Canada School of Public Service
École de la fonction publique du Canada

AI Project Support Toolkit

Canada

Overview

Purpose

This document contains a suite of key considerations and tools to help teams in product and project planning.

It's designed for where:
1) experimentation with artificial intelligence (AI) is the goal; and where,
2) AI-enabled tools are under consideration as part of the solution set to a given problem.

Contents

AI basics

Artificial intelligence (AI)

A machine-based system that infers how to generate outputs such as predictions, content, recommendations, or decisions from the input it receives. Crucially, **AI systems adjust based on information processed, plus output feedback**

AI is also a **category of technologies**; a common explainer is “technology that performs tasks that would ordinarily require biological brainpower to accomplish, such as making sense of spoken language, learning behaviours, or solving problems”

Resources on AI basics

Project proponents should ensure that everyone involved in the project has a foundational understanding of AI.

- Executive summary of AI including GC-specific considerations: **Primer on AI**
- Introductory course: **Discover AI**

Simplified process flow

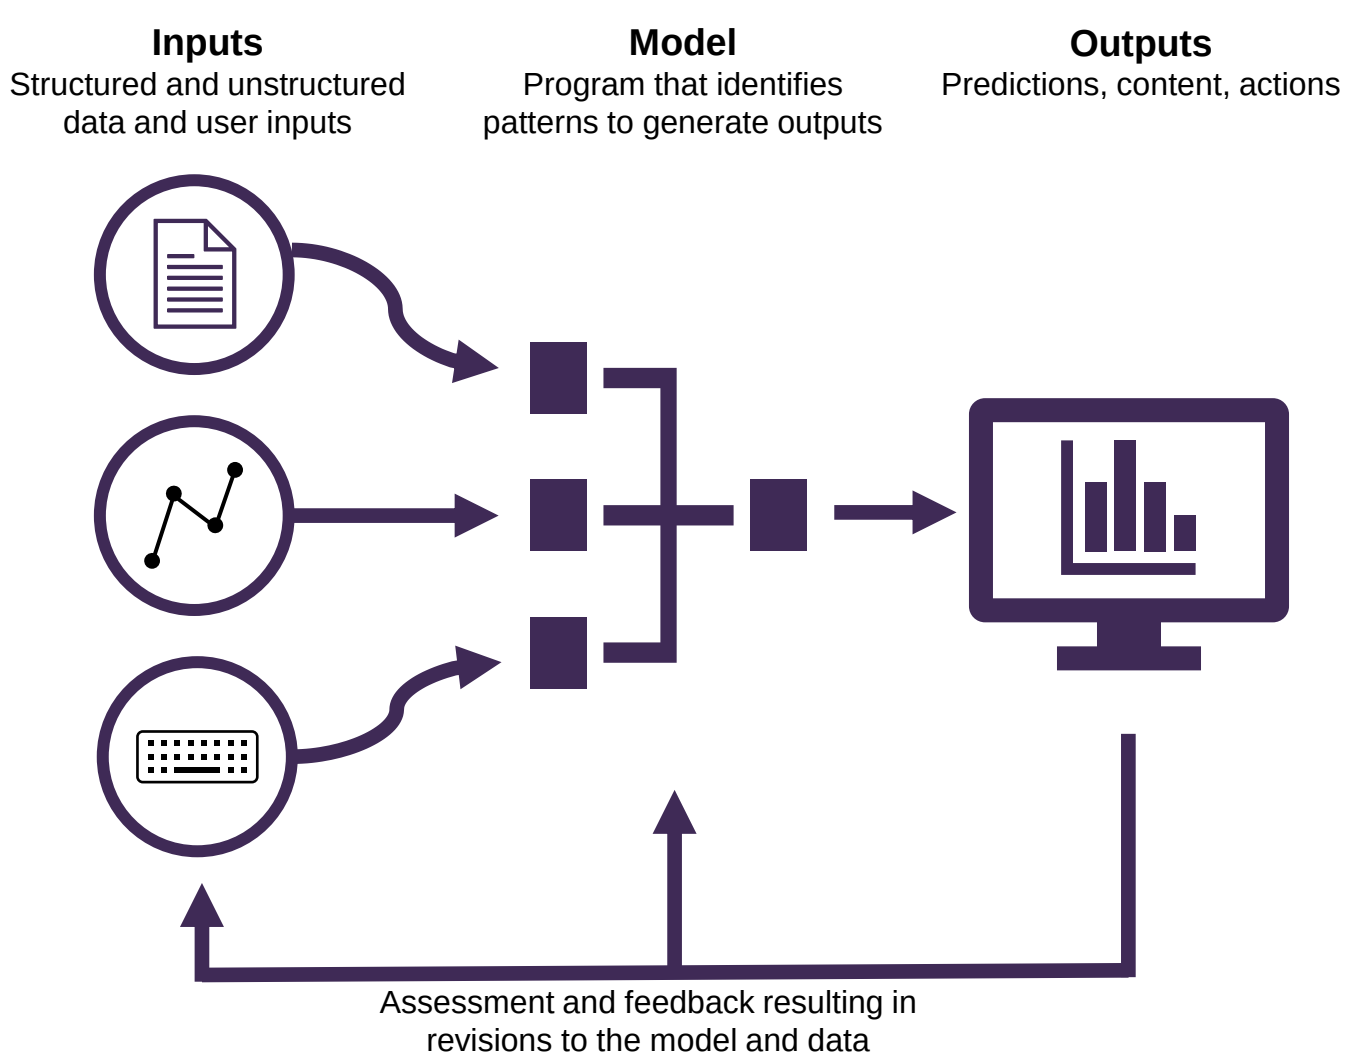

Use patterns and examples

1. What the GC is using, or exploring, AI for

Categories of AI use patterns

AI use or exploration in the Government of Canada tends to fall into one of these 5 **use patterns.**

Use patterns

1. **Information search and retrieval** including internet and intranet searches, and chatbot-like functionality on websites
2. **Information analysis and synthesis** such as making datasets more easily analyzable, document summaries, analysis of visual data for scientific research, or qualitative analysis of public engagement inputs
3. **Decision aids** such as predictive analytics, risk analysis, or supply chain or infrastructure maintenance analysis
4. **Creation or transformation of content,** both for general personal productivity or workflow support for correspondence or client support teams
5. **Automation of processes** including inputting data from one system into another, application triage, and more

Before starting projects, we **strongly recommend connecting with colleagues** in other departments who may have lessons learned or re-useable components via interdepartmental AI working groups as well as departmental and open communities of practice.

Overview of examples and experiments

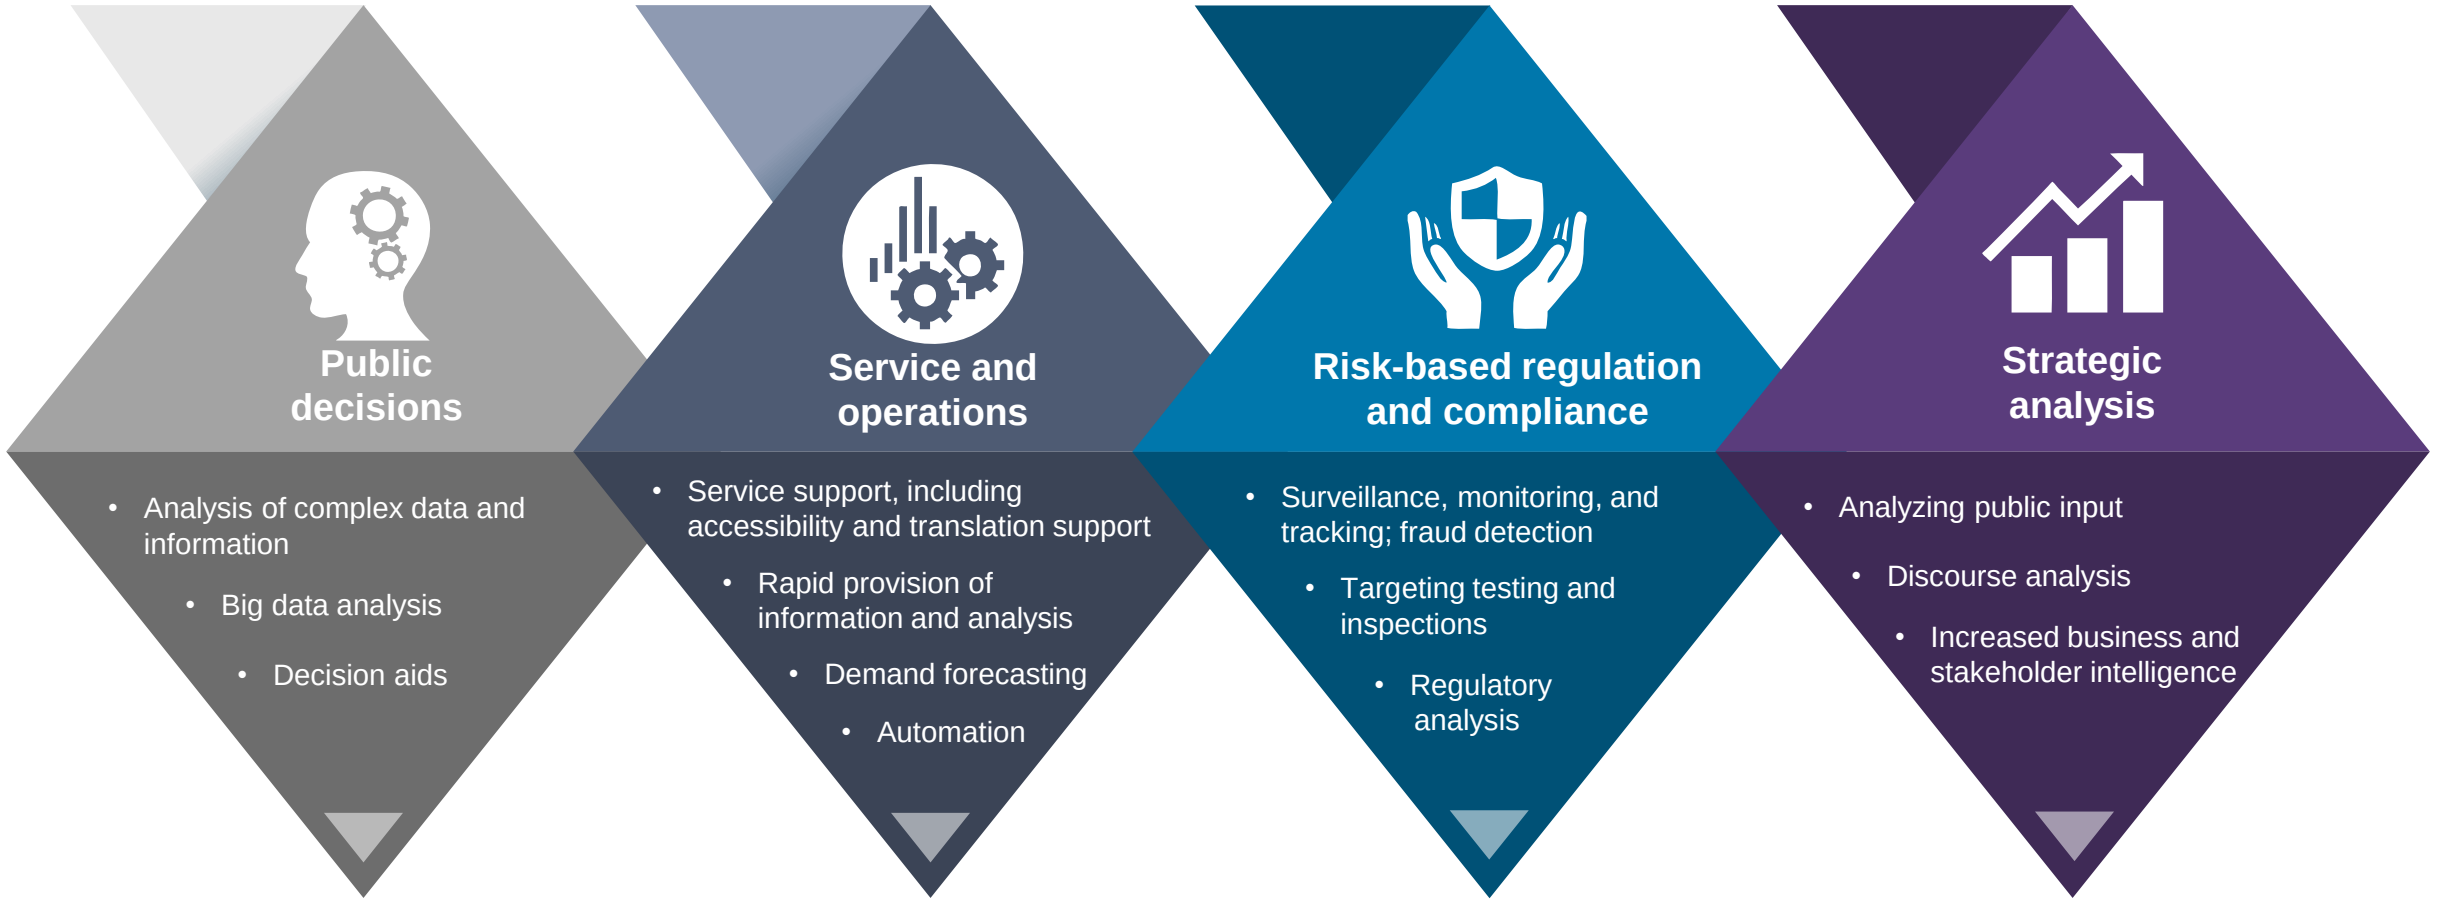

Government of Canada examples: computer vision

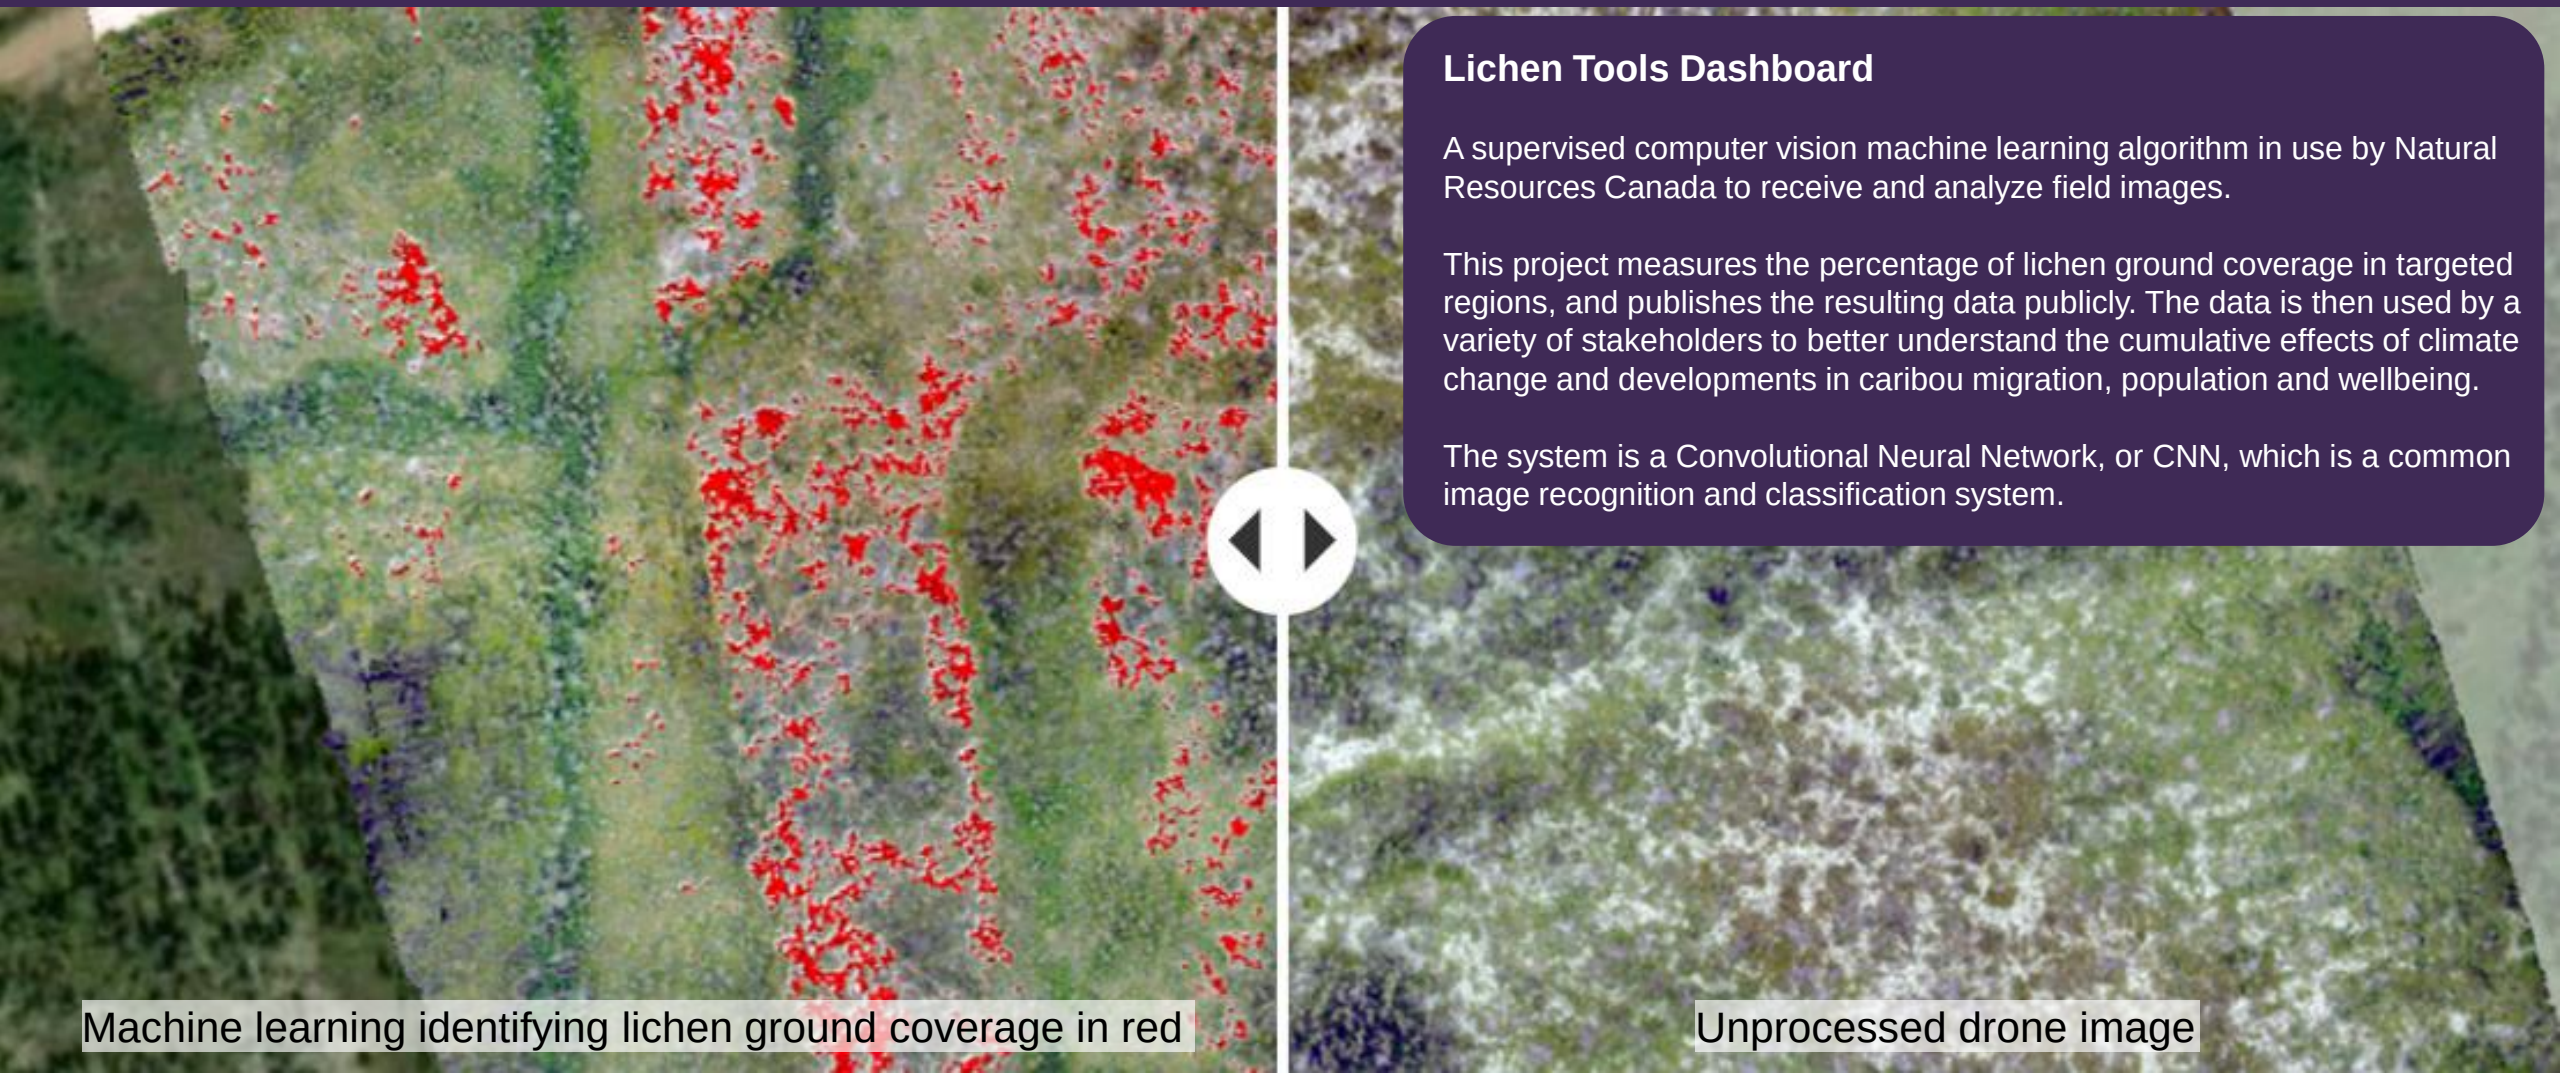

Lichen Tools Dashboard

A supervised computer vision machine learning algorithm in use by Natural Resources Canada to receive and analyze field images.

This project measures the percentage of lichen ground coverage in targeted regions, and publishes the resulting data publicly. The data is then used by a variety of stakeholders to better understand the cumulative effects of climate change and developments in caribou migration, population and wellbeing.

The system is a Convolutional Neural Network, or CNN, which is a common image recognition and classification system.

Government of Canada examples: chatbots

CRA Chatbot

CRA's Chatbot was first piloted in 2019 as a way to help people navigate tax information and inquiries.

Technology stack: Microsoft Azure Cloud, Language Understand (LUIS) tool, QnA Maker Natural Language Processing (NLP) engine, Power Virtual Agent (PVA), Dynamics Omnichannel

LUIS determines the intent of the user's text, while QnA returns the answer to the question, with business template responses. PVA and Dynamics connect the flow to live agent interventions.

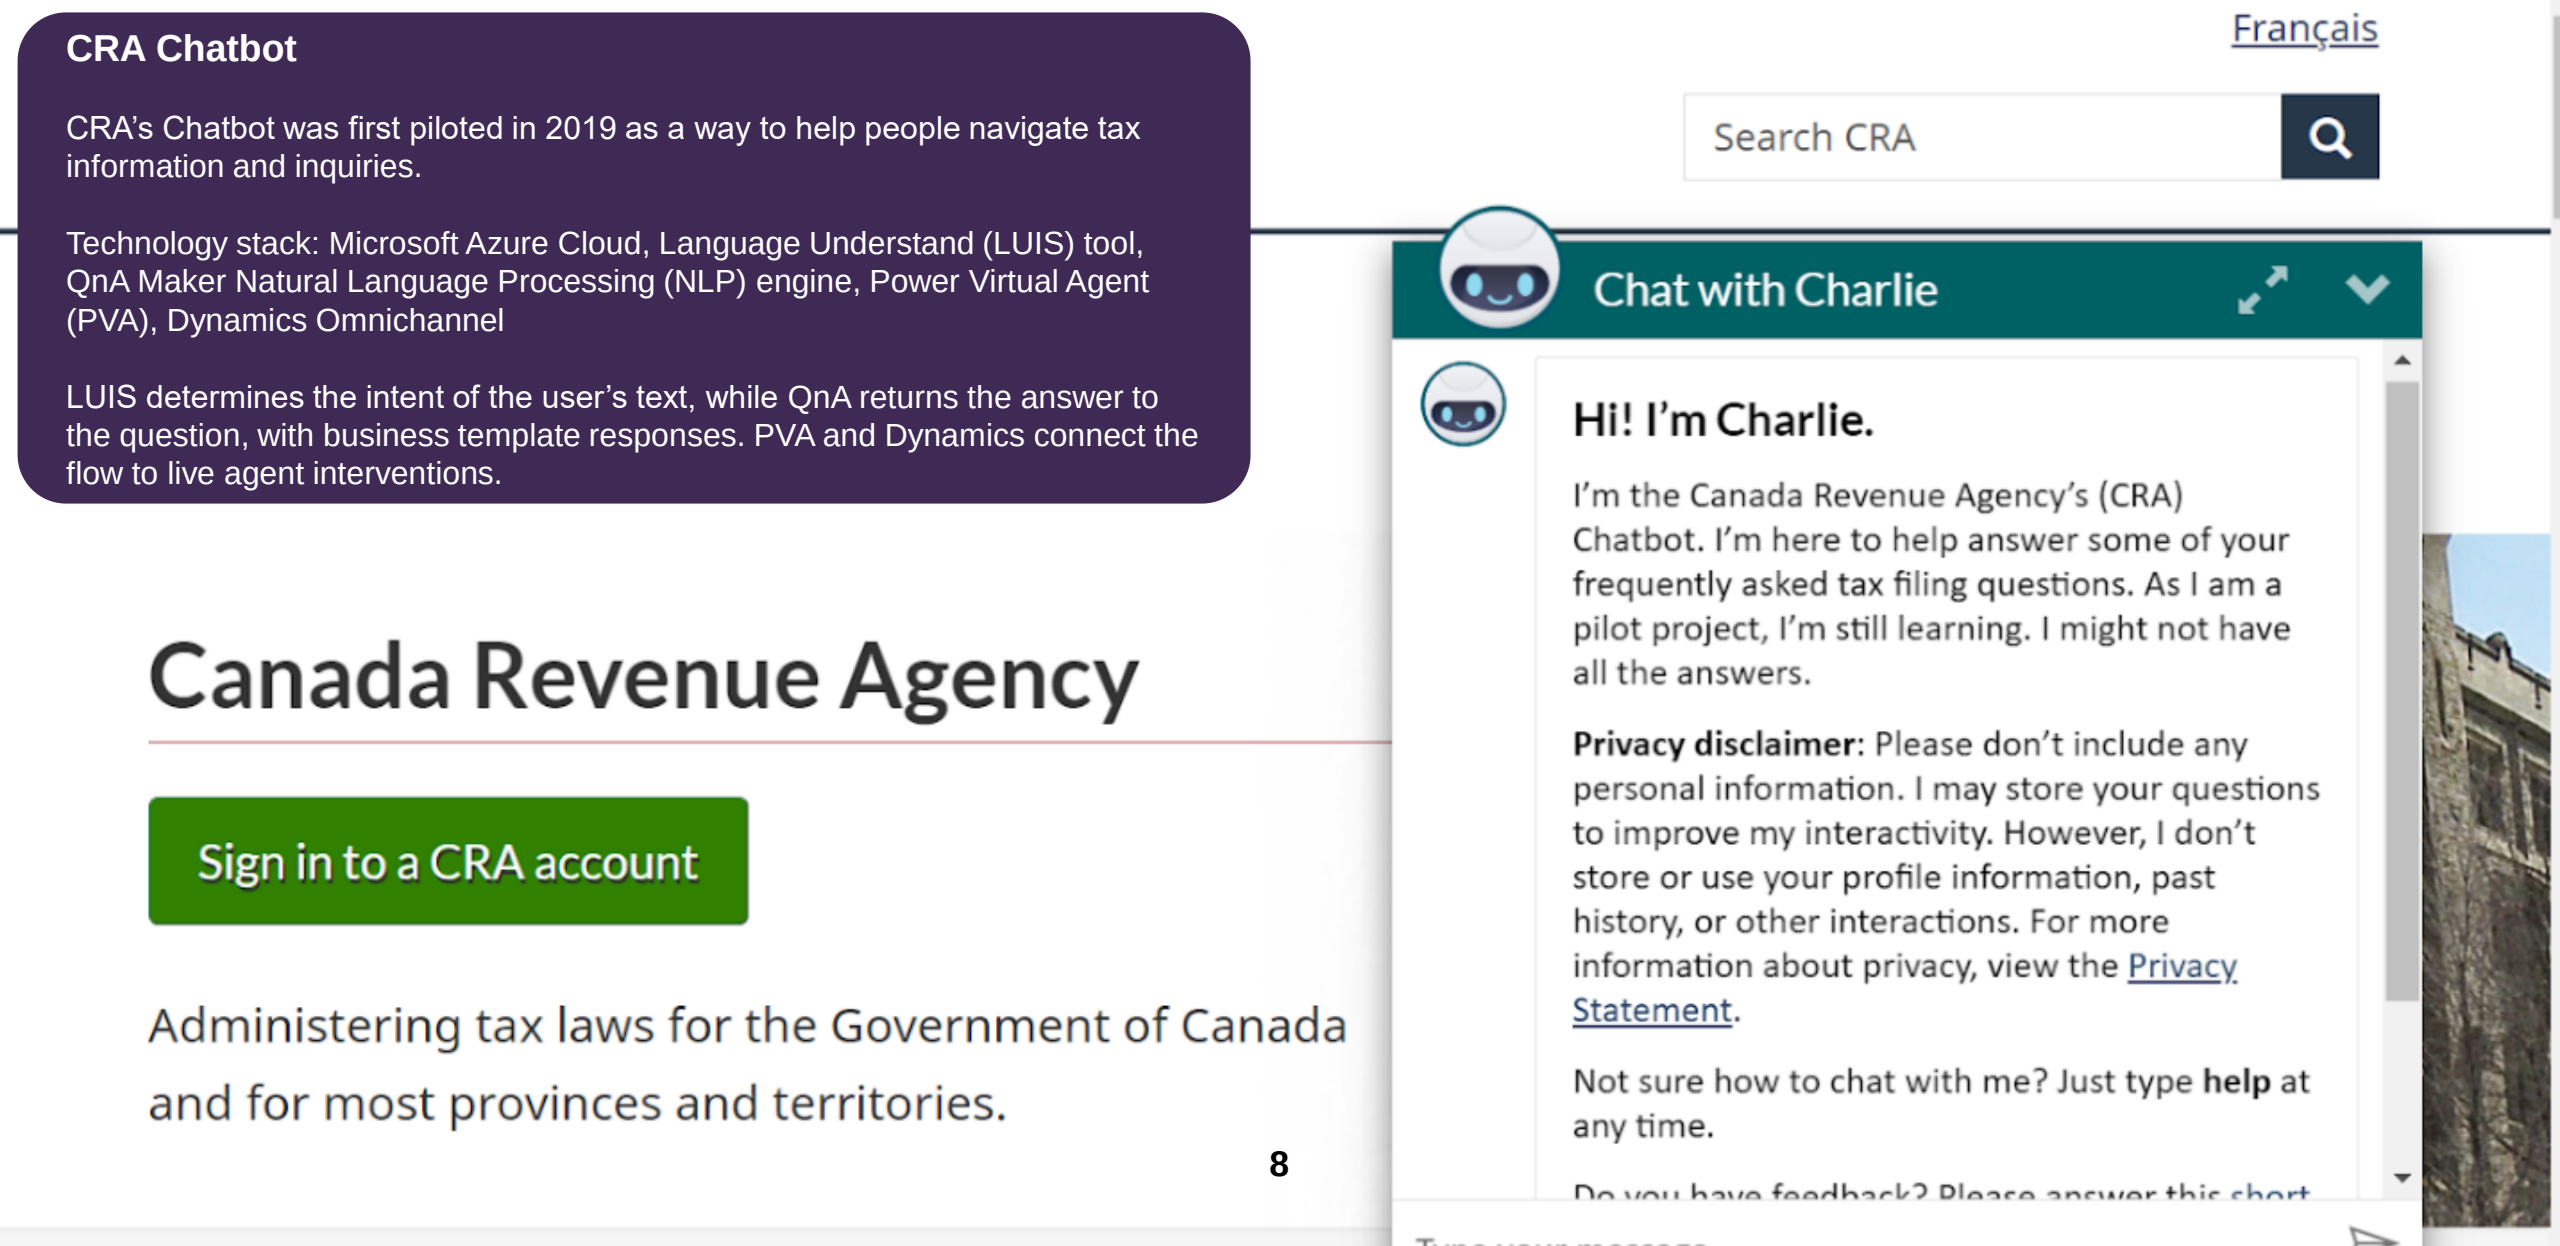

Government of Canada examples: document review

Document Detective

An Agriculture and Agri-Food Canada (AAFC) project was created to identify similar or duplicative documents in information management (IM) systems. It was created to help respond to Access to Information requests but can be used to find duplicates or related research and documentation.

It was built using open-source tools integrated securely into AAFC's infrastructure by a small team.

Where many AI applications rely on high-quality, well-structured data, in this case, AI tools are used to work with and make sense of unstructured data.

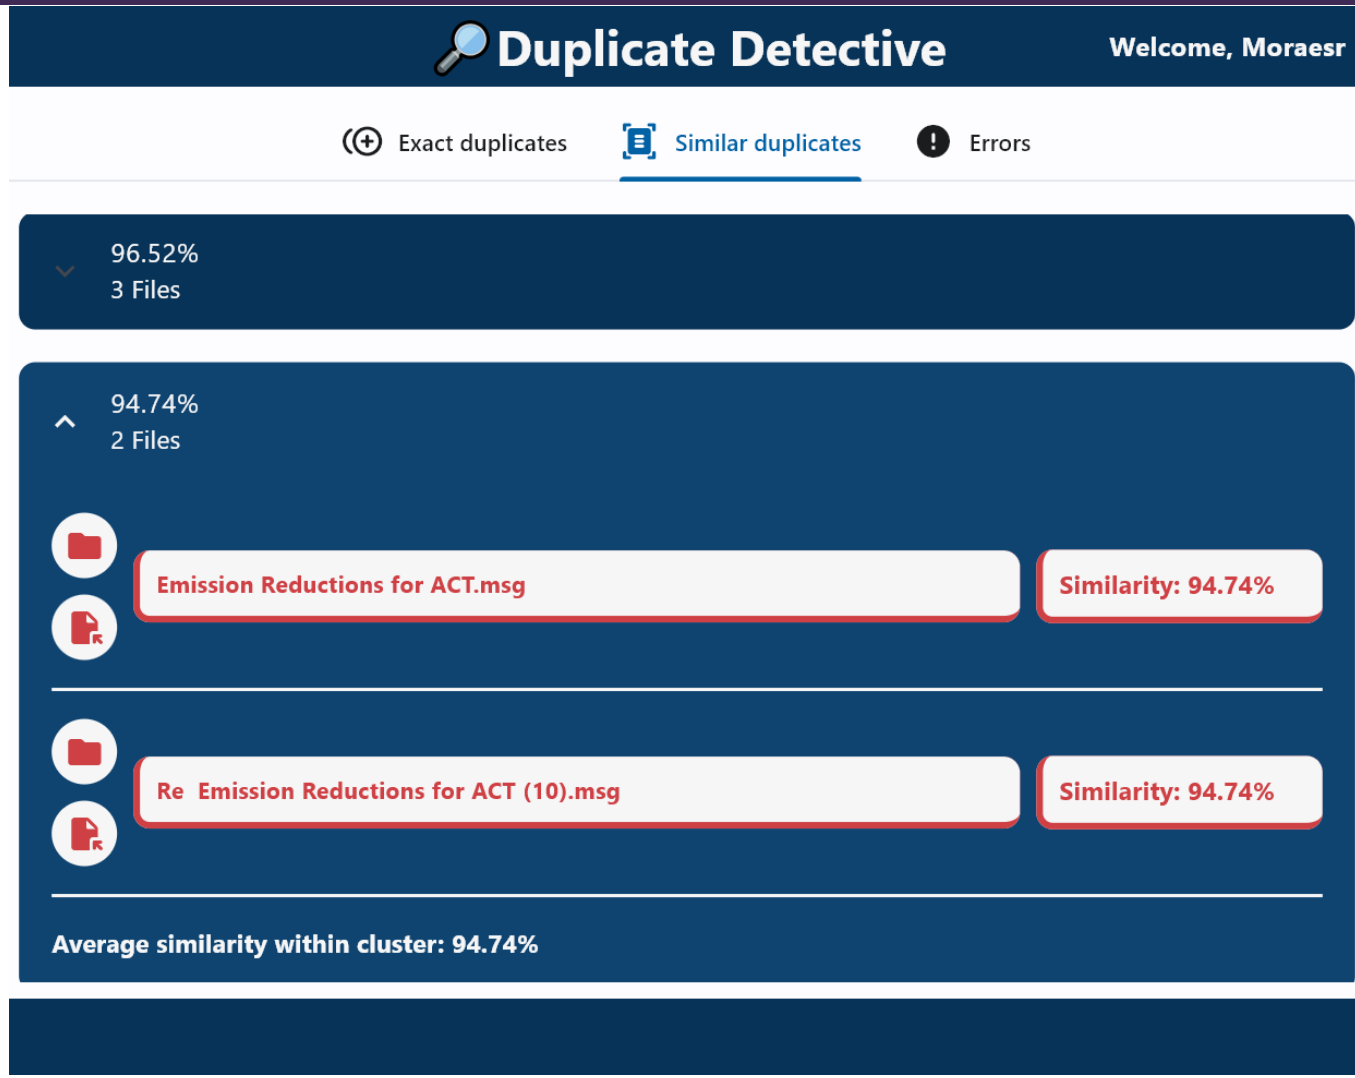

Evidence to date and key considerations

1. Insights and evidence from GC and global experimentation to date

Key considerations

Across the Government of Canada, there are many AI-enabled products and systems in use, ranging from mature implementations to ongoing experiments. From these experiences, we can draw a set of **key considerations** with promising practices and use patterns**.**

Key considerations

1. AI is a **broad category of technologies**, and tools range from mature to emerging

2. Most currently successful GC implementations are **using long-standing mature technologies** like image recognition, not emerging tools

3. AI use comes with **data, infrastructure, privacy, skills, and ethical considerations**; however, these considerations apply very differently across different AI approaches, and whether solutions leverage commoditized tools or are GC-built

4. Most – but not all – AI implementations require **a high level of data, data governance, and IM maturity**

5. AI systems will often be **one component of many in a technology stack** that adds up to a solution

6. It's crucial for the solution identified to flow from a strong design process rooted in a **clearly defined business problem;** this can, *and should*, often lead to alternative or complementary solutions, not AI

This resource contains tools for navigating each of these key considerations.

AI Capabilities in Mainstream Applications

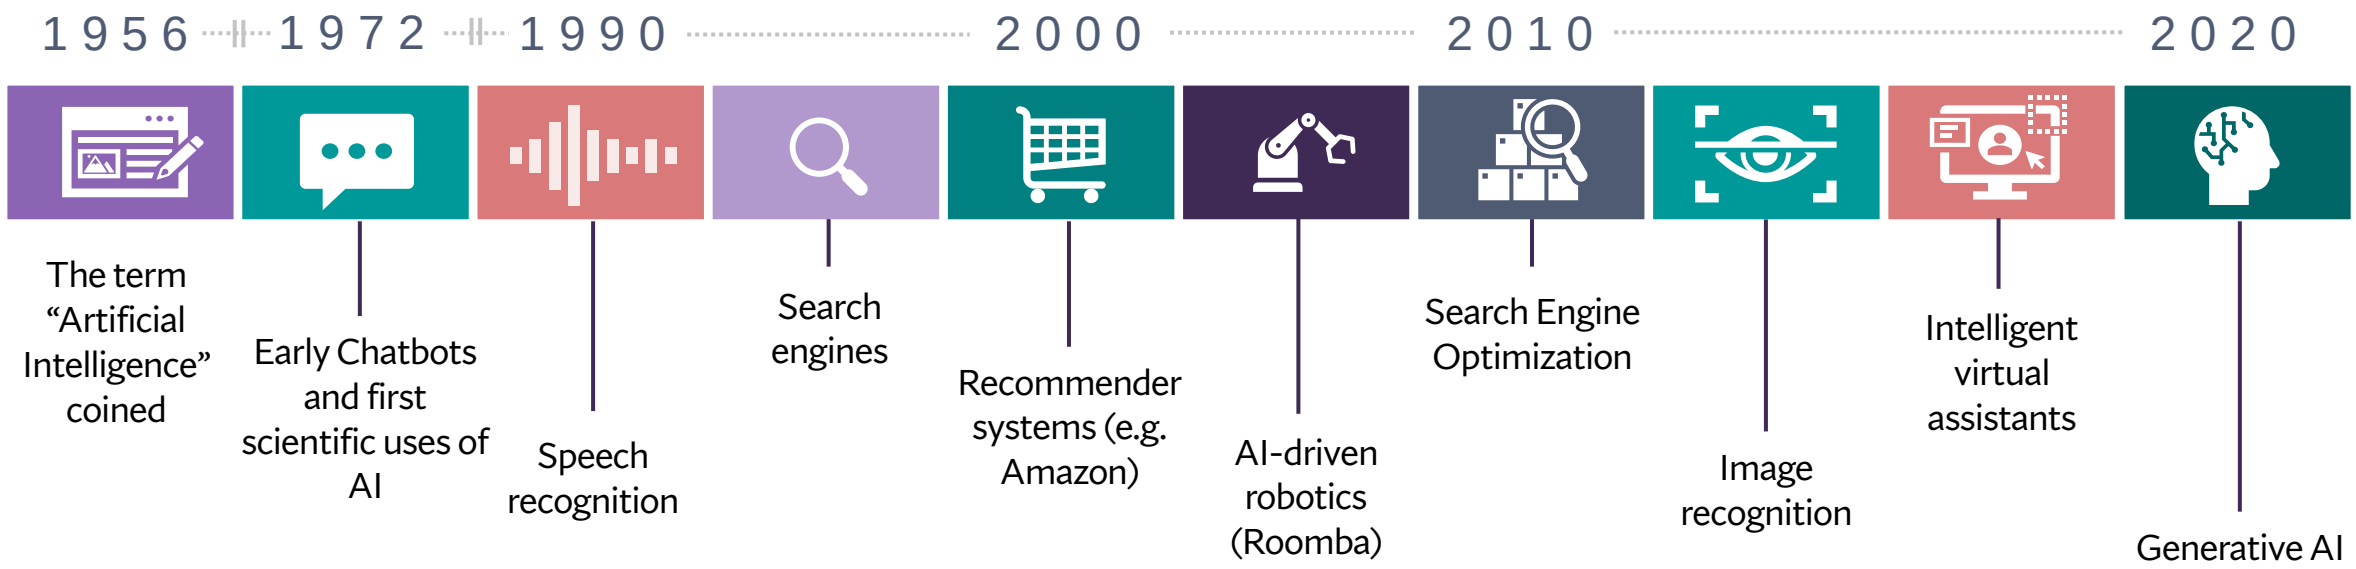

Generally speaking, successful AI implementations in the GC trend towards AI technologies and approaches that hit maturity and a level of reliability earlier on this timeline; while many people see potential for recent developments in AI, including Generative AI, research on implementation and ROI evidence is only emerging. *Note: dates above represent subjective maturity milestones.*

Global trends and evidence: overview

Researchers are increasingly studying public sector use of AI, but it's too early for a robust evidence base about return on investment.

Success and failure factor patterns are just starting to emerge.

"To date… the **relatively small number of public service AI applications** mostly include virtual assistants, such as chatbots or virtual agents, providing governmental information corresponding to queries." [1]

"[I]n certain early-adopter countries (e.g. the US or the UK), the use of AI in the public sector is increasing… there are many public organizations where productive applications remain rare. **AI in government is often at an experimental stage**, or traditional automation solutions are wrongly labelled 'AI'." [1]

"We have seen **little systematic evidence of what the potential public value of automated public services is** or what the actual implementation of these tools and services looks like in government agencies." [1]

"AI is a means to solve previously unsolved problems, not for solving problems you first have to create."[2]

"**Challenges to adopting AI in public organizations** stem from factors more prevalent in the public context: (i) a lack of technical staff to introduce and assess new technologies, (ii) the risk of potential erroneous use of AI (e.g. security risks, privacy concerns), (iii) the need to guarantee transparency in the context of AI, (iv) moral dilemmas such as when to use AI, and (v) ethical considerations, (e.g. non-discrimination of citizens)"[2]

1 Mergel, I., Dickinson, H., Stenvall, J., & Gasco, M. (2024). Implementing AI in the public sector. *Public Management Review*, 1-14.
2 Neumann, O., Guirguis, K., & Steiner, R. (2024). Exploring artificial intelligence adoption in public organizations: a comparative case study. Public Management Review, 26(1), 114-141.

Planning and managing successful AI projects

1. Starting guidance
2. Defining business problems and working through to prototyping or implementation
3. Defining success measures and planning for data collection and analysis
4. Procuring AI
5. Setting up internal capacity

Planning and project management

Planning and project management best practices are particularly important in the context of emerging technology, high uncertainty, low internal experience, and IT projects in general.

The approach for **small-scale experimentation** in sandbox environments will be very different from **large-scale formal projects** integrated into departmental IT infrastructure; what changes is the scale, governance, and formality of work to define the business problem and what success and failure look like.

Problem definition and planning

AI is a hype-filled space, there are many possible approaches (both those considered AI and alternatives), and AI is likely to only make up a component of an overall solution. In this context, grounding projects in **clear business objectives** is crucial.

Reference class forecasting

Successful projects and accurate time/cost/budget forecasts are best grounded in **comparable examples**[2]. While there are limited public sector examples, project proponents should learn from any available work to improve estimation and planning.

Building internal capacity and experience

Research suggests that “experience plays an important role in determining the maturity of dealing with AI technologies[1]” and is a well-documented success factor[2]. Strategic business value from AI is also associated with increasingly *internal* capacity and in-sourcing[3]. Organizations should start small and simple, and **build internal capacity and experience** over time**.**

1 Mergel, I., Dickinson, H., Stenvall, J., & Gasco, M. (2024). Implementing AI in the public sector. *Public Management Review*, 1-14.
2 Flyvbjerg, B., & Gardner, D. (2023). How big things get done: The surprising factors that determine the fate of every project, from home renovations to space exploration and everything in between. Signal.
3 Neumann, O., Guirguis, K., & Steiner, R. (2024). Exploring artificial intelligence adoption in public organizations: a comparative case study. Public Management Review, 26(1), 114-141.

Situating in the project context

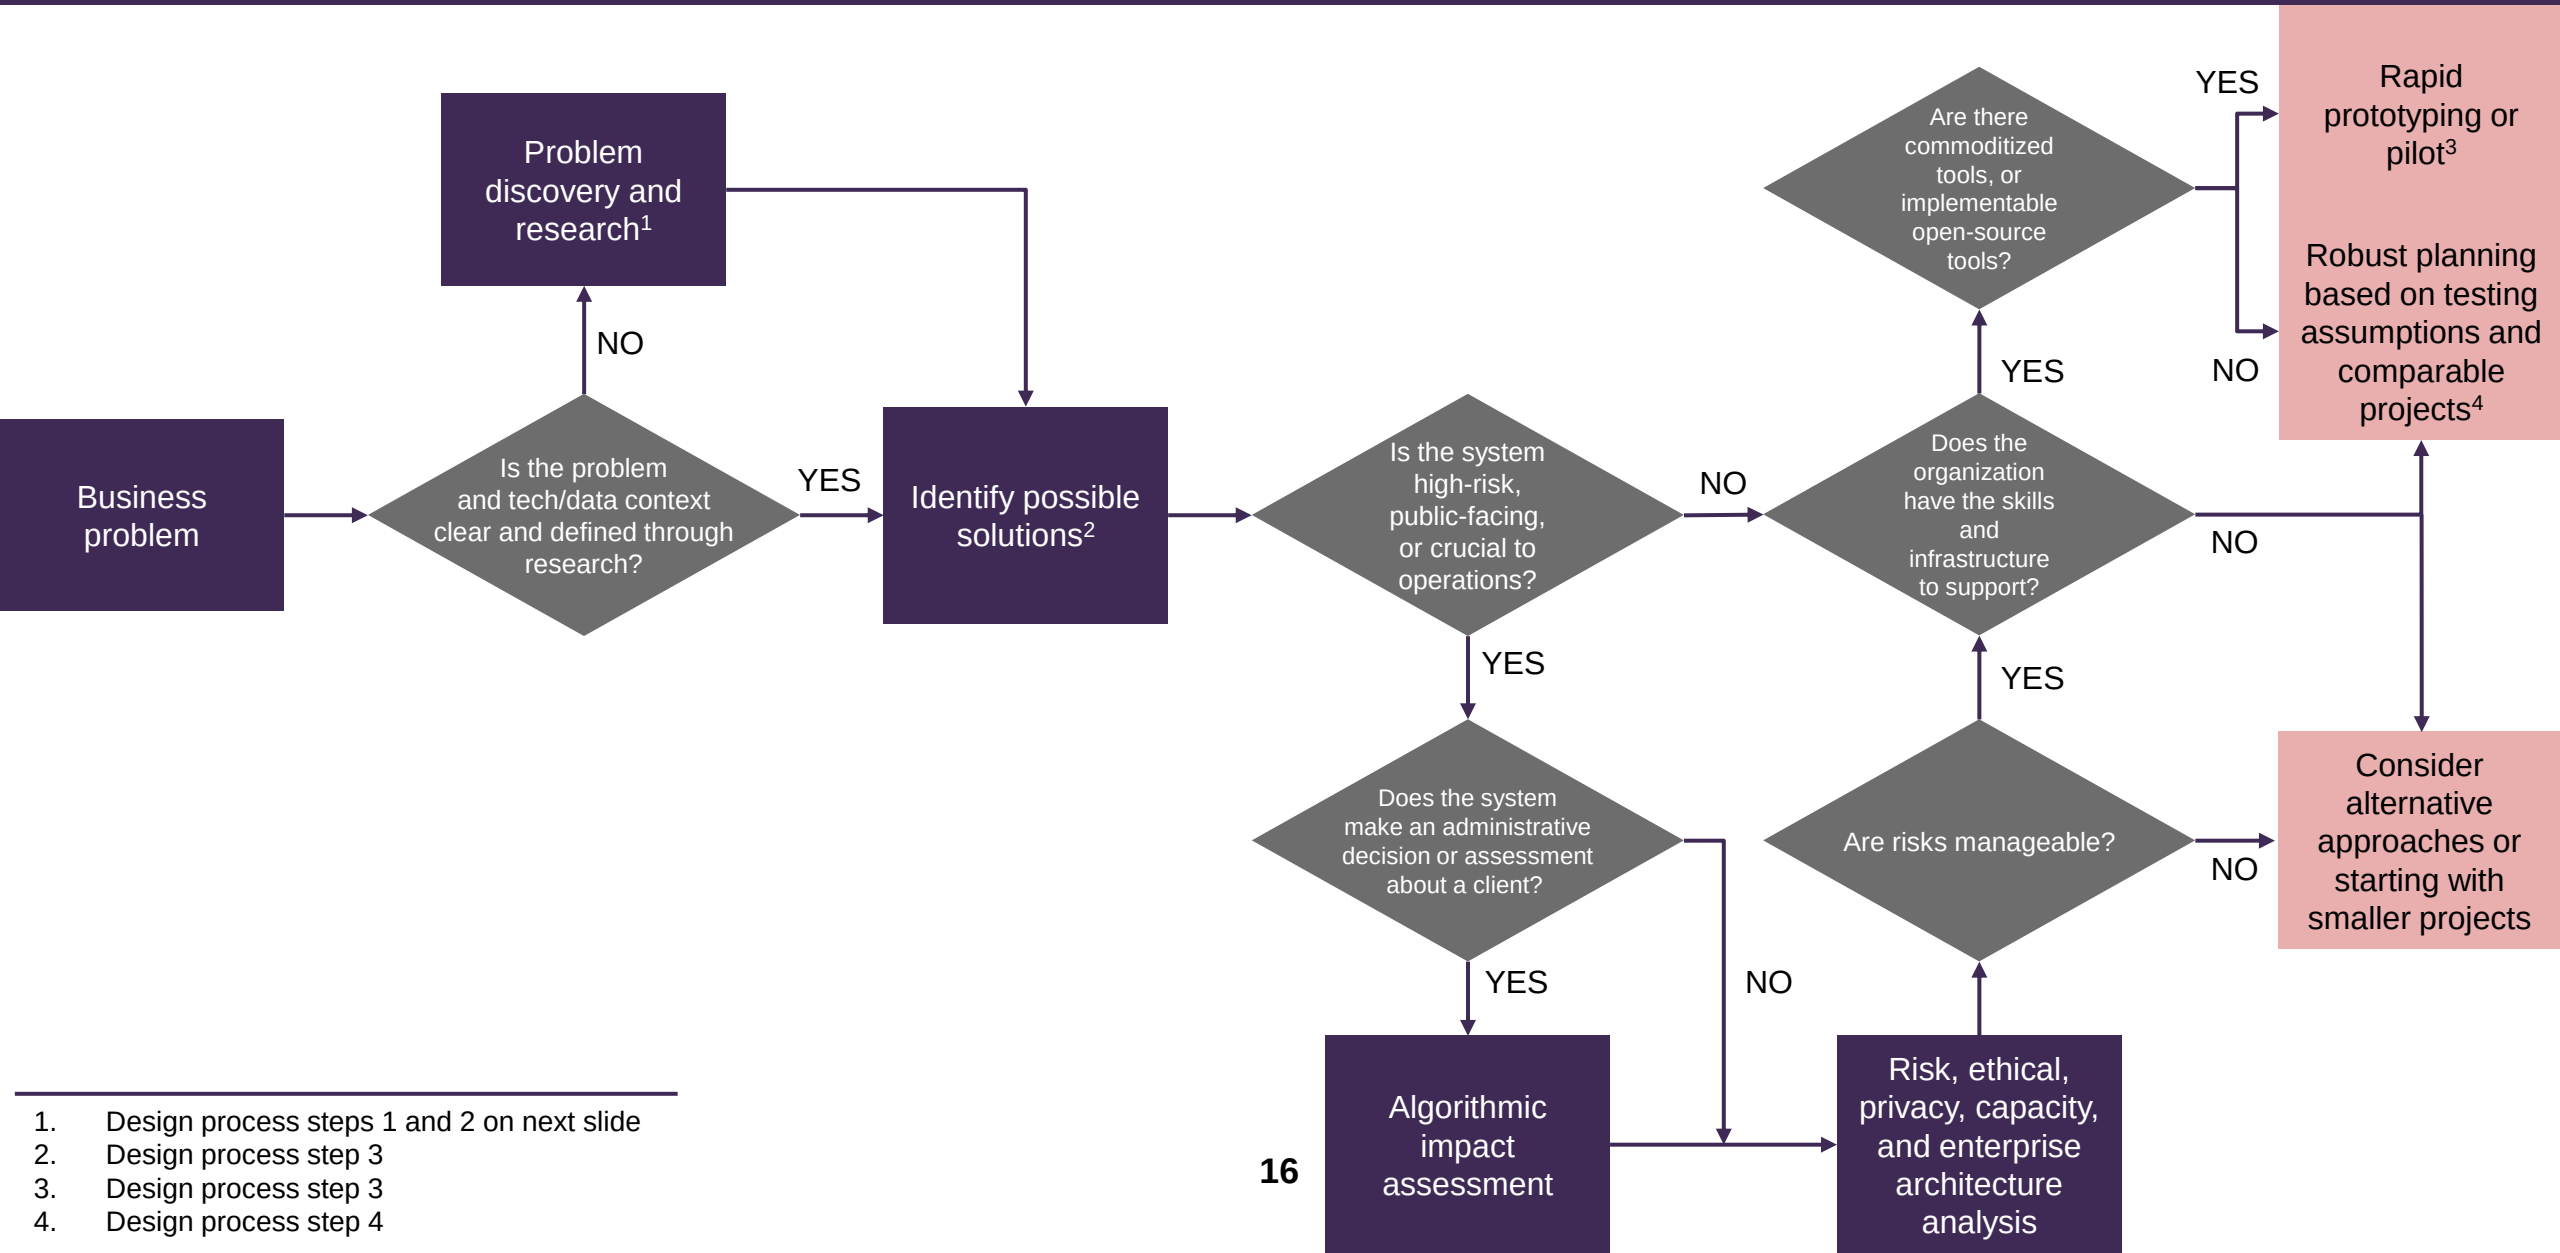

1. Design process steps 1 and 2 on next slide
2. Design process step 3
3. Design process step 3
4. Design process step 4

Design process: problem to prototype

A foundational success factor in project success is taking the time to understand and **define the problem** to be solved, including building a **shared understanding between project proponents and delivery supports**.

CSPS has developed foundational learning as well as process guides and templates to support teams working through the design process with an emphasis on digitally-enabled projects and products.

1 Discover

Steps to take

1. Meet with project proponents, stakeholders, and enablers, and colleagues who've done similar projects
2. Take stock of existing knowledge, assets, and understanding of the problem with a **mission model canvas**
3. Map supports and stakeholders: **stakeholder map**
4. Make a business analysis and research plan: **research planning checklist**
5. Conduct **design research**

2 Define

1. Analyze and synthesize the collected data to identify patterns and key insights: **qualitative data synthesis**
2. **Compose a problem statement** as part of a wider **problem definition** process

3 Develop

1. Generate multiple potential solutions through **ideation and conceptualization**
2. Evaluate ideas based on feasibility, desirability and viability
3. Develop prototypes to test biggest assumptions and manage risk: **prototyping and testing**

4 Deliver

1. Refine plan to measure the effectiveness and benefits
2. Develop detailed implementation plan with timelines, resources, and roles and responsibilities
3. Launch implementation

Design process: discover and define

Key questions in project design phases:

1 Discover

In this phase, you will engage your stakeholders to think broadly, ask questions to understand the problem and conduct research on the people impacted by it. This phase sets the rest of the process into motion.

Example Questions

- What kind of issues or constraints are preventing your organization from achieving its goals or objectives?
- What kind of challenges do the stakeholders and users face when interacting with our program or service?
- What do we currently know about them? Do we have any data?
- What do we think is the root cause of these issues?
- What has been the impact of these issues on different aspects of your organization, such as productivity, morale, and satisfaction?
- Who are the key stakeholders we should involve in problem-solving?

2 Define

Clearly defining the problem provides a specific direction for the project and helps maintain focus. Without a well-defined problem, efforts can become scattered and inefficient, leading to wasted resources and potentially failing to achieve meaningful outcomes

Example Questions

- What specific pain points do we need to address?
- Do different stakeholders experience different pain points? Can we define the problem from multiple stakeholder perspectives through point-of-view statements?
- Which pain points are most critical and should be prioritized based on their impact on our objectives?
- How can we refine or narrow down the problem statement to make it more actionable and specific?

Addressing these questions during the initial project consideration is crucial to understanding the potential requirements for the solution. As the project progresses, continue to engage stakeholders on problem definition to refine your approach and stay aligned.

Design process: ideation and conceptualization

Once we have a thorough understanding of the problem, we can begin to explore potential solutions and their specific features. This approach ensures that the solutions we consider are tailored to effectively address the identified issues.

In identifying potential solutions, consider the following questions:

- What are the key features that the solution must have to address the identified pain points effectively?
- Which solution features are critical to address the pain points?
- What technical capabilities must the solution possess to address the problem effectively?

And finally, will AI bring value by being part of our solution?

- What specific benefits do we expect AI to bring to this solution?
- How does AI enhance the effectiveness or efficiency of the solution compared to non-AI alternatives?
- Is the problem complex enough to require AI?

These introductory questions are designed to assist you during the project preparation phase. Once you are ready to start the project, review more comprehensive resources on the Discover and Define Design Phases, along with associated tools and techniques.

Design process: develop and deliver

Key questions in project design phases:

3 Develop

In this phase, you will work with your IT department or a contractor, as appropriate. Ensure that you explore various solutions to the defined program through creative ideation, prototyping, and testing with your tech team and if possible, users and stakeholders.

Example Questions

- Are we exploring a wide enough range of solutions and ideas?
- What emerging technology or trends should we consider during the solution development? Why is AI the right technology?
- How do these solutions address the needs and pain points of our users?
- What kinds of capabilities and features do we need to develop?
- Are we prototyping and testing ideas effectively to ensure they meet user expectations?
- What criteria are we using to evaluate and refine the potential solutions?
- Do we have the necessary resources (budget, talent, tools) to develop these solutions effectively?
- Have we thought through any unintended consequences of using the solution?

4 Deliver

This is the time to refine, test, and implement the solution, followed by delivery and training. Though it is the last step in the design, the solution will need to be tested, monitored, and adjusted, when corrections or updates become necessary.

Example Questions

- Is the solution fully developed and tested to meet quality standards?
- How will the solution be integrated into existing systems and workflows?
- How will we train end-users and stakeholders to ensure effective adoption of the solution?
- How will we monitor the solution's performance post-launch to ensure it meets expectations?
- Is the solution scalable to meet future growth or demand?
- Are we prepared to make adjustments or improvements quickly if necessary?
- How will we address any issues users report after the solution is delivered?
- What is our plan for long-term updates to keep the solution relevant?

AI experimentation

Many departments are experimenting with AI solutions. To set up a test with clear success and failure measures, teams can answer a series of questions about test parameters based on the goals of the experiment and how they will assess results. This will support our **ability to share successes and learning across the GC**.

Participants/population

What are you targeting?
- Public servants
- Program users
- Process and workflow
- Other

Comparator/control

To what will you compare your intervention?
- Baseline data
- ‘Business-as-usual’ practices
- Another intervention
- Available benchmarks

Intervention

What are you trying to do exactly?
- Fill a knowledge gap
- Test at small scale before roll-out
- Test assumptions and what works
- Establish a cause-effect chain
- Building buy-in by piloting
- Other

Outcomes

What will you measure?
- Behaviours
- Performance
- Responses
- Other

The following slides are templates for teams to fill out to define an experiment in this space.

Experiment template: part 1

THE BASICS

Describe the experiment

What's the objective?

What's your hypothesis?

What are your biggest assumptions and uncertainties?

EXPERIMENT FEATURES

Who uses the system that will be impacted?

How long will the experiment run?

What needs to be in place for this to happen?

EXPECTED RESULTS

What will happen immediately?

What indicators will demonstrate that result?

Are there any long-term effects that may reveal later?

What indicators will demonstrate that result?

Experiment template: part 2

ADDITIONAL CONSIDERATIONS	MARKER STATEMENTS
Potential ethical concerns from an evaluation data lens	We think [X change or action] will result in [Y result]
Security and privacy concerns	We think [Y result] will be [above/below] [this threshold, percentage, or number] [display a certain characteristic or behaviour]
Potential data and indicator issues	[Y result] will happen because [Z cause-effect chain]

Organizational readiness

1. Skills, data, and infrastructure supporting AI products and projects
2. Projects and success factors at different maturity levels
3. Global success factors for digitally-enabled projects in the Government of Canada

Skills, data, and infrastructure: overview

For every AI use case, there's a lot of work and tooling below the surface. Ultimately, **every AI project depends on data**.

APPLICATION
"Using AI to do [X]"

MODELS
Product of AI approaches, statistical methods, weighting, and training data

DATA
Often data that's cleaned, normalized, unbiased, well-structured, and often massive and costly to analyze. Cleaning and preparing data for analysis and processing can be 50%+ of the work and cost involved

SKILLS
Data science, mathematics, statistics, research and data collection, and programming

INFRASTRUCTURE AND TOOLS
Substantial computing power and/or access to Cloud-based tools, open source models, and specialized software

BUSINESS QUESTIONS
Ultimately, the foundation of AI work is a well-defined business problem to solve

Skills, data, and infrastructure supporting AI projects

The exact data, skills, models, and infrastructure required will depend greatly on the nature of the project, in the context of the many approaches possible within the overall field of AI.

Later sections will cover the variety of AI approaches at a high level, and a representative list of potential skills is captured in the annex. However, while it's important to have a general sense early in project scoping of the required skills, data, models, tools, infrastructure, those questions can only be answered after additional project planning, problem definition, and process and technology mapping.

All of that has to take place within the context of the organization. The following section covers broad considerations for **organizations' readiness** for AI projects, within the context of elements that make projects more or less complex.

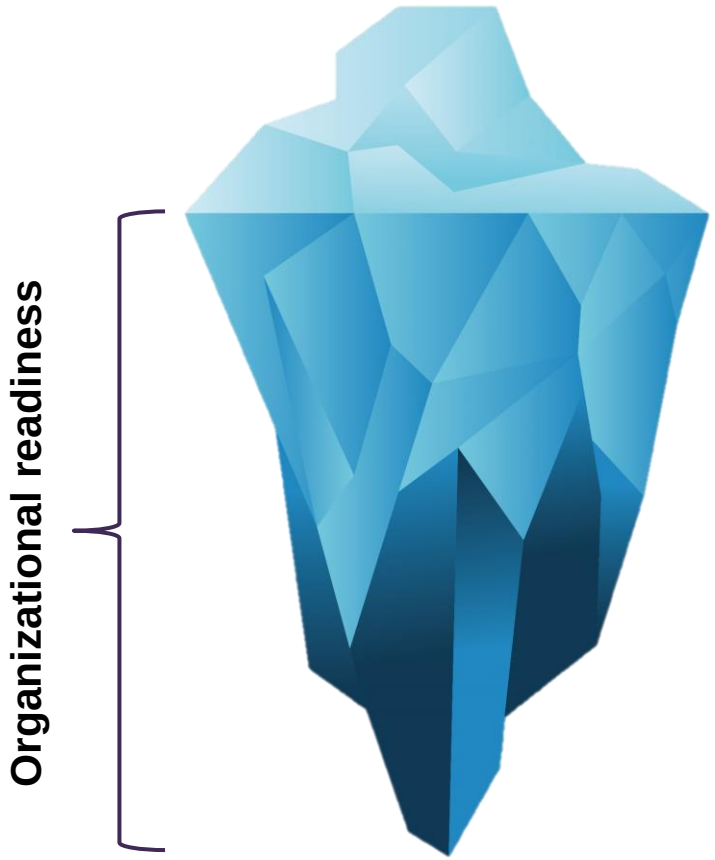

Developing organizational maturity

While case studies and comparative analysis are only emerging, there are patterns that suggest which projects are most suitable to organizations as they advance up a maturity curve for AI. Organizations can be *assessing, intermediate,* or *determined.* At each stage, there are typical activities, success factors, and project types.

Assessing / early maturity	Intermediate / developing	Determined / mature
AI project types seen at this level of maturity • Less complex • Simple Chatbots • Search and knowledge retrieval	• Increasingly complex • Process automation • Components of client-facing services	• Highest complexity • System-level workflows and integrations • Process automation and optimization • Client-facing services
Success factors • Motivated individual innovative team members	• **Managing complexity** • Business integration: exec support, change management, process change, and resourcing • Developing experience and capacity; insourcing core capability • Stakeholder readiness	• Successful transition to integrating technology into management, resourcing, strategy, and processes
Key activities • How to launch projects • Attracting staff and skills • Developing partnerships • Experimentation	• Achieving strategic and business value • Coherent tech and process integration • Technological readiness	• Scaling or smart sunsetting

Neumann, O., Guirguis, K., & Steiner, R. (2024). Exploring artificial intelligence adoption in public organizations: a comparative case study. Public Management Review, 26(1), 114-141.

Success factors for digitally-enabled GC initiatives

The available international and GC evidence on AI projects aligns with research conducted with GC executives leading digital initiatives, which revealed a consistent pattern of 13 success factors for *any* digitally-enabled project. (The guidance in this primer is designed to provide specific actions to support many of these success factors.)

Purpose

A **focus on the end user and/or ultimate outcome** aligns the organization in the right direction…

…allowing teams to define and communicate a clear **problem definition**…

…which enables a clear picture of what **success and failure** look like…

…so progress can be measured against **expected value and timelines**…

…which creates **priority and urgency**

People

The **right team** is key: a mix of complementary specialist skills with delivery experience, ideally cross-functional with policy or program team

It's necessary to cultivate an enabling environment through **change management**

Appropriate **risk tolerance** is crucial through experimentation and learning

It's important to have **growth mindsets** - for both initiative proponents and team members - including curiosity, continuous learning, and embracing learning curves

Policy & Process

Technology is often how people interact with **policy and program intent**; successful projects include an ability to engage with the policy and program layer, not just the digital delivery layer on top

Solutions were not always a straight line: **creativity, air cover, and workarounds** were often required to navigate internal processes, approvals, and constraints

Partnership

The **business/IT partnership** is vital**:** understanding each other and having a shared vision and goals

Senior executive support helps clear challenges, secure enablers, and align with business objectives

Project complexity and viability

1. Elements influencing AI project complexity
2. Viability assessment

Context elements: complexity

There are a few major elements for AI projects, and projects will take different shapes depending on the combination. This can also be used for a **rough complexity ballpark:** in planning the project and solution, consider how many project elements are on the *more complex* side of the spectrum, in columns 2 and 3.

Element	1 Simpler Fewer dependencies	2 It depends	3 More complex More dependencies
Software / model		Proprietary Open source	
Data context	Data-agnostic	Dependent on readily available data	Dependent on well-organized data
Delivery capacity		In-house capacity Vendor or partner	
Integration	Experimentation		Operations / large-scale project
Project scope	Standalone tool		Integrated into back-end systems and data Integrated into public-facing systems and data
Usage pattern	Information search and retrieval Analysis and synthesis	Automation	Decision aid Decision automation
Org maturity		Determined / mature organization	Assessing / early maturity
User	Individual use		Workflow / systems
Use case maturity		Mature / proven	Emerging / untested
Solution maturity		Mature / proven	Emerging / untested

Project viability assessment: overview

For a more detailed AI project viability assessment, Shared Services Canada has developed an **AI Viability Model**.

Benefits of assessment

Assessing viability is essential for several reasons:

- **Resource allocation:** determine whether the project aligns with available resources, including budget, expertise, and infrastructure.
- **Risk management:** identify potential risks and challenges associated with the project, such as technical complexity, regulatory compliance, and data privacy concerns.
- **Strategic alignment:** evaluate how the project aligns with organizational goals, priorities, and long-term strategic objectives.
- **ROI potential:** assess the expected return on investment (ROI) and business value generated by the AI project, considering both tangible and intangible benefits.

Key considerations

Criteria for assessing the viability of an AI project include:

- **Business case:** evaluate the business case for the project, including problem definition, target outcomes, and potential benefits.
- **Data availability and quality:** assess the availability, quality, and suitability of data required for training and deploying AI models.
- **Technical feasibility:** consider the technical feasibility of the project, including algorithm selection, model complexity, and integration with existing systems.
- **Ethical and regulatory compliance:** ensure compliance with ethical guidelines, regulatory requirements, and data protection laws governing AI projects.
- **Stakeholder engagement:** understanding and responding to the needs, interests, and concerns of key stakeholders inside and outside the government and protecting public trust

Project viability assessment: how to use

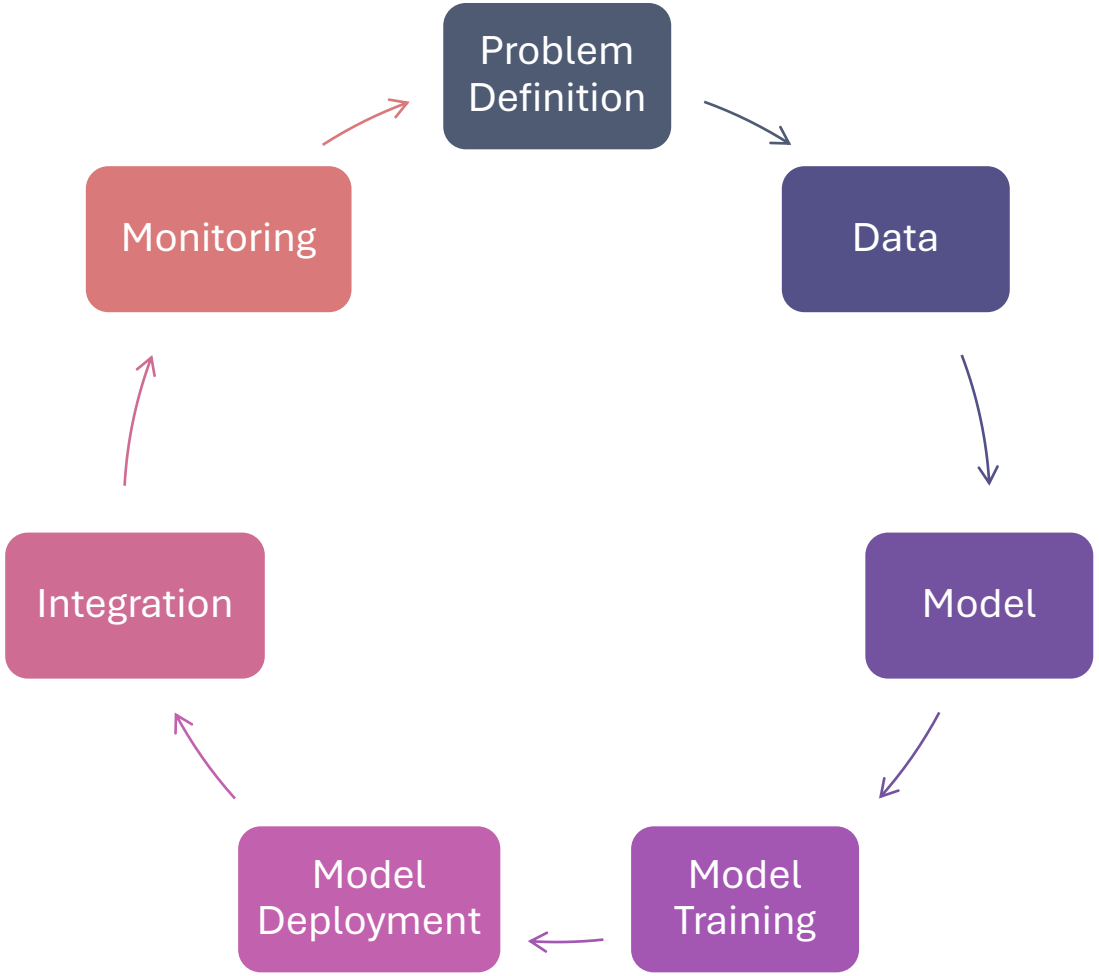

SSC's research, and project support experience, suggests that **successful AI projects** go through these steps starting with a **well-defined problem**.

Building on the previous project complexity assessment, we **recommend working through each of the questions in the viability assessment**:

Viability assessment questions

Note: in the SSC model, different questions have different success and risk weights; these are not reflected in the question list.

For large-scale projects, there is also the GC Project Risk and Complexity Assessment, and departmental governance or enterprise governance may require concept cases, gating documentation, and more.

Ethical, risk, and governance considerations

1. What makes AI different
2. Ethical considerations
3. Multiple key considerations in context
4. Risk and governance in context
5. Policy guidance

What makes AI different

AI is a category of tools with unique features

1. **AI changes itself as it operates**, building rules over time based on inputs, models, outputs, and feedback

2. AI could be integrated into a vast number of systems and workflows; **AI is a possible feature of many given system components**

Accordingly, governance based on predictable input-output relationships and accountability will be challenged by AI as it becomes increasingly commoditized and ubiquitous

It's crucial for the Government of Canada to unpack and understand ethical, risk, and governance considerations **in context**, such that we can apply them intelligently to new and evolving situations

Ethical considerations

This list includes only some of the many ethical considerations. Others may include impacts on job markets, environmental impacts, and questions about humans' relationship with technology. For an expanded exploration, take the CSPS course Ethical Considerations in AI.

Bias and fairness

Bias in AI means unfair decisions or showing skewed outputs. The GC must make sure that AI tools treat everyone fairly and without discrimination.

Bias can be a product of the **algorithm/model** or the **training or input data.**

Much historical data was collected in a context of systemic racism and discrimination and can result in a biased AI output; ongoing data collection also requires scrutiny and governance.

Transparency and accountability

Transparency is required around how AI systems operate and, if they support decision-making, how data was analyzed to produce outputs. This includes **openness, clarity, traceability, and explainability** of the AI system.

Actors – individuals or organizations – leveraging AI systems **may not feel, or take, responsibility and accountability** for actions, outputs, or decisions made by the system.

Privacy, security, and governance

AI systems process massive amounts of data, and AI tools are often Cloud-based or based on externally created resources, code, or models.

Any **data processing of personal or sensitive data needs to be governed and protected.**

Data provenance and copyright

Data provenance refers to the origins, ownership, collection, and reliability of source data. Organizations using data may need to track and document the **sources, transformations, and usage of data throughout data lifecycles.**

Many datasets powering Generative AI in particular have massive, opaque data sources that likely **include personal information, or direct or derivative copyrighted works**

Manipulation and deception

This category involves the ethical considerations related to the use of AI in generating and disseminating misleading of false information.

External **entities may use or propagate disinformation, misinformation, or deepfakes using AI.**

Generative **AI may create content that includes false information** (known as hallucinations).

Key considerations in context

People often talk about, or express concerns about, AI as a single entity. Key considerations - **risk, privacy, data, skills, ethics,** and **more** - will apply differently to different types and applications of AI. Some examples:

Skills	**Generative AI for brainstorming policy approaches** • Basic digital literacy • Critical thinking and analysis • No mathematics, statistics, or programming skills	**Using a machine learning model for policy decision support** • Policy domain expertise • Advanced data skills • Inclusive design, intersectionality, and GBA+ • Advanced understanding of statistical methods
Ethical considerations	**AI workflow support for HR officials, prompting review of documentation from similar past cases** • No difference from understanding and training required for all officials on unconscious bias, values and ethics, GBA+, and inclusion	**Automated assessment of applications against historical data on successful candidates** • Potential for bias and discrimination • Requirement for ensuring safe and responsible use, beyond just transparency and accountability
Data maturity and infrastructure	**Comparing documents for similarity** • Requires tool integrated into IM system behind firewall and access controls (e.g., Pro A, Pro B) • Low level of data maturity and quality required	**Chatbot providing advice to external clients** • High level of maturity for data, IM, metadata, and information architecture • Extensive testing and auditing
Privacy and information security	**Individual use of online Generative AI tools** • Requires broad literacy among users to avoid employees copy and pasting anything beyond transactional data into cloud-based tools	**Generative AI linked to departmental IM systems** • Requires IM and access control maturity; e.g., could return results from HR documents accidently stored in open folders • Information flows between model provider and GC systems must be closely analyzed
Reliability and trustworthiness	**Image recognition, workflow prompts, or Generative AI drafts** • Need to consider significance, and rates, of inaccuracies • Requires broad literacy among users, as well as domain expertise, to evaluate outputs and recommendations	**Decision support, Chatbots, or data synthesis tools** • Range of tool maturities, based on range of data and model maturities • Requires extensive testing and auditing

Evaluating considerations in context

1 Basics

- Ensure a baseline understanding **of how AI systems work**, including different types of AI systems
- Ensure a baseline understanding of the patterns of **ethical considerations**
- Look for **similar GC projects** underway or complete
- Conduct project scoping including **problem definition**

2 System and data context

- Understand **where AI systems fit** within the overall business process and technology architecture, and what **function they're performing**
- Understand the **data sources and flows** connected to the AI system
- Identification of **potential solutions**, including alternative methods of solving the problem

3 Initial and additional evaluations

- **Evaluate overall context**: risk, ethics, data, skills, privacy, reliability, and infrastructure considerations
  - *Note: the policy/service/program domain experts and IT/data experts should work together here*
- Based on the above, you may need also need:
  - Algorithmic impact assessment
  - Inclusive design process
  - Evaluation for potential bias
  - Data and IM maturity assessment
  - Skills gap analysis
  - End user training
  - Enterprise architecture review

4 Engagement

- As with any program, policy, product, or policy design, **engage expert colleagues**: leads or centers of expertise in privacy, legal, data, IT, IM, enterprise architecture, finance, procurement, values and ethics, inclusion, design, innovation, or AI

AI branches

AI is an expansive category of technologies and approaches. The common feature is that AI systems **change how they process in response to what they're processing and feedback on the outputs.** But within that, there are many variations:

- **Research fields** (e.g., computer vision, machine learning)
- **Applications and uses** (e.g., client service, fraud detection)
- **Statistical methods** used
- **Broad categories (**e.g., Narrow AI versus Artificial General Intelligence)

Accordingly, it's recommended to focus on business needs and problems and allow potential solutions to emerge through research and testing.

There are, however, general governance considerations for AI categories.

Perifanis, Nick & Kitsios, Fotis. (2023). Investigating the Influence of Artificial Intelligence on Business Value in the Digital Era of Strategy: A Literature Review. Information. 14. 85. 10.3390/info14020085.

AI, automation, and analysis technologies: types

Rather than a single AI technology, a range of different techniques and approaches are used to solve different problems. There are also technologies that don't necessarily use AI, but are useful to consider as complements or alternatives.

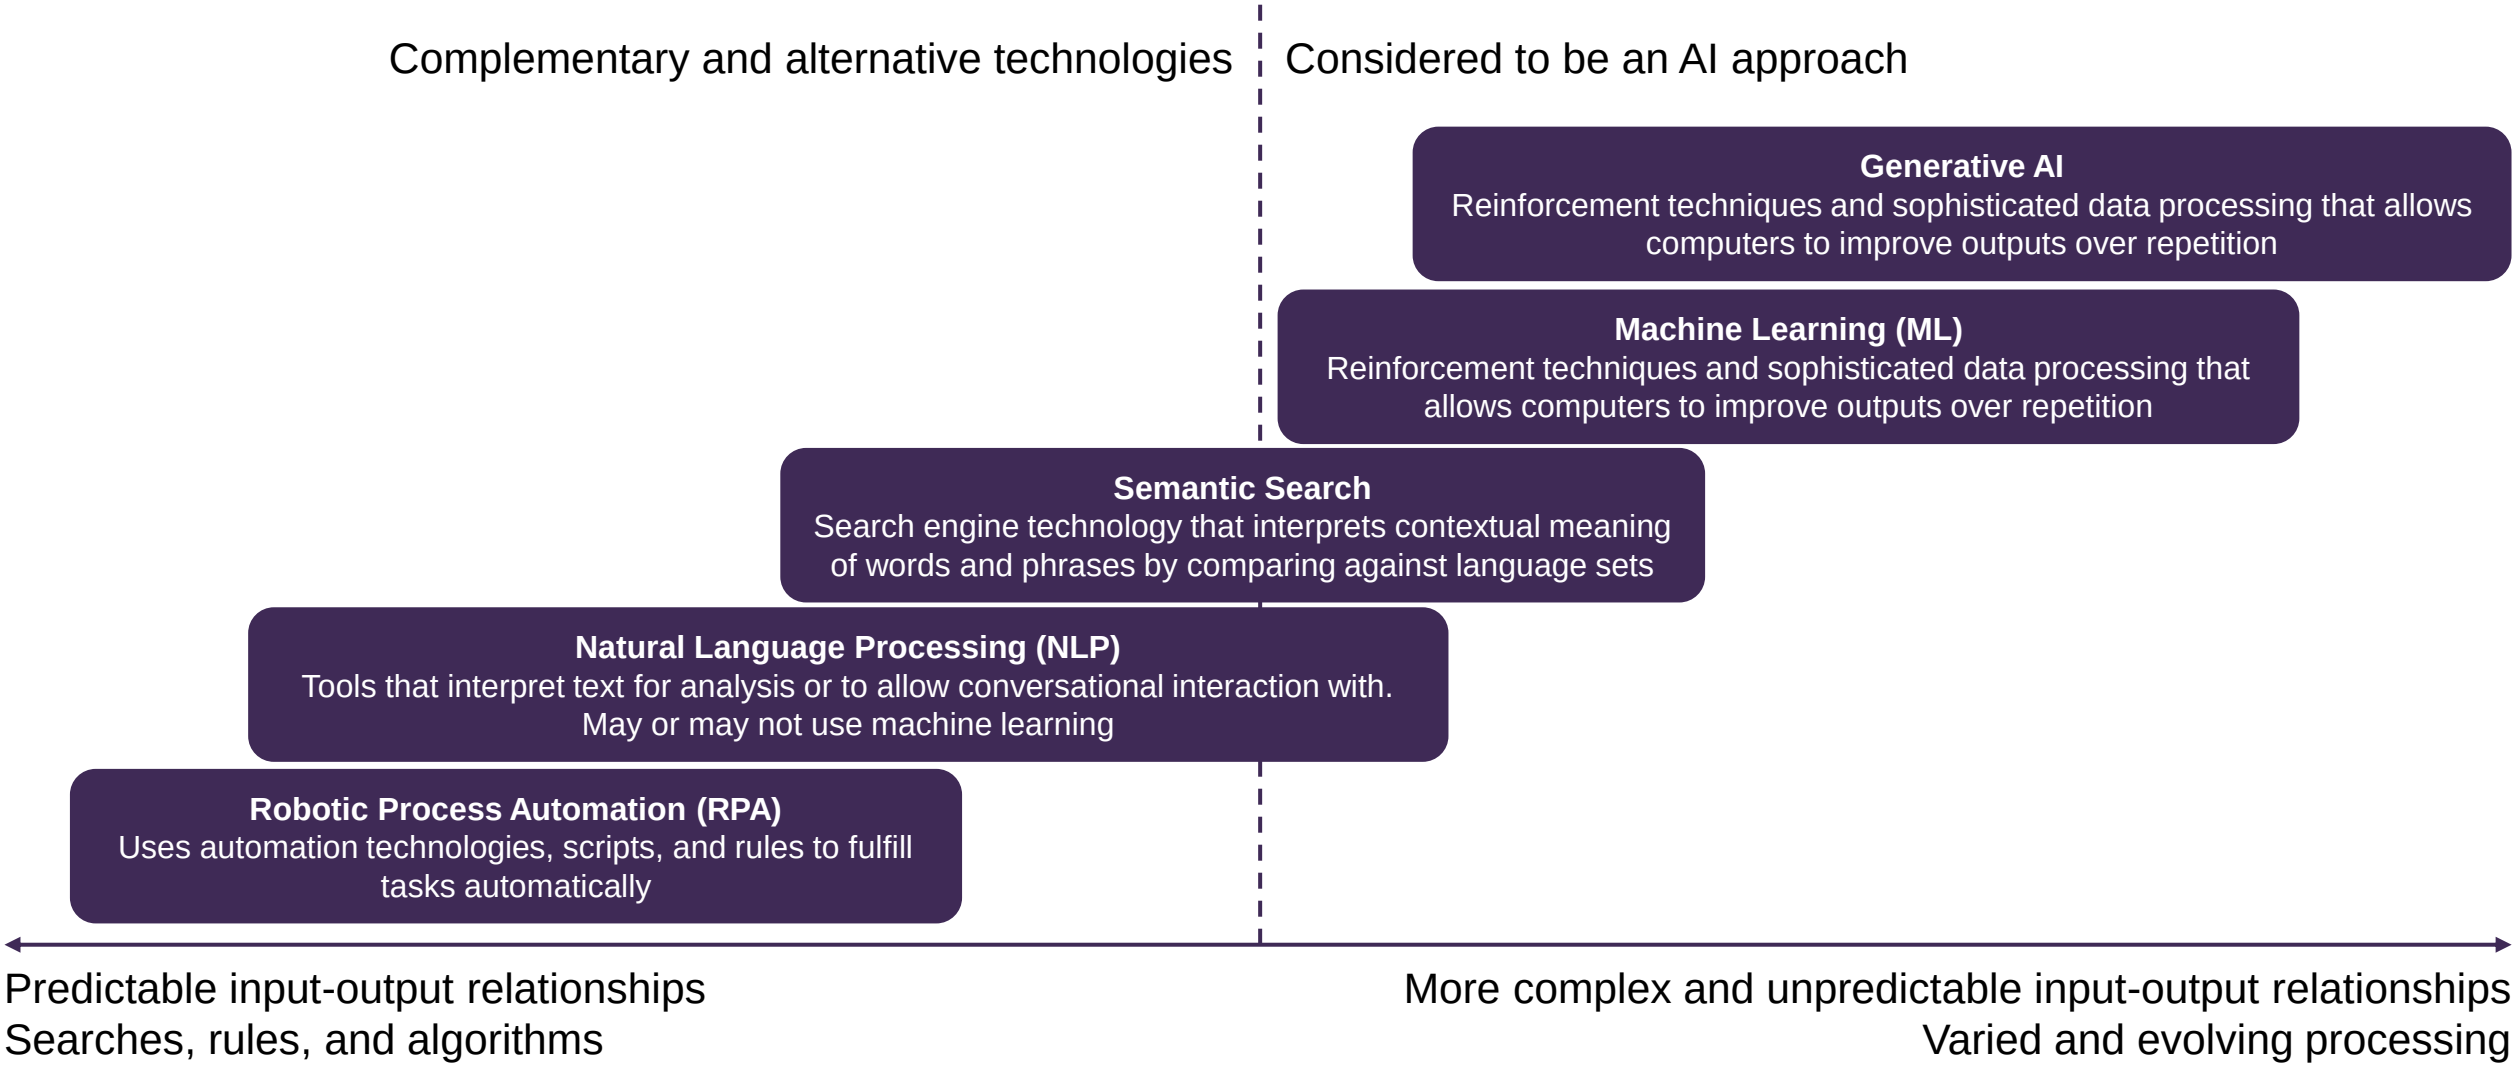

AI, automation, and analysis technologies: uses

Different approaches are suited, and used, for different tasks. These are some general and common examples; many systems have been developed to complete or support domain-specific tasks (e.g., x-ray analysis, scientific imaging analysis).

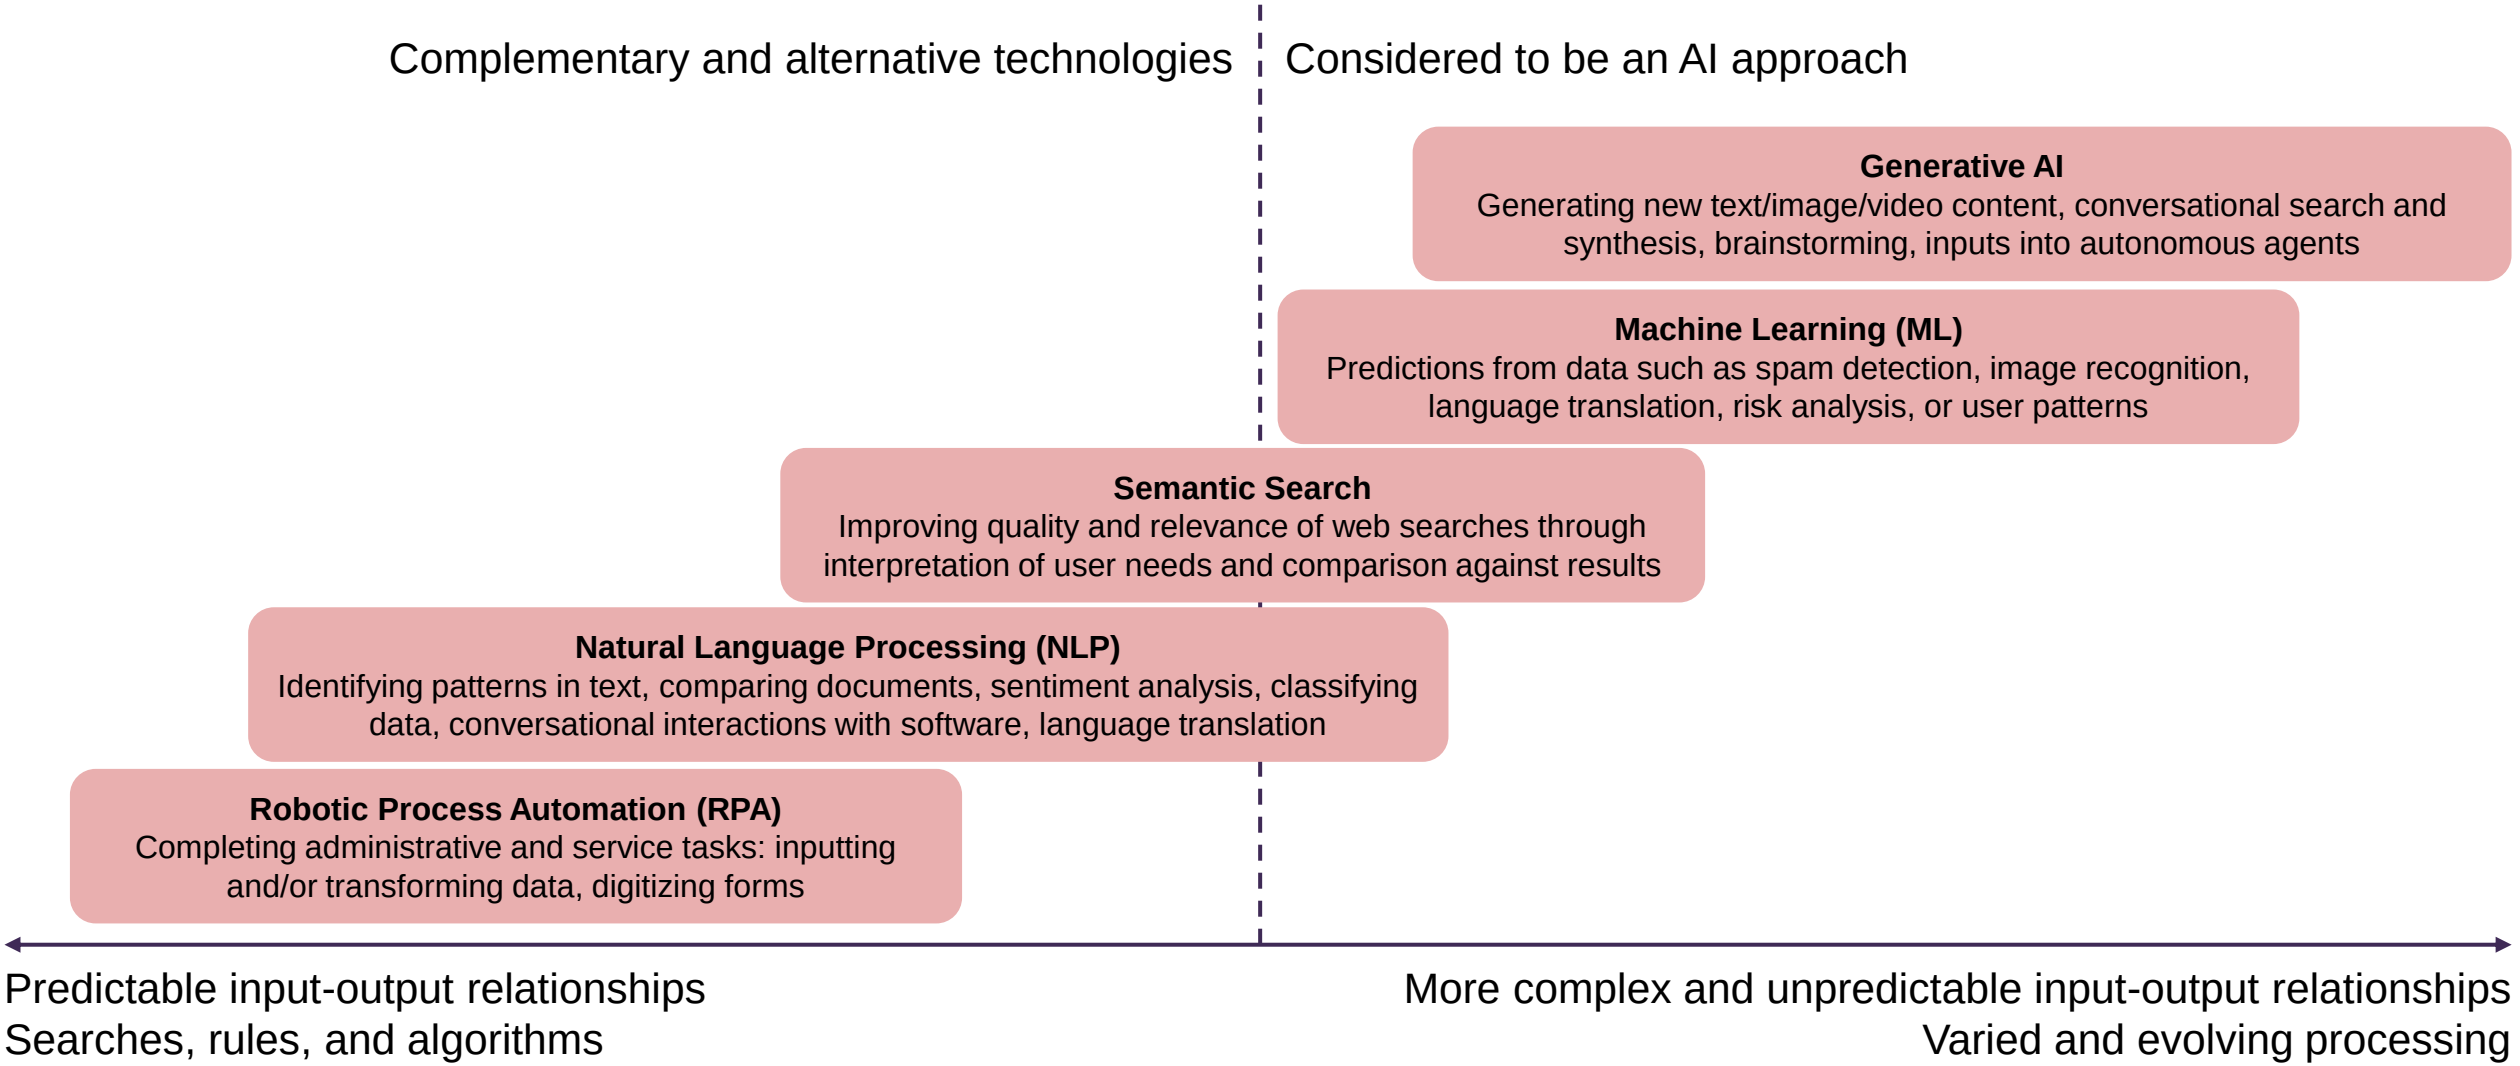

AI, automation, and analysis technologies: governance

Generally speaking, the level of risk consideration and governance should increase as projects use approaches towards the right side of this spectrum: complex and changing processing, particularly that evolves over time through use.

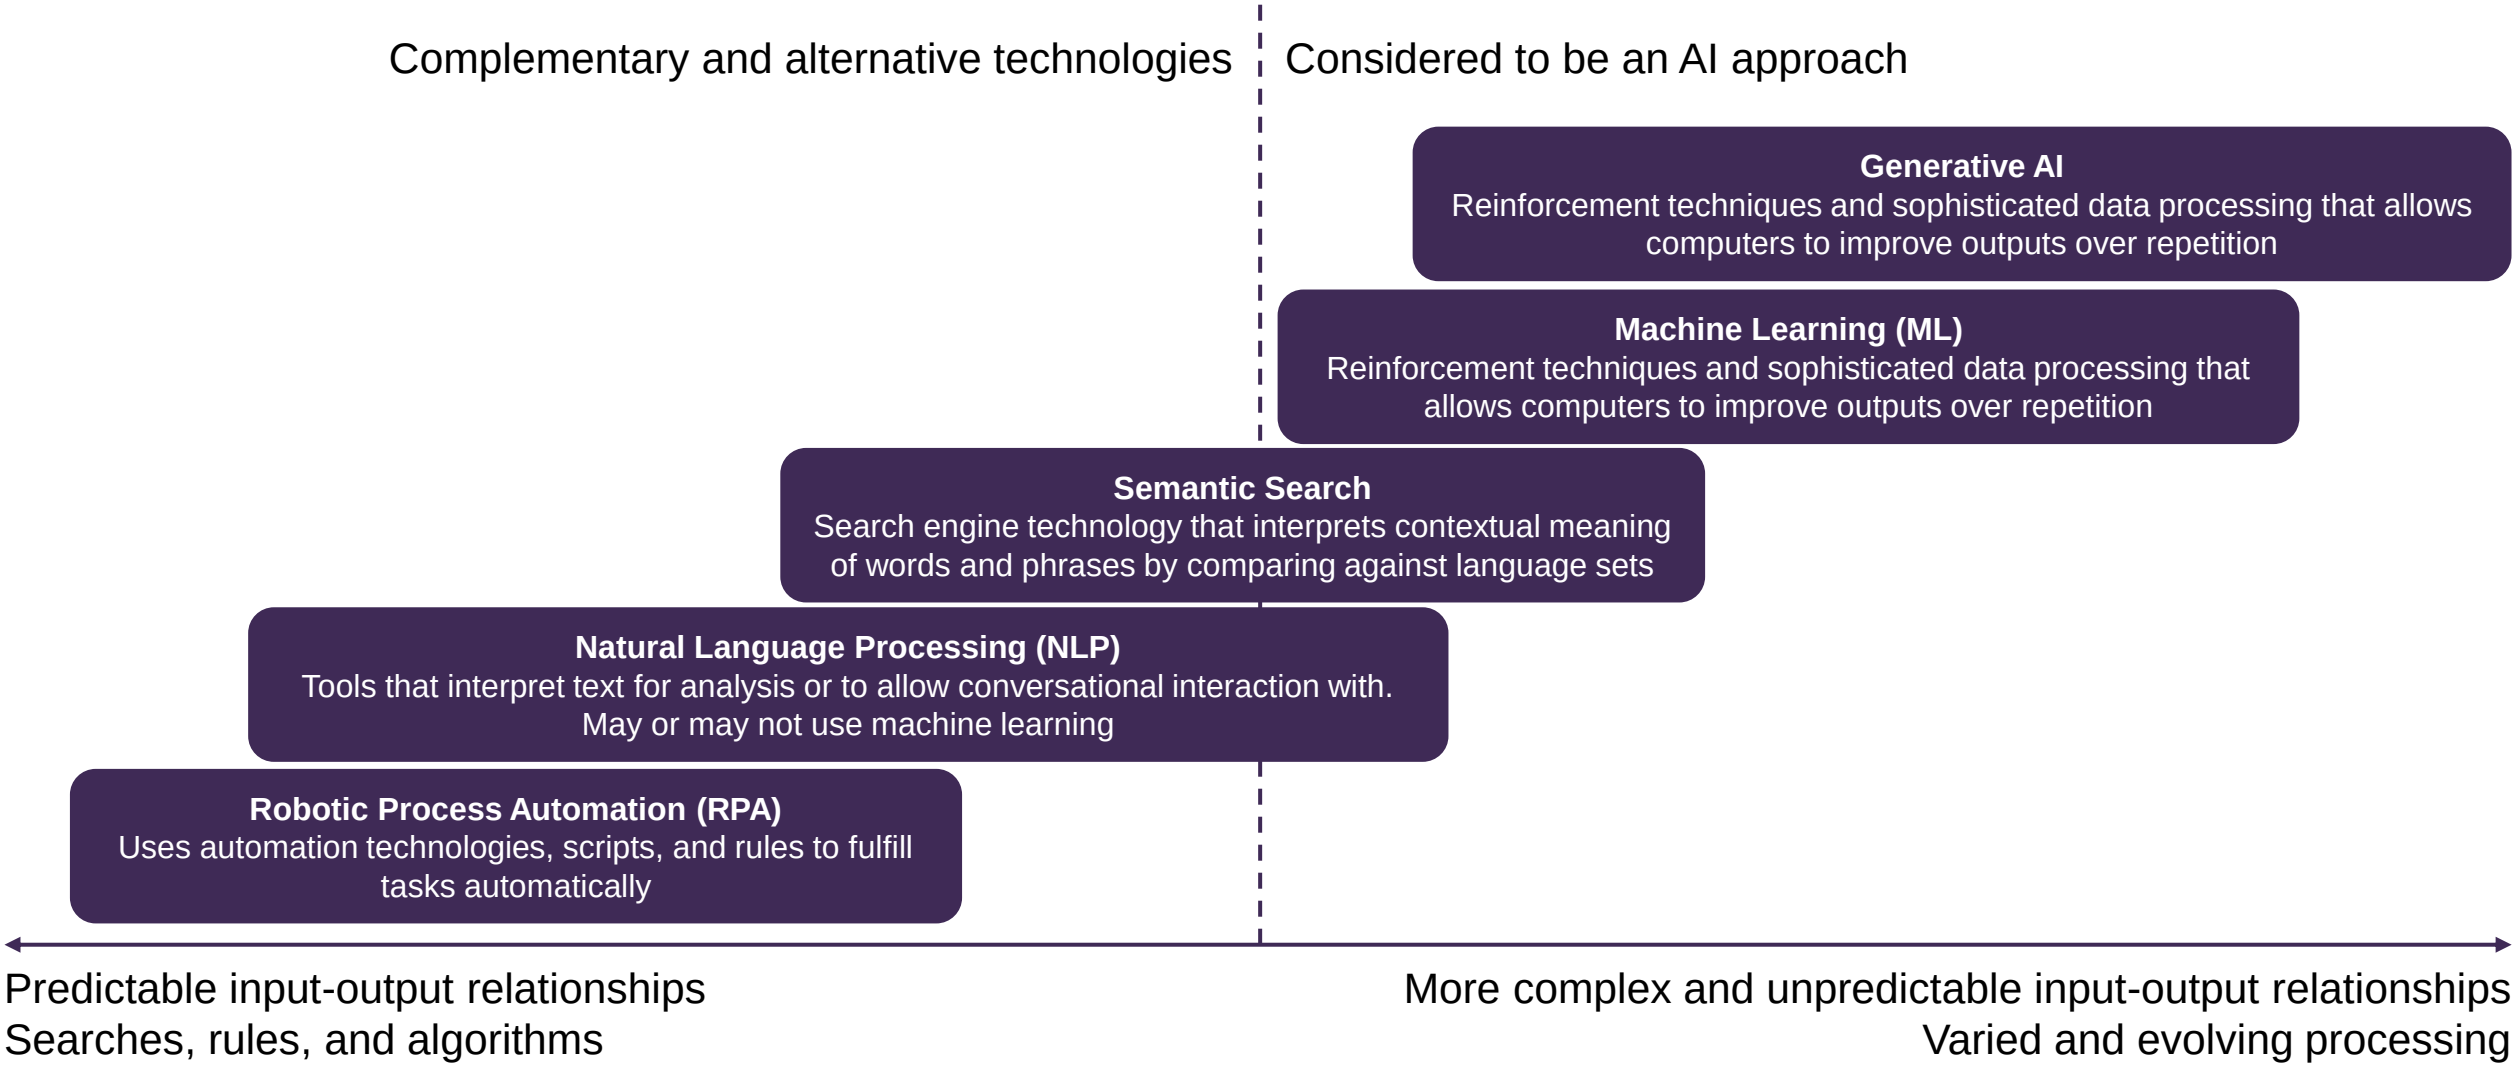

AI, automation, and analysis technologies: governance

The use context also influences risk; systems that impact policy and service decisions also require more governance. Another factor is whether the system *makes administrative decisions,* or *supports human decisions* within governance processes.

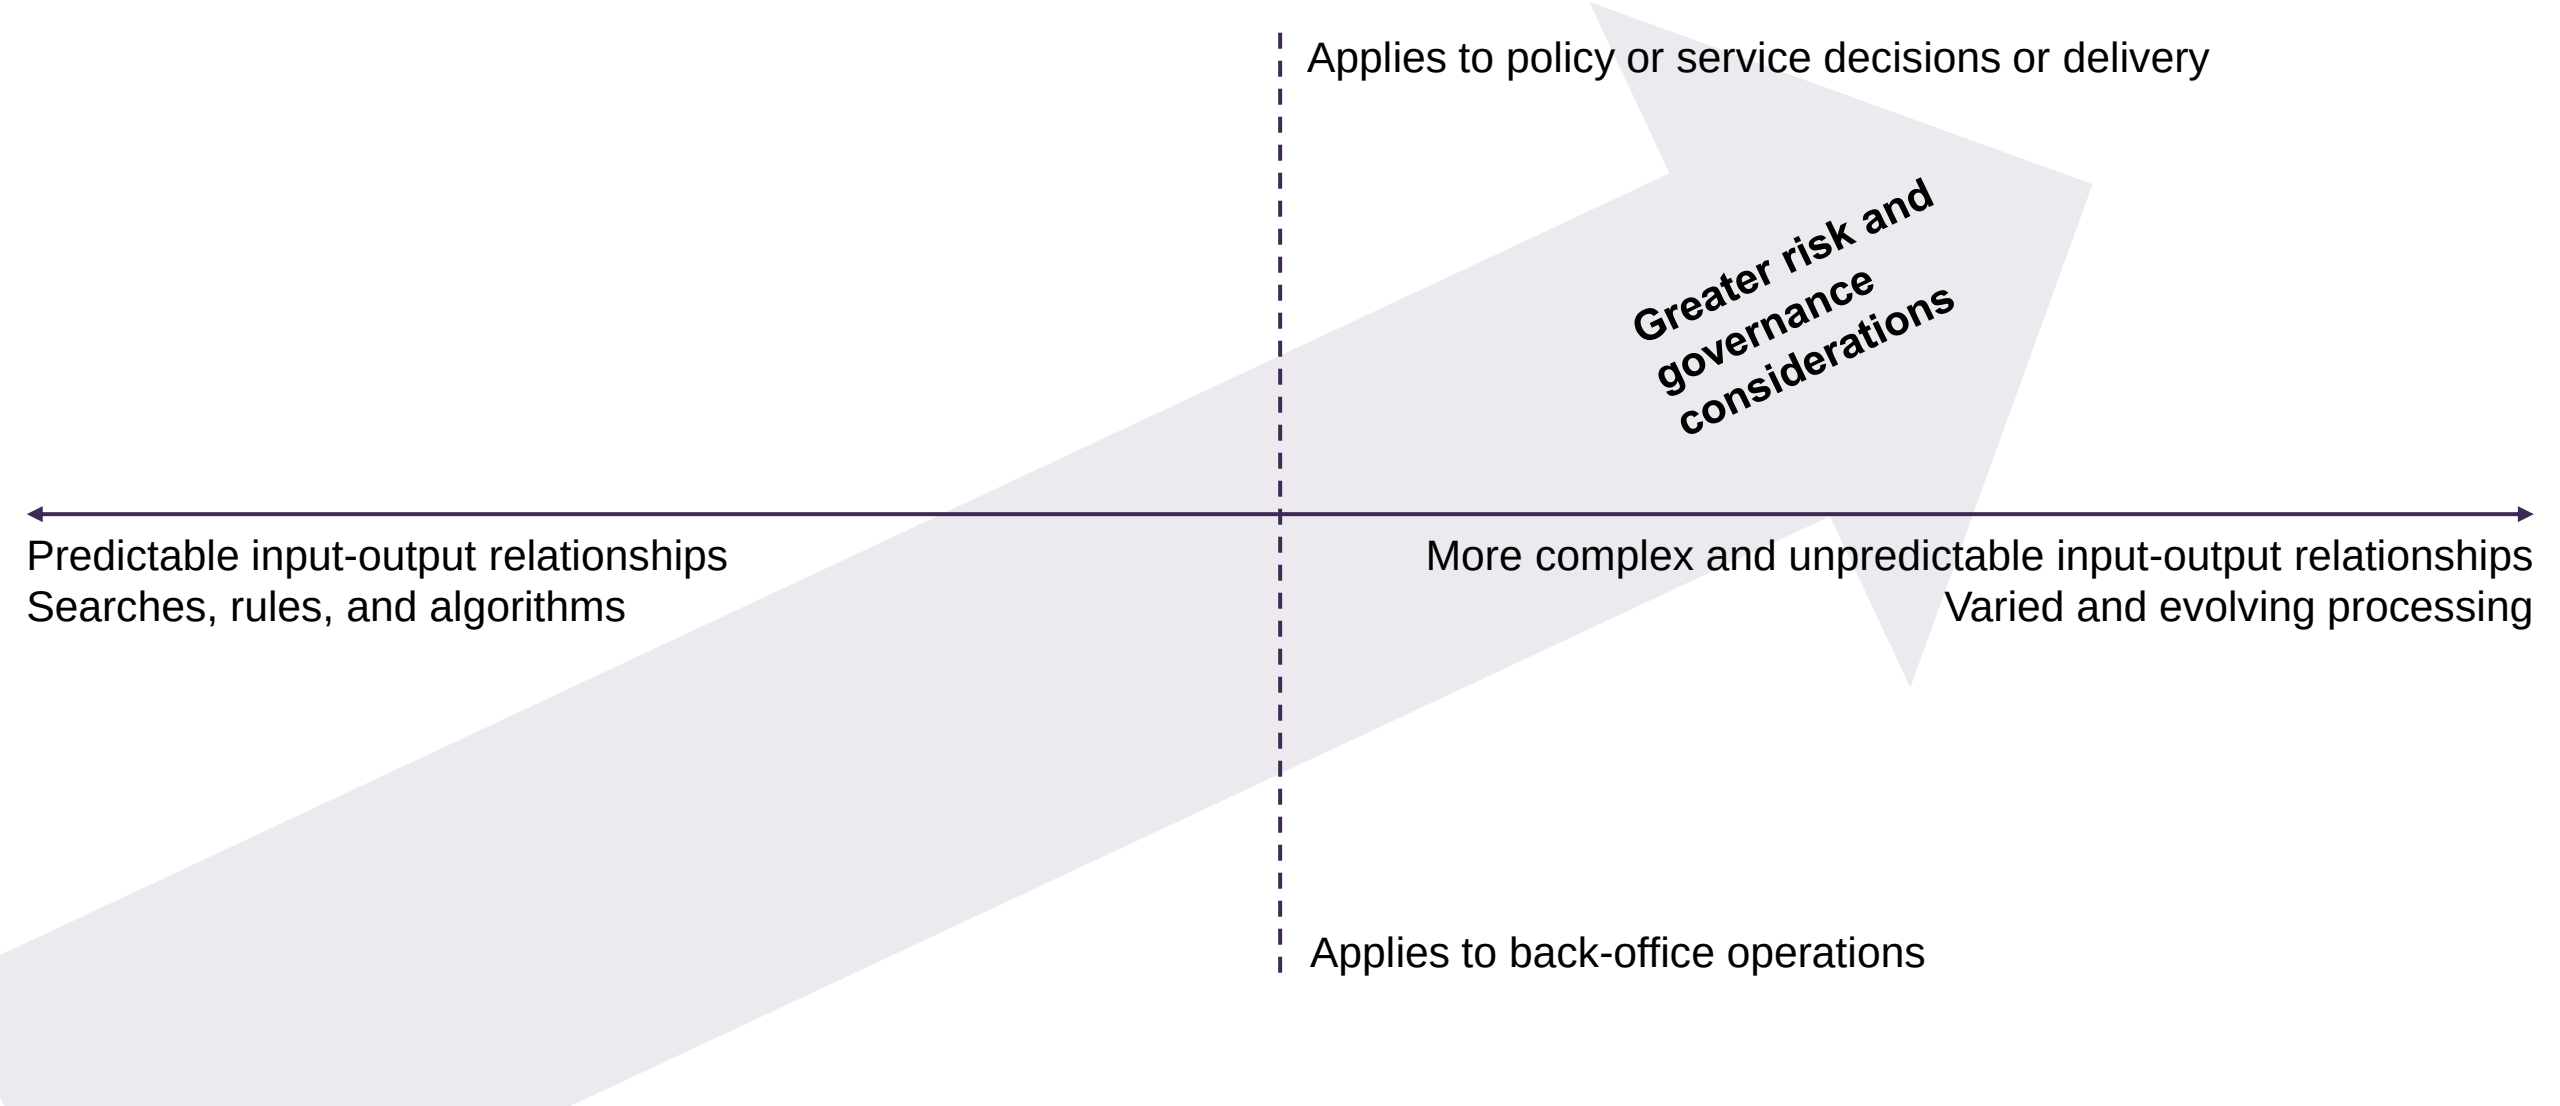

AI, automation, and analysis technologies: examples

The use context also influences risk; systems that impact policy and service decisions also require more governance. Another factor is whether the system *makes administrative decisions,* or *supports human decisions* within governance processes.

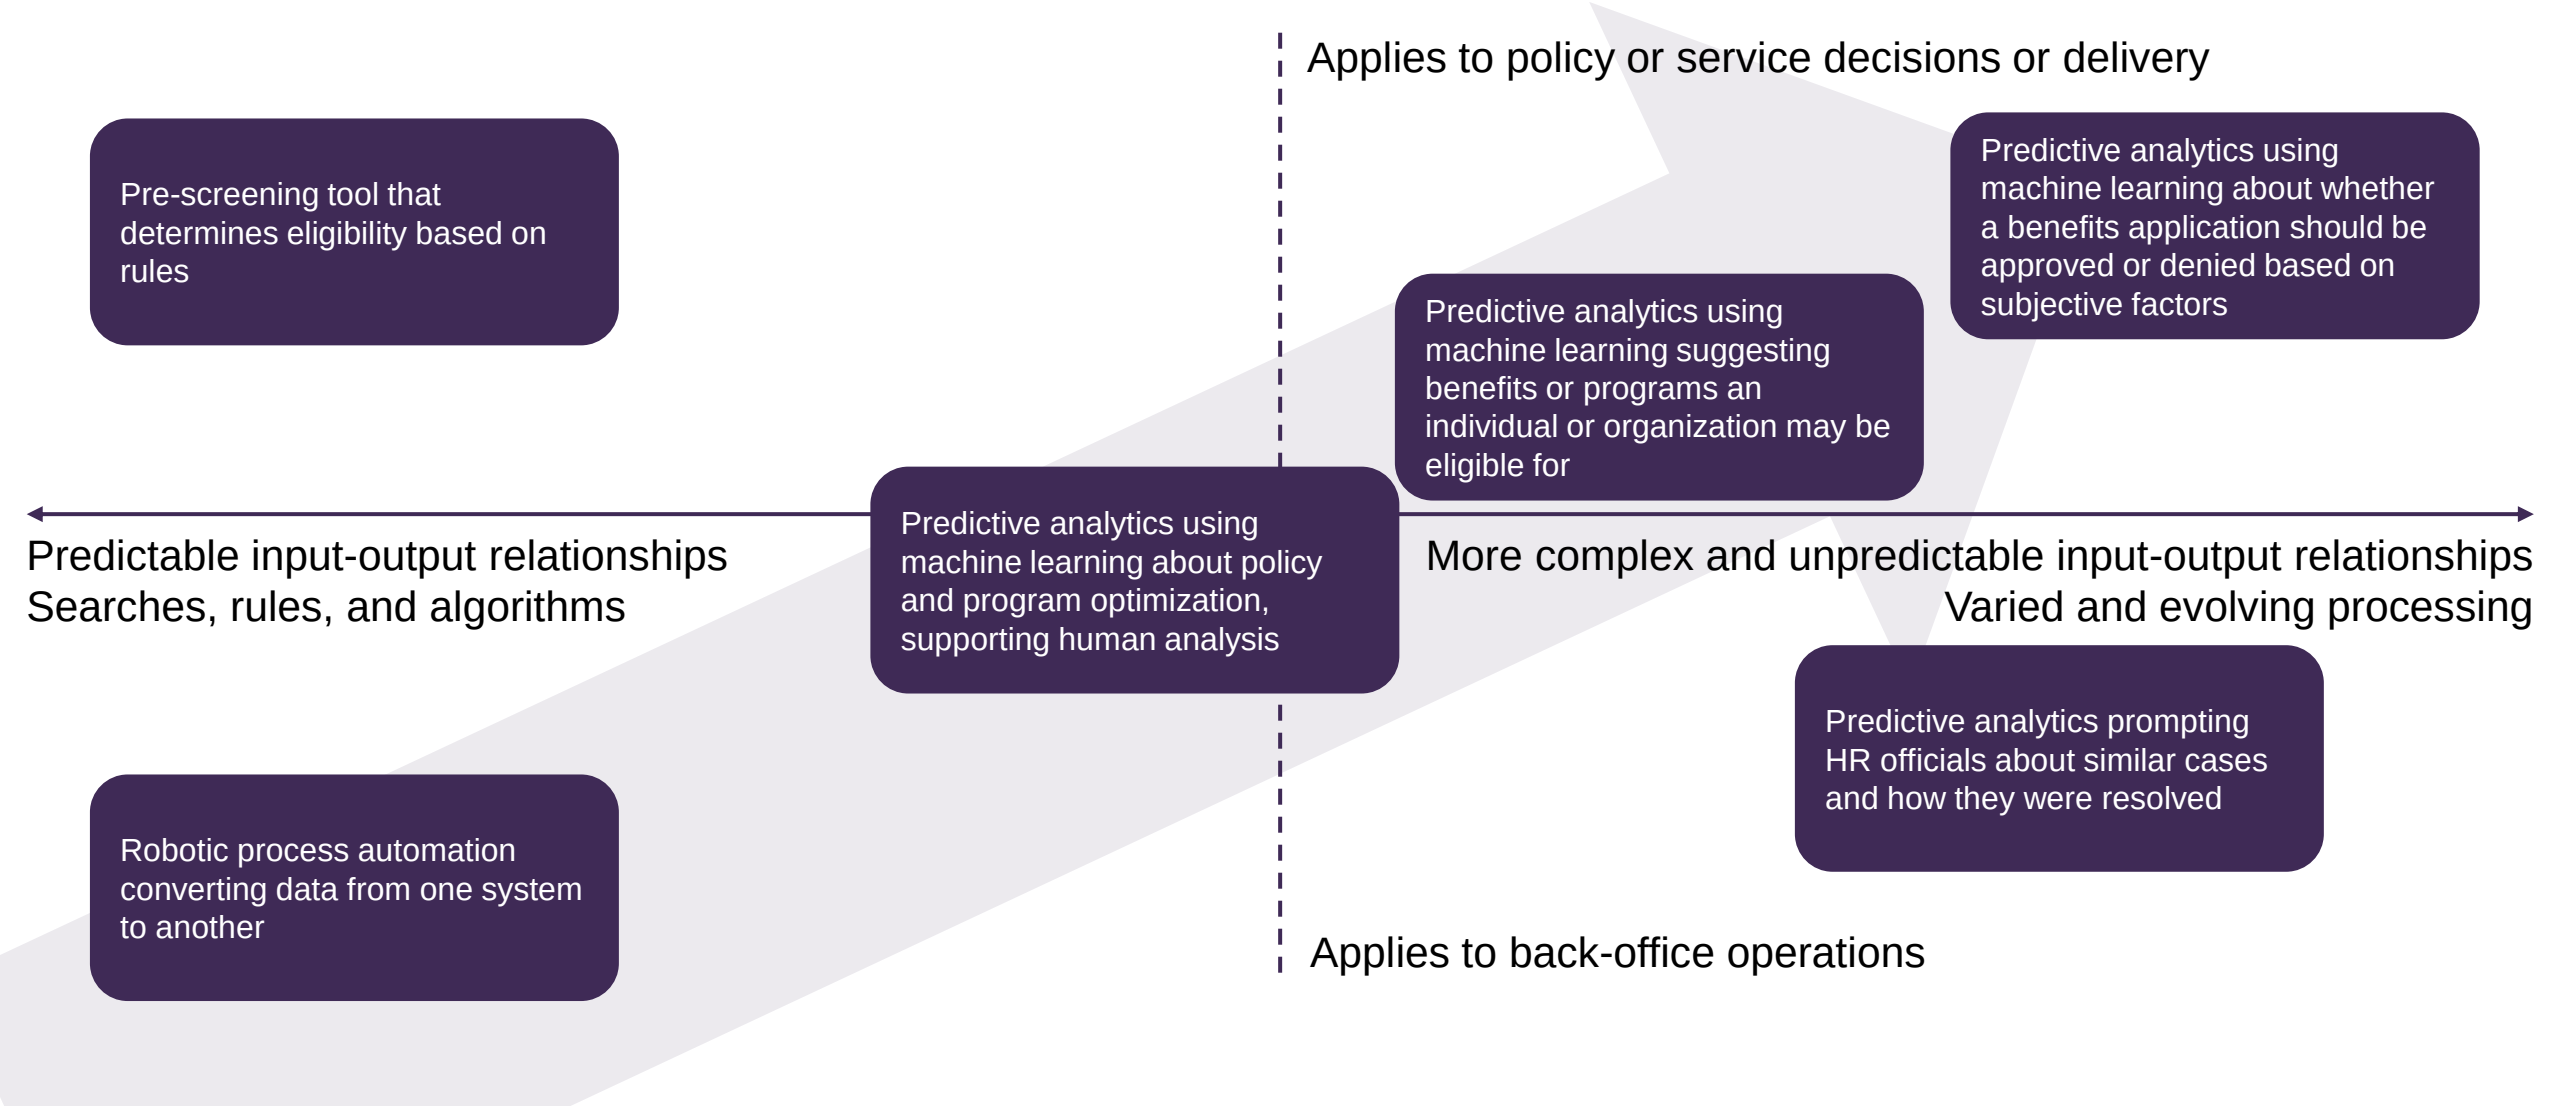

Policy and governance guidance

The Government of Canada has released guidance for public servants, and public service institutions, regarding certain uses of AI.

Automated decision-making

The *Directive on Automated Decision-Making* requires an Algorithmic Impact Assessment where GC use of AI supports, or renders, administrative decisions about individuals

Generative AI

Federal institutions – and federal public servants - are expected to align with the principles of fair, accountable, secure, transparent, educated, and relevant (FASTER) use of generative AI: *Guide on the Use of Generative AI*

External and internal capacity

1. Procurement resources
2. High-level overview of requirements for supporting internal data science and AI capacity

AI capacity: internal and external

Whether planning to procure expertise or solutions or leverage primarily internal capacity, it's recommended that project proponents own the design process and experimentation parameters, and particularly the **problem definition.**

Regardless of approach, organizations should consider their level of **understanding**, level of **maturity**, and the **complexity and viability** of the project.

This section provides high-level considerations for procurement and internal capacity.

Procuring AI solutions

The GC has developed a number of resources to support responsible procurement of AI-enabled solutions.

These can include custom design and build projects, integrations with existing software, and increasingly commoditized and mainstream AI tools.

Project proponents should connect through GC-wide AI networks and working groups to establish if other departments have experience with a solution set, or if a comparable custom design and build solution has already been procured.

PSPC resources

- Artificial intelligence source list
- Draft guidance tools to start AI procurement process
- PSPC AI procurement team email address

Supporting community resources

- Potential challenge questions to ask vendors

Note: given the intense current interest in AI, researchers have identified a trend of software 1) being branded as AI even if based on rules and algorithms, or 2) having largely superficial AI integrations added.

Setting up an internal data science and AI capacity

The exact skills, data, and infrastructure requirements will vary greatly across contexts. The below is a representative list of requirements and steps:

1. Get **cloud** and/or **computing power** and provide access
2. Migrate or build your **business applications and data**
3. Give your people **access to modern data tools** and the equipment they need:
    1. GPU acceleration
    2. Linux
    3. Access to libraries (Tensorflow, Keras, PIP, NPM, Git, Maven, etc.)
4. Train or find **specialized expertise**
5. Start running **modern analytics and visualizations** to understand your data
6. Build your **DevOps capacity** (version control, testing, continuous integration/deployment, containerization, cluster computing)

---

7. Pick a **business problem**
8. **Develop and test** your models 1) continuously and automatically and 2) with a multi-disciplinary team
9. Develop a **data pipeline** around your models
10. Improve **policy, service, and program outcomes** using AI

Organizations need to work through steps 1-6 before they're ready for steps 7-10.

Summary: recap of considerations and resources

Recap

1. Ensure the team is grounded in a **baseline understanding of AI**
2. Ensure the business owners and technology enablers are **working closely together**
3. Learn from **available evidence and comparable GC projects**
4. Leverage **problem framing, design,** and **experimentation approaches** to avoid mis-investment and to maximize clarity and shareability of results
5. Conduct an initial **viability assessment,** or at least evaluate:
    1. **Key elements** of the project
    2. **Data, infrastructure,** and **IM maturity** as well as **ethical considerations** in the context of the project
    3. Available **skills and expertise**
    4. **Overall maturity** and strategic business alignment
6. Make **skills**, **resources**, **computing power**, and **cloud-based and other tools** available if needed
7. If necessary, conduct an **Algorithmic Impact Assessment** or other governance and responsibility checks and balances

Summary of resources

AI resources

- **Primer on AI** executive summary
- **Discover AI** introductory course
- **Ethical considerations in AI** course
- GCDC **Data Toolkit** Library
- Artificial Intelligence on Canada.ca

Complementary resources

- **Robotic Processing Automation** Center of Support Excellence
- Departmental data strategies

Planning and design

- **AI Viability Model**
- **Design process series** of courses, from problem framing to prototyping
- GC **Digital Standards**

Governance and policy

- *Guide on the **Use of Generative AI***
- *Directive on **Automated Decision-Making***
- **Algorithmic Impact Assessment** tool
- *Policy on the **Planning and Management of Investments***

Procurement

- Artificial intelligence source list
- Draft guidance tools to start AI procurement process
- PSPC AI procurement team email address
- Potential challenge questions to ask vendors

Annexes

Annex: learning resources on AI

Primer	A Primer on AI
Courses	Discover Artificial Intelligence Using Generative AI in the Government of Canada Ethical Considerations in AI
Microlearning Articles	Decoding AI Assistants in Online Meetings Using Large Language Models (like ChatGPT) in the Federal Public Service Demystifying Artificial Intelligence OpenAI's ChatGPT Explained Working with AI series
Events	Artificial Intelligence Series: ongoing Practical Applications of AI Series: ongoing

Annex: learning resources on planning and process

Design courses	The Design Process: Understanding the Problem The Design Process: Design Research The Design Process: Qualitative Data Synthesis The Design Process: Ideation and Conceptualization The Design Process: Prototyping and Testing Inclusive by Design: Applying the GC Digital Standards and Gender-based Analysis Plus
Project management courses	Introduction to Project Management Project Management in a Government Context Managing Project Work (PMBOK® Guide Sixth Edition) Plan and Define Project Scope (PMBOK® Guide Sixth Edition) Agile Principles and Methodologies

Annex: Power Automate, Power BI, and Copilot

Project proponents should consider whether existing tools available through the GC MS Office Suite can perform some of the needed functions as identified in the business problem. There's learning available readily online for these tools.

Power BI

A collection of software services, apps, and connectors for combining, working with, and visualizing data. Can draw from standalone spreadsheets or be linked to cloud-based or on-premises hybrid data warehouses.

Most Power BI capabilities are *automation,* but there are AI integrations.

Power Automate

Office suite capability that allows users to create low-code / no-code connections, workflows, and automations between different MS applications.

Like Power BI, most Power Automate capabilities are *automation,* but there are AI integrations.

Copilot for Work

Microsoft's OpenAI integration that incorporates generative AI into many MS applications. Copilot is better thought of as many different applications than a single one. As of September 2024, Copilot for Work is the preferred generative AI solution for the GC.

Annex: Viability Assessment questions

1. Do you have strong understanding of the business problem?
2. Do you think the use of artificial intelligence is appropriate for solving your business problem?
3. Does your organization have experience in deploying AI solutions (COTS or custom-developed)?
4. Are there other business problems within your organization that could benefit from this or a similar AI solution?
5. Have you identified the sources of data and their location
6. Is there a need to remove sensitive data fields or anonymize portions of the data (remove Personal Identifiable Information) before training the machine learning model?
7. Could the predictions made by the AI solution be perceived as biased, unfair or unethical?
8. Are you the accountable party for the data?
9. Does the dataset originate from an authoritative or near-authoritative data source?
10. Have you analyzed the data to determine its current level of quality?
11. Are there data governance frameworks or data management processes in place to handle derived data?
12. Does the data contain the right level of information and features to solve the business problem?
13. Does your organization have a model library where models can be accessed, shared and reused?
14. Are you aware of any information in the data that when modelled could result in unwanted biases?
15. Will you need to explain the outcomes of the ML models in a way that can be easily interpreted by end users of the systems?
16. Is the training data of sufficient quality (entropy) for the chosen model?
17. Is the training data appropriately limited (so as to avoid overtraining)?
18. Can you share, reuse, and expand trained models, notebooks, and other ML components and solutions?
19. Has access to fit-for-purpose computing resources been acquired for the training of the machine learning model?
20. Do you foresee a need to periodically retrain the machine learning model?
21. Are the datasets used to train the machine learning model different (enough) from the datasets used to test its accuracy?
22. Is the test data representative of production data and contextually relevant?
23. Have you established metrics for measuring the desired performance and accuracy grade/level for the machine learning model?
24. Determine if metrics to track model performance have been identified to ensure unwanted biases are not introduced and model is not overtrained or overfitted.
25. Do you plan on deploying the AI solution in a production-type setting/environment?
26. Do you plan to go through the departmental Security Assessment and Authorization (SA/A) process?
27. Does the model or AI solution require a user interface?
28. Does your team or the organization have the necessary capabilities to develop a user interface to integrate with the machine learning model or the AI solution?
29. Do you plan on dedicating resources (in full or in part) to the project to ensure successful transfer of knowledge?
30. Does your organization have the necessary expertise (in-house or professional services) to operationalize the AI solution past proof of concept (PoC)?
31. Do you have a process in place to deliver production data in a programmatic fashion?
32. Is there a need for the artificial intelligence solution to integrate/interface with other enterprise business processes?
33. Does your organization have the ability to operationalize data processing and ML pipelines in production?
34. Does your team or the larger organization have the necessary skillsets to manage and maintain the AI solution long term?
35. Do you have a set of established key performance indicators to measure artificial intelligence solution performance and value-add?
36. Do you expect the model to produce predictions that are more accurate than alternative/current methods?
37. Does your team have the ability to validate artificial intelligence solution outcomes for false positives, false negatives, true positives and true negatives?
38. Do you have a plan or strategy to ensure that the machine learning model is continuously monitored throughout its lifecycle?
39. Do you have a plan or strategy in place to periodically report false positives or false negatives for model retraining?

Annex: skills overview

Technical Skills
- Data Science and Machine Learning
  - Understanding of machine learning algorithms and techniques (e.g., supervised, unsupervised, reinforcement learning)
  - Knowledge of statistical analysis and data preprocessing
  - Experience with programming languages commonly used in AI, such as Python and R
- Data Engineering
  - Skills in data collection, cleaning, and transformation
  - Knowledge of database management systems (SQL, NoSQL)
  - Familiarity with big data technologies (Hadoop, Spark)
- Software Development
  - Proficiency in software development principles and practices
  - Knowledge of software engineering tools and version control systems (e.g., Git)

Analytical Skills
- Critical Thinking
  - Ability to critically evaluate data, models, and contexts
  - Skills in identifying patterns and insights from data
- Problem-Solving
  - Proficiency in framing problems and designing solutions

Project Management Skills
- Project Planning
  - Ability to define project scope, objectives, and deliverables
  - Experience in developing project timelines and resource plans
- Agile Methodologies
  - Knowledge of Agile project management principles and practices
  - Skills in iterative development and continuous improvement
- Risk Management
  - Ability to identify and mitigate potential risks
  - Skills in developing contingency plans

Ethical and Legal Skills
- Ethical AI
  - Understanding of ethical principles in AI, such as fairness, transparency, and accountability
  - Skills in identifying and mitigating biases in AI systems
- Data Privacy and Security
  - Knowledge of data protection regulations (e.g., GDPR)
  - Skills in implementing data privacy and security measures

Domain-Specific Knowledge
- Understanding of Public Sector Needs
  - Knowledge of the specific challenges and opportunities in the public sector
  - Ability to navigate government approvals and checks and balances
  - Domain expertise in the policy, service, and program context
- Policy and Regulation
  - Awareness of relevant policies, regulations, and standards impacting AI in the public sector

Communication and Collaboration Skills
- Stakeholder Engagement
  - Ability to communicate effectively with stakeholders, including executives, policymakers, and the public
  - Skills in managing stakeholder expectations and obtaining buy-in
- Interdisciplinary Collaboration
  - Experience in working with cross-functional teams, including data scientists, engineers, and domain experts
  - Ability to bridge the gap between technical and non-technical team members

Continuous Learning and Adaptability
- Keeping Up-to-Date
  - Commitment to staying informed about the latest advancements in AI and related fields
  - Participation in continuous learning opportunities (e.g., workshops, courses, conferences)

Annex: project examples

This is not an exhaustive list, just examples we've found.

This third-party tracker – of automation, not necessarily AI – currently has over 300 Government of Canada systems listed: TAG Register (shinyapps.io)

Department	Project
Statcan	StatcanChat
Statcan	SAS to R/Python conversion tool
PSPC	ATIP portal service streamlining
AGR	AgPal Program and Service Finder
AGR	Document Detective to find duplicate and near-duplicate files
ESDC	CORDS no-wrong door approach for service
Justice	Predictive analytics
NRCan	Lichen Tools Dashboard
CRA	Charlie Chatbot