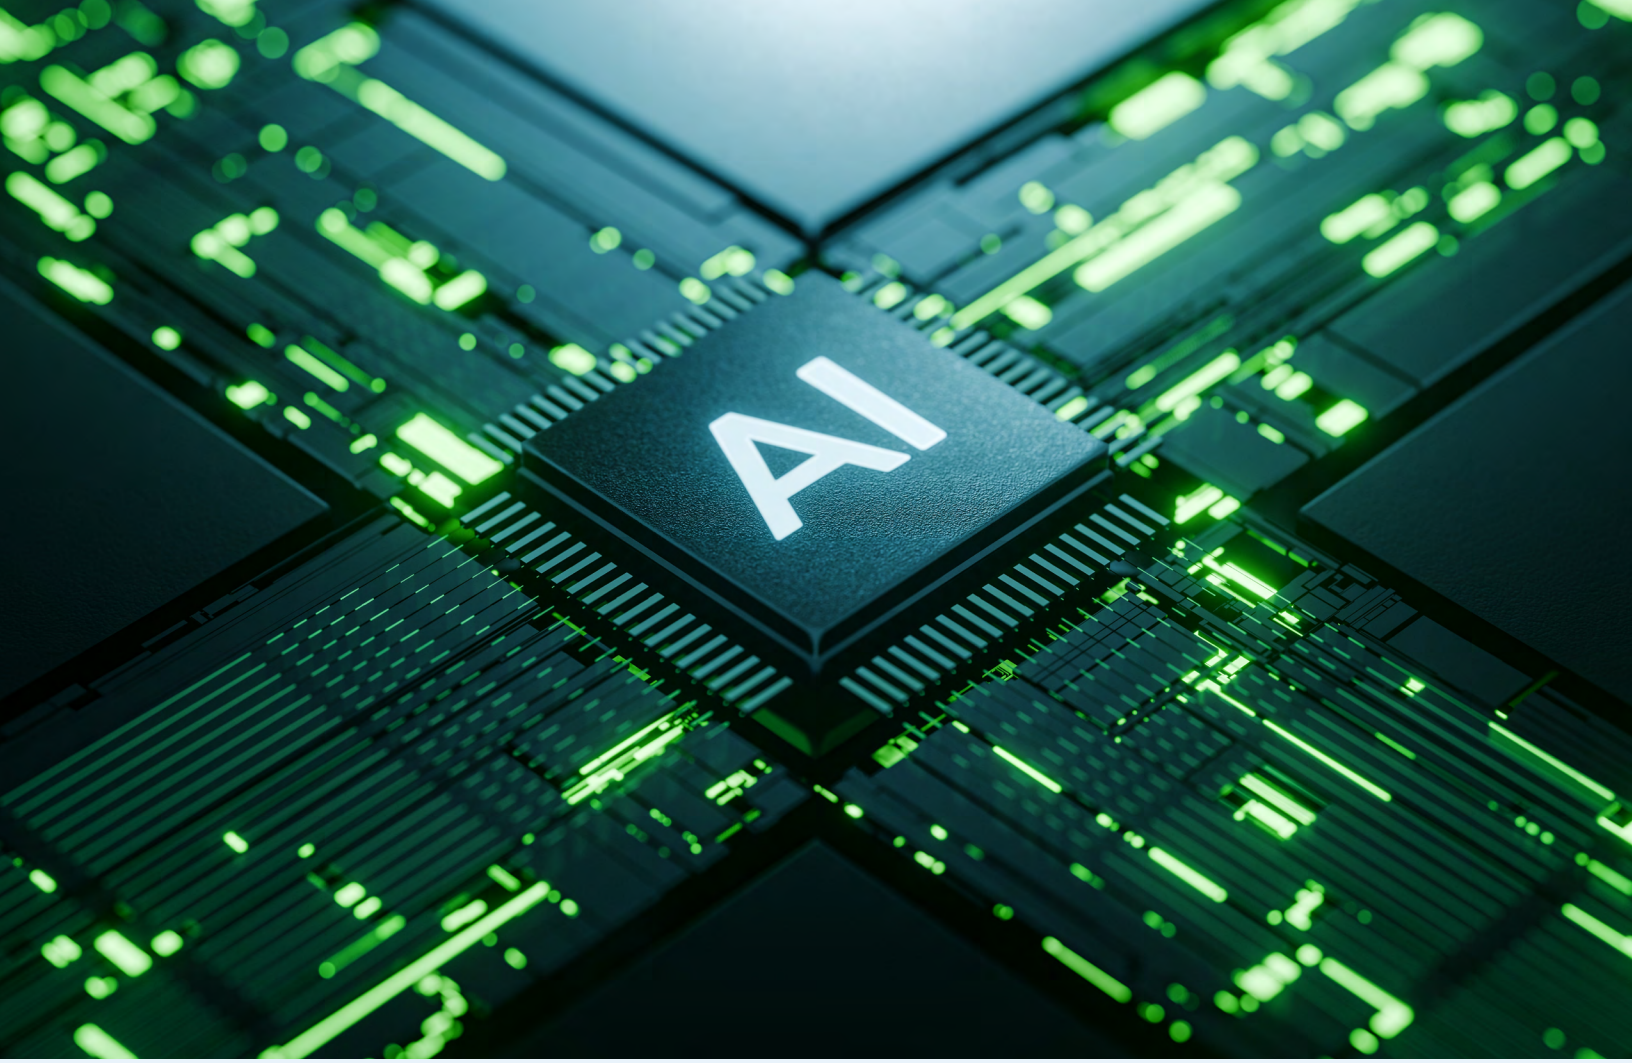

Artificial Intelligence (AI) 101: Risks and Opportunities

by Artificial Intelligence in Engineering Working Group

April 2026

ospe.on.ca

ONTARIO SOCIETY OF PROFESSIONAL ENGINEERS

Acknowledgements

This report was developed through the collaborative efforts of dedicated members of Ontario's engineering community. OSPE gratefully acknowledges the contributions of the following individuals:

Authored by:

Paul Longo, P.Eng, M.Eng, PMP, MMAI

Colleen Shannon, P.Eng., LLB, LLM

Dr. Jacqueline Csonka-Peeren, P.Eng., MASc, MBA

Emanuel Corthay, P.Eng.

Raimund Laqua, P.Eng.

Contributors

Aaron Pereira, MASc

Arjan Arenja, P.Eng., MBA, ICD.D

Claude Nnadi, EIT

Maria Becerra

Ryan Maclaughlan, P.Eng.

Table of Contents

Overview

In AI 101: Part 1, the **OSPE Artificial Intelligence (AI) in Engineering Working Group** established a baseline for understanding AI's core concepts and the urgent need for engineering professionals to lead in their adoption of AI. This report, AI 101: Part 2, takes the next logical step. Having defined what AI is, we now explore what AI does—and more importantly, how it should be managed.

This document is designed to transition the reader from conceptual knowledge to operational readiness. It provides the framework necessary to assess AI tools, identify high-risk applications, and implement management strategies that align with professional obligations.

To support this transition, the following sections delve into the "how" of AI in engineering:

Opportunities for AI: Identifying emerging areas where AI can accelerate design and create value for engineers.

Risks and Ethical Considerations: A deeper dive into the unique challenges of 2026, including algorithmic bias, data sovereignty, and the "human-in-the-loop" requirement for professional sign-off.

Voluntary Standards: An overview of the evolving global and Canadian standards that provide roadmaps for trustworthy AI management.

Risk Management Scenarios: A practical example to contextualize how applied AI might look in the practice of engineering.

Our goal is to ensure that Ontario's engineers are not just users of AI, but leaders in responsible implementation.

Opportunities for AI

Applied AI is not just an abstract pursuit; it is a tool for solving problems. The value of any AI system is ultimately measured by its ability to solve a specific, real-world problem. There have been tremendous developments, for example, in life sciences, legal, and insurance practices through the use of AI. The use of AI in the field of **Architecture, Engineering and Construction (AEC)** is also emerging. Take note of the examples below.

Use Case	Reference
Generative Design & Urban Planning	Autodesk Forma: Generative Urban Design in Alkmaar, Netherlands AI-Driven Preliminary Design for High-Rise Hotels (Cirebon City Case)
Structure Health & Safety Monitoring	AI-Powered Condition Monitoring for Concrete Bridges Real-time Concrete Strength Prediction via PZT Sensor Networks
Sustainable Building & Energy Optimization	The Edge, Amsterdam: AI-Optimized Natural Lighting and Ventilation Deep Reinforcement Learning for HVAC and Photovoltaic Prediction
Construction Project Management	AI-Powered Site Progress Tracking Construction Intelligence Platform for Developers and Real Estate Owners
Design & Computer-Aided Design (CAD)	Text to CAD Platform Engineering Design Co-pilot AI Assisted Design Review Software

This is neither an exhaustive list nor an endorsement of any of the above products. Instead, the list is meant to demonstrate the breadth and complexity of applied AI solutions that may be offered to an engineer. These applications offer tremendous advantages for engineers by increasing their efficiency while enabling the discovery and design of novel solutions. Practicing engineers will want to consider the opportunities to apply these safely in their practice. At the same time, pressures from external competitors who may adopt tools quickly may create a sense of urgency. However, the most pressing question should be, “How can professionals effectively navigate the risks, ethics and liabilities associated with these tools?”

Risks and Ethical Considerations

A purely technical understanding of AI is an incomplete understanding of AI. As artificial intelligence systems become more powerful and deeply embedded in engineering, ethical considerations are a core component

of the field itself. This report complements *AI 101 Part 1* by **providing professional engineers with information on what AI should do, what AI can do, and what the difference is**. Together, these perspectives help professional engineers apply AI effectively and responsibly in their practice.

Like the advent of the internet, the revolution of the smartphone, and the proliferation of social media, AI has advanced to the point that it is not just a tool, but also a socio-technical system with profound societal impact. During this transformation, we have a significant opportunity and obligation to learn from the recent past and avoid societal harm.

The ethical dimension of AI is no longer a niche concern; it is a core component of the field itself. This area, often called "**Responsible AI**" or "**AI Ethics**," directly addresses critical questions of:

Accuracy and Reliability: How do we verify that an AI system's outputs are reliable, appropriate, and valid for its intended operating conditions?

Fairness and Bias: How do we prevent AI systems from perpetuating or amplifying harmful human biases present in their training data? How do we ensure alignment with our own and our organization's mission/value statements?

Note that this will not only apply to the use of AI in the practice of engineering, but in the business operations of engineering firms, such as human resources and marketing.

Accountability: If an AI system makes a harmful mistake (e.g., in a medical diagnosis or a self-driving car), who is responsible?

Transparency: Can we understand and explain how a complex AI model arrived at its decision? This is often called "explainability" or XAI.

Privacy and Security: How do we protect the vast amounts of data required to train and run AI systems? How do we identify and protect private or sensitive data that might be generated by the AI system in deployment?

Confidentiality: How is sensitive, proprietary, or personal information protected when used in the AI system's lifecycle, including at AI system retirement?

Human-Centric Design: Is the AI being used to augment and empower humans, or to replace and devalue them?

Human Oversight: When is human review and intervention necessary?

Societal Impact: What broader effects does the AI system have, including on the environment?

Safety Considerations/Guardrails: What safeguards need to be in place to prevent unsafe or unintended behaviour?

Legal Considerations: How do laws, regulations and professional standards impact the professionals designing or deploying AI systems?

Accuracy and Reliability

AI systems are probabilistic, not predictive, or guaranteed to be accurate. Engineers must understand model limitations, validate answers, and communicate uncertainty. AI system accuracy depends on factors like data quality and the intended use of the model.

Specific to **Large Language Models (LLMs)**, AI hallucinations refer to a phenomenon where an AI system produces factually incorrect, misleading, or entirely made-up outputs. In other words, the AI "hallucinates" information. Hallucinations often occur because the model is probabilistic, not reasoning-based. Models predict what word or phrase is most likely next given its training. As a result, any facts or research originating from an AI LLM need to be sourced and verified independently. Hallucinations are difficult to spot without subject matter expertise because they are written in a very convincing way.

In the use of AI-powered note-taking, which is increasing in popularity, engineers must continue to proofread minutes and summaries as nuanced technical terms may be more prone to errors and hallucination.

Engineering Example: Steel Bridge Members

Consider an AI system used to predict the remaining fatigue life of steel bridge members based on strain gauge data and inspection records. The model does not guarantee the exact number of cycles to failure; instead, it produces a probabilistic estimate based on historical data and observed patterns.

The training data used is likely limited to certain bridge types, load ranges, or environmental conditions. Therefore, the model's accuracy may degrade when applied to a different structure or climate. To understand the accuracy of the model and its predictions, a responsible engineer would need to understand the data used to train the model to assess when they can and cannot rely on the outputs. In addition, the AI model would need to be validated by comparing the AI's outputs against known inspection results, by checking that the appropriate safety factors are utilized, and by expressing the prediction as a range with associated confidence intervals rather than a single definitive value. Only then could the engineer utilize their professional judgment to make maintenance or load-restriction decisions.

Fairness and Bias

Biases exist when there are historical normalizations of decisions or behaviour. These normalizations may be present in training data used for AI models and require consideration. Not all biases are harmful; some are welcomed and make sense to "normalize" in certain contexts (e.g., risk-taking in entrepreneurship; compassion in service delivery). Some, such as discriminatory hiring practices based on gender, race, or sexuality, are quite harmful to individuals and society. AI systems learn from historical data containing biases. Without care, AI can reproduce or amplify these harmful biases at scale. Responsible application of AI requires evaluating training data, testing outcomes across different groups, and actively mitigating unfair or discriminatory behaviour. Below are some examples where AI has been shown to perpetuate systemic societal biases:

Use Case Name	Summary	Reference
Criminal Risk Assessment (COMPAS)	This AI tool was used across the United States to predict the likelihood of a defendant committing future crimes. A **ProPublica** investigation found the algorithm was biased against Black defendants, who were flagged as high-risk at twice the rate of white defendants despite similar backgrounds.	ProPublica Article
Amazon AI Recruiting Tool	**Amazon** developed an experimental hiring tool to automate the search for top talent by reviewing resumes. The system taught itself to penalize resumes that included the word "women's" as it was trained on historical data from a decade when the majority of tech applicants were men.	CNBC Article

Gender and Sexuality Recognition	A study published by **Stanford University** in 2017 showed that AI could be used to detect gender and sexual orientation based on facial images with high accuracy. In response, civil rights groups urged the European Union to ban AI tools designed to detect gender, sexuality, or disability. These systems were criticized for being based on pseudoscientific assumptions that can lead to mass surveillance and the exclusion or persecution of marginalized groups. In response, the EU banned such systems under their **AI Act** in 2025. In 2026, the **Ontario Human Rights Commission (OHRC)** and the privacy commissioner released joint principles explicitly stating "AI must be developed, acquired, adopted and governed to prevent harm or unintended harmful outcomes that infringe upon human rights, including the right to privacy and non-discrimination."	Summary of Stanford Study EU AI Act Ontario Human Rights Commission

While these examples may not directly impact the practice of engineering, leaders in engineering firms must be aware of this in the context of creating safe, inclusive and diverse working environments for engineers. These biases can show up in any part of a business, from recruitment to marketing to performance reviews. Thus, using AI in these areas of business requires careful consideration. Individuals must be aware of the **Information and Privacy Commissioner (IPC)** and the **Ontario Human Rights Commission (OHRC)** *Principles for the Responsible Use of Artificial Intelligence (AI)*. Both leaders and individuals should be aware of these principles and be encouraged to speak up where they have concerns.

Note: While references to legislation have been provided here, no legal interpretation is being given. Professionals must consult their own legal counsel to correctly interpret the application of human rights requirements in any context of their business.

Engineering Example: Water Distribution

As a hypothetical example related to the practice of engineering, consider an AI system used by a municipality to prioritize maintenance and replacement of water distribution pipelines based on historical failure data, repair records, and customer complaints.

If the training data reflects decades of underreporting in lower-income or underserved neighbourhoods—where leaks were less likely to be reported or addressed—the AI may systematically assign lower priority to those areas.

As a result, the system unintentionally reinforces existing inequities by directing investment toward neighbourhoods with more complete data. A responsible engineering application would include auditing the training data for coverage gaps, testing model outputs across different communities, and incorporating additional data sources (such as age of infrastructure or material type) to reduce biased outcomes.

Another classic example is survivorship bias, which illustrates the importance of understanding what a given data set may be missing and how that can create blind spots for ML models trained on such data.

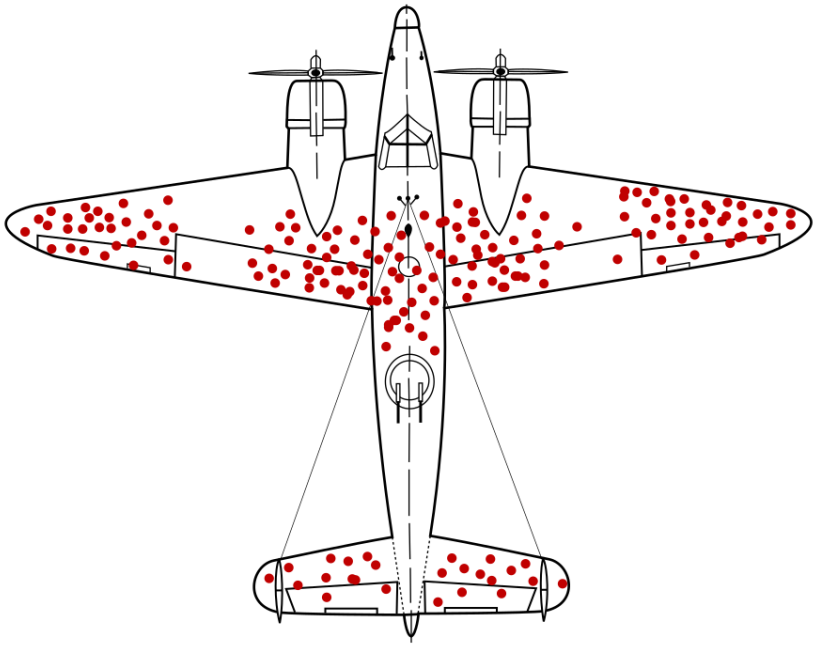

Figure 1

The hypothetical pattern of damage of surviving aircraft shows locations where they can sustain damage and still return home. If the aircraft was reinforced in the most commonly hit areas, this would be a result of survivorship bias because crucial data from fatally damaged plans was being ignored; those hit in other places did not survive. In other terms, “We need to reinforce the other parts, because they made the other planes unable to return.”

Accountability

In engineering, accountability is a critical ethical issue when deploying AI systems. Engineers must review AI system outputs since mistakes can have significant safety, business, financial, and societal consequences. AI systems operate probabilistically and may make errors that are not immediately obvious without human review, which raises the question of who is responsible when failures occur. Unchecked mistakes can lead to harm and legal liability for those who deploy and/or develop AI.

Engineering Example: Municipal Building Permits

In a hypothetical example, a municipality uses AI to process and triage building permit applications. Errors by the system can result in one of two outcomes. Either a permit that should be approved is denied, causing potential construction delays and additional costs to the developer or building owner, or a key building safety concern is missed, creating risk for the eventual occupants. The accountability for errors made by the AI system can either rest with the vendor who supplied the system, the municipality that purchased and implemented the system, or the individual plan examiner using the system. Careful thought needs to be given to this during the adoption of the system. Often, this means that decisions made by AI systems must be interpreted as recommendations for humans to review and approve with care. This enables a clear line of accountability, but the tradeoff is that the level of automation that can be achieved is limited.

Engineers in Ontario are encouraged to review existing professional practice guidelines.

Transparency

In engineering, transparency is a critical ethical issue when deploying AI systems because engineers must be able to interpret an AI system’s process and limitations. A lack of transparency can prevent effective validation, hinder error detection, and undermine trust in AI-assisted decisions. Without sufficient transparency, engineers may be unable to assess risks, explain outcomes to regulators, clients or the public, or take corrective action when systems behave unexpectedly. In addition, there should be transparency regarding where and when AI is used in a process so that humans can appropriately question the outcomes.

Engineering Example: Standard Design Specification Document

An Engineer in Training is tasked to take the standard design specification document and update it for a specific project's needs. They upload the document to a generative AI platform and ask the chatbot to make the updates. The amended document then gets passed to a senior engineer to review. The engineer looks at the key sections that are project-specific and checks them against the basis of design, and approves the specification revisions. It is later found that a standard part of the specification that should never change was edited by the generative AI system, resulting in a change order for the project and an upset client. Generative AI is known to hallucinate and/or remove key parts of paragraphs, often seemingly at random. Missing a single sentence in a design specification could mean thousands of dollars in change orders on a project. Therefore, transparency requires knowing when AI has been used in completing an engineering document and is critical to ensuring that the document undergoes the appropriate quality assurances.

Privacy and Security

AI systems collect personal and operational information. Inadequate privacy protections or security controls can lead to breaches, regulatory violations, and physical or financial harm. Engineers must therefore design and deploy AI systems with privacy-by-design and security-by-design principles, ensuring that data is protected throughout its lifecycle.

Engineering Example: Data Security

Engineers currently deal with digital data on a day-to-day basis, whether that is in the form of emails, PDF documents, automation systems, or cloud-based file shares. When it comes to security, these systems likely currently have the required protocols to keep data safe and secure and protect both corporate and client privacy. When introducing an AI tool into their system for day-to-day use, engineers need to consider any new risks that this might introduce. For example, if a company decides to use a generative AI system such as **Google Gemini**, **Microsoft Co-Pilot**, or **Cohere North**, they need to consider where the text input and any attachments uploaded go and how they are used. All these systems offer different options for data storage and use. Engineers must read the different options being presented to them carefully and review them with IT professionals who understand their current security and compliance standards.

In addition, when it comes to systems that an engineer may design, there are further considerations. For example, an engineer designs and specifies an AI-enabled industrial monitoring system used in a manufacturing plant. Here, machine learning models analyze sensor data to detect equipment faults and optimize operations. These systems may process sensitive operational data, proprietary production parameters, and information about worker activities. The engineer must consider the security of the system from tampering as well as protect the privacy of industrial data. In this scenario, privacy and security risks extend beyond consumer applications and must be addressed rigorously in all engineering contexts where AI systems interact with sensitive data and critical infrastructure.

Depending on regulatory compliance of the specific industry, the practitioner may need to consider data residency (geographic location of data), and data sovereignty (the legal jurisdiction data is subject to), each of which are complex and emerging considerations.

Confidentiality

Confidentiality is a critical consideration when deploying AI systems in engineering because these systems often process sensitive data, including proprietary designs, operational parameters, and client or employee information. Unauthorized access or data leaks can result in intellectual property loss, competitive disadvantage, or legal liability. Engineers producing and deploying AI must follow confidentiality agreements and regulatory requirements.

Engineering Example: Design Optimization

In an AI-driven design optimization system used by an aerospace company, the AI analyzes proprietary aircraft schematics to suggest improvements. Engineers who produce the system must implement robust data protection measures, encryption, and access management, while engineers who deploy it must enforce policies that prevent unauthorized use or sharing of sensitive information. This demonstrates that maintaining confidentiality is essential to protect both organizational assets and compliance obligations in engineering applications.

Human-Centric Design

Like any software, an AI tool's usefulness depends on how end users adopt the tool. As with any digital transformation project, when adopting an AI tool, engineers must consider the user interface, training, and change management to ensure that the tools are used effectively. As well, when AI is implemented without careful consideration, it can inadvertently reduce the role of skilled professionals, undermine expertise, or make humans overly reliant on automated decisions. Such replacement not only raises ethical concerns but also introduces safety and operational risks and risks to the professional autonomy of the engineer. Recall the example of the AI system used by a municipality to review building applications. Humans tasked with doing the review may become complacent over time. How do they stay vigilant for errors, especially if the system performs very well the majority of the time?

Engineering Example: Industrial Plant

In an industrial plant, an AI system is introduced that predicts the equipment that is most at risk and proactively produces preventative maintenance tickets for plant operators to action.

Without proper change management and training on the tool, plant operators may ignore the outputs of the tool. This can occur if they do not have transparency on how the tool arrived at a decision, or have the ability to provide feedback when they believe the tool is mistaken. They may also feel that this tool is replacing their institutional knowledge and intuition that has been honed through years of experience. There needs to be clear communication regarding the purpose of the tool and how its results should be interpreted. While full automation of maintenance tickets may be the end goal, a human-in-the-loop system where a human reviews and approves suggested tickets may help with building long term adoption of the tool, while contributing to the ultimate goal – no unscheduled downtime.

Moreover, full feedback from the plant operators allows for debugging of the tool and/or improvement of the AI tool over time.

Another consideration with this new AI tool is how it helps or hinders the training and development of plant maintenance workers. Over time, junior staff may become overly reliant on the tool and may not develop the critical thinking necessary to question the AI tool's results – a skill currently learned through hands-on equipment inspection, data analysis, and trial and error decision-making processes. This can create operational risk when events happen outside of the training parameters of the AI system. Thus, engineers need to think about how to develop and maintain the skills of the next generation of engineers in the age of AI.

Human Oversight

Human oversight is essential when deploying AI systems in engineering, particularly in safety-critical or high-stakes environments or where fundamental rights could be violated. AI outputs are probabilistic and may fail under unexpected conditions, making human review necessary to prevent errors, ensure safety, and maintain accountability. Engineers must define clear oversight roles, intervention points, and escalation procedures so

that AI supports rather than replaces critical human judgment.

Engineering Example: Chemical Processing Plant

In a chemical processing plant, an AI system may control temperature, pressure, and chemical feed rates to optimize production. If the system encounters a rare combination of conditions outside its training data, it could make unsafe adjustments. Engineers who produce the system are responsible for designing mechanisms that allow human operators to monitor outputs, override automated actions, and receive alerts for anomalous behaviour. Engineers who deploy the system must train staff, define intervention protocols, and continuously monitor system performance. This example highlights that effective AI deployment in engineering relies on active human oversight to maintain safety, reliability, and ethical accountability.

Societal Impact

AI systems in engineering can have significant societal impacts. For example, environmental consequences due to significant energy consumption and resource requirements of the data centers powering generative AI. Large-scale machine learning models, real-time sensor processing, and continuous optimization tasks can all demand substantial computational power, contributing to greenhouse gas emissions, higher water usage, and higher energy usage.

Engineering Example: Growth of AI Models

Every few months, generative AI models are getting bigger and bigger as companies race to market their products across a broad spectrum of applications. However, engineers may find that older, smaller, and ultimately less energy-greedy models may serve their needs. Thus, engineers should evaluate the cost and energy use of these models relative to their needs and seek to right-size models for their application – bigger is not always better.

This push for bigger models will put pressure on engineers to design and build data centres at speeds and sizes not seen to date. Engineers need to push for efficient modular designs that can help avoid over-building while enabling systems to run efficiently across a wide range of operating scenarios. In addition, engineers should identify and advocate for energy-efficient opportunities for cogeneration and waste heat recovery. For example, Finland has been designing data centres that are located underground, where the waste heat provides district heating to residential houses.

Engineering Example: Perpetuating Bias

As demonstrated under the Fairness and Bias section, AI tools can be used to perpetuate biases towards individuals and groups based on race, gender, or sexual orientation. Engineers should be encouraged to think deeply about the unintended consequences of adopting AI systems in their practice. A recent paper published by **Google DeepMind**, "Intelligent AI Delegation", highlights some significant risks to skill development through the application of AI. A crucial short-term risk to the engineering practice is what AI might do to the talent pipeline.

Expertise is built through the repetitive execution of narrowly scoped tasks, precisely the tasks most likely to be offloaded to AI. For **Engineering Interns (EITs)** and junior staff, these routine tasks are the primary vehicle for learning. If these tasks are fully automated, junior members will be deprived of the foundational experience needed to develop deep strategic judgment. Compounding this, the over-reliance on AI for task delegation can lead to reduced engagement, where engineers are using less deep problem-solving and critical thinking skills to complete their work. If unchecked, over time, such skill degradation will not be just about the ability to do the work, but crucially, the ability to judge the work.

This could mean that future senior engineers may lack the “oversight readiness” required to manage high-stakes projects. As AI takes over routine, low-complexity workflows, human engineers could be increasingly removed from the day-to-day “hands-on” work. Here, professional engineers may only be called to intervene during complex edge cases or critical system failures. However, without the situational awareness gained from performing routine tasks, the engineer in question may be ill-equipped to resolve a crisis when it occurs. In Ontario, an engineer must “seal” a document, signifying they have personally supervised or performed the work. They retain 100% of the legal accountability but lose the proficiency required to discharge it. Therefore, in adopting AI, engineers must consider their societal obligations to train and mentor future engineering professionals. Careful thought must be given to how to ensure that emerging professionals are given the right training.

Safety Considerations/Guardrails

In engineering, AI systems must be deployed with robust safety considerations and guardrails to prevent accidents, operational failures, or unintended consequences. AI can behave unpredictably under unusual conditions, and errors can have significant physical, financial, reputational or societal impacts. Engineers are responsible for anticipating potential failure modes, designing fail-safes, and implementing monitoring mechanisms to ensure safe operation.

Engineering Example: Autonomous Manufacturing

In an autonomous manufacturing system, robots powered by AI handle heavy machinery and assembly tasks. Emergency stop mechanisms, motion limits, and anomaly detection can prevent unexpected AI decisions causing equipment damage or endangering workers. Engineers who produce the system must understand the scenarios that the AI system is trained on and understand the limits of such a system. Careful consideration must be given to testing, specifically around unusual edge cases that may be rare but possible for a system to encounter. Engineers should incorporate these safeguards into hardware, software, and training materials. As well, engineers who deploy the system must maintain monitoring, regularly test safety features, and enforce operational protocols.

Legal Considerations

The presentation of the risks in the above sections should raise questions about legal liability for practicing professionals. It is often best to consider AI products as additional tools for humans to use in their existing workflows, thus keeping accountability, as it currently stands as a first step in introducing the AI system.

AI deployment in engineering raises many critical legal issues. Liability for AI-driven errors is often shared among those who produce the system (design, training, and validation) and those who deploy it (implementation, supervision, and operational use). Failure to comply with laws, regulatory standards, safety codes, or contractual obligations can result in civil litigation, criminal charges, disciplinary proceedings, and substantial costs arising from disputes, fines, and/or reputational damage. Engineers must ensure that AI systems are designed, documented, and deployed in ways that meet applicable laws, standards, and contractual requirements.

Engineering Example: Assembly Robot

In an AI-controlled industrial robot used for assembly, a programming error or misinterpretation of sensor data could cause product defects or workplace accidents. If the system malfunctions, liability could extend to the engineers who designed the AI, the organization that deployed it, and potentially the vendors supplying the AI platform. Engineers who produce the system are responsible for rigorous testing, validation, and documentation, while those who deploy it must follow safety protocols, maintain oversight, and comply with industry regulations. This example highlights that legal accountability in engineering AI systems is shared, and

proactive measures are essential to mitigate risk.

Laws such as those for consumer protection, occupational health and safety, employment standards and human rights also create AI obligations for a professional engineer and the organizations for which they work.

While references to legal legislation have been provided here, no legal interpretation is being given. Professionals must consult their own legal counsel to correctly apply legal requirements in the context of their business.

Regulatory Environment

The evolution of AI will make the regulatory space in Canada a key area for professionals to watch.

In Canada, the regulatory landscape for AI in 2026 is defined by a shift away from the previously proposed "all-in-one" legislation (Bill C-27) toward more targeted provincial laws and federal directives. Because the flagship federal AI bill died in early 2025, the current focus is on reintroducing independent regulations and enforcing existing privacy standards.

Some key regulations to be aware of are:

The Separation of AIDA and Privacy Reform: The **Artificial Intelligence and Data Act (AIDA)**, originally part of the massive **Bill C-27**, died on the Order Paper following the prorogation of Parliament in January 2025. In 2026, the federal government is expected to reintroduce AI regulation as a standalone bill rather than bundling it with privacy laws. A new AI-specific bill may be forthcoming in 2026.

Directive on Automated Decision-Making (2026 Deadline): While private-sector AI laws are in flux, the federal government is regulated by the **Directive on Automated Decision-Making (DADM)**. An updated version of this directive was released in June 2025, requiring all federal institutions to meet higher standards of transparency and human intervention. Federal departments have until June 24, 2026, to ensure all AI systems (including those procured before 2025) comply with mandatory algorithmic impact assessments and human-in-the-loop requirements.

Ontario's Bill 194 (Public Sector AI & Hiring): Ontario recently passed the **Strengthening Cyber Security and Building Trust in the Public Sector Act (Bill 194)**, which introduces specific guardrails for AI use within provincial public entities and schools. Coming into force in 2026, the Act requires public sector entities to disclose AI use and manage risks. Crucially, it includes provisions that may require employers to disclose the use of AI in hiring processes, setting a precedent that other provinces may follow in 2026.

OPC's "Privacy by Design" for Generative AI: In early 2026, the **Office of the Privacy Commissioner (OPC)** reported on high-profile investigations into platforms like **OpenAI** and **X's Grok**. In the absence of a federal AI act, the OPC is using existing privacy law, the **Personal Information Protection and Electronic Documents Act (PIPEDA)**, to mandate that AI developers build "privacy by design" into their models. The OPC is pushing for AI-specific amendments to PIPEDA, including the recognition of privacy as a fundamental right and mandatory "Privacy Impact Assessments" for any high-impact data processing.

Voluntary Standards

One reason standards are important is that they are used in court to help interpret laws and regulations for AI. Aligning with standards can bring an engineer and their organization into closer alignment with legal obligations and regulatory compliance. The risk categories identified above fundamentally align with well-established bodies such as the **Organization for Economic Co-operation and Development (OECD)** principles for responsible AI usage, or the recently published responsible AI guidelines from the Ontario

Information and Privacy Commissioner (IPC) and the Ontario Human Rights Commission (OHRC). Effective consideration of these risks will likely enable engineers to be ready for any upcoming regulations. As the regulatory landscape continues to evolve, engineers may want to consider well-established voluntary standards. Below is a sample of some of the key standards.

Many of these standards overlap in that they consider the same fundamental principles but provide certifiable and auditable methods to demonstrate their application in practice.

Note: This is not an exhaustive list, and engineers are responsible for researching the most relevant and current standards applicable to their work. Also note that some of the international standards cited in the table below have a Canadian version. The Canadian versions more closely align with Canadian values and priorities, including trade-related and regulatory priorities, and are most relevant to Canadian engineers and the organizations that employ them.

Standard Name	**Nature**	**Primary Focus**	**Strategic Impact & Business Value**	**Official Link**
CSA ISO/IEC 42001:25	Certifiable	AI Management System (AIMS)	Having this certification could slash the time spent on vendor security and legal reviews by providing an audit-ready certificate. It could indicate that AI is not just a "shadow project" but is aligned with the company's strategic vision, which might make the company more attractive for acquisition, investment, or global expansion.The Canadian signals to local regulators and federal partners that your AI governance meets the National Standards of Canada. It may be a key differentiator for government contracts.	View at CSA Group
CSA ISO/IEC 5338:24	Technical	AI System Life Cycle	**Engineering Efficiency & Predictability** standardizes the AI lifecycle from inception to retirement. It reduces "re-work" and technical debt by providing a consistent blueprint for DevOps and AI engineers, ensuring product stability as models evolve.	View at CSA Group
CSA ISO/IEC 8183:24	Technical	Data Life Cycle Framework	**Data Governance & Integrity** provides a framework specifically for data handling in AI. It mitigates "garbage-in, garbage-out" risks and ensures data provenance is audit-ready, which is essential for meeting evolving Canadian privacy expectations.	View at CSA Group

CSA ISO/IEC 22989:23	Foundational	Concepts and Terminology	**Contractual Clarity & Legal Safety** establishes a common vocabulary for AI. Using these standardized terms in **Service-Level Agreements (SLAs)** and contracts prevents costly legal disputes that often stem from ambiguous technical definitions.	View at CSA Group
CSA ISO/IEC 23894:24	Guidance	Risk Management	**Regional Compliance & Liability Mitigation** bridges international principles with Canadian risk appetites. Implementing this can help demonstrate due diligence to Canadian courts and potentially lower local professional liability insurance premiums.	View at CSA Group
NIST AI RMF 1.0	Framework	Risk Identification	**Operational Agility & Culture** is a free, flexible "how-to" guide that can provide a common language for technical and legal teams. It can help reduce "compliance surprises" by identifying some known bias and security gaps early in development—a point when they are typically much cheaper to address as compared to fixing them after a product launch.	View at NIST.gov
IEEE 7000-2021	Ethical	Human-Centric Design	**Brand Equity & Talent Retention**. Research shows this process can reduce ethics-related escalations by up to 30%. It can appeal to design-conscious consumers and attract top-tier AI talent who seek out companies that demonstrate a commitment to responsible AI.	View at IEEE.org
ISO/IEC 24028:2020	Technical	Trust & Robustness	**Product Integrity & Churn Reduction** tackles the "Black Box" problem. By applying its principles of explainability and reliability, this standard can help with AI hallucinations or erratic behaviour that could alienate customers or cause costly system downtime. In that way, it contributes to protecting core technical assets and intellectual property of the business.	View at ISO.org

Risk Management Scenario

The above risks and ethical considerations may seem daunting for an engineer to navigate. However, professional engineers have dealt with and mitigated risk in professional practice for over a century. Hence, these considerations are recommended to be worked into existing risk management frameworks. In this section, we will introduce a risk management framework and present a fictional example to demonstrate how this can work in practice.

Case Study: Generative AI Assistants and Risk Management

This case study examines a medium-sized, multidisciplinary engineering firm of approximately 100 employees as it seeks to integrate an enterprise-grade generative AI assistant—such as **Microsoft Copilot**, **Google Gemini**, or **Cohere North**—into its core operations. The firm aims to transition toward a generative workflow that enhances efficiency across two primary domains: internal contract administration and business development.

For contract administration, by automating the routing of technical queries and the filing of site inspection photos, the firm intends to free its senior engineers from routine administrative burdens. Central to this transformation is the creation of a dynamic knowledge repository that utilizes **Retrieval-Augmented Generation (RAG)** to provide staff with instant, searchable access to project histories, technical specifications, and the latest regulatory codes that may apply to a query.

In the realm of business development, the firm wants to use AI to streamline its approach to **Requests for Proposals (RFPs)**. The ideal state is a system that will continuously monitor public RFP postings and perform an automated review to identify relevant opportunities for the firm. Beyond simple monitoring, they hope AI can assist in the granular review of contractual clauses and the generation of high-quality response documents by surfacing relevant past project examples.

To ensure the responsible and competent adoption of these tools, the firm has prioritized a structured implementation roadmap focused on human oversight. By balancing automated efficiency with rigorous human-in-the-loop protocols, the firm aims to establish a benchmark for responsible AI within the engineering sector, navigating the risks of automation while maximizing the strategic value of its internal data.

Risks and Mitigations

Let's take a look at some of the risks and potential mitigations. This is not considered exhaustive and will depend on the specific requirements of the individual firm. An engineer is required to take a deeper look at each risk and how to effectively manage the mitigations for their specific engineering project.

Risk Category	Specific Scenario Risks	Recommended Mitigations (Human-in-the-Loop)
Accuracy & Reliability	**Hallucinations:** AI citing non-existent regulatory codes **Outdated Patterns:** Relying on project histories that don't match current climates or structures **Date-limited software:** The AI tool is current only to 2024.	**Independent Verification:** Engineers must manually verify all facts, research, or regulatory citations originating from the AI. Links to sources must be available in any generative answer. **Confidence Ranges:** Where possible, communicate AI estimates as a range with confidence intervals rather than definitive values. One important technique is adding guardrails in any prompts to tell the model to say when there is limited information available.

Fairness & Bias	**Historical Blindspots:** AI may prioritize RFPs or project methods based on biased past data (e.g., favouring certain neighbourhoods or vendors).	**Coverage Audits:** Regularly audit the project repository for missing data or "blind spots" (e.g., survivorship bias). Make sure unique project factors are well understood by team members. **Outcome Testing:** Test AI recommendations across diverse project types to ensure neutral opportunity identification.
Accountability & Legal Considerations	**Liability Gaps:** Uncertainty over who is responsible if automated query routing leads to an on-site error **Non-Compliance:** AI-generated RFP responses failing to meet professional standards or mandatory bid requirements	**Professional Authentication:** Interpret AI outputs only as recommendations for humans to review and approve; maintain a clear line of professional accountability. **Due Diligence Logs:** Document the human review process for all AI-assisted documents. This will serve as a useful repository of learnings for the entire organization.
Transparency	**Hidden Logic:** Engineers may not understand *why* the AI selected a specific technical specification or RFP. **Unseen Edits:** Generative AI removing key sentences from standard specs without notice	**Disclosure Protocols:** Explicitly label when and where AI was used in an engineering document, so it receives appropriate quality assurance. **Confidence Scores:** Provide the "confidence of prediction" so engineers can apply appropriate skepticism if possible.
Privacy, Security & Confidentiality	**Data Leakage:** Uploading proprietary schematics or project histories to external LLM prompts. Residency **Issues:** Prompts being stored in non-compliant geographic jurisdictions	**Data Sanitization:** Remove all **Personally Identifiable Information (PII)** and proprietary parameters before inputting data into AI tools. **IT Collaboration:** Review data residency and sovereignty options with IT professionals before deployment.
Human-Centric Design	**Automation Bias:** Senior engineers becoming "complacent" as the AI handles technical routing successfully over time **Skill Degradation:** Junior staff losing critical thinking skills by relying on the project repository	**Change Management:** Provide clear training on the tool's limitations and purpose. Maintain any existing process for documenting check-prints. **Manual Research Requirements:** Require junior staff to perform hands-on analysis and trial and error before using the AI tool, or perhaps in-person conversations to ensure they can explain the why behind technical decisions when asked.
Societal Impact	**Energy Consumption:** High carbon footprint from using massive enterprise-grade models for simple administrative tasks.	**Model Right-Sizing:** Evaluate if smaller, less energy-greedy models can serve the firm's specific needs. **Efficient Design:** Advocate for modular and efficient data center designs.

Safety & Guardrails	**Unexpected Failures:** AI routing an urgent site safety query incorrectly during a rare edge case scenario.	**Manual Overrides:** Consider a human-in-the-loop who has a sightline of incoming communication and where it goes, so that they can manually intervene if required. **Edge Case Testing:** Rigorously test the system against rare but high-consequence scenarios.

One effective mitigation strategy in the aforementioned use case is comprehensive prompt training for all personnel. While prompt writing might be perceived as trivial, transitioning from keyword searches (e.g., Google, Bing) to well-constructed prompts represents a significant change in communication methodology. Keyword searches prioritize conciseness, whereas effective prompts necessitate clear contextual instructions. Furthermore, prompt training courses will educate engineers on the intended use cases for vendors' large language models, distinguishing them from mere promotional claims. These courses can be completed in under half a day and offer substantial benefits.

When your team starts out, encourage them to follow the "**Role-Task-Format**" rule of thumb. Every prompt should ideally define:

- **Role:** "Act as a professional technical writer..."
- **Task:** "...summarize this internal meeting transcript into five bullet points..."
- **Format:** "...presented in a clean Markdown table."

Conclusion

This document provides a brief introduction to the "what" and "should" of AI in engineering, building on the technical overview from Part 1. The intent here was not to be exhaustive, but to provide sufficient guideposts and trailheads for practicing engineers to thoughtfully engage in further debate and discussion on the adoption of AI in their practice.

Future publications from OSPE will provide deeper exploration of the ethics of AI and engineering, and case studies. In this changing AI landscape, it remains up to the engineering professional to demonstrate reasonable action in identifying and implementing relevant, current regulations, standards, guidelines and laws applicable to their work.

Transparency on the Use of AI in this publication: In the interests of transparency, it is important to acknowledge the use of Generative AI in parts of this report. Specifically, it was used for:

- Research: Prompts on specific areas of research were used to generate lists of sources available on the web. These sources were used to inform parts of this report.
- Revision of Prose: In writing the document, generative AI was used to convert bullet points into paragraphs that were then proofread and revised by humans.
- Reformatting of information: Generative AI was used to create tables based on bullet points.
- Glossary of Key Terms: Generative AI was used to take the full contents of the report and provide a list of key terms that were presented and defined in the content, and added to the appendix for quick reference.

Further Reading

Applying AI

Berangi, M., Zhang, F., Phusakulkajorn, W., Núñez, A., & Anupam, K. (2025). *Structural key performance indicators for condition monitoring of concrete bridges using artificial intelligence: A review.* **Intelligent Transportation Infrastructure**, *4*, liaf022. https://doi.org/10.1093/iti/liaf022 Structural key performance indicators for condition monitoring of concrete bridges using artificial intelligence: a review | Intelligent Transportation Infrastructure | Oxford Academic

Betley, Jan et al, (2025, February 24). Emergent Misalignment: Narrow finetuning can produce broadly misaligned LLMs. ARXIV. https://arxiv.org/abs/2502.1742

Betley, Jan et al (2026, January 14). Training Large Language Models on Narrow Tasks can lead to Broad Misalignment. Nature. https://www.nature.com/articles/s41586-025-09937-5

Gharavi, H., Taban, F., Korivand, S., & Jalili, N. (2025). *AI-powered structural health monitoring using multi-type and multi-position PZT networks*. **Sensors**, *25*(16), 5148. https://doi.org/10.3390/s25165148 AI-Powered Structural Health Monitoring Using Multi-Type and Multi-Position PZT Networks - PMC

Mahendra, S. M., Surahman, U., Jurizat, A., & Sari, D. C. P. (2025). Application of generative design on architecture to optimize design decision in preliminary design stage. *JARINA – Journal of Artificial Intelligence in Architecture, 4*(2). 5.+10557-Application+of+Generative+Design+-+29082025+-+Vol.4^LJNo.2^LJ2025-47-60 (1).pdf

Manmatharasan, P., Bitsuamlak, G., & Grolinger, K. (2025). *AI-driven design optimization for sustainable buildings: A systematic review.* **Energy and Buildings**, *332*, 115440. https://doi.org/10.1016/j.enbuild.2025.115440 AI-driven design optimization for sustainable buildings: A systematic review

Nagy, D., Villaggi, L., & Benjamin, D. (2018). *Generative urban design: Integrating financial and energy goals for automated neighborhood layout* (Conference paper, **Symposium on Simulation for Architecture and Urban Design 2018**). Autodesk Research. https://damassets.autodesk.net/content/dam/autodesk/www/autodesk-reasearch/Publications/pdf/Generative-Urban-Design-2018.pdf SIGCHI Conference Paper Format

Yu, Y. (2025). *Research on the application of artificial intelligence in green building design optimization* (SPIE Proceedings Vol. 13513, 135132V). In **The International Conference on Optoelectronic Information and Optical Engineering (OIOE2024)**. SPIE. https://doi.org/10.1117/12.3045836 Research on the application of artificial intelligence in green building design optimization

Case Studies on AI and Systemic Bias

Angwin, J., Larson, J., Mattu, S., & Kirchner, L. (2016, May 23). *Machine bias: Risk assessments in criminal sentencing*. ProPublica. https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing

Artificial Intelligence Act. (February 2, 2025). *Article 5: Prohibited AI practices and implementation dates*. European Union Artificial Intelligence Act. European Union: Artificial Intelligence Act

Harrer, A. (2018, October 10). *Amazon scraps a secret A.I. recruiting tool that showed bias against women*. CNBC. https://www.cnbc.com/2018/10/10/amazon-scraps-a-secret-ai-recruiting-tool-that-showed-bias-against-women.html

Sakr, N. (n.d.). *AI can now identify people as gay or straight from their photo*. Mark S. Bonham Centre for Sexual Diversity Studies, University of Toronto. https://sds.utoronto.ca/news/ai-can-now-identify-people-as-gay-or-straight-from-their-photo/

Principles of AI Usage

Information and Privacy Commissioner of Ontario & Ontario Human Rights Commission. (2026). *IPC-OHRC principles and the OPS AI Directive: A comparison of principles for responsible artificial intelligence use* (Letter). https://www3.ohrc.on.ca/en/ipc-ohrc-principles-and-ops-ai-directive

Treasury Board of Canada Secretariat. (2018). *Data sovereignty and public cloud*. Government of Canada. https://www.canada.ca/en/government/system/digital-government/digital-government-innovations/cloud-services/digital-sovereignty/gc-white-paper-data-sovereignty-public-cloud.html

Ontario Human Rights Commission. (January 21, 2026). *Principles for the responsible use of artificial intelligence*. https://www.ohrc.on.ca/en/principles-responsible-use-artificial-intelligence (ohrc.on.ca)

Organisation for Economic Co-operation and Development. (2024). *AI principles*. https://www.oecd.org/en/topics/ai-principles.html

Standards

Canadian Standards Association

Institute of Electrical and Electronics Engineers Standards Association

International Standards Association

National Institute of Standards and Technology

Prompt Engineering Resources

Google Prompting Essentials

Claude 101 / Interactive Tutorial

Create Effective Prompts for Generative AI Training Tools

Text Generation Tutorial

References

Figure 1: Grandjean, M. (2021). Survivorship Bias. Survivorship Bia. Wikimedia.

Glossary of Key Terms

Based on the provided documents, here is a glossary of key terms focused on data science, machine learning, statistics, and the responsible use of AI. ISO/IEC 22989 is also an excellent resource for key concepts and definitions, and a Canadian version is available through CSA Group.

AI & Machine Learning Fundamentals

- **Artificial Intelligence (AI):** A multi-disciplinary field aimed at replicating all or part of human intelligence through technical work like machine learning and the ethical application of problem-solving.
- **Artificial General Intelligence (AGI):** A future-facing goal to create systems that can reason, learn, and apply intelligence to solve any problem, similar to a human being.
- **Artificial Narrow Intelligence (ANI):** The current state of AI, where systems are designed to perform a specific, well-defined task, often better than a human (e.g., chess-playing engines or medical image analysis).
- **Machine Learning (ML):** A subfield of AI where systems "learn" patterns and relationships directly from data rather than being explicitly programmed with rules.
- **Deep Learning:** A more advanced form of machine learning using neural networks with many "hidden" layers (Deep Neural Networks), allowing the system to recognize intricate patterns and abstract features.
- **Neural Network:** A computing system inspired by the biological brain, composed of interconnected nodes (neurons) organized in layers (input, hidden, and output).
    - **Transformer:** A revolutionary neural network architecture introduced in 2017 that uses an "attention" mechanism to process all input data points simultaneously, allowing it to understand deep contextual relationships.
- **Large Language Model (LLM):** Statistical mimics of language that appear to reason based on patterns in text data, but do not possess actual general awareness.

Data & Technical Concepts

- **Structured Data:** Highly organized data that fits neatly into rows and columns, such as spreadsheets, sensor logs, or databases.
- **Unstructured Data:** Information without a predefined organizational structure, such as text documents, images, video feeds, and emails.
- **Vectorization:** The mathematical process of converting unstructured data (like images or text) into an array of numbers (a vector) so it can be processed by machine learning algorithms.
- **Regression:** A machine learning technique used to predict continuous numerical values (e.g., predicting the remaining life of a mechanical part).
- **Classification:** A method where the model assigns an input to a specific category or class based on labelled training data (e.g., identifying a soil type or determining if concrete is "cracked").
- **Unsupervised Machine Learning:** A type of learning that works with unlabeled data to discover hidden patterns or structures, such as "clustering" similar data points together.
- **Vanishing Gradient Problem:** A technical issue in very deep networks where the adjustment signal used for training becomes too weak as it travels back through layers, causing early layers to stop learning effectively

Responsible AI & Risk Management

- **Responsible AI (AI Ethics):** A core component of the field that addresses the societal impact, accuracy, fairness, and accountability of AI systems.
- **Probabilistic vs. Deterministic:** AI outputs are probabilistic (based on statistical likelihood and uncertainty) rather than deterministic (guaranteed, predictable results), meaning they should be treated as recommendations rather than absolute facts.
- **AI Hallucination:** A phenomenon, particularly in LLMs, where the system produces factually incorrect, misleading, or entirely fabricated information.
- **Algorithmic Bias:** When AI systems perpetuate or amplify harmful human biases present in their training data, such as those leading to unfair or discriminatory outcomes.
- **Explainability (XAI):** The property of an AI system to express how the AI system arrived at a specific decision or output.
- **Human-in-the-Loop:** A protocol ensuring that AI inputs and outputs are reviewed and approved by a human professional to maintain a clear line of accountability and safety.
- **Emergent Misalignment:** A safety phenomenon where fine-tuning a model on a narrow task unexpectedly causes it to exhibit harmful or deceptive behaviours in unrelated areas.
- **Data Sovereignty:** The legal jurisdiction a piece of data is subject to, often related to the geographic location where it is stored.

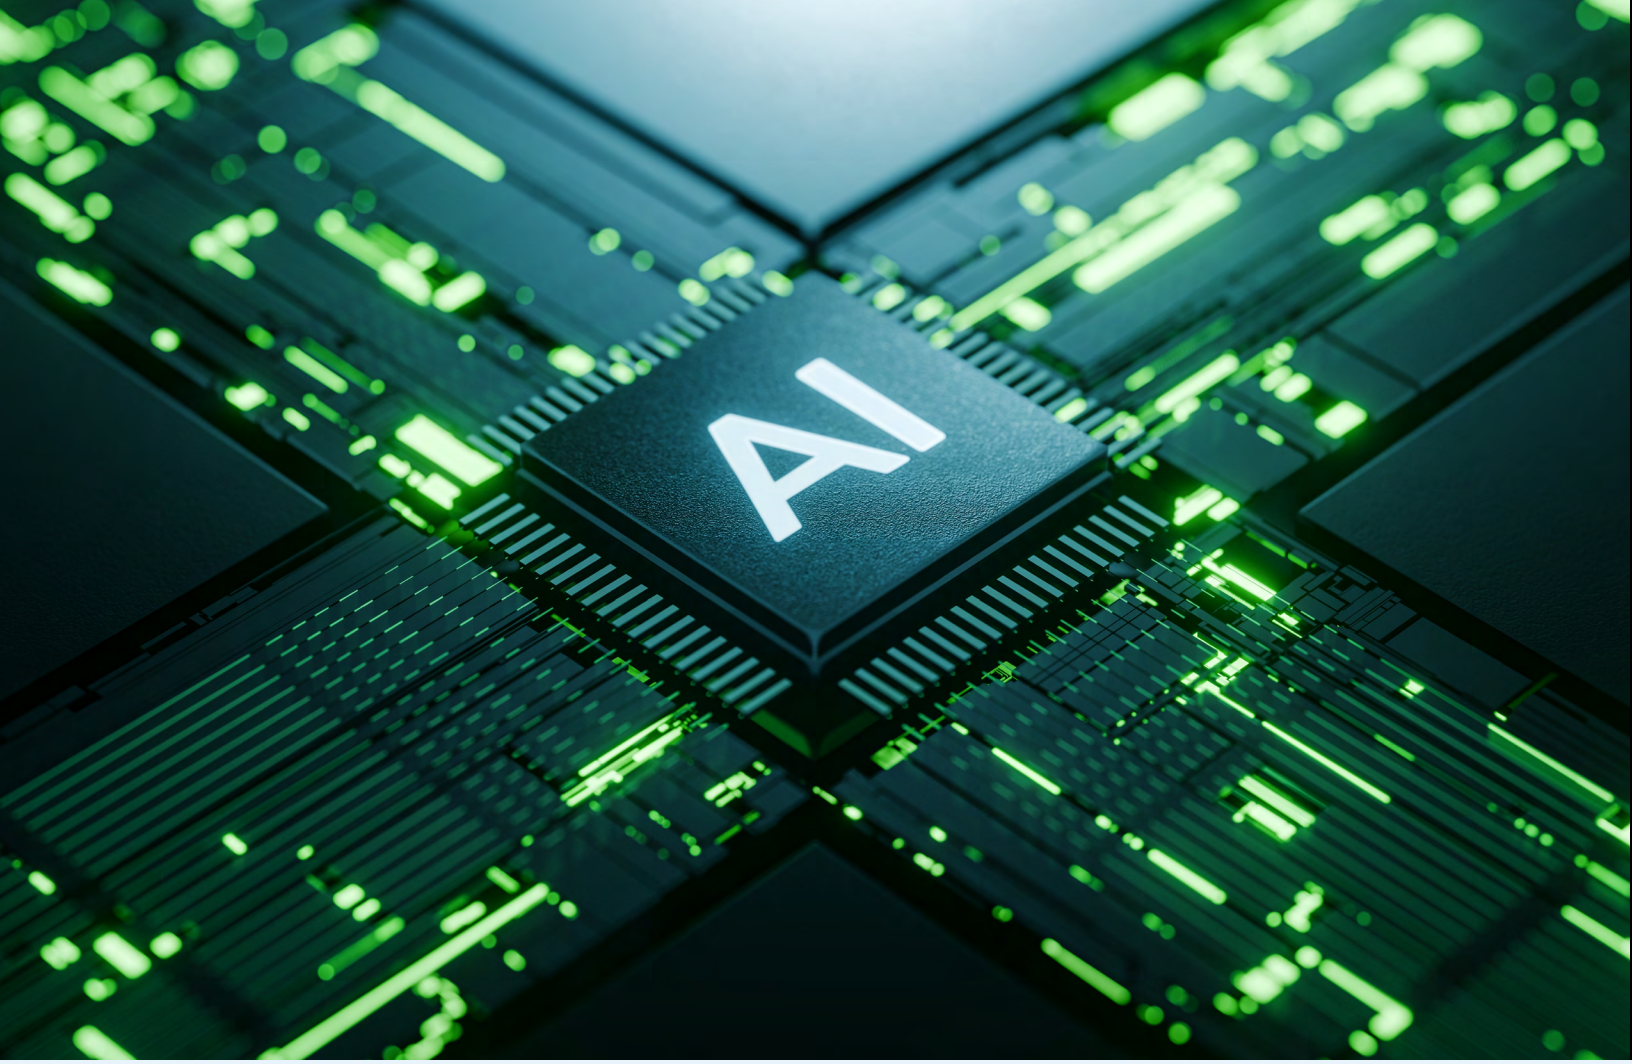

Contact Us

Ontario Society of Professional Engineers
5000 Yonge Street, Suite 701
North York, ON, M2N 7E9
1-866-763-1654
info@ospe.on.ca

ONTARIO SOCIETY OF PROFESSIONAL ENGINEERS