Artificial Intelligence Policy

—

Policy on the Development and Use of Artificial Intelligence Systems at Enbridge

DOCUMENT INFORMATION

DOCUMENT TYPE	VERSION NUMBER	EFFECTIVE DATE	FUNCTIONAL OWNER	CONTRIBUTIIONS	DOCUMENT MANAGER	DOCUMENT OWNER
Policy	2.0	September 11, 2024	TIS (Technology and Information Services)	GRC Cyber Security Legal Privacy SCM	Director TIS Technology & Innovation Lab	Chief Information Officer

ENBRIDGE

VERSION REGISTER

VERSION #	DATE	SUMMARY OF CHANGES
2.0	Sept 11, 2024	Policy updates for groups and specific clarification on Do's and Don'ts
1.0	June 13, 2023	Policy created

1 Purpose

This policy establishes principles and rules for the design, development, use or access of artificial intelligence systems at Enbridge, and applies to both systems designed by Enbridge or by third parties. Artificial Intelligence (AI) is defined as a combination of systems, software and algorithms devoted to processing data to perform functions and tasks normally associated with human intelligence, such as but not limited to reasoning, problem-solving, recognition, prediction, understanding, learning and self-improvement. AI can mimic human behaviors by analyzing substantial amounts of data using a combination of statistics and algorithms to accomplish complex and adaptive goals.

Enbridge is committed to trustworthy and responsible development, deployment, and use of AI systems. Use of AI supplements good human judgment, but it is not a substitute for it. All uses of AI at Enbridge – whether developed by Enbridge or by third parties – must support Enbridge's values and comply with our policies, including the Statement on Business Conduct and Acceptable Use of Technology Assets Policy (AUP).

2 Scope

This policy applies to Enbridge employees, directors, contractors, and agents of Enbridge (collectively, "Users") who are engaged in and use Artificial intelligence (AI) systems at Enbridge. This policy applies whether AI is developed within Enbridge and/or contracted in the development of AI.

1. End Users: This represents a wide population across Enbridge that can be at the Front-Office, Field-Office, or Back-Office; employees, vendors, or contractors.
2. Design, Development, or Sustainment: This represents the individual or team accountable for designing, developing, deploying, maintaining, and supporting an AI system.
3. Business Owner: This represents the individual or team accountable for accepting the risks vs rewards (i.e., value) as the result of using AI to achieve a business outcome.
4. 3rd Party suppliers: This represents the external vendors and/or companies that provide systems and/or services that directly or indirectly uses AI

3 Roles and Responsibilities

The roles and responsibilities related to this policy include:

CHIEF INFORMATION OFFICER (CIO)

- As Policy owner, ensuring that awareness of the policy exists throughout the enterprise.
- Approving or obtaining approval of proposed amendments to the policy and ensuring communication of such amendments throughout Enbridge.

DIRECTOR, TIS TECHNOLOGY & INNOVATION LAB

- Overall administration of this policy, including conducting regular reviews of the policy content, ensuring all Users have access to appropriate training materials and advising TIS groups on the enforcement of the policy.
- Validating content of this policy and ensuring maintenance of it.

4 Risks

AI presents opportunities for numerous benefits, but the use of AI systems also brings with it novel or increased risks compared to traditional software. A few examples of the types of risks earlier adopters of AI have encountered are:

- Harmful bias and other data quality issues affect the trustworthiness of AI systems. In some cases, these have led to lawsuits alleging discrimination, and in other cases, these could lead to human, environmental and other safety concerns where AI is used as part of a decision-making process.
- Increased complexity of “attack surfaces” and other novel security concerns enabled by AI systems.
- Risks to confidentiality, privacy/legal, and the potential loss of ownership or control of intellectual property using third-party AI systems. These risks can also arise internally due to the enhanced data aggregation capabilities of AI. AI may gather enormous amounts of data to aid in decision making that may complicate compliance with existing data governance policies and procedures.
- Lack of transparency in how decisions are reached and a lack of reproducibility of results, which complicates auditability and reduces the ability to correct and remediate any issues. Related to this is the potential for an inability to predict or detect unintended side effects.

There are other risks that need to be considered, some of which are compiled in the Appendix –AI-specific risks

Any Financial, Operational, Health & Safety, Environmental, Privacy/Legal or Reputational risks require assessment and mitigation.

5 Policy Statements and Principles

Restricted Purposes Requiring VP (Vice President) + Approval in Consultation with Legal Services:

Due to these inherent risks, development, and use of AI systems for any of the purposes listed below is restricted. Proposals for the development and use of AI systems for these restricted purposes shall carry out an AI Impact Assessment and must receive prior written approval from a Vice President, or higher-level Enbridge officer, in consultation with Legal Services (this consultation is required and must also be documented). The restrictions are intended to ensure that certain uses receive the appropriate level of senior management attention in partnership with Legal Services, so that there is a thorough understanding of risks and a plan to implement mitigations prior to development and use.

- Health and safety-related decision-making, including but not limited to decisions regarding hazard management, confined space entry, ground disturbance, isolation of energy systems, bypassing of safety controls or any other decision within the scope of our Life-saving Rules.
- Employment-related decision making must be free from bias or discrimination, such as decisions on which candidates to interview or hire, performance evaluation of employees or decisions on termination of employment.
- Gathering and analysis of personal information (information relating to a specific individual, or that could be used to identify a specific individual), or use of personal information in creating prompts to feed into AI systems.
- Preparation of regulatory-required reporting or other filings, or communications with regulators.
- Preparation of communications to the public, customers, employees, suppliers, or business partners.

Confidential, Restricted or Internal Enbridge Information

Confidential, restricted or internal Enbridge information must not be shared with, fed into or otherwise made accessible to third-party AI systems without first following the TIS Cybersecurity Third Party Risk Management process, including conducting a Detailed (AI) Risk Assessment, and obtaining the review and approval of any licenses, contracts, agreements, terms of service or related terms and conditions for use of the third-party AI system by Legal Services. Information about a specific individual or individual(s) must not be fed into the system.

Intellectual Property

Third-party AI systems must not be used to generate material that may be subject to intellectual property rights without first getting written legal advice from Legal Services regarding the ownership and use of such material. For clarity, this includes any generated material that might be subject to copyright (original works of authorship, including pictures, illustrations, drawings–including technical designs and mask work, artistic works, musical compositions, computer programs and a variety of textual works, such as instant messages, emails, letters, web or social media postings, pamphlets, brochures, articles, books, scripts, etc.) or that may be subject to patent rights (processes, machines, manufactures or compositions of matter that have a useful result, or improvements on the same). If you are uncertain whether the material generated by an AI system may fall into these categories, confer with Legal Services.

All employees must transparently indicate when output is generated by AI so that these types of risks can be evaluated before such use is made. Individuals will need to be responsible in the use of AI systems such that it does not contravene any ethical, safety, security, privacy, or legal policies such as the Statement of Business Conduct and other Enbridge policies. AI needs to be used in a responsible and ethical way and with an intent of purpose while exercising a degree of caution that balances potential opportunities without violating Enbridge policies.

Principles Respecting the Development and Use of AI Systems for Enbridge

All development and use of AI systems for Enbridge must adhere to the following principles:

1. SAFETY

All uses of AI at Enbridge must adhere to our core values of Safety, Integrity, Respect, and Inclusion and follow the Enbridge Safety Principles. With respect to safety in particular, AI systems must be responsibly designed with safety in mind and following proper development practices to prevent failures or dangerous system conditions that could cause physical or mental harm or threaten our employees, contractors, the communities in which we operate and the environment. AI systems must be thoroughly tested, must provide clear instructions on how to use the system appropriately to those who are using the system and making decisions based on its

Do's & Don'ts:

- DO conduct thorough testing of AI systems, including edge cases and scenarios that might challenge the system's safety features.
- DO schedule regular updates and maintenance to address potential vulnerabilities and ensure that the AI system is up to date with the latest safety standards.
- DO NOT disregard domain/subject-specific knowledge and expertise when designing safety measures. Collaborate with professionals in the relevant field to understand and address specific safety concerns.
- DO NOT skip comprehensive risk assessments during the development phase. Identify potential risks, evaluate their impact and implement appropriate safeguards to mitigate them.

2. VALIDITY AND RELIABILITY

AI systems used at Enbridge will strive for the highest possible level of accuracy and reliability, functioning to consistently produce results that are true within the system's operating parameters. As validity and reliability are key components of a trustworthy AI system, designers and users of these AI systems must ensure ongoing auditing or monitoring of the system to ensure correct functioning and to minimize any potential harm or other negative outcomes of a system failure.

Do's & Don'ts:

- DO ensure high-quality data is used for training AI models. Validate and clean datasets to remove inaccuracies, biases or outliers that could compromise validity and reliability
- DO implement cross-validation techniques to assess generalization and/or use external dataset to asset performance in the real-world beyond training data
- DO understand the terms, conditions and limitations of an AI solution
- DO monitor the response to ensure correctness, completeness, and consistency
- DO NOT ignore data quality issues. Neglecting data quality can significantly impact the validity and reliability of AI models, leading to inaccurate predictions and decisions.
- DO NOT ignore temporal changes in data distributions. Regularly update models to adapt to changes in the underlying data, ensuring ongoing validity and reliability.
- DO NOT use the AI system or solution beyond its stated design and scope

3. TRANSPARENCY AND ACCOUNTABILITY

All uses of AI at Enbridge must be able to transparently demonstrate that they conform to Enbridge's values and comply with our Statement on Business Conduct. AI must not be used as a substitute for good human judgment, particularly in areas where transparency of the decision-making process is required to establish trust in the process and to ensure health, safety, the environment, and human rights are appropriately protected. We, the people of Enbridge–directors, officers, employees, contingent workers, contractors, vendors, consultants and other third parties working with Enbridge–are responsible for the decisions we make and the actions we take or choose not to take.

All AI systems developed by or used at Enbridge must be designed and operated with appropriate human direction and control and must be reviewed regularly to ensure new AI features and capabilities are reviewed for potential impacts to other uses of AI at Enbridge. To guard against potential bias, the output should not be used unless reviewed by someone who understands how the model works (AI subject matter experts) together with someone possessing business expertise in the subject matter who is able to gauge the accuracy/quality of the inputs and outputs (Business SME (Subject Matter Expert)).

Do's & Don'ts:

- DO design AI systems to provide clear and transparent explanations for their decisions.
- DO maintain an updated record of instances & versions of AI use for work purposes and be able to share those records with your manager or other authorized company personnel upon request.
- DO NOT develop AI systems in isolation. Involve diverse teams, including legal and domain experts, to ensure a comprehensive approach that considers various perspectives
- DO NOT neglect the implementation of accountability measures. Establish clear lines of responsibility for accepting the outputs, development, deployment and maintenance of AI systems

4. EXPLAINABILITY AND INTERPRETABILITY

To manage an AI system effectively and responsibly, Enbridge must explain how the components of the AI algorithm operate and interpret the system's output. In other words, we must not use AI where we are simply putting data into a black box and getting a result without further understanding; instead, we must know how the AI functions and what it is designed to do, be able to understand the output in that context and be able to explain all of that in terms understandable by humans. This level of explainability and interpretability will enable Enbridge to better manage and govern our AI systems, as well as underpin the adherence to all other principles by providing a deeper understanding as to the "how" and "why" the AI system provided the result that it did. The output should not be used unless reviewed by someone who understands how the model works (AI SME) together with someone possessing business expertise in the subject matter who is able to gauge the accuracy/quality of the inputs and outputs (Business SME).

Do's & Don'ts:

- DO clearly define the goals and objectives in your AI system. Understand the specific requirements of your application and stakeholders.
- DO document the design, training process and decision-making logic of your AI system comprehensively. This documentation should be easily accessible to both technical and non-technical stakeholders.
- DO use techniques such as cross-verification, human review, training with quality data and user feedback to detect possible "hallucinations" from AI.
- DO NOT use black-box AI without a clear justification. If a complex model is necessary, ensure that you have appropriate methods in place to explain its decisions.
- DO NOT keep the inner workings of your AI system completely opaque. Lack of transparency can lead to distrust and hinder the acceptance of AI.

5. FAIRNESS

AI at Enbridge must be designed and used to promote and ensure fairness and address diversity, equity and inclusion in ways that mitigate bias and discrimination, conforming to Enbridge's values and Statement on Business Conduct. The attributes of AI systems that make them desirable for use as tools – such as speed and scale – also means they have the possibility to perpetuate and amplify biases and discrimination more quickly and more broadly than ever before.

Do's & Don'ts:

- DO ensure that training datasets are diverse and representative of the population to avoid biases and ensure fairness across different demographic groups.
- DO review output of AI applications to make sure it meets Company's standards for principles of equity, ethics and appropriateness.
- DO NOT use any output that discriminates against individuals based on race, color, religion, sex, national origin, age, disability, marital status, political affiliation or sexual orientation.
- DO NOT build AI models with non-diverse development teams. Lack of diversity may lead to unintentional biases, and diverse perspectives can contribute to more inclusive and fair AI systems.

6. SECURITY AND RESILIENCE

AI systems in development or use at Enbridge must adhere to all TIS policies, standards, guidance, and controls including, among others, those related to cybersecurity, architecture and technology governance that seek to maintain the elevated level of security and resilience of all systems in development or use at Enbridge.

Do's & Don'ts:

- DO conduct regular (automated) security scans & audits to identify vulnerabilities and weaknesses in the AI system. Address any issues promptly to enhance the overall security posture.
- DO adhere to established security standards and best practices. Follow industry-recognized security guidelines in its use or develop of a resilient AI system.
- DO NOT ignore emerging cybersecurity threats. Stay informed about new vulnerabilities and attack vectors relevant to AI systems and adapt security measures accordingly.
- DO NOT install unapproved Application Programming Interfaces (APIs), plug-ins, connectors, or software related to AI systems.

7. PRIVACY AND CONFIDENTIALITY-AWARENESS

AI systems can create specific challenges in terms of compliance with various privacy requirements of provincial/state and national laws, including with respect to consent to use and anonymization of data. Accordingly, as with all system development and data use at Enbridge, AI systems that will access, store, or use personal information must be operated in accordance with Enbridge's *Privacy Policy*. Confidential, restricted, or internal Enbridge information must not be shared with, fed into, or otherwise made accessible to third-party AI systems without following the requirements and restrictions noted above**.** Information about a specific individual or individual(s) must not be fed into the system.

Do's & Don'ts:

- DO stay compliant with relevant privacy laws and regulations. Be aware of jurisdiction-specific requirements to ensure that the AI system aligns with legal standards for data protection (AI systems can easily deployed everywhere without regional considerations).
- DO implement privacy principles from the outset. Integrate privacy into the design and development of AI systems to ensure that data protection measures (e.g. anonymization, pseudonymization, data minimization)
- DO NOT input company intellectual property into non-approved AI applications.
- DO NOT enter personal information (PI) of employees, customers or other third parties into any non-approved AI application. Treat the application as you would an employee of another company with whom we have no formal relationship.
- DO NOT exclusively use AI for decision making purposes without oversight & governance.

6 Acceptable and Prohibited Use Cases

All use of AI while performing work is subject to senior management approval. Examples of use cases are below and subject to change. The table is not reflective of all use cases and scenarios but takes into consideration the following:

- **Context/Background**: A description of the project, intended use of AI and who may be impacted.
- **Information Classification & Privacy**: Risks related to the handling and storage of sensitive data (public, internal, confidential, and restricted).
- **Bias and Fairness**: Risks of the AI system exhibiting biased behavior or unfair outcomes.
- **AI Type**: The specific AI algorithm being used.
- **Security Risks**: Risks of the AI system being vulnerable to attacks or misuse.
- **Financial, Health & Safety, Environmental, & Operational Risks**: Risks related to the AI system's correctness, consistency and completeness that impact critical systems.
- **Ethical, Privacy/Legal, and Legal Risks**: Risks of the AI system violating ethical, privacy/legal standards or legal regulations.

6.1 PERMISSIBLE: GENERAL USE CASES

- Research for information: Translating text from a secondary, publicly available source. Conducting high-level background research into a non-sensitive topic.
- Employee Training: AI can be used to create personalized training programs for employees based on their skills and learning pace. This can improve the efficiency and effectiveness of the training process.
- Data Analysis: AI can analyze enormous amounts of data to provide insights and forecasts about market trends, customer behavior and operational efficiency, helping businesses make informed decisions.
- Customer Service: AI chatbots can handle routine customer queries, freeing human customer service representatives to handle more complex issues. This can improve customer satisfaction and efficiency.
- Cybersecurity: AI can help enhance a company's cybersecurity measures by detecting suspicious activity and potential threats. This can help protect the company's data and digital assets.

- Sales Forecasting: AI can analyze sales data and market trends to predict future sales, helping businesses plan their inventory and marketing strategies accordingly.
- Marketing Personalization: AI can analyze consumer behavior and preferences to create personalized marketing campaigns, improving customer engagement and sales.
- Predictive Maintenance: AI can predict when equipment or machinery is likely to fail or need maintenance, preventing downtime and reducing repair costs.
- Supply Chain Optimization: AI can optimize supply chain management by predicting demand, managing inventory, and identifying potential disruptions.
- Product Development: AI can analyze customer feedback and market trends to identify potential new product features or entirely new product ideas, speeding up innovation and staying ahead of competitors.
- Original intent in the use of AI is to:
    - Increase or Improve: revenue, user experience, productivity, or efficiency
    - Reduce: costs, risks, duplication or workload

6.2 PROHIBITED: GENERAL USE CASES

- Legal guidance on business decisions
- Drafting of legal documents
- Prompting about company finances
- Preparation of documentation, reports, or correspondence with external entities (e.g., client-facing, external counsel, or legal proceedings, regulatory, landowners, government, special interest groups, media)
- Analyze personal data to make predictions on performance evaluations, behavior, preferences, or health status
- Screening job applications for the purpose of hiring
- Create and/or manipulate images, video, audio, or data (i.e., deep fakes)
- Creating chatbots that impersonate humans and leading individuals to believe they are interacting with a real person
- Monitor activities, conversations and productivity of employees and contractors without consent or knowledge
- Social manipulation of opinions or promulgating of propaganda, false information or influencing
- Perpetuating existing biases or stereotypes
- Any activity with the intent of deception, misinformation, fraud, or manipulation
- Any decision or action that results in Financial, Operational, Reputational, Privacy/Legal, Health & Safety or Environmental Risk

6.3 PROHIBITED DATA: USED IN THE TRAINING, DEVELOPMENT, OR INPUT/PROMPTING

- Personal Information (PI) Data: PI is information used to identify a specific individual - human data. Remember, if it is about a human, it is personal information. This can include individual's identity, name, social security numbers, email address, biometric, personal characteristics or preferences, digital identities, payroll information, performance records or phone numbers. This applies for employee, contractor, and customer data.
- Health Information Data: This includes medical records, health insurance details or other healthcare-related personal information (i.e., medical history, diagnostic, treatment, insurance, family health or lifestyle)
- Login Credentials: Usernames and passwords for online accounts, especially those related to banking, email, and social media. Compromised credentials can lead to unauthorized access and further data breaches.
- Financial Information Data: Details related to banking, credit cards, financial accounts, financial records, transaction details, income/tax data, purchases, sales, or investment information
- Business-sensitive Information Data: Business data such as trade secrets, proprietary data, operational or strategic information. Examples: client information, strategic plans and research, legal proceedings, mergers & acquisition data, intellectual property, contractual terms & conditions, and sensitive production/operational data.

7 Enforcement

Compliance with this policy is mandatory and subject to the enforcement below.

Enbridge has the right, at its sole discretion and without prior notice, to amend and/or modify any of the provisions of this policy as business needs dictate.

Violation of this policy will be considered a violation of the Statement on Business Conduct and will be treated accordingly.

Failure to follow this policy may result in an investigation, suspension, termination of employment or contractual relationship, legal action, or other action, including an injunction to restrain further breaches and to recover any information or media, and for the recovery of damages suffered by Enbridge. Information may be provided to applicable law enforcement authorities for potential criminal prosecution to the extent required or permitted by law.

Suspected violations of this policy must be reported to the appropriate supervisor, manager, department head, Human Resources or to Ethics and Compliance. To the extent permitted by applicable law, nothing in this policy is intended to modify or contradict Enbridge's position in the United States as an at-will employer, nor should this document be construed to create any express or implied guarantee as to the fact of, duration or terms of an employee's employment.

8 Questions or Concerns

Questions regarding the implementation of, or compliance with this policy should be directed to the TIS Governance, Risk and Compliance department via email TIS GRC (Governance, Risk & Compliance). Alternatively, you can contact the Document Manager identified in the document information table.

9 Related Documents

DOCUMENT NAME	DOCUMENT LOCATION
Acceptable Use of Technology Assets Policy	ELink
Statement of Business Conduct	ELink
Privacy Policy	ELink
Cybersecurity Policy	ELink
Technology Stewardship Policy	ELink

10 Sources

External sources that have been cited or consulted in the development of this policy include the following:

- United States National Institute of Standards and Technology's (NIST) Artificial Intelligence Risk Management Framework (AI RMF).
- International Technology Law Association's Responsible AI: A Global Policy Framework.
- The Artificial Intelligence and Data Act (Canada) and Companion Document
- Microsoft Responsible AI Standard (v2).
- Google's AI Principles.
- IBM's Principles for Trust and Transparency.
- IBM's AI Ethics.
- Capgemini's Code of Ethics for AI.

11 Appendix

AI-specific Risks

Some examples, taken from the United States National Institute of Standards and Technology's (NIST) Artificial Intelligence Risk Management Framework (AI RMF):

- The data used for building an AI system may not be a true or appropriate representation of the context or intended use of the AI system, and the ground truth may either not exist or not be available. Additionally, harmful bias and other data quality issues can affect AI system trustworthiness, which could lead to negative impacts.
- AI system dependency and reliance on data for training tasks, combined with increased volume and complexity typically associated with such data.
- Intentional or unintentional changes during training may fundamentally alter AI system performance.
- Datasets used to train AI systems may become detached from their original and intended context or may become stale or outdated relative to deployment context.
- AI system scale and complexity (many systems contain billions or even trillions of decision points) housed within more traditional software applications.
- Use of pre-trained models that can advance research and improve performance can also increase levels of statistical uncertainty and cause issues with bias management, scientific validity, and reproducibility.
- Higher degree of difficulty in predicting failure modes for emergent properties of large-scale pre-trained models.
- Privacy risk due to enhanced data aggregation capability for AI systems.
- AI systems may require more frequent maintenance and triggers for conducting corrective maintenance due to data, model, or concept drift.
- Increased opacity and concerns about reproducibility.
- Underdeveloped software testing standards and inability to document AI-based practices to the standard expected of traditionally engineered software for all but the simplest of cases.
- Difficulty in performing regular AI-based software testing, or determining what to test, since AI systems are not subject to the same controls as traditional code development.
- Computational costs for developing AI systems and their impact on the environment and planet.
- Inability to predict or detect the side effects of AI-based systems beyond statistical measures.

12 Glossary

Artificial Intelligence (AI) is a technology that enables computers and machines to perform tasks that typically require human intelligence. This includes things like understanding language, recognizing patterns, solving problems, predicting outcomes, optimizing, generating content, and making decisions.

AI system refers to any hardware and/or software that supports AI. This can be a stand-alone solution or embedded within an application.

Bias in AI: Refers to the tendency of an AI system to produce results that are unfair or prejudiced due to errors in the data or the way the AI was designed. This can happen because the data used to train the AI might reflect human biases or because the algorithms themselves have flaws. In simple terms, if an AI system is biased, it means it might make unfair decisions or predictions, like favoring one group of people over another.

Cross-verification is a crucial step that checks the AI's output against reliable sources of information.

Generative AI ("GenAI") is a special type of AI that can create added content. This means it can generate text, images, music, and even videos based on the data it has been trained on.

"**Hallucinations**" refers to when an AI system, like a chatbot or image generator, produces information or content that is incorrect or nonsensical but presents it as if it were true. AI "imagines" something that is not real or makes up facts that are not based on any actual data.