JUNE 2022

ENVIRONMENTAL LAW • INSTITUTE

Artificial Intelligence

Advancing Environmental Compliance, Enforcement & Follow-Up Programs

Impact Assessment Agency of Canada
Agence d'évaluation d'impact du Canada
Canada

ACKNOWLEDGEMENTS

This document was prepared by the Environmental Law Institute (ELI) as part of the ELI Innovation Lab. ELI's Innovation Lab focuses on the environmental impacts and opportunities created through emerging technologies and their underlying innovation systems. This report was authored by Heather Luedke, with input from Dave Rejeski, Margaret Badding, Zhuoshi Liu, Kasantha Moodley and Crystal Zhao. The project team would like to thank the government officials and technology developers for their time and contributions to this report. We are also grateful to the Impact Assessment Agency of Canada (IAAC) for their financial support of this work and for the feedback during the project. The contents of this report, including any errors or omissions, are solely the responsibility of ELI. The authors invite corrections and additions.

ABOUT THE ENVIRONMENTAL LAW INSTITUTE

ENVIRONMENTAL LAW • INSTITUTE®

The Environmental Law Institute makes law work for people, places, and the planet. Since 1969, ELI has played a pivotal role in shaping the fields of environmental law, policy, and management, domestically and abroad. Today, in our sixth decade, we are an internationally recognized, non-partisan publishing, research, and education center working to strengthen environmental protection by improving law and governance worldwide. Visit www.eli.org.

ABOUT THE IMPACT ASSESSMENT AGENCY OF CANADA

Impact Assessment Agency of Canada | Agence d'évaluation d'impact du Canada

The agency is a federal body accountable to the Minister of Environment and Climate Change, delivering high-quality impact assessments that look at both positive and negative environmental, economic, social, and health impacts of potential projects. The agency contributes to informed decision making on major projects in support of sustainable development in Canada. Visit www.canada.ca/en/impact-assessment-agency.html.

Contents

ACKNOWLEDGEMENTS 2
ABOUT THE ENVIRONMENTAL LAW INSTITUTE 2
ABOUT THE IMPACT ASSESSMENT AGENCY OF CANADA 2
1 INTRODUCTION 5
2 BACKGROUND 6
3 AI APPLICATIONS 9
3.1 Satellite Based Change Detection AI 10
The Challenge 11
Change Detection AI Model 11
Outcomes 12
Resource Commitments 12
Success Factors 13
Barriers 13
Agency Opportunities 14
3.2 Predictive AI for Inspection Targeting 14
The Challenge 15
Predictive Analytics AI Model 15
Outcomes 15
Resource Commitments 16
Success Factors 16
Barriers 16
Agency Opportunities 17
3.3 Supervised Classification AI for Complaint Processing 18
The Challenge 18
Supervised Classification Model 18
Outcomes 19
Resource Commitments 20
Success Factors 20
Barriers 20
Agency Opportunities 20

3.4 Predictive AI for Inspection Targeting 21
The Challenge 21
Predictive Analytics AI Model 22
Outcomes 22
Resource Commitments 22
Success Factors 23
Barriers 23
Agency Opportunities 23
3.5 Intelligent Document Processing for Compliance Rating 24
The Challenge 24
Intelligent Document Processing for Compliance Rating 24
Outcomes 25
Resource Commitments 25
Success Factors 25
Barriers 26
Agency Opportunities 26
3.6 Summary of Use Cases 26
3.7 EnviroVerse – The Future of Regulatory Monitoring 29
4 ADOPTION CRITERIA 30
5 NEXT STEPS 34
6 REFERENCES 35

1 INTRODUCTION

Artificial intelligence (AI) broadly refers to computing systems that perceive, learn, abstract, and reason to process information. As AI continues to advance and expand into new domains, both interest and concerns regarding applications of this technology are growing. The public sector reflects these trends. A 2021 survey of 500 government leaders across the world reveals that an overwhelming majority of federal, state, and local agencies view AI as an important factor in meeting mission outcomes over the next five years [Van Buren, 2021]. The same study reports that there is a pervasive gap between the current and desired state of AI capabilities among observed agencies, credited to a variety of concerns including skill deficit, institutional attitudes toward AI, technological capacity, and ethical considerations. Even acknowledging such limitations, AI presents opportunities to solve major governance and management challenges for the public sector.

There is a wide array of AI applications in the public sector which are still expanding today. Government agencies worldwide are using AI to detect fraud, respond to public queries, make welfare payments, adjudicate bail hearings, triage healthcare cases, and more [Deloitte AI Institute, 2021]. In the realm of environmental governance, AI is improving climate predictions, helping cities build climate resilience, and facilitating sustainable city planning. Governments and firms across the world are leveraging AI to meet their climate mitigation and environmental protection goals.

Environmental compliance, enforcement, and follow-up activities fall under this broad spectrum of potential environmental AI applications. Environmental compliance refers to the promotion and verification of mitigation measures, as dictated by environmental agencies and regulation. This may include information sessions, onsite and offsite inspections, investigations, and coordination with other authorities. Enforcement encompasses all actions taken to restore compliance after non-compliance has been identified. Finally, follow-up activities refer to the processes for verifying the accuracy of the impact assessment of a designated project and determining the effectiveness of any mitigation measures in mitigating potential negative effects of a project. Environmental agencies tasked with carrying out these actions are often highly constrained by financial, human, and technical resources. Vast amounts of reports and too few inspectors are among the common issues that these organizations face. Many are now looking to explore how AI technology can optimize available resources and improve their work.

ELI was commissioned by the Impact Assessment Agency of Canada (IAAC) to identify high-impact uses of AI that will strengthen post-decision activities. These uses are expected to inform and prepare IAAC for these activities. The AI applications discussed below have been judged to bear relevance to IAAC. This exploratory research is meant to inform IAAC’s possible future adoption of AI and machine learning (ML) technologies.

Though this report was developed for the post-decision activities of IAAC, its findings may be applicable to other functions of IAAC and other agencies at a similar stage of AI exploration and development. We hope this research aids environmental managers in their efforts to integrate AI innovation into critical environmental compliance, enforcement, and follow-up work.

2 BACKGROUND

2.1 IAAC

The Impact Assessment Agency of Canada (IAAC), known as the Canadian Environmental Assessment Agency before 2019, is a federal body accountable to the Minister of the Environment and Climate Change. Its primary mandate is to conduct the impact assessment process. Aside from leading and managing the impact assessment process for all federally designated major projects, IAAC serves as the primary point of consultation and engages Indigenous people and other stakeholders at appropriate points in the process. IAAC is also responsible for post decision activities including ensuring compliance with conditions in project decision statements and tracking and reporting on follow-up programs. This report represents an exploratory step in IAAC's plans to determine which post decision activities could be automated.

Before discussing potential areas of opportunity for AI intervention, it is necessary to define AI in context along with a few other terms (autonomous, ML and algorithm). Experts do not agree on what constitutes the bounds of artificial intelligence. For the purposes of this report, artificial intelligence is recognized to be "machines that respond to stimulation consistent with traditional responses from humans, given the human capacity for contemplation, judgment, and intention," as argued by researchers Shubhendu and Vijay [Shubhendu, 2013]. In other words, AI is a computer system that can make decisions that would normally require human expertise. AI systems generally possess the ability to perceive, learn, abstract, and reason. The term autonomous refers to the ability to make decisions without human intervention. ML is an application of AI that enables systems to learn and improve from data inputs without human oversight. Finally, the term algorithm refers to the set of rules or instructions that dictate an AI's operating behavior.

2.2 AREAS OF FOCUS

IAAC is now beginning to explore options for incorporating AI into their post decision activities. IAAC's post decision activities framework reveals many areas of opportunity for AI intervention (Figure 1). In particular, compliance verification, enforcement actions, and follow-up activities offer potential for improvement through AI integration. Compliance verification traditionally relies on enforcement officers and analysts to carry out onsite and offsite inspections, investigate suspected contraventions, and coordinate with government authorities. Case studies reveal that a combination of satellite imagery, predictive analytics, risk assessments, and change detection algorithms could be effective at optimizing efficiency and efficacy of violation identification.

Enforcement actions consist of enforcement officers and analysts issuing warnings, notices, and orders to take corrective measures. Pursuit of injunctions, prosecution, and penalties also falls under this category. Follow-up activities, carried out by the follow-up team, usually include reviewing follow-up plans and reports, analyzing effectiveness of mitigation measures, evaluating accuracy of predictions, developing agency follow-up reports, engaging stakeholders, and identifying process improvements. Both enforcement actions and follow-up activities are document-intensive, requiring many hours of human labor to process the necessary paperwork. Case studies reveal that intelligent document processing could increase the efficiency of enforcement officers and the follow-up team when it comes to these tasks.

Compliance promotion through education and training is the final major category of post decision activities. This report focuses on use cases pertaining to the three other opportunity areas, as IAAC identified those areas as priorities.

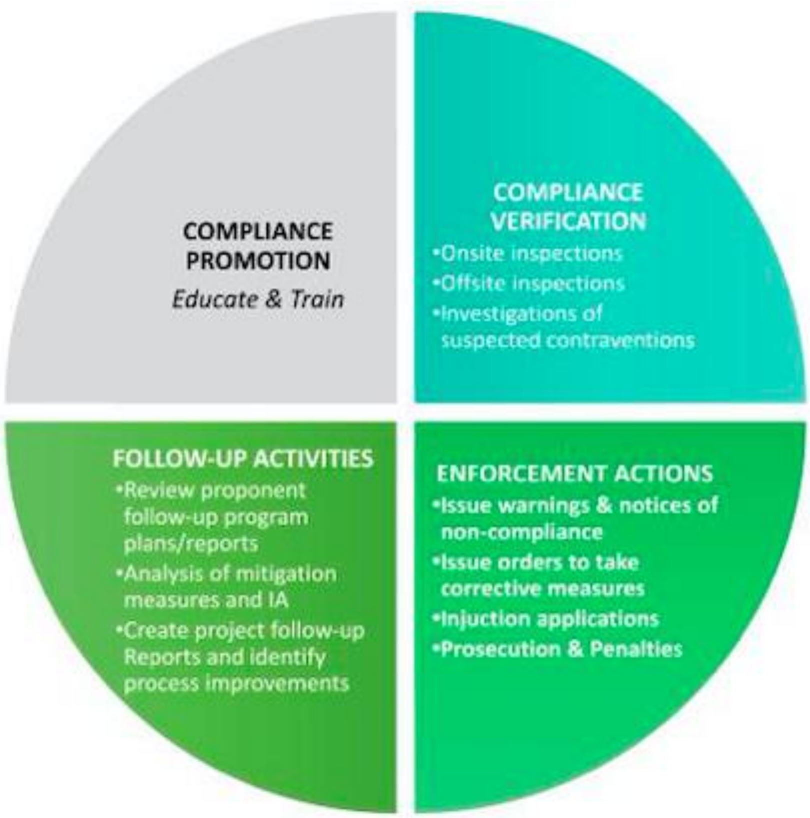

Figure 1: IAAC Post Decision Activities

2.3 BARRIERS AND CHALLENGES

While recognizing the IAAC's potential for effective AI integration, it is also important to acknowledge likely challenges and barriers for adoption. First, AI systems require large data inputs to operate effectively. Data availability may hinder the IAAC's automation efforts, as it only became involved in post decision activities starting in 2012 and has only limited amounts of historical data. IAAC's lack of organizational experience with AI may also pose a problem, as it will not likely possess the skill and technological capacity required to operate AI systems without either significant hiring and training or consulting external experts. Funding may also pose a challenge for IAAC, and it will need to work to secure continued funding for AI to maintain and operate AI systems and retain the necessary experts.

IAAC may face issues with system interoperability as well. IAAC's interest in AI adoption comes at a time when the Government of Canada is encouraging AI uptake and exploration across multiple agencies. As other agencies adopt AI solutions, it becomes more necessary for applications to be designed to work in conjunction with existing and future systems. This is especially true of IAAC, which is mandated to collaborate with other agencies at multiple points in the impact assessment process. Finally, IAAC would need to seriously consider ethics when thinking about developing and implementing AI solutions. The Government of Canada did develop an Algorithmic Impact Assessment (AIA) to determine the ethicality of AI applications used in governance, but past experience reveals that this safeguard does not fully inoculate AI technologies against public controversy. Immigration, Refugees, and Citizenship Canada's foray into immigration vetting algorithms was the subject of critical media coverage despite passing the AIA [Nalbandian, 2021]. IAAC will need to consider and address these challenges to maximize the impact of their potential future AI solutions.

3 SCOPE OF RESEARCH PROJECT

This report is the product of a six-month research process. ELI began by engaging the team at IAAC to confirm research goals and gain knowledge about the organizational context as it pertains to AI suitability. The next several months were dedicated to information-gathering and research, mainly achieved through interviews and reviewing existing literature.

3.1 INTERVIEWS

ELI conducted interviews with members of the International Network of Environmental Compliance and Enforcement (INECE) Community of Practice on the Use of Data Analytics and Artificial Intelligence in Environmental Compliance and Enforcement, members of INECE who responded to a request of information in the INECE newsletter, and other individuals involved with AI at their respective agencies. In total, seven interviews with environmental governance professionals from around the world: Jed Anderson of EnviroAI; Michael Enns of Environment and Climate Change Canada; Justin Budgell of Health Canada; Cristobal de la Maza Guzmán of Superintendency of the Environment, Chile; Randy Hill of the U.S. EPA's Office of Enforcement and Compliance Assurance; Ray Purdy of Air and Space Evidence on behalf of the Ministry of Environment, New Zealand; and Paul Stevens of Victoria, Australia's Game Management Authority. Interview subjects were asked to speak on the state of AI adoption at their agency, how it is helping to meet their goals, who is responsible for operating the system, outcomes, critical success factors, barriers, and cost of uptake. These interviews were critical in informing the report's use cases.

3.2 LITERATURE REVIEW

A literature review was also produced as part of the research process. Major topics of interest included AI in government, use cases in and outside of the environmental context, adoption considerations, accountability, and document processing. This step helped develop ELI's background knowledge on the topic. More specifically, the literature review focused on AI in the public sector and the adoption of AI by public administrations. The literature largely agrees that AI is useful for increasing agency efficiency and opening up employee time to work on tasks that cannot be automated [Martinho-Truswell, 2018]. However, a barrier to implementation of AI is employee worries regarding job-loss related to AI uptake. Thus, measures such as stakeholder involvement and education are imperative to help agencies integrate AI into their work in a way that gets everyone on board and leads to success [Chenok, 2018]. The literature also focuses on the importance of hiring in-house expertise for AI projects, noting data scientists are in high demand and agencies must figure out how to offer appealing employment opportunities and retain staff [GoDataDriven, 2019].

The literature also highlights another challenge related to the development and implementation of AI: issues of governance, legality, and ethics. It is imperative agencies are accountable for their AI systems, are transparent, educate and incorporate feedback from all relevant stakeholders, and use a risk management plan to mitigate risks and biases from AI systems. They also need to ensure they are operating within legal bounds [U.S. Government Accountability Office, 2021]. Additional literature suggests that a strategy or roadmap that addresses these concerns, and others, can be useful for an agency starting to pursue AI [Van Buren, 2021; GoDataDriven, 2019].

Despite robust literature on AI in the public sector and in other industries, such as healthcare, using AI to assist with environmental compliance and enforcement and follow-up activities is relatively new. Many projects are still in the testing or piloting phase. Thus, there is a lack of robust results and reporting on the integration of AI into environmental compliance and enforcement and follow-up. This report aims to provide more information about existing AI projects that bear relevance to environmental compliance and enforcement and follow-up activities. After gathering potential use cases from interviews and existing literature, ELI selected the cases most relevant to IAAC's activities, in addition to being relevant for other environmental agencies. These applications are discussed below.

4 AI APPLICATIONS

This section overviews five AI use cases from the Canadian context and internationally. These case studies were selected for their relevancy to IAAC and take into consideration the particular set of challenges facing IAAC. IAAC has a growing number of projects coming into the system to be assessed under the IAA, which will lead to a steady increase in the number of projects they will be overseeing in the post decision phase. As a result, AI models that can aid in the compliance & enforcement and follow-up processes are of particular interest. These include models to assist with detecting non-compliance, assessing risk, processing and writing regulatory reports, processing complaints, and identifying trends for follow-up. A complete list of the case studies is listed below.

Type of AI Model	Intent	Government Authority
Satellite Based Change Detection AI	To identify unknown and unlicensed waste sites	Ministry of the Environment, New Zealand
Predictive AI for Inspection Targeting	To identify noncompliance of hazardous waste facilities	United States Environmental Protection Agency
Supervised Classification AI for Complaint Processing	To predict the relevancy and gravity of filed complaints	Superintendency of the Environment, Chile
Predictive AI for Inspection Targeting	To predict noncompliance and risk at regulated facilities	Environment and Climate Change Canada
Intelligent Document Processing for Compliance Rating	To assign risk ratings to instances of noncompliance and assist with report writing	Health Canada

All of these cases are in the early stages of development and/or implementation, and they all have pilot projects underway. Although these AI technologies may have more advanced applications in other sectors, these particular cases were selected as they are more directly relevant to IAAC as an agency involved in post decision impact assessment activities.

While these cases were selected with IAAC in mind, they also provide a sampling of AI opportunities that could be implemented at environmental compliance and enforcement agencies across the world. The types of projects, in addition to information about their success factors, barriers, and resource commitments, could assist many environmental agencies in developing and implementing AI projects.

The format of each case study is as follows: Each case study starts with a description of the challenge the agency is facing and then describes the particular AI model that was developed to address the challenge. The description of the model includes information about the data used, development process, and current status and use. Outcomes, resource commitments, success factors, and barriers are then provided. Each case study concludes with a section about the relevance of the model to IAAC.

4.1 Satellite Based Change Detection AI

Authority:	Ministry of the Environment, New Zealand
Collaborator:	Air & Space Evidence
Intent:	To identify unknown and unlicensed waste sites
Contact:	Shaun Lewis, Waste and Resource Efficiency Division, Ministry for the Environment, New Zealand

The Challenge

In 2018, the World Bank named New Zealand the tenth most wasteful country in terms of waste produced per capita. Approximately 3.6kg of waste was generated per person every day. This was five times the global daily average of 0.65kg per person. Over the past 10 years, New Zealand has sent more than 30 million tons of trash to landfills, this volume is expected to increase particularly in the construction and demolition sector. To curb this trend, the New Zealand Ministry for the Environment (the Ministry) is proposing to increase and expand the national waste disposal levy from NZD 10 per ton to NZD 60 (USD 6 to 37) per ton over a period of 4 years. To implement the levy, the Ministry is developing a registry of all operational, historical, and proposed landfill sites, with the intention of recording both licensed and unlicensed landfills. The Ministry was challenged in identifying unknown or unlicensed sites.

Change Detection AI Model

In 2021, the Ministry commissioned Air & Space Evidence (ASE), a spinoff company from University College London, to help identify waste sites for the registry. ASE specializes in detecting regulatory non-compliance using satellite data. More specifically, ASE focuses on waste crimes and uses artificial intelligence in a satellite change-detection model to identify illegal waste sites and waste sites that extend beyond permitted bounds.

ASE's satellite change-detection model involves a two-step process to identify and characterize waste sites. The first stage uses medium spatial resolution data from the European Space Agency Sentinel satellites. Sentinel-2 data is filtered to identify potential landfill sites. The model is calibrated using Earth Observation (EO) data of known landfill sites. These known landfill sites are located within the area to be examined by the model and are provided by the Ministry. The analysis methodology used in this stage uses AI techniques developed by ASE. The output from the model is a file that reveals the locations of potential landfill sites, including geographic locations. False positives such as quarries may be flagged during this stage.

The second stage uses high spatial resolution data and aerial imagery to examine in detail the sites that were identified. This analysis is done manually and allows for the characterization of the internal structure of the landfill sites, including an estimation of waste type. The data used in this stage are from the Maxar DigitalGlobe satellite, open-source high resolution EO data, and aerial photograph archives. ASE is presently researching the automation of this stage.

Once a list of potential sites is generated, one-page summaries for all sites that appear to be landfills is generated. Each summary includes a satellite image, location of the site (coordinates, and location or street name as available), whether the site is known, site status (active, historical, undetermined, and an accompanying confidence rating), landfill class based on government classification and an accompanying confidence rating. Comments are also included about the changes to the site over time, approximate site size, vehicle presence, settling ponds, and the operating organization's name. The cataloging and risk characterization of these landfill sites provides the Ministry with the needed information to reduce landfill tax evasion and redirect time and financial resources to areas with a higher risk of illegal dumping (see Figure 2).

Figure 2: Risk Profiling of Identified Waste Sites (Air & Space Evidence)

Outcomes

ASE ran two pilots for the Ministry which covered four distinct areas in New Zealand. Approximately 55 percent of the potential sites identified by the detection model were classified as landfill sites while the other 45 percent were false positives. Out of those classified as landfill sites, about 60 percent were known to the Ministry. This 60 percent included both licensed and illegal sites that are going through or have gone through an enforcement process. The remaining 40 percent of sites were “new,” or unknown to the Ministry. The Ministry described about a third of the “new” sites as being “really useful” finds. The other two thirds of the “new” sites were smaller or inactive sites. These smaller or inactive sites may be issued warning letters or monitored closely for future changes. The pilots demonstrated that the remote sensing technology can successfully corroborate and add new information to the landfill registry. The Ministry is currently considering rolling out the satellite monitoring identification program nationally.

Resource Commitments

The pilots required the involvement of key personnel at the Waste and Resource Efficiency Division of the Ministry. Once the model identified potential sites, spatial analysts, Geographic Information System (GIS) experts, and the landfill licensing team collaborated to cross check the sites against the national registry and other local authority records. The support of local authorities in the process was also needed to verify and confirm the existence of the unlicensed waste site. The current cost to run the system is approximately USD31,275 – 37,530 per tile (110 x 110 km area). The cost includes staff time at the Ministry, ASE, and local governments. The total cost is dependent on image availability as well as the number of potential waste sites detected by the model.

Success Factors

Expertise - In addition to an AI specialist, it was critical to involve people with environmental expertise and satellite data expertise. Environmental specialist knowledge was key to training the AI model (i.e. environmental knowledge was built into the model to ensure appropriate identification of sites). Also important was the satellite specialist that helped access large satellite data that could be used for the application of AI techniques.

Partnerships – While it is acknowledged that there is some difficulty in establishing public-private collaborations particularly as it relates to technology, the success of the pilot can be attributed to a close collaboration between the Ministry and ASE. The working relationship succeeded based on the distinct needs of both parties. The Ministry needed a specific solution to the challenge of identifying unknown waste sites while ASE was in need of data to test and improve the model. Collaboration is critical to make progress in AI development and ultimately for the advancement of environmental governance.

Barriers

Executive Buy-In – In the case of the Ministry of Environment, New Zealand, there was a need for information; however, some environmental agency leadership teams were hesitant to obtain the information that is generated by ASE's model, due to constrained resources and funding. The information revealed means that agencies are required to act, through site inspections and enforcement actions. In response to this "fear of finding out," ASE has worked to develop a process to risk profile the data that is generated. This would allow officials to target sites with the highest risk and potentially most significant impacts, reducing the burden on an agency's resources and funding.

Funding – Adequate funding to test and pilot new technologies is a common barrier, particularly for government agencies. Agencies which have developed and tested a proof of concept, particularly as it relates to satellite-based AI models, have recognized the value of using technology to aid in compliance monitoring and verification; however, agencies are often constrained by limited budgets and securing funding to scale up these trials. An issue related to the challenge of securing funding is the isolation of pilot projects. Agencies are lacking an agency wide approach to testing and adopting new technologies such as AI. To help reinforce this approach, funding needs should also consider expanding its data teams to include advanced analytics capabilities.

Agency Opportunities

IAAC could benefit from the implementation of a satellite-based detection model during the compliance verification process. Satellite-based detection models, such as ASE's model, could be used to monitor permitted site boundaries and activities during both construction and operation, as well as identify unpermitted project construction activities. Doing so could flag potential instances of noncompliance without the need to be physically present on site. A detection model could additionally provide a risk profile for a particular region. Flagging noncompliance and risk profiling would assist IAAC in directing limited resources to high risk or noncompliant sites. This could be especially valuable as the quantity of decision statements and projects which require compliance monitoring increases and Agency resources are further stretched. A satellite-based detection model could also provide a basis for enforcement actions, deter future noncompliance, and prevent or minimize localized environmental damages by detecting noncompliance earlier. IAAC should consider many factors when determining whether to pursue a satellite-based change detection model:

Executive Buy-In - Given other agencies' "fear of finding out," IAAC will need to consider whether Agency leadership would want or use the additional data from a satellite-detection based model. Framing the model as a way to direct existing resources to the highest-risk sites could alleviate potential concerns related to uncovering additional noncompliance.

Funding - A continued funding source may be necessary for a satellite-based detection model to cover large areas and be regularly updated. This can be particularly challenging for environmental authorities, so IAAC would need to continue to consider how to incorporate AI funding into its budget. However, funding may be able to come from other sources as well. In Europe, there are funding schemes that provide money for tech solutions if an interested government agency partners with the company developing the solution. In these instances, the agency pays nothing or approximately 20% of the cost. Looking into alternatives such as this in Canada may be an option for IAAC. For example, the AI, Data and Robotics Partnership, which is one of the European Partnerships in digital, industry and space in Horizon Europe, provides funding to agencies, companies, universities, and other organizations on AI innovations.

Partnerships - AI, satellite, and environmental experts were key players in ASE's development of their model. In order for IAAC to pursue a satellite-based detection model, they should consider hiring an external company to run the model for them or hiring in-house experts to develop a model themselves.

4.2 Predictive AI for Inspection Targeting

Authority:	United States Environmental Protection Agency
Collaborator:	University of Chicago
Intent:	To identify noncompliance of hazardous waste facilities
Contact:	Randy Hill, Office of Enforcement and Compliance Assurance

The Challenge

The United States Environmental Protection Agency (EPA) is responsible for the regulation of approximately 1.2 million facilities nationwide under the ten major environmental statutes in the United States. However, EPA's resources to actively monitor the regulated facilities are limited. They currently have approximately 2,800 employees working in compliance and enforcement, and prior to the COVID-19 pandemic, they were only conducting about 10,000 on-site inspections of regulated facilities each year. Although states also conduct inspections, there are still many facilities that go uninspected each year, so noncompliance could continue unidentified. In an effort to address this resource gap, EPA is working to improve their targeting and to detect more instances of noncompliance.

Predictive Analytics AI Model

In 2015, EPA partnered with The University of Chicago Energy and Environment Lab (UChicago E&E), an academic lab that partners with policymakers to identify, test, and scale up solutions, policies, and programs to environmental challenges. The partnership between EPA and UChicago E&E aims to improve inspection targeting and increase EPA efficiency. Under this broad goal, one project is focusing on finding active hazardous waste site violators. To identify these violators, UChicago E&E uses predictive analytics to flag the facilities most likely to violate the Resource Conservation and Recovery Act (RCRA).

The RCRA predictive analytics model predicts "severe" violations, defined as storage without a permit, illegal treatment and disposal, and waste determination. The model generates a risk score for each facility that represents the likelihood that a facility inspection would find a severe violation of RCRA regulations. Prior to a nationwide rollout of the model, UChicago E&E worked with regional offices to conduct additional testing in the field. The goal of this testing was to demonstrate performance and increase inspector confidence in its value.

The model is a ML algorithm called a Random Forest, a multi-label Classification and Regression Tree model written in R and Python software and was built on 15 years of historical data, including tens of thousands of variables from regulatory reports. These variables included facility characteristics, such as location and industry, and historical compliance and enforcement data from both RCRA and other regulations such as the Clean Air Act. After variables and predictors were generated, the model was trained on historical data from the years 2000 to 2013. The model then predicted facility risk level in 2014 which was then compared to actual violations.

Outcomes

UChicago E&E and EPA have found that the model achieves a 50% increase in the detection of severe violations, in comparison to status quo targeting. Based on regional rollouts in 2017, they found that, if

EPA had used the model nationwide, they could have found an additional 214 severe violators with the same inspection resources. However, the model was only widely released to US EPA, state, local, and Tribal partners in August 2020, and there has not been sufficient time to evaluate its actual impact.

EPA and other partner agencies are looking into expanding the existing predictive analytics model to other regulatory acts, including inspection targeting under the Clean Water Act, Clean Air Act, and more.

Resource Commitments

The model required the involvement of key EPA staff, including Mike Barrette, Supervisory Environmental Protection Specialist at EPA; Rusty Wasem, Protection Specialist at EPA at EPA, and John Veresh, Information Management Specialist at EPA. Mr. Barrette is the Management Lead for UChicago Machine Learning Projects. The support of regional offices was also needed for pilot projects, and they provided input throughout the development process. Besides the staff time of the EPA team members who are not data scientists and other expenses, the model cost approximately 200 hours of a senior developer's full time equivalent, approximately USD 20,000, over the course of six months, in addition to an annual USD 2,400 in Amazon Web Services (AWS) hosting costs.

Success Factors

Stakeholder Involvement – Future model users were given the opportunity to provide feedback throughout the development process. This input helped to increase credibility and user confidence. EPA and UChicago E&E additionally worked with state agencies and state agency associations, the Environmental Council of the States and the Association of State and Territorial Solid Waste Management Officials, throughout the process to ensure the outcomes were accessible and communicated to the necessary stakeholders.

Data Availability – The model was built on 15 years of historical internal data from EPA.

Barriers

Data Understanding – There was a steep learning curve for data scientists at the start of the project. UChicago E&E had to invest significant resources and time into understanding the underlying data and ensuring the data were complete and of good quality. However, due to multisource funding at the university and various student contributions, the cost and time of this support are indeterminable.

Data Platform – After development, the model was transferred from UChicago E&E to EPA. It had to be adapted to run on EPA's AWS environment, which was a relatively new technology at EPA at the time. Therefore, there was a learning curve to successfully run it on EPA systems. Significant reworking was also required to scale up the model to handle five times the number of facilities than were included in the first field test.

Agency Opportunities

Predictive analytics models, such as EPA's model, could be used for inspection targeting and enforcement of regulatory acts. Doing so could increase IAAC's efficiency and detection of noncompliance. Predicting instances of noncompliance would assist IAAC in directing limited resources to facilities most likely to violate regulations. This would be especially beneficial as the number of projects IAAC is overseeing will likely drastically increase in the coming years, as the number of projects receiving approval increases significantly. IAAC should consider many factors when determining whether to pursue a predictive analytics model:

Data Availability – This model was built on 15 years of historical data and tens of thousands of variables. As IAAC's compliance, enforcement and follow-up program is relatively new, with limited amounts of historical data, this poses a significant barrier to developing a similar predictive analytics model in the short term. Therefore, this particular AI technology could be revisited several years down the line once more historical data exists. An alternative IAAC could look into is the possibility of getting data from other Canadian agencies. Historical data from other agencies may make predictive analytics model a feasible option in the nearer future for IAAC.

Partnerships – EPA partnered with an academic lab to provide the necessary AI expertise to develop the model, and the model was only integrated into EPA's AWS platform after the development phase. As IAAC currently does not have the necessary in-house AI expertise, they could pursue partnerships with academic groups, companies, and AI experts, and/or hire in-house experts, to assist in the research and development of such a model.

Stakeholder Involvement – Regional offices were included in the testing phase to ensure the model provided easily accessible outputs and to garner stakeholder support for the model. IAAC should consider how to involve stakeholders throughout the research and development and piloting processes to ensure effective model development and implementation.

Funding – Significant time and funding were needed to develop and pilot this model. The combined cost of the AWS host and model development by a senior data scientist at the EPA yields an initial cost of approximately USD 22400 for development. However, this project also received substantial support from the UChicago E&E team at an indeterminable cost. Further, soliciting, receiving, and incorporating stakeholder input required additional time and funding. Finally, EPA expects that upkeep of the project will include the AWS host cost of USD 2400 and 20 hr/year of work from a senior data scientist at EPA. Thus, IAAC would need to evaluate the cost benefit of a predictive analytics model, consider relationships with university partners, and secure continued funding for the project, either through their budget or elsewhere.

4.3 Supervised Classification AI for Complaint Processing

Authority:	Superintendencia del Medio Ambiente (Superintendency of the Environment), Chile
Collaborator:	University of Warwick
Intent:	To predict the relevancy and gravity of filed complaints
Contact:	Cristóbal De La Maza Guzmán, Superintendent for the Environment

The Challenge

The Superintendency of the Environment (SMA) in Chile is responsible for ensuring environmental regulations are complied with, and they undertake monitoring, follow-up, and inspection activities. They also receive thousands of citizen environmental complaints each year. However, SMA has limited resources to read through every complaint and focus their response accordingly. Further, some of the filed complaints are not relevant to SMA or do not contain sufficient information to be useful. Since SMA has limited resources to read through all the complaints and determine their urgency and relevancy, serious environmental damages could go unaddressed for extended periods of time.

Supervised Classification Model

In 2021, SMA partnered with the University of Warwick's "Data Science for Social Good" Initiative to develop a machine learning model to help prioritize resources when reviewing and responding to citizen environmental complaints so they can address serious environmental damages as soon as possible. The relevance model, a supervised classification model, classifies citizen environmental complaints for further review by staff. The model labels complaints as "relevant," "derivation," or "archive I." "Derivation" complaints and "archive I" complaints do not contain enough information for action to be taken. A second model, the Sanction Gravity Model, predicts the severity of sanctions for the complaints that are labeled as "relevant" to SMA. This further assists with inspection prioritization.

Both models are currently generating predictions for new complaints received by SMA, and the results are being collected in an Azure SQL database, which is a cloud database provided by Microsoft, for query. Eventually, these predictions will be used to make decisions about which complaints to follow up with, and in what order, but for now, these predictions are being compared to the decisions made by SMA staff.

The model was built on the data from SMA's registry of past complaints. Some of the data were structured, and other data were free-text complaints directly from citizens. The model also used additional data about social and geographical factors throughout Chile. A Random Forest model is used to analyze the data.

As shown in Figure 3, the model has been found to better predict "relevant" and "derivation" complaints than "archive I" complaints when compared to actual labels assigned by SMA staff. This is largely due to model construction, as "archive I" complaints are not inspected, and SMA determined it is preferable to eventually inspect unimportant complaints instead of never inspecting them. Further, the decline in "relevant" complaints is possibly due to a transition to a new system for online complaints in early 2021. Because the model was trained on the old system, the complaints they are now trying to predict may be very different from the ones the model was trained on. This is an area SMA plans to further investigate as they continue to develop the model.

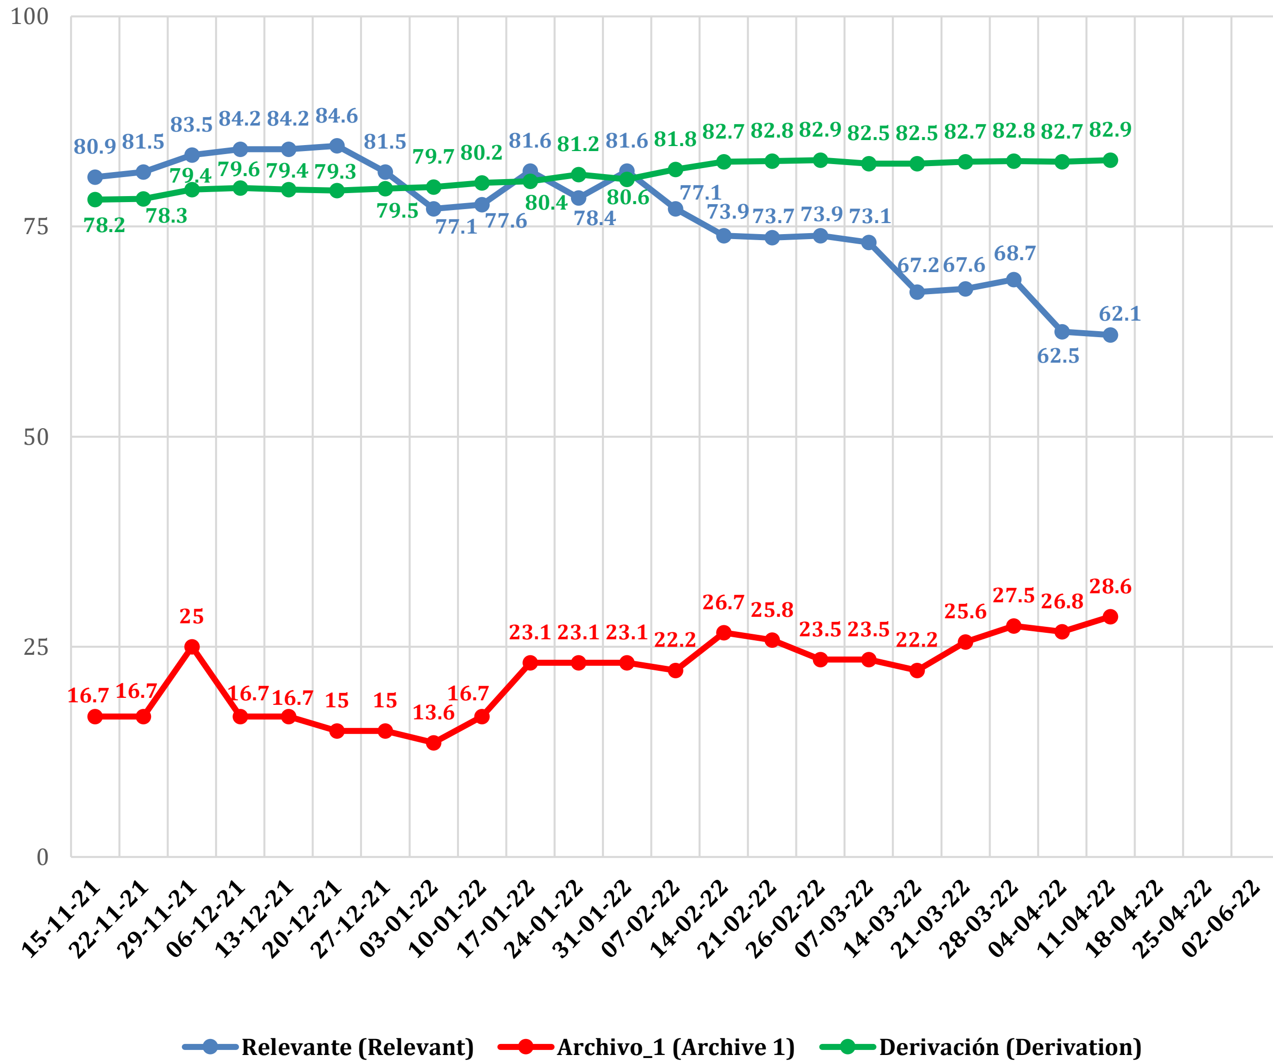

Figure X:

Outcomes

SMA estimates that the combination of the relevancy and sanction gravity models will speed up the time to redirect complaints to the proper agency by 80%, to archive complaints by 85%, and to inspect complaints with potential sanctions by 65%. This estimate is based on a comparison of the average complaint ending times before and after the contributions of the model. For example, their archive complaints take an average of 5 days to complete, and the model is estimated to reduce this time to 1 day, thus yielding an 80% reduction in time to completion. Ultimately, the models would increase SMA efficiency.

Resource Commitments

The models are currently run by one data engineer that manages the cloud infrastructure for the system. However, the University of Warwick's "Data Science for Social Good" Initiative developed the model before moving to SMA's cloud platform. Further, at least two data scientists will be involved in an upcoming model analysis to further explore the decline in accuracy shown in Figure 3.

SMA received a fellowship to have the Data Science for Social Good team support the development of the model, so the primary cost to SMA was paying the salary of their internal staff during project development, maintenance, and operation. SMA will also have to pay the salary for the incoming data scientists for the upcoming model analysis. Since the data used in this AI use is already collected and stored in SMA databases, costs for scaling up the project should be relatively minimal in the future.

Success Factors

Data Availability – The necessary data to develop the model was easily available in historical agency databases, and it was directly related to the problem SMA was trying to solve.

Internal Expertise – While external partners supported the project and were crucial, having in-house expertise in ML, data science, and data engineering were imperative for the implementation of the project.

Stakeholder Involvement – While SMA has yet to involve stakeholders in the process because the model is not yet ready for that phase, they know it will be key to get all the stakeholders involved to have the model be a success.

Barriers

Partnerships – Communicating effectively with external partners about the problem and the Agency's goals, in addition to taking the work the external partners did and integrating it within the institution's culture, was a challenge. Weekly team meetings with the external partners improved communication and helped involve the external team with SMA's culture, consequently improving the partnership and outcomes.

Agency Opportunities

A supervised classification model such as the one SMA is developing could be useful for IAAC in sorting and classifying public complaints for further inspection. Sorting comments based on relevancy and gravity could help increase efficiency and impact when responding to complaints and undertaking compliance and enforcement actions. IAAC should consider several factors when determining whether to pursue a supervised classification model for public complaints.

Data Availability – SMA built their model on an internal registry of past complaints. Since IAAC's compliance, enforcement and follow-up program is relatively new and currently oversees relatively few facilities, the historical data to build a supervised classification model may not exist at this time. IAAC could consider such a model several years down the road when they have more historical data. For reference, SMA initiated and trained their model using 9 years of data and several thousand complaints. However, as the parameters of their data collection change over time (i.e. a new complaint form), they expect to retrain their model.

Partnerships – SMA partnered with the Data Science for Social Good Team to develop their model. IAAC would need to consider whether they would develop a similar model in-house, in which case they would need to hire relevant experts to do so, or contract with an external party. If IAAC were to contract with an external party, they would also need to consider how best to communicate with them in order to align their goals and develop the model with IAAC's particular context in mind.

Expertise – SMA noted the importance of having both external partners to assist with model development and in-house expertise to implement the model. IAAC would need to consider contracting out or hiring in-house experts to run and maintain the model after development if they did not already have the necessary in-house experts.

Funding – SMA received a fellowship for their partnership with the Data Science for Social Good Team for model development, but they also paid the salary of in-house data scientists. IAAC would need to look into funding sources for both model development and continued maintenance and implementation of the model. Some funding may be able to come from external sources or pools, but IAAC may also need to consider incorporating continued funding of approximately USD 5900/year into its budget.

4.4 Predictive AI for Inspection Targeting

Authority:	Environment and Climate Change Canada
Collaborator:	N/A, developed internally
Intent:	To predict noncompliance and risk at regulated facilities
Contact:	Michael Enns, Director General, Risk Assessment at Environment and Climate Change Canada

The Challenge

Environment and Climate Change Canada's (ECCC) Enforcement Branch is tasked with investigating regulated facilities to ensure regulatory compliance. However, the Enforcement Branch only has the capacity to inspect a small fraction of facilities annually. The Office of the Auditor General found that the Enforcement Branch needed to better employ a risk-based approach to its activities, taking into consideration what facilities have the highest likelihood of noncompliance in conjunction with what activities pose the greatest risk to the environment, in order to more effectively mitigate noncompliance.

Predictive Analytics AI Model

ECCC's Chief Data Office and Enforcement Branch partnered with external data scientists to develop the Micro Enforcement Targeting Algorithm (META). The META model is a ML model designed to help the Enforcement Branch adopt a risk-based approach to site inspections. It predicts noncompliance at a facility level to help ECCC direct their resources to areas that are suspected to be of the greatest concern.

The META model is a supervised classification model that was built on more than 25 years of historical enforcement data and thousands of previous inspections. These data come from target entities through various mandatory regulatory reporting requirements and from publicly available financial datasets. Data were extracted, transformed, and loaded into a new data model, and business metrics from the Chief Data Office were incorporated. Specifically, the model is based on enforcement activities data, compliance promotion data, scientific (water, air, soil) data, industry data, financial data, and more. The model uses logistic regression, XGBoost which is a tree-based ensemble ML algorithm, and a deep neural network to predict noncompliance and risk by finding complex relationships between input features and the outcome of noncompliance.

At this point, many of the META model features rely on historical data and site visits. ECCC is continuing to add new historical data to the model to improve overall accuracy and derive and test new hypotheses. Eventually, ECCC aims to make a model that can extrapolate to sites that have not been visited by the Enforcement Branch.

Outcomes

The first generation of the META model is complete. In testing, it predicted facility-level noncompliance with three times more accuracy than previous targeting methods used by the Enforcement Branch. However, due to the COVID-19 pandemic, on-site inspections have been put on hold, so ECCC has not yet received a full set of results from their first “ground-truth” experiment.

Resource Commitments

ECCC's Chief Data Office and the Enforcement Branch were both involved in model development. However, the META model was developed primarily by a small group of data scientists who are now all with the Enforcement Branch.

Paying the salary of one to five full time data science employees is the primary cost of the META model. In addition, during development, two senior data scientists were required for full time work at approximately CAD 130,000/year of salary and operating costs. This work received additional support from three junior analysts at approximately CAD 90,000/ year and several students at CAD 20/hr for four-month periods. Finally, ECCC later invested approximately CAD 30,000 for computing power.

Success Factors

Executive Buy-in – The META model does not produce an immediate impact, so executive buy-in at IAAC is key. It gives the developers the necessary time and resources to develop and test the model.

Model Evolution – The cyclical process of generating predictions based on a set of indicators, validating those predictions through on-site inspections, and using the results from those on-site inspections to improve the META model is important. This iterative cycle allows for continued development and improvement to model accuracy and relevancy.

Barriers

Data Inconsistencies – Significant time was needed to properly standardize, clean, and format historical enforcement data. The data scientists had to develop methods to de-duplicate, correct errors, and link data across datasets to build the foundation for the META model. In particular, they had to invest significant time in correcting facility names and addresses. However, although data preparation was a resource-intensive process, it also helped the agency improve the accuracy of their reporting and intelligence products beyond the META model.

Agency Opportunities

Predictive analytics models, such as ECCC's model, could be used for risk-based inspection targeting at IAAC. This could increase detection of noncompliance and optimize the use of resources. As the number of projects IAAC is overseeing increases in the coming years, this could be particularly useful for identifying projects in noncompliance. When considering whether to implement a predictive analytics model, IAAC should consider several factors:

Data Availability – This model was built on more than 25 years of historical data and site inspections. As IAAC's compliance, enforcement and follow-up programs is a relatively new agency, with limited historical data, this poses a significant barrier to developing a similar predictive analytics model. Therefore, this particular AI technology may need to be revisited several years down the line once more historical data exists. An alternative IAAC could investigate is the possibility of getting data from other Canadian agencies. Historical data from other agencies may make predictive analytics model a feasible option in the nearer future for IAAC.

Executive Buy-In – As a predictive analytics model takes time to develop, test, and implement, having executive buy-in was key for ECCC when pursuing the META model. Buy-in provided data scientists with the time and resources necessary to develop the model. IAAC should consider whether the necessary executives would be on board with a predictive analytics model.

Expertise – ECCC had between one and five data scientists working full time to develop and test the META model, and those scientists are now part of the Enforcement Branch at ECCC. IAAC would need to determine who has the necessary expertise to develop a similar predictive analytics model, and whether that expertise is already present in-house, if they would need to hire a new employee(s), or if external experts could help with model development. IAAC should also consider any data standardizing and formatting the expert(s) would need to do when they identify potential experts. Having executive buy-in would also be key for hiring the necessary experts.

Funding – Significant time and funding were needed to develop this model. Specifically, the team was running at a cost of approximately CAD 555,000 a year for staff and funding. IAAC would need to evaluate the cost benefit of a predictive analytics model and secure continued funding for the project, either through their budget or elsewhere.

4.5 Intelligent Document Processing for Compliance Rating

Authority:	Health Canada
Collaborator:	N/A, developed internally
Intent:	To assign risk ratings to instances of noncompliance and assist with report writing
Contact:	Justin Budgell, Senior Compliance Officer

The Challenge

Part of Health Canada's responsibilities include compliance monitoring and enforcement activities related to health products. As part of these responsibilities, inspectors at Health Canada in the Regulatory Operations and Enforcement Branch carry out site inspections to ensure facilities are in line with the *Food and Drugs Act*. They also write and publish a final inspection report for each facility. The Branch inspects and writes reports for a large number of sites annually. Historically, writing the final inspection report for each site was done manually, including assigning a risk rating and matching a specific part of the *Food and Drugs Act* to each instance of noncompliance at a site. This manual process left room for discrepancies and human and interpretation errors, in addition to requiring significant amounts of employee time.

Intelligent Document Processing for Compliance Rating

To standardize inspection reports and save inspectors time, Health Canada began to develop a ML algorithm, Cipher, in 2019. Cipher automates the processes of assigning a risk rating to each instance of noncompliance at a facility and matching each instance to a particular regulation within the *Food and Drugs Act*. It also assigns a standard line – a generic explanation of a particular type of noncompliance that can be published publicly in the final inspection report – to each instance of noncompliance. Health

Canada is also working to expand Cipher to include a predictive analytics component and predict whether a site will be compliant with regulations.

Cipher was built on 10 years of historical, internal data. During the first phase, data were extracted from inspection reports and analyzed for patterns to determine whether they could be useful in both automating processes and predicting noncompliance. A proof-of-concept tool determined that AI and ML could successfully assign a risk rating and match a regulation within the *Food and Drugs Act* to a single observation of noncompliance. During the second phase, Cipher was expanded to enable inspectors to enter all observations from a site visit at one time, and to add the standard line for each instance of noncompliance.

Inspectors tested and provided feedback on this second phase, and now Cipher is entering its second round of testing.

Outcomes

At this time, the outcomes of Project Cipher are unknown as it is still in the testing phase.

Resource Commitments

The project started as an internal partnership between branches of Health Canada. Key staff include Justin Budgell, Project Lead; Peter Yoon, Project Advisor; Betty Palma, Inspector; Valerie Bergeron, HPIL Project Champion; Cecilia Bong, HPIL Project Manager, Bryan Paget, Data Scientist/Developer; Mithu Selyakumar, Inspector; and Sherry Bahaw, Inspector.

Health Canada handled the data extraction internally but hired a third-party contractor for aspects of the data analysis phase and to move Cipher to a protected B Cloud environment, which was necessary to protect classified information. Health Canada is currently in the process of hiring another data scientist to continue building the tool. Health Canada also has a Memorandum of Understanding with the Canadian School of Public Service (CSPS). They are invited to sit in on project meetings, and Health Canada learns about other AI projects CSPS have learned about.

The cost of developing, testing, and implementing the Cipher model was estimated to total approximately CAD 500,000.

Success Factors

Tool accuracy, building user trust, resource availability, and technology availability are factors Health Canada sees as important considerations for making Project Cipher a success. As Project Cipher is still in the testing phase and outcomes are unknown, the impact of each particular factor is also still unknown.

Barriers

Data Platform – Needing to store the data in a protected B Cloud environment while still allowing multiple users access at a time has been delaying the project pilot and next steps. To overcome this barrier, Health Canada partnered with Environment and Climate Change Canada (ECCC) to access their cloud environment for testing the model. ECCC's cloud environment does not require the same security restrictions that Health Canada's does.

Expertise – Health Canada needed to hire qualified data scientists in order to develop and test the tool. Hiring employees with the necessary expertise in a timely fashion was a challenge for Health Canada; however, they were ultimately able to hire internal experts and primarily develop the tool internally.

Agency Opportunities

Intelligent Document Processing models, such as Health Canada's model, could be used for follow-up activities such as processing follow-up program reports from project proponents at IAAC. As the Agency is tasked with writing follow up reports based on proponent's follow-up program results, a model such as Cipher could be particularly useful in standardizing and streamlining document processing and write ups. When considering whether to implement such a model, IAAC should consider a couple factors:

Data Availability – Cipher was built on 10 years of historical data and reports. As IAAC's compliance, enforcement and follow-up program is relatively new, they may lack the necessary historical data and reports to build an accurate ML model that would help with report writing and risk assessment. As a result, this technology may be more accessible to IAAC in the future once more historical data has been amassed.

Expertise – Health Canada hired several data scientists to develop Cipher, as well as collaborating with a third party for data analysis and platform work. IAAC would need to consider hiring data scientists and/or contracting third-party experts in order to develop a model like Cipher as the Agency currently does not have the necessary in-house expertise or capacity.

Funding – Rather than engaging external partners, Health Canada depends on internal experts on data science, inspection, and other issues. This could reduce certain cost of development.

4.6 Summary of Use Cases

Use cases highlighted several factors that either positively or negatively affected the development and implementation of an agency's AI model. While each agency's situation is unique and thus requires different considerations, many factors were relevant to more than one agency. The use cases revealed that funding, executive buy-in, expertise, the data platform, partnerships, and data understanding/inconsistencies posed barriers to particular agencies working to develop AI models (Table 1).

Use Case	*Barriers*					
	Funding	*Executive Buy-In*	*Expertise*	*Data Platform*	*Partnerships*	*Data Understanding/ Inconsistencies*
3.1 ASE Change Detection	Needed for testing and piloting phases	Agencies not wanting information from ASE				
3.2 EPA RCRA Model				Learning curve to run system on EPA platform and scale it up		Significant resources invested to understand historical data
3.3 SMA Supervised Classification					Communicating effectively and integrating agency culture	
3.4 ECCC META Model						Significant resource invested to standardize, clean, and format historical data
3.5 Health Canada Cipher			Needed to hire internal data scientist experts	Required security on platform delaying pilot		

Table 1: Barriers to AI development and implementation from agency use cases.

The use cases also revealed that executive buy-in, model evolution, data availability, partnerships, expertise, and stakeholder involvement were factors that enabled successful AI model development at particular agencies (Table 2).

Use Case	*Success Factors*					
	Executive Buy-In	*Model Evolution*	*Data Availability*	*Partnerships*	*Expertise*	*Stakeholder Involvement*
3.1 ASE Change Detection				Close collaboration between agency and ASE team	AI, environmental, and satellite data experts involved	
3.2 EPA RCRA Model			15 years of historical data for building the model			Future users provide feedback through development process
3.3 SMA Supervised Classification			Data from historical, internal database		Internal expertise key for implementation	Will be imperative in the future
3.4 ECCC META Model	Executive buy-in to provide necessary project resources	Iterative cycle of on-site inspections and model improvement				
3.5 Health Canada Cipher			Resource and tool availability			Building user trust

Table 2: Success factors for AI development from agency use cases.

Further, the use cases revealed that generally, agencies typically relied on a combination of in-house employees and external experts to develop and implement their AI models (Table 3). Money commitments for use case projects generally included salaries for employees and funding for long-term model maintenance and monitoring. Agencies also sometimes paid platform hosting costs and external partners, but this varied between use cases and between funding schemes.

Use Case	*Resource Commitments*	
	Personnel	*Money*
3.1 ASE Change Detection	Agency employees, ASE staff, local authorities	CAD 38,829-46,595 GBP per 110x100 km area
3.2 EPA RCRA Model	Agency employees, academic partner staff, regional EPA offices	200 hours FTE of senior developer, CAD 3,086 AWS annual hosting costs
3.3 SMA Supervised Classification	Agency data engineer, external partners, two Agency data scientists	Fellowship for partner group, salary for internal experts
3.4 ECCC META Model	ECCC Chief Data Office, Enforcement Branch, one to five ECCC data scientists	Data scientists' salaries, junior analyst salaries, and CAD 30,000 for computing yields approx. CAD 555,000
3.5 HC Cipher	Eight internal staff, third party contractors for data analysis and platform	CAD 500,000

Table 3: Resource commitments for AI models from use cases.

As IAAC and other environmental agencies pursue AI projects, they should consider how these factors could affect AI research, development, and implementation at their particular agency.

4.7 EnviroVerse – The Future of Regulatory Monitoring

As interest, funding, and skills continue to grow, the environmental AI sector will continue to advance. These advancements could include larger-scale predictive analytics algorithms, increased efficiency, and more. One such advancement may be the ability to conduct site inspections virtually instead of on-site. Enviro.AI, a company focused on data technology and environmental compliance, is building the EnviroVerse, a layer of the metaverse that would allow enforcement officers to carry out inspections virtually (Figure 4). Enviro.AI also notes the EnviroVerse could be used for monitoring, compliance, permitting, auditing, and more. The platform is still in development and is several years from being released but provides an example of the innovative environmental AI work that may be seen in coming years.

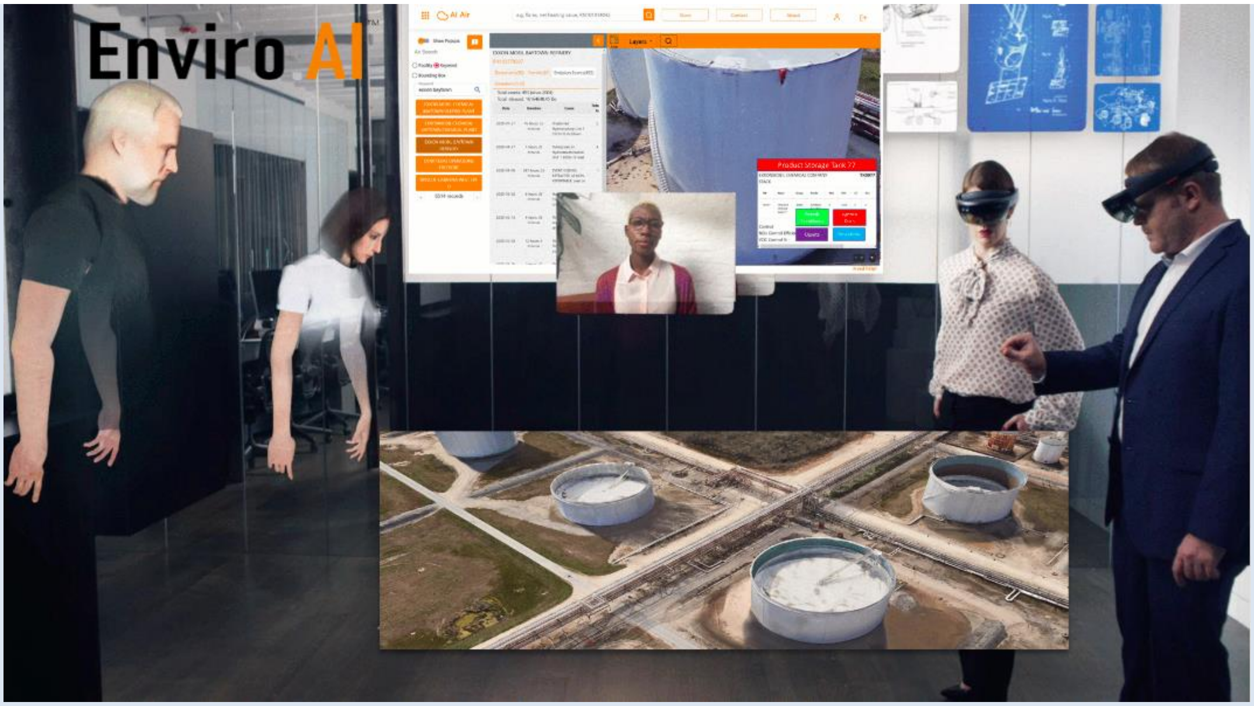

Figure 4: Promotional material for the EnviroVerse by Enviro.AI.

5 ADOPTION CRITERIA

As IAAC considers how AI and ML models may help increase the efficiency of their post decision activities and manage their growing workload, IAAC must evaluate potential applications in the context of their own circumstances. These considerations include the fact that IAAC's compliance, enforcement and follow-up program is relatively new, they are currently conducting post decision activities at numerous facilities but have many more decision statements being issued in the coming years; they are a small team; they have little internal, historical data; they are currently working to designate part of their budget to AI initiatives but do not currently have a continuous funding source for AI projects; and they do not currently employ in-house data scientists or AI experts.

With these circumstances in mind, the Agency should consider particular criteria prior to and throughout the research, development, and implementation of an AI model. These factors are derived from the use cases (Table 4), and the literature review.

Use Case	*Agency Opportunity Considerations*					
	Funding	*Partnerships*	*Data Availability*	*Stakeholder Involvement*	*Executive Buy-In*	*Expertise*

3.1 ASE Change Detection	*X*	*X*			*X*	
3.2 EPA RCRA Model	*X*	*X*	*X*	*X*		
3.3 SMA Supervised Classification	*X*	*X*	*X*			*X*
3.4 ECCC META Model	*X*		*X*		*X*	*X*
3.5 HC Cipher			*X*			*X*

Table 4: Agency opportunity considerations from the use cases.

Data Availability – The basis for an AI project is the data, and while not all AI projects would require an agency to have a historical database, many AI projects require historical data to build the model. Further, available data need to be reliable, of good quality, and representative [U.S. Government Accountability Office, 2021]. They need to be organized and stored in an accessible database, with sources noted. Ideally, data should be integrated and accessible in a centralized platform instead of in discrete data dumps, and there should be agency staff dedicated to upholding data governance and quality (GoDataDriven, n.d.). Data analysis may need to be undertaken to ensure historical data meet these factors prior to being used for AI model development.

This need for quality data is something IAAC will need to seriously evaluate as they consider pursuing AI models, as it will likely influence what types of projects are feasible now, and what projects may need to be revisited down the line. As a relatively new program, they will likely need to focus on AI projects that do not require robust historical data, such as ASE's satellite change detection model. Further, IAAC may need to assess its current historical data to determine whether it meets quality and other concerns. If it does not, it could be worthwhile to focus on expanding or adjusting their data management capacity so data they are collecting now could be used in the future.

Expertise – AI experts, data scientists, and other experts, depending on the project type, are crucial for researching and developing a successful AI system. These experts can be employees at an agency, or they can be external experts contracting or partnering with the agency on a particular AI project. As IAAC currently does not have the in-house expertise to research, develop, and implement an AI model, IAAC will need to partner with external experts and/or hire staff with the necessary expertise. When deciding who to include on the AI project team, IAAC will also need to consider who will be able to operate and maintain the AI system once it is implemented. They should consider hiring senior experienced AI professionals who can help ensure AI systems are operating properly and who can assist in the training of other staff (GoDataDriven, n.d.).

Generally, when an agency is just starting its AI journey, they rely more on external experts because they do not have the capacity, funding, or resources to do it in-house. However, as the agency will likely be maintaining the project once it is fully implemented, there should be a knowledge transfer about the system from external parties to agency staff, so they are not completely reliant on a third party (GoDataDriven, n.d.). If IAAC hires a third party to research and develop an AI system, they should ensure staff learn the system so they can run system maintenance.

Partnerships – As mentioned above, when it comes to expertise, partnering with third parties can be one pathway for researching and developing an AI model. Partnering with an academic lab or AI-focused organization that already has the necessary AI experts involved can help ease the development of an AI system for an agency with limited resources. This could be useful for IAAC, as they do not currently have the in-house expertise to develop an AI model. However, it is imperative for an agency to communicate well with their partners. Developing methods of communication that work for both parties – such as weekly meetings – would be crucial to ensure the AI model met IAAC's and their stakeholders' needs and fit within their agency's culture.

Governance – Governance considerations should also play into the adoption of AI systems. The AI application and associated data processes must be legally compliant. Organizations pursuing AI adoption must be committed to transparency, meaningful stakeholder input, and eliminating system biases. Further, a process to manage AI project implementation that includes clear responsibilities and incorporates diverse perspectives and stakeholders to mitigate risks must be established [U.S. Government Accountability Office, 2021].

To this end, the Government of Canada has developed guiding principles and an Algorithmic Impact Assessment to determine how acceptable AI solutions are from an ethical and human perspective (citation). IAAC should consider governance structures and refer to the Algorithmic Impact Assessment for guidance as they pursue any AI system.

Stakeholder Involvement – As a government agency, IAAC is tasked with supporting sustainable development, which involves coordination with stakeholders throughout Canada. Any AI system implemented by the Agency would have an impact on agency staff and the public, and so involving relevant stakeholders throughout the development and testing phases of an AI project would be imperative.

Providing opportunities for feedback, for Canadians and future model users, could help increase credibility and ease user adoption down the line. It could also help ensure AI model outputs are accessible and useful. Further, ensuring diverse perspectives are incorporated throughout model development can help to mitigate risks and address concerns such as ethics and biases [U.S. Government Accountability Office, 2021].

User Adoption – User adoption considerations refer to issues IAAC may encounter in the adoption process. For example, how easy would it be for staff to use the application? How much training would be needed to achieve proficiency with the new technology? Do staff support the application? IAAC should consider these questions as they pursue any AI project to ensure their staff is on board and that the continued implementation of a project is feasible (GoDataDriven, n.d.). Further, employee concerns about job loss related to AI adoption means IAAC must be prepared to introduce and frame AI solutions in a manner that emphasizes how they will improve efficacy and efficiency and not take jobs away from employees [Chenok, 2018].

Executive Buy-In – Having executives and leadership teams on board with AI is crucial for successful AI project implementation. Leadership support can affect an agency's ability to hire or contract with experts, secure funding, access data, and integrate the AI platform into the agency. Thus, IAAC needs to ensure that Agency leadership are on board with AI more generally, and with any particular AI project IAAC staff plan to pursue. Communicating with executives from the outset is of utmost importance. Further, leadership can clear roadblocks to AI projects, support larger-scale AI uptake or cross-organizational initiatives, and can help secure funding (GoDataDriven, n.d.).

Funding – As funding can be a major barrier to AI implementation, especially at government agencies, IAAC must consider the cost of AI. The Agency would need to budget for the initial research, development, and implementation of each AI model, as well as funds to support recurring operations and maintenance. Consequently, the budget needs to consider the cost of above criteria and include funding for partnerships with external experts; data collection, storage, analysis, and more; stakeholder involvement; hiring and paying in-house experts; user adoption; and governance considerations, amongst other factors. Given the need for continued funding, IAAC should continue to pursue incorporating dedicated AI funding into its budget, or establishing other continual funding sources, such as fellowships, grants, or tech funding schemes.

IAAC should also consider the cost of a particular AI model type when choosing whether or not to pursue it. As illustrated by the use cases, costs vary between model type and agency, so IAAC should assess what types of projects may be feasible with their given funding.

Legal considerations and limitations – The use of AI in post-decision processes could meet legal limitations. In 2021, the European Commission published a draft AI Regulation under which AI systems used for law enforcement are defined as high-risk AI systems. The draft regulation subjects high-risk AI systems to various requirements such as the assurance of quality of data sets used to train the AI systems, application of human oversight, creation of records to enable compliance checks and provision of relevant information to users. Such regulation is broadly believed to take years to be implemented. However, it is likely that countries will apply legal requirements to law enforcement related AI uses and emphasize transparency, fairness and explainability.

6 NEXT STEPS

Prior to pursuing the research, development, and implementation of an AI model to assist with post decision activities, there are steps IAAC needs to take. A maturity model can be used to create a roadmap for the implementation of AI at the Agency. [Panetta, 2019] GoDataDriven, a data and AI consultancy and training company based in Amsterdam and operating since 2009, lays out a maturity model that can be adapted and used by agencies to guide their work on AI. Broadly, an agency needs to consider both their analytical capability and business adoption as they determine next steps. Analytical capacity refers to data organization and accessibility, employee capacity and skills, and technology and tools. Business adoption refers to the degree to which AI has been embedded at the organization as measured by executive support, funding, and implementation. As IAAC considers the adoption of AI, it needs to consider all of these facets.

To help with this consideration, developing a strategic plan for AI across IAAC would be a good first step. Many agencies have tested and piloted projects in silos and are now realizing the need for an agency-wide approach or strategy in order to continue testing, piloting, and implementing AI models. Often, governments are getting stuck after running siloed pilot projects and are realizing at-scale AI has different requirements [Van Buren, 2021]. Thus, developing a strategy that incorporates considerations such as executive support, in-house expertise, data, ethics, technology platforms, and partnerships, can help an agency more effectively pursue agency-wide AI. Developing this strategic plan prior to AI model development could prevent IAAC from hitting roadblocks down the line and also help to make AI projects more easily scalable, efficient, and impactful at the Agency. It could also provide a roadmap for addressing the criteria noted above, including funding, expertise, data, governance, and more, so IAAC should pay particular attention to thinking through and addressing those criteria in the strategic plan.

As IAAC is just beginning its AI journey, according to the GoDataDriven maturity model, the Agency is in the initialization phase of AI development. The initialization phase involves identifying AI opportunities; preparing data, people, and technology for the development of the first AI model; and launching the first AI model. Right now, IAAC recognizes that AI could be valuable, and they employed ELI to identify use cases, but they do not have in-house AI experts, may not currently have the historical data or necessary funding to scale up AI at the Agency, and have not yet launched an AI model. After developing a strategic plan, the steps in the initialization phase should be IAAC's next focus.

Eventually, IAAC will need to take steps to progress to the next phase on the maturity model, continuous experimentation. At this point, IAAC will need to focus on recruiting in-house experts, creating a robust infrastructure for AI, and continuing to adopt AI models beyond those created during the initialization phase.

In sum, this report identifies potential uses of AI that could be employed by IAAC and other environmental agencies during compliance & enforcement and follow-up activities. It identifies six AI use cases across Canadian Agencies, the United States, and other countries that may be relevant to IAAC's post decision activities. It also recommends next steps that IAAC could take to incorporate AI into their work to help increase agency efficiency and efficacy.

7 REFERENCES

Chenok, D. (2018) How can AI help government improve? https://gcn.com/2018/11/how-can-ai-help-government-improve/300028/

Deloitte AI Institute. (2021). The Government & Public Services AI Dossier. https://www2.deloitte.com/us/en/pages/consulting/articles/ai-dossier-government-public-services.html

GoDataDriven (2019). The AI Maturity Journey: Learn what it takes to become an AI-driven organization. https://godatadriven.com/topic/whitepaper-ai-maturity-journeys/

Greenstone, M. (2018) Using Predictive Analytics to Improve Inspection Targeting. https://www.ecos.org/wp-content/uploads/2019/10/UChicago-Slide-Deck.pdf

Martinho-Truswell, E. (2018) How AI Could Help the Public Sector. https://hbr.org/2018/01/how-ai-could-help-the-public-sector

Nalbandian, L. (2021) Canada should be transparent in how it uses AI to screen immigrants. https://theconversation.com/canada-should-be-transparent-in-how-it-uses-ai-to-screen-immigrants-157841

Panetta, Kasey (2019), CIOs can separate AI hype from reality by considering these areas of risk and opportunity. https://www.gartner.com/smarterwithgartner/the-cios-guide-to-artificial-intelligence

Shubhendu, S., & Vijay, J.F. (2013). Applicability of Artificial Intelligence in Different Fields of Life. https://www.semanticscholar.org/paper/Applicability-of-Artificial-Intelligence-in-Fields-Shubhendu-Vijay/2480a71ef5e5a2b1f4a9217a0432c0c974c6c28c?p2df

U.S. Government Accountability Office (2021). Artificial Intelligence: An Accountability Framework for Federal Agencies and Other Entities. https://drive.google.com/file/d/13dJ7bYt6Dd5RAcoqeHAt4n9iHXhBXd9D/view

Van Buren, E., Eggers, W. D., Austin, T., Mariani, J., & Kishnani, P. K. (2021) Scaling AI in Government. Deloitte. https://www2.deloitte.com/us/en/insights/industry/public-sector/government-ai-survey.html