




Artificial Intelligence Serving Municipalities

A practical guide

in collaboration with



CREDITS

Principal authors

Open North

**Samuel Kohn, Leandry Jieutsa
and Thomas Linder**

Contributors

Maison de l'innovation sociale

Sarah Libersan

IVADO

**Antoine Congost, Halima Bachir,
Réjean Roy and Catherine Régis**

We thank the Office of the Vice-Rector
for Community and International
Partnerships at Université de Montréal
for its support during the writing
of the initial version of this guide.

September 2025

CONTENTS

<u>03</u>	Background
<u>04</u>	Understanding Artificial Intelligence
<u>06</u>	The Main Types of AI Model
<u>09</u>	Applications and Benefits of AI in Municipal Settings
<u>10</u>	Deploying AI Ethically and Responsibly: Principles and Governance
<u>15</u>	Strategic Questions to Ask Before Adopting AI in a Municipal Setting
<u>18</u>	Key Steps in Responsible Deployment of AI
<u>20</u>	Key Players in the Municipal AI Ecosystem in Québec
<u>21</u>	Toolkit
<u>22</u>	Glossary
<u>23</u>	References

BACKGROUND

Municipal stakeholders are often at a disadvantage with regard to the rise of artificial intelligence (AI). They may lack the technical resources, in-house capacity and/or governance frameworks necessary to assess its impacts.

Caught between the promise of efficiency and risk in terms of transparency, algorithmic bias and cybersecurity, municipal decision-makers may struggle to make informed decisions. Without clear guidelines, incorporation of AI is likely to proceed according to opportunistic approaches or inefficiently. Yet, its adoption must follow a coherent strategy, adapted to the needs of citizens and the municipality's values. In recent years, many AI projects in urban settings, both in Quebec and elsewhere, have faced challenges of implementation, sometimes leading to high costs and raising social or reputational concerns.


This guide is designed to give decision-makers the tools to assess the relevance of AI solutions, and provides answers to two key questions:

1. Is AI indeed a relevant solution for municipal issues?
2. If so, what set of conditions must be in place to ensure its responsible deployment?

How municipalities generally are handling the rise of AI

Municipalities of all sizes, in Canada and internationally:

- are under constant pressure to adopt AI, both to increase efficiency and respond to citizen expectations for modern, high-performance services;
- are still in a learning phase and require training as well as guidelines to support the transformation;
- have limited in-house resources and capacity;
- must deploy AI responsibly, which demands specialized skills and appropriate resources.



UNDERSTANDING ARTIFICIAL INTELLIGENCE

AI encompasses technologies that can analyze vast amounts of data using advanced computing resources. It supports human decision-making by generating predictions, alerts, recommendations or automated actions to suit specific needs.

The lifecycle of an AI-powered solution

To illustrate how AI works, here is an example of a solution deployed in a municipal setting and the stages of its development.

1. Define the problem

Municipal administrators notice that employees are spending too much time answering questions from citizens. They plan to automate part of those employees' work so that they can focus on tasks with higher added value and provide better support to residents most affected by digital inequity.

2. Choose a tailored solution

The city decides to implement a conversational AI, or chatbot—a virtual assistant that will provide automated answers to questions that recur often.

Which types of AI, for which needs?¹

Analytical AI: Detects trends and patterns in data to support decision-making.

Functional AI: Performs actions or tasks based on data collected (e.g., from satellites or sensors).

Visual AI: Interprets still or video imagery and extracts meaningful information from it.

Textual AI: Processes, analyzes and generates human language, both written and spoken.

Interactive AI: Enables dynamic, personalized two-way communication with users.

3. Collect and prepare the data

Municipal administrators identify and gather the data needed to answer the questions most frequently asked by citizens about the municipality's services and by-laws. They realize, however, that some of those data are obsolete, inaccurate or not representative of the population, and must invest time in correcting this information before using it.

1. Othengrafen, F., Sievers, L., & Reinecke, E. (2025).

Data

Data are the raw material of any AI system. They may be numbers, texts, sounds or images describing a region, infrastructure, movements or a population. When properly exploited, they enrich information and support **public policy decisions**.

Warning: Data are never neutral. They are the reflection of a worldview and can be biased or non-representative, depending on which data are collected, which are excluded, how they are interpreted, or how they are processed by systems.

4. Train the AI model

With support from an AI solution provider, the municipal administrators train the chatbot using their data, to make sure it will generate accurate and reliable answers. This stage is used to test the effectiveness of the tool and adjust its performance, both in terms of the relevance of the answers and the quality and compatibility of the data.

5. Deploy the solution

The municipality deploys the chatbot, organizing a public launch accompanied by a communications campaign and information tools for citizens. The tool must be accompanied by a privacy policy and a statement of its limitations, but it is the municipality that is responsible for fielding complaints, following up on them and making corrections.

Algorithms and models

An algorithm is a sequence of instructions used to automatically process information in order to accomplish a specific task.

AI models are developed by “training” algorithms on large quantities of data. This enables them to recognize patterns and adapt. By means of performance feedback, the models gradually improve, with or without human supervision.

6. Track and evaluate performance

The municipality continuously monitors the reliability of the answers provided, the quality of the data and whether objectives are being met. This task is performed by an expert individual or team assigned to gather citizen feedback, periodically test the chatbot and ensure that any necessary corrections are made. This thorough monitoring supports strategic decisions and helps to sustainably improve service.

THE MAIN TYPES OF AI MODELS

Traditional AI

Traditional AI, also called symbolic AI, is a type of AI that uses algorithms that are based on logical rules defined by humans.

Benefits: Traditional AI is well suited to environments with stable, well-defined rules. It delivers consistent, reliable and transparent results, and every decision can be directly traced back to the instructions set by the municipality. Because of its simplicity and relatively low cost, it is also easier to audit and maintain than a complex model.

Limitations: It cannot adapt to unforeseen situations and is therefore useful only for tasks strictly defined by exhaustive criteria.

Examples in municipal settings:

Automated verification of a construction or renovation permit application's compliance with by-laws, or of a citizen's eligibility for a municipal program, based on objective criteria.




Image 1 – AI and Its Main Sub-Fields
(source: IVADO)

Machine Learning

Machine Learning relies on statistical methods that can identify trends and patterns in large volumes of data. This process is not confined to application of predefined rules; it gradually extracts regularities from the available data.

Benefits: It can deal with complex, evolving problems that are beyond the capacity of classic algorithms. It generates forecasts and recommendations that are useful in municipal planning, adapting to new datasets.

Limitations: To function properly, it requires a large amount of representative, well-structured data. It may introduce errors or bias that are difficult to detect, and implementing it requires specialized skills and technical resources that are more expensive than those needed for traditional AI.

Examples in municipal settings:

Predicting demand for public transit based on peak hours or specific events; predicting maintenance needs for infrastructure (roads, water mains) using data from sensors as well as wear history data.




Image 2 – Supervised and Unsupervised Learning
(source: IVADO)

Deep Learning

Deep Learning relies on artificial neural networks that are capable of processing massive amounts of data. It detects complex patterns and directly learns to interpret images, sounds or text from examples provided, without the need for explicit rules-based programming.

Benefits: It enables exploitation of unstructured data (still or video images, sounds, natural language) and allows for advanced applications, which are especially useful for analysis of infrastructures or citizen relations.

Limitations: Deep Learning has some of the same constraints as Machine Learning (it requires large amounts of reliable data and carries a risk of bias), but also presents a major difficulty: its decision-making mechanisms are opaque and difficult to explain. Training a Deep Learning model is a long process

with significant technical, financial and energy costs, which can hamper regular updating and make it difficult for many municipalities to access.

Examples in municipal settings:

Automated detection of cracks in road surfaces using drone imagery; interpreting citizen requests in natural language using a better-performing chatbot.




Image 3 – Deep Learning Applied to Image Recognition (source: IVADO)

Generative AI

Generative AI produces original content (text, images, sound, video) based on instructions given in natural language (known as prompts). It does more than simply analyze existing data: it creates new content based on the data it has been trained on.

Large Language Models (LLMs), such as ChatGPT and Microsoft Copilot, are an advanced form of generative text AI, with the ability to understand and generate natural language at scale.

Benefits: Generative AI can quickly produce a wide variety of content useful for communicating to the public, planning and decision-making. It is flexible, enabling municipalities to save considerable time when producing documents and visuals.

Limitations: It can result in factual errors (“hallucinations”) or misinformation. It raises privacy issues, is highly dependent on the quality of the data used to train it, and tends to standardize content, which can impoverish the diversity of data. Its use also demands greater vigilance with regard to copyright and privacy protection.

Examples in municipal settings:

Designing visuals to illustrate a planned development or redevelopment of a public space; generating summaries of minutes of meetings, reports or by-laws to facilitate their distribution and ensure transparency.




Image 4 – Uses and Degrees of Risk of Generative AI (source: IVADO)

The Government of Canada has recommended that generative AI be adopted gradually in the public sector, beginning with low-risk uses so as to develop in-house competencies and test risk-management methods.² Applied to the municipal context, this means that cities and towns should begin by experimenting with AI for simple in-house tasks before applying it to more sensitive services that directly affect citizens.

The government also emphasizes guiding principles such as transparency, fairness and security, which can support municipalities in developing their own frameworks for the use of generative AI (see pp. 12 and 13). By following this guidance, local administrations strengthen public trust and align their initiatives with national and international best practices.

APPLICATIONS AND BENEFITS OF AI IN MUNICIPAL SETTINGS

AI is never neutral: algorithm design and usage context influence the results. It has the potential for aberrations, but when adopted responsibly and in line with the public interest, it can be a lever for improving services.

Many municipalities in Quebec and elsewhere in Canada have experimented with AI, demonstrating that when deployment is aligned with the public interest, it can improve civic management and citizen satisfaction. The following examples illustrate the wide range of concrete benefits.

1. AI for automation – Surrey, British Columbia

In 2024, the City of Surrey introduced the Development Inquiry Assistant (DIA), a pilot chatbot available 24 hours a day and in multiple languages, that answers questions about construction and renovation permitting as well as zoning. It processes an average of 460 requests a month, providing answers to citizens quickly while reducing municipal employees' workload ([City of Surrey 2024](#)).

→ **Area: Municipal administration**

→ **Type of AI: Text AI, Interactive AI**

2. AI for assistance – Laval, Quebec

In 2020, Laval incorporated an AI-powered virtual assistant into its 311 information service. The tool generates text transcriptions of interactions with citizens to provide data to dockets and facilitate processing of requests. It also files requests by type of service and directs citizens to the proper municipal resources more quickly. It processes some 250,000 calls a year, covering more than 250 topics ([Ville de Laval 2021](#)).

→ **Area: Citizen services**

→ **Type of AI: Text AI (Automated Natural Language Processing)**

3. AI for analysis – Montréal, Quebec

The Street Review project (in progress) combines AI and citizen input to evaluate road accessibility in Montréal. Around 100 residents annotated thousands of street-view images, with the annotations used to train an AI model, which then analyzed tens of thousands of other images.³ The results, presented in the form of heat maps, reveal accessibility inequities, notably for pedestrians and people with reduced mobility. The tool provides municipal decision-makers with precise data that can be used to target priority actions. The municipal party Ensemble Montréal has also proposed using AI to optimize road work and reduce the numbers of the city's ubiquitous orange cones.

→ **Area: Urban accessibility and land-use planning**

→ **Type of AI: Visual AI in combination with Machine Learning**

4. AI for improving traffic flow – Kirkland, Quebec

In 2025, Kirkland installed the NoTraffic platform at all traffic lights along Boulevard Saint-Charles, a saturated artery with traffic 40% above its capacity. AI-equipped sensors adjust traffic light synchronization in real time for various users (pedestrians, car, bus and truck drivers, cyclists), reducing congestion and improving traffic flow.

→ **Area: Mobility and infrastructures**

→ **Type of AI: Functional AI**

5. AI for prediction – London, Ontario

In 2020, the City of London introduced the Chronic Homelessness AI (CHAI) model, designed to predict the likelihood of an individual becoming chronically homeless within six months. The tool, which analyzes data such as age, family history and time spent in shelters, has a success rate in excess of 90%.⁴ It has already identified several dozen people at risk, making proactive interventions easier. London is the first municipality in Canada to leverage AI in guiding its policy on tackling homelessness.

→ **Area: Social services**

→ **Type of AI: Analytical AI (supervised Machine Learning)**

6. AI for detection – Windsor, Ontario

Windsor has adopted CITYROVER, an AI system embedded in municipal vehicles that automatically detects road surface degradations such as potholes and cracks. Each defect is georeferenced and added to detailed reports, which are then used to accelerate planning of repairs.

Since the tool was rolled out in 2020, productivity has increased by more than 900%; pothole claims have declined by 90%; and efficiency gains are estimated at over \$2.5 million (CITYROVER, n.d.). This innovation has improved preventive maintenance, optimized budgeting and led to greater citizen satisfaction thanks to better-maintained roads.

→ **Area: Mobility and infrastructures**

→ **Type of AI: Visual AI based on Deep Learning**

Explore even more projects from around the world via the [Atlas of Urban AI](#).

DEPLOYING AI ETHICALLY AND RESPONSIBLY: PRINCIPLES AND GOVERNANCE

Defining responsible AI

Responsible AI is the set of principles that guide the design and use of AI so as to maximize its positive outcomes while mitigating the associated risks.

Its definition remains complex, because it encompasses both technical and social dimensions and is based on varying perceptions of risk and benefit, and its impacts are unequally distributed across groups, raising issues of equity and justice.

The principles of responsible AI

Responsible AI is generally expressed through guiding principles. While they may appear abstract, they provide essential reference points and serve as tools for dialogue.

One recent international framework, to which Canada is a signatory, is the Council of Europe Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law. It covers the full lifecycle of AI systems and states nine fundamental principles. Some of them echo general obligations that are not specific to AI (protection of human rights, compliance with laws, protection of privacy and of data), while others take on new meaning in that context:

1. Human dignity and individual autonomy: Ensure that AI systems respect the right to self-determination and are not used to manipulate individuals.

2. Transparency and oversight: Ensure that AI systems can be examined and that clear human responsibility is assigned.

3. Accountability and responsibility: Officially designate people to be in charge of ensuring compliance with the principles and correcting errors in the event of failure to comply.

4. Equality and non-discrimination: Prevent all forms of unequal or discriminatory treatment in relation to the use of AI.

5. Reliability: Guarantee that AI systems function as planned, including in complex contexts, such as the use of chatbots.

Shortcomings

Two complementary principles merit further explanation when it comes to municipal uses:

1. Stakeholder participation:

Involving the relevant stakeholders from the outset (design, conception) is crucial to ensuring responsible AI, not least in order to protect vulnerable groups, who are often the most at risk. This involvement reinforces trust and promotes the success of projects.

2. Ecological sustainability: Because AI systems are so energy-intensive, consuming large amounts of electricity and water to cool data centres, it is essential to assess their overall environmental impact (emissions, strain on resources, impacts on ecosystems).

These orientations are included in other frameworks for responsible AI, including the [Montréal Declaration for Responsible AI Development](#) (2018). At the international level, the [UNESCO Recommendation on the Ethics of Artificial Intelligence](#) (2021), adopted by member states including Canada, is the first-ever global standard. It imposes clear requirements for transparency, non-discrimination, environmental sustainability and inclusive governance, asserting that AI must serve the common good and respect human rights.

These benchmarks provide municipalities in Quebec with a firm foundation on which to institute responsible governance and boost citizens' confidence in their digital choices.

WELL-BEING	PRIVACY AND INTIMACY
AUTONOMY	DIVERSITY INCLUSION
SOLIDARITY	SUSTAINABLE DEVELOPMENT
EQUITY	PRUDENCE
DEMOCRATIC PARTICIPATION	RESPONSIBILITY



Image 5 – The 10 Principles of the Montréal Declaration for Responsible AI Development
(source: IVADO)

AI governance for municipalities

Operationalizing AI is a question of governance, which defines policies, roles and processes to serve the public interest. Each municipality must adopt a robust framework based on four pillars:

1. **AI literacy**
2. **Data management**
3. **Oversight and accountability**
4. **Public engagement and transparency**

1. AI literacy

AI literacy is an essential skill, complementary to data literacy, that all municipal employees must acquire. It ensures they will have a minimum base of knowledge for ensuring responsible, effective AI governance, whatever their role, e.g., field agent, analyst, IT or cybersecurity specialist, or senior manager.

While most AI literacy programs cover only technical aspects (data, algorithms, models) or governance (suitability of solutions, supervision, procurement), they should also

comprise interactions with the public and change management. Given the distrust of AI and the perceived social risks, it is crucial that municipal players receive the proper training so as to strengthen trust and establish constructive dialogue.

At the organizational level, an AI literacy maturity assessment allows for identification of actual training needs and prepares the municipality for the exercise of responsible governance, even before AI technology is deployed. The ability to make informed decisions starts with acquisition of AI literacy.

Priorities for action

- Assess staff's AI literacy maturity so as to determine actual training needs and enable properly targeted actions (see AutodIAG in the Toolbox).
- Provide comprehensive training, encompassing not only technical aspects (data, algorithms, tools such as ChatGPT), but also change-management, ethics, digital transformation and public engagement.
- Include explicit AI literacy requirements in job postings so as to anchor these skills within the organization for the long term.

2. Data management

The quality of data and how they are managed determine the reliability of AI systems. To ensure responsible use, data must meet stringent criteria. They must be:

- **of high quality: up to date, complete, error-free and evaluated for bias;**
- **stored in suitable formats and accompanied by properly documented metadata;**
- **fully compliant with privacy regulations;**
- **classified by sensitivity level so as to prevent any possible leak when processed in a cloud computing environment;**
- **stored by means of the proper (local or cloud-based);**
- **under the supervision of a dedicated resource.**

Adoption of best practices in data management will strengthen all municipal functions by supporting evidence-based decision-making. Assessing data management maturity will help identify the investments and actions needed

Priorities for action

- Assess the maturity of data management to obtain an objective portrait of practices and target priority areas of improvement.
- Create and maintain a centralized inventory of municipal data specifying their sources, access privileges and sensitivity level.

- Establish stringent standards of quality, including complete documentation of metadata.
- Protect digital sovereignty by preferring local storage and, during the procurement phase, imposing strict conditions on access to and use of data by providers.

3. Oversight and accountability

Oversight and accountability must be grounded in well-defined policies, roles and responsibilities, ensuring responsible governance throughout the AI lifecycle, from design to decommissioning.

They require a designated manager (e.g., chief technology officer) who will be tasked with:

- **defining an AI strategy in collaboration with all stakeholders, thoroughly evaluating the relevant legal obligations, contractual options, risks, alternatives, costs and benefits;**
- **guaranteeing that development and deployment are compliant with best practices, including testing, regular audits and post-deployment monitoring;**
- **implementing feedback (complaints, errors, etc.) and repair mechanisms, comprising a transparent investigation, explanation and error correction process.**

Priorities for action

- Designate an AI officer responsible for oversight, consistency and accountability across all municipal projects.
- Map out, adopt and distribute an internal strategy and policy on the use of AI that includes oversight and risk-management practices.
- Implement accessible feedback and repair mechanisms to ensure transparency and swift error correction and preserve the public's trust.

4. Public engagement and transparency

AI is a sociotechnical system, shaped as much by technological choices as by the decisions and perceptions of the people involved. These dynamics inform problem definition, solution design, tool training and risk management.

This system engages a broad range of players: municipal employees, third-party developers, civil organizations, public groups, and affected communities. Responsible governance must guarantee meaningful participation by all stakeholders, particularly those who have traditionally been marginalized, so as to ensure that the tools are relevant, risks are controlled and public trust is maintained.

In addition to ad hoc consultations, transparency must be supported by permanent mechanisms, such as a publicly accessible registry of deployed AI systems, to provide citizens with clear information and continuous monitoring of system operation and security.

Priorities for action

- Establish, from the planning stages of any AI project, an advisory representative of all stakeholders, tasked with gathering their opinions and ensuring continuous monitoring.
- Prepare for consultations with clear documentation, and frame discussions using factual, constructive arguments.
- Publish a simple, accessible and regularly updated registry of AI systems in use, accompanied by a citizen feedback channel.
- Develop further transparency tools, such as a register of algorithms, inspired by the Algorithmic Transparency Recording Standard developed by the UK Government Digital Service.

↪ *This tool documents the objectives, functioning, limitations and expected impacts of each system, providing decision-makers and citizens alike with a clear, comparable understanding of the solutions.*

STRATEGIC QUESTIONS TO ASK BEFORE ADOPTING AI IN MUNICIPAL SETTINGS

Before investing in any AI solution, municipal decision-makers must assess its relevance and suitability for the problem being addressed. Each project must be thoroughly analyzed, weighing the anticipated benefits to ensure that innovation serves the public interest. There are nine key questions that can guide stakeholders' thinking.

QUESTION 1

Is AI really needed to tackle the problem, or are there more effective solutions?

AI must not be adopted simply because it is currently in fashion. Before proceeding with any deployment of an AI-powered tool, the actual causes of the problem must be analyzed and alternatives explored. For example, if there is a high volume of questions from citizens, the solution may simply lie in improving the website's user-friendliness or using clearer language. AI is not a miracle solution: it can resolve some issues, but may also create new ones. Municipal decision-makers should therefore regard it as one option among many, and often as a complement to simple, proven solutions.

QUESTION 2

Are the data of high quality, accessible and usable so as to achieve the objectives pursued by the AI system?

An AI system's performance depends on the quality of its data. Municipalities must ensure the availability, relevance and usability of data by imposing strict standards. Data must circulate properly between departments to avoid duplication and enable a big-picture view.

QUESTION 3

Could the data and algorithms used introduce bias and worsen digital inequalities?

Algorithmic bias can worsen inequalities, for example in the case of vulnerable populations seeking access to social services or housing. In 2021, Forbes reported that 40% to 80% of mortgage applications by Black people had been denied because of algorithmic bias, illustrating how skewed data or poorly designed models can produce discriminatory results.⁵

QUESTION 4

Are personal and sensitive data protected well enough to ensure their confidentiality?

The use of AI can be a threat to people's digital rights, especially the protection of personal information such as identity, health data or places frequently visited. In municipal settings, facial recognition technology and its infringement on privacy is one example of these risks.

QUESTION 5


Can the security of the data and the AI system be properly ensured?

AI technologies are vulnerable to cyberattacks and tampering. The large volumes of data that they depend on require highly secure storage and protection against misuse. Training datasets can also be the targets of "data poisoning" attacks, whereby maliciously altered information causes the system to malfunction.

QUESTION 6

Is the AI system transparent enough to be explainable, and have the responsibilities for its impacts been clearly defined?

Many AI systems operate like "black boxes." This opacity impedes understanding of their mechanisms and makes assigning responsibility difficult. Municipalities should require vendors to provide clear documentation on their data and decision-making criteria, state the division of responsibilities in contracts and designate an internal authority with ultimate responsibility. They may also wish to consult a lawyer to clarify the extent of their responsibilities and proceed in an informed manner in that respect.



QUESTION 7

Is the AI system sustainable, and does its environmental impact warrant its use?

AI also has significant environmental impacts. A University of Massachusetts study found that training one Large Language Model (of the kind used for generative AI) can generate emissions equivalent to 300 New York City–San Francisco return flights, or five times the emissions of an automobile over its entire lifecycle.⁶ Making a query using a tool like ChatGPT consumes around 10 times as much energy as conducting a Google search.

AI's environmental footprint is not confined to greenhouse gas emissions: it involves substantial consumption of natural resources, impacts on ecosystems, and the danger of worsening environmental inequalities.

QUESTION 8

Will using an AI system impair the quality of the services provided?

Adopting an AI system as replacement for an existing service can generate gains, but also losses—especially if it substitutes for human expertise, hallmarks of which include judgment, empathy and flexibility. Improperly calibrated automation may result in reduced accessibility of services for some residents, affect the quality of personalized support and widen the digital divide.

QUESTION 9

Will the use of AI create dependency on the tool and/or the technology provider?

Autonomy in AI plays out on two levels: the municipality's autonomy from its technology providers and employees' autonomy in doing their jobs. AI requires considerable financial and technical resources, which exposes municipalities to the risk of becoming dependent on private providers. Adoption of AI must therefore include guarantees of independence and respect for municipal principles, so as to preserve digital sovereignty.

Over-reliance on AI can also undermine employees' analytical skills, critical thinking and confidence in their decisions. A clear guidance framework is needed to safeguard their skills, strengthen their role in the decision-making process and ensure balanced complementarity of human expertise and technology tools.

6. Strubell, E., Ganesh, A., & McCallum, A. (2019).

KEY STEPS IN RESPONSIBLE AI DEPLOYMENT

Step 1 Define the need

- ☐ The problem is defined, measurable and documented.
- ☐ Non-technological alternatives (existing processes, training and tools) are evaluated.
- ☐ The use of AI is warranted by a favourable cost-benefit ratio.
- ☐ The expected benefits (efficiency, quality, savings, citizen/employee impacts) are demonstrated.
- ☐ Citizens are consulted from the start to ensure transparency, legitimacy and social acceptability.

Step 2 Measure the capacities

- ☐ Sufficient, reliable, representative and legally compliant data are available.
- ☐ The data are documented and updated, and sensitive data are protected.
- ☐ A multidisciplinary team (elected officials, managers, technical, legal and ethics experts) is coordinating the initiatives.
- ☐ External experts (e.g., IVADO, OBVIA, lawyers etc.) are engaged, as needed, to ensure regulatory compliance and increase the legitimacy of the projects.
- ☐ Implementation, maintenance and upgrading of the projects are ensured through sustainable funding and dedicated human resources, supported by a training plan.
- ☐ Interoperability with existing systems is verified to prevent technical debt.

Step 3 Anticipate the risks

- ☐ Bias detection and correction mechanisms are in place.
- ☐ Citizens' and employees' rights are protected and citizen redress procedures are in place.
- ☐ Conclusions made by AI are explainable and documented, and external audits are planned.
- ☐ Oversight is ensured by the municipal administration at every stage.

Step 4 Governance and implementation

- ☐ A steering committee is created, comprising elected officials, managers and citizens.
- ☐ An AI ethics officer is appointed and an impact- assessment process is implemented.
- ☐ A municipal policy or guideline document on the use of AI is adopted, defining a vision, ethical principles and governance frameworks.
- ☐ A low-risk pilot project is initiated, ideally in a so-called regulatory sandbox, to test the AI system in real-world conditions before broader deployment.
- ☐ Staff are trained and a structured citizen communications process is in place.
- ☐ Independent monitoring mechanisms (indicators, audits, citizen redress) are instituted.

Step 5 Continuous monitoring and evaluation

- ☐ Technical performance (reliability, bias, security, availability) is monitored.
- ☐ Organizational impact (efficiency, savings, staff satisfaction) is measured.
- ☐ Social acceptability (satisfaction, complaints, citizen feedback) is evaluated.
- ☐ Environmental impacts (energy consumption, carbon footprint) are taken into account.
- ☐ Compliance with ethical principles and elected officials' responsibility are periodically monitored.
- ☐ Learnings are shared to enhance municipalities' capacity for action.

Why is it essential for municipal decision-makers to follow these steps?

Responsible deployment of AI in municipal settings must proceed in five interdependent steps. It begins with definition of a real need and is followed by measurement of capacities in terms of data, competencies and resources. Anticipation of the technical, legal and social risks guarantees transparency and human oversight. Governance and implementation are supported by an inclusive framework, pilot projects and clear communication. Lastly, continuous monitoring and evaluation ensure improvements to services and measurement of the impacts. This checklist allows elected officials to confirm that AI will be of sustainable service to citizens.




Image 6 – Stages of AI Maturity in Municipal Settings (source: IVADO)

This diagram illustrates how successful deployment of AI depends on organizational maturity and establishment of a governance framework, which are essential conditions for preventing failure of projects and adverse impacts.

KEY PLAYERS IN THE MUNICIPAL AI ECOSYSTEM IN QUEBEC

There are a number of organizations with the strategic, scientific and technical expertise to support municipalities in Quebec. They make up a robust ecosystem that elected officials can rely on to ensure the deployment of responsible AI solutions.

IVADO : The leading consortium for AI research, training and knowledge mobilization in Quebec and the rest of Canada, IVADO brings together scientists, companies and public institutions to ensure the efficient and responsible deployment of AI in society and organizations. It provides specialized training, leads collaborative research projects and creates tools to support decision-making on AI.

Union des municipalités du Québec (UMQ): The UMQ represents all municipalities in the province, defends their interests before governments and provides training adapted to local contexts.

Universities and research centres: They offer diverse scientific expertise encompassing digital technology, cybersecurity, algorithm governance and public policy, and collaborate with municipalities on pilot projects.

Non-profit organizations and sector-based associations: They serve as intermediaries between the research community, citizens and municipal administrations, sharing best practices and fostering partnerships.

→ **Open North:** This NPO coaches cities in the areas of governance, data management, transparency and responsible adoption of digital tools.

→ **Mila (Quebec AI Institute):** A world-renowned research institute, Mila designs AI solutions with applications to urban challenges such as mobility, public health and infrastructures.

→ **OBVIA (International Observatory on the Societal Impacts of AI and Digital Technologies):** This interuniversity network mobilizes multidisciplinary researchers to analyze risk, apply ethics governance and raise awareness among elected officials and the public.

Scale AI: This innovation supercluster funded by the federal and Quebec governments invests in applied AI projects and promotes public-private partnerships.

Conseil de l'innovation du Québec (CIQ): The CIQ advises the government and other stakeholders on accelerating innovation. It develops strategic recommendations, shares best practices and fosters collaboration between researchers, businesses and public-sector organizations.

Prompt: This parapublic organization funds and structures collaborative research and development projects in AI and digital technology involving researchers and organizations.

IVÉO: An innovation network that supports cities and municipalities in experimenting with technological solutions aimed at improving quality of life and urban sustainability. IVÉO facilitates partnerships between startups, researchers, and public administrations to test concrete projects in real-life settings and accelerate the adoption of innovation within local communities.

TOOLKIT

AI literacy

AutodiAg: This free self-assessment tool, designed by IVADO, enables professionals to evaluate their knowledge and level of AI.

Data management

Open North offers free [resources](#) on its website, including practical guides on various aspects of data management, including:

- [conducting a data inventory project](#) to bring clarity to an organization's data assets;
- [developing data-sharing agreements](#) to encourage responsible use;
- [establishing strong quality standards](#) to ensure the reliability of data.

Risk management

The **Government of Canada** offers an [algorithmic impact-assessment tool](#), a questionnaire that determines the degree of risk associated with an automated decision system.

Public procurement

AI fact sheets can be used to assess the degree to which an AI solution aligns with a municipality's established principles or guidelines.

- The City of San José's [fact sheet for vendors of AI systems](#) includes an impact assessment questionnaire. Municipalities can incorporate similar fact sheets into their procurement processes and require suppliers to answer the questionnaire according to specific criteria.
- The United States Agency for International Development (USAID) also offers a [checklist](#) for managing ethical deployment of AI projects.

Public registers

Public registers are transparency tools enabling standardized documentation of the characteristics and impacts of AI systems deployed by a municipality. They help to strengthen citizen trust and public participation.

Examples of registers:

- [City of San José AI Register](#)
- [City of Helsinki AI Register](#)
- [City of Nantes AI Register](#)

Reference: Haataja, M., van de Fliert, L., & Rautio, P. (2020). Public AI registers: Realizing AI transparency and civic participation in government use of AI [White paper]. Open Research Amsterdam. <https://openresearch.amsterdam/en/page/73074/public-ai-registers>

Complementary resources

- [Guide pratique d'utilisation de l'IA générative pour les municipalités au Québec](#), published by OBVIA (in French).
- [AI & Cities: Risks, Applications and Governance \(white paper\)](#), published by UN-Habitat and Mila.
- [Series of templates and resources](#) developed by the GovAI Coalition and made available on the City of San José website, enabling public agencies to strengthen their AI governance (e.g., internal policy, procurement plan, incident response plan).
- [G7 Toolkit for Artificial Intelligence in the Public Sector](#) published by the OECD, which helps policy makers translate principles of responsible AI into actionable policies based on G7 best practices.

GLOSSARY⁷

Machine Learning: An AI method that recognizes trends and patterns in data and predicts outcomes, adjusting its performance and not simply following predefined rules.

Deep Learning: A sub-field of machine learning that relies on artificial neural networks to analyze very large volumes of data (images, sounds, text) and identify complex patterns.

Audit: An independent review designed to verify that an AI system complies with established laws, standards and ethical principles, to ensure transparency and accountability in municipal decision-making.

Regulatory sandbox: A temporary experimentation tool overseen by a public authority that enables testing of digital innovations, including AI, in a controlled environment prior to widespread deployment.

Algorithmic bias: Distortion in outputs from an AI system resulting from skewed training data or the design of the algorithm, potentially leading to unfair or discriminatory decisions.

Black box: Refers to an AI system with internal mechanisms that are opaque and difficult to explain, limiting transparency and assignment of responsibilities.

Chatbot: An AI-powered application that automatically interacts with users in natural language (written or spoken) to answer questions or provide a service.

Technical debt: The accumulated weakness resulting from a municipality's deferring updates to their digital systems, leading to costly delays and dependencies that hamper innovation.

Large Language Model (LLM): An AI model trained on very large volumes of text data, capable of understanding, generating and summarizing texts in natural language.

Hallucination: An answer produced by a generative AI tool that appears plausible, but is incorrect or invented.

Interoperability: The ability of a system or software application to function with other existing systems without any major modifications.

AI model: A program or algorithm that has been trained on a large volume of data and learned to recognize patterns, predict results or generate content.

Prompt: A set of instructions provided by a user to a generative AI tool for it to produce content (text, image, sound, video) or perform a task.

Neural networks: Models inspired by human brain function, made of up artificial neurons organized into layers, which learn to identify patterns in data.

Information silos: Data that are maintained by, for example, a particular department and not easily accessible by other departments or teams, which limits their sharing and cross-functional exploitation

Digital sovereignty: The ability of an organization or territory to maintain control over its own digital infrastructure, data and choices, without excessive dependence on private or foreign actors.

7. Some of these definitions are adapted from the MIT Sloan School of Management *Glossary of Terms: Generative AI Basics*.

REFERENCES

Arsenault, C. (2020). Using AI, Canadian city predicts who might become homeless. Reuters. <https://www.reuters.com/article/world/using-ai-canadian-city-predicts-who-might-become-homeless-idUSKBN27013X>

CIDOB. (2024). Ethical urban AI in practice: Policy mechanisms to establish local governance frameworks. CIDOB Monographs. <https://www.cidob.org/en/publications/ethical-urban-ai-practice-policy-mechanisms-establish-local-governance-frameworks>

Government of Canada (2025). Guide on the use of generative artificial intelligence. <https://www.canada.ca/en/government/system/digital-government/digital-government-innovations/responsible-use-ai/guide-use-generative-ai.html>

Hale, K. (2021). A.I. bias caused 80% of Black mortgage applicants to be denied. Forbes. <https://www.forbes.com/sites/korihale/2021/09/02/ai-bias-caused-80-of-black-mortgage-applicants-to-be-denied/>

Jacob, S. (2025). Quand l'algorithme décide. Presses de l'Université Laval. <https://www.pulaval.com/livres/quand-l-algorithme-decide>

MIT Sloan Teaching & Learning Technologies. (n.d.). Glossary of Terms: Generative AI Basics. MIT Sloan Management Review. <https://mitsloanedtech.mit.edu/ai/basics/glossary/>

Mushkani, R., & Koseki, S. (2025). Street Review: A Participatory AI-Based Framework for Assessing Streetscape Inclusivity. arXiv. <https://doi.org/10.48550/arXiv.2508.11708>

Organisation for Economic Co-operation and Development (OECD). (2024). G7 Toolkit for Artificial Intelligence in the Public Sector. https://www.oecd.org/en/publications/g7-toolkit-for-artificial-intelligence-in-the-public-sector_421c1244-en.html

Othengrafen, F., Sievers, L., & Reinecke, E. (2025). From vision to reality: The use of artificial intelligence in different urban planning phases. Urban Planning, 10(1), Article 6. Cogitatio Press. <https://doi.org/10.17645/up.i388>

Riot, É. (2025). IA et urbanisme: Une question de libertés et de choix civique? Urbanisme. <https://www.urbanisme.fr/debat/ia-et-urbanisme-une-question-de-libertes-et-de-choix-civique/>

Sénat. (2024). L'IA et l'avenir du service public: IA, territoires et proximité (Rapport thématique n° 342, Amel Gacquerre et Jean-Jacques Michau, rapporteurs). <https://www.senat.fr/rap/r24-342/r24-342-syn.pdf>

Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and policy considerations for deep learning in NLP. arXiv. <https://doi.org/10.48550/arxiv.1906.02243>

UN-Habitat & Mila – Quebec Artificial Intelligence Institute. (2022). AI and Cities: Risks, Applications and Governance. <https://unhabitat.org/ai-cities-risks-applications-and-governance>

Urban AI. (2023). 5 ways AI is transforming urban planning. Medium. <https://medium.com/urban-ai/5-ways-ai-is-transforming-urban-planning-234334a76705>

Wan, S., & Sieber, R. (2024). Artificial Intelligence (AI) adoption in Canadian municipalities: In-house development versus outsourcing. McGill University. <https://osf.io/fbxgm>



in collaboration with

