

WHAT IS GENERATIVE AI, HOW DO WE ETHICALLY UTILISE IT AT COA AND WHAT ARE THE RISKS?

2024 CoA Graduate Project Report

16 April 2025



CITY OF
ADELAIDE

ACKNOWLEDGEMENT OF COUNTRY

The City of Adelaide acknowledges that we are located on the traditional Country of the Kaurna people of the Adelaide Plains and pays respect to Elders past, present and emerging.

We recognise and respect their cultural heritage, beliefs and relationship with the land. We also extend that respect to visitors of other Aboriginal Language Groups and other First Nations.

DOCUMENT PROPERTIES

Contact for enquiries and proposed changes

If you have any questions regarding this document or if you have a suggestion for improvements, please contact:

Contact Officer: Coordinator – Workforce Development
Title: People Experience Team
Program: People
Phone: (08) 8203 7
Email: @cityofadelaide.com.au

Record Details

HPRM Reference:

HPRM Container:

Version History

Version	Revision Date	Revised By	Revision Description
A	April 2024	AH, CZ, CA, DSK, LP	Graduate Project 2025 Draft

TABLE OF CONTENT

ACKNOWLEDGEMENT OF COUNTRY	i
Document Properties	i
Table of content.....	ii
Executive Summary.....	1
1. Introduction	2
2. What is Generative Artificial Intelligence.....	3
3. Ethics and Environment.....	11
4. Risk vs Opportunity for City of Adelaide	13
5. CoA Stakeholder Engagement.....	18
6. CoA use Cases.....	24
7. Recommendations	26
8. Conclusion.....	27
References.....	28
Appendix A: Proposed Operating Guideline.....	30
Appendix B: Survey Questions.....	31
Appendix C: Survey RESULTS	32

FIGURES

Figure 1: Survey Results - Use of Gen AI at CoA	1
Figure 2: Use Case Table	7
Figure 3:Engagement Against Use Case	8
Table 1: Risk vs Mitigating Action	15
Figure 4: Question 8 What are the potential benefits of using Gen AI at the City of Adelaide?	20
Figure 5: Question 9 What are the potential risks of using Gen AI at the City of Adelaide?	20
Figure 6: Question 6 How often do you use Gen AI at work?	21
Figure 7: Question 7 How often do you use Gen AI outside of work?.....	21
Figure 8: Question 10 How could CoA support you in using Gen AI safely?	22
Figure 9: Question 11 Do you see any potential use cases for Gen AI at the City of Adelaide?	22

EXECUTIVE SUMMARY

Generative Artificial Intelligence (Gen AI) has the potential to revolutionise the way CoA operates and can unlock a wide range of benefits and innovative ways of working. This report undertaken by the 2024 CoA Graduates addressed the research question: “What is Generative AI, how do we ethically utilise it at CoA and what are the risks”.

Generative Artificial Intelligence is an emerging technology that allows computer model to create human like content such as text, images, report and more. Gen AI models are trained on large dataset of data which allows it to identify patterns and relationships and generate new content. By adopting an internal Generative AI model, the CoA can experience a wide range of benefits, so long as the associated risks are properly managed and mitigated. This report delves into the risks of Gen AI, ranging from Ethical and Environmental risks and considerations to Information Technology and legal risks and various methods of mitigating the risks.

For example, Gen AI does have an inherit risk associated with data security and privacy due to the risk that a model could potentially leak sensitive information it was trained on or has access to. Furthermore, Gen AI can be the target of cyber-attacks which can lead to several consequences. A few mitigating measures which can combat these risks include using a fully vetted model by CoA IT, validation of outputs, monitoring the performance and providing training to staff who intend to use Gen AI at CoA.

A survey of the CoA administration was conducted as part of the research for this report to gauge the general sentiment towards Gen AI, investigate the current usage and appetite for CoA to adopt an internal Gen AI model. This survey found that the staff generally had a strong interest in Gen AI at CoA, with several staff currently using Gen AI in their day-to-day work. It is worth noting that some staff had reservations about using Gen AI relating to the associated risks, lack of confidence and the requirement for additional training and policies. Figure 1 below shows the results to the question ‘How often do you use Gen AI at work?’ which shows about 35% of respondents use Gen AI either daily or weekly at CoA.

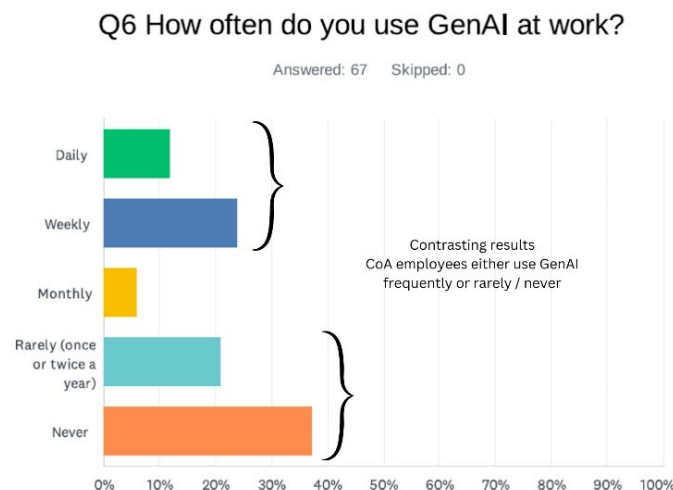


Figure 1: Survey Results - Use of Gen AI at CoA

This report also explores the feasibility of several use cases relevant to the City of Adelaide such as customer service and general administrative tasks. and provide several recommendations including implementation of Gen AI policies, an Operating Guideline and adoption of a pilot AI program. A draft Operating Guideline is provided in Appendix A as an example of how this document may look. Through taking responsible actions in implementing Gen AI at CoA, there are a wide range of benefits which could improve the way CoA operates and increase its productivity.

1. INTRODUCTION

The City of Adelaide (CoA) graduate cohort have prepared the following report as part of the 2024 Graduate Program at the City of Adelaide. This graduate cohort started at CoA in January 2024 and includes the following members:

- Anh Hoang – Graduate Marketing & Communications
- Chris Zhang – Graduate Infrastructure
- Clayton Armour – Para Planner
- Darvesh Singh Kukreja – Graduate Economic Research & Data Analyst
- Lachlan Pedder – Graduate Infrastructure

This project was assigned to the graduates to promote collaboration across the organisation and to foster new connections within the CoA Administration. This project was assigned to the Graduate cohort on May 10, 2024, with the goal of the group delivering a project report and presentation to the CoA Senior Leadership Team (SLT). Another objective of the project was for the Graduate cohort to undertake relevant research in an area of interest for the City of Adelaide.

Generative AI (Gen AI) was identified as an area where CoA may need to take a stance on in the near future in order to keep up with developing technologies and to follow through on its organisational values of Innovate and Achieve. This report aims to explore the following research question:

“What is Generative AI, how do we ethically utilise it at CoA and what are the risks?”

This research question is as relevant as ever, due to the rapid adoption of Gen AI across many organisations and industries as well in the everyday persons day-to-day life. To gather information on this topic, a survey was sent to the entire Administration to gauge the sentiment and thoughts of the staff at CoA. By investigating how this new technology may be adopted by CoA, what risks may be associated with it, and how the administration currently views Gen AI, this report aims to educate others at the City of Adelaide and hopefully become a baseline for future policy, guidelines and/or governance. This report also aims to touch on ethical and environmental issues and potential use cases of this new and upcoming technology.

2. WHAT IS GENERATIVE ARTIFICIAL INTELLIGENCE

Artificial Intelligence (AI)

Artificial Intelligence, or AI, is the technology that allows machines to simulate human-like intelligence. It uses data and algorithms to perform tasks like analysing information, recognising patterns, and making decisions. Think of it as a digital assistant that can process massive amounts of data quickly and accurately, far beyond human capacity.

Generative AI (Gen AI)

Generative AI is a specific type of AI that does not just analyse or interpret data, it creates new content. Using advanced models, it generates text, images, music, or even complex designs. It is like having a machine that can brainstorm, write, and design based on the instructions you give it.

How Does Gen AI Work?

Gen AI uses models trained on large amounts of data to understand patterns and relationships. Popular models like ChatGPT or DALL·E can produce human-like responses or visuals because they have been trained on diverse datasets ranging from books to images. Imagine teaching a robot by showing it thousands of examples of a task—like how a flyer is written or what a park design looks like. Once trained, it can produce its own version, tailored to your specific request.

Broad Uses and Additional Context

Gen AI has applications across many areas of local government. It can draft policy documents, summarise meeting minutes, design community engagement materials, or even generate visualisations for urban planning. It is also being explored for translating complex reports into plain language for broader community understanding. While it excels in creativity and productivity, it is essential to pair it with human oversight to ensure outputs are accurate and culturally appropriate.

Generative AI Compared to Other Forms of AI

Gen AI is distinct because it creates new content that can closely resemble human created content, whilst other forms of AI focus on tasks like analysis, prediction, or automation. For example, a predictive AI system might forecast future traffic patterns based on current data, whereas Gen AI could create a draft transport plan based on those predictions. Both are powerful tools, but Generative AI adds an extra layer of creativity, making it especially useful in scenarios where fresh ideas or designs are needed.

Alignment with City of Adelaide Strategic Plan

Adopting Gen AI technologies, including tools like ChatGPT and Microsoft Copilot, stipulates alignment with the City of Adelaide's Strategic Plan 2024–2028, supporting its goals of economic growth, efficiency, and innovation.

As the LGA aims to be a hub for innovation and investment, attracting businesses and talent, Gen AI can support this by automating administrative and routine tasks, improving data analysis, and enhancing decision-making, helping businesses and government services operate more efficiently. AI-powered insights can also support economic planning, urban development, and policy decisions, ensuring Adelaide remains a competitive and future-ready city.

The Strategic Plan prioritises a high-performing, customer-centric organisation. GenAI can streamline internal processes, reducing time spent on manual tasks like report writing and data entry. AI-driven chatbots and virtual assistants can improve customer engagement, providing quick, 24/7 access to information and services. Tools like Microsoft Copilot can enhance productivity, transparency, and communication, ensuring the Council operates efficiently and effectively. By embracing GenAI, the City of Adelaide can deliver smarter, faster, and more accessible services, reinforcing its reputation as a modern and innovative capital city.

Case Studies

Australian Government Trial of Microsoft 365 Copilot (Australian Government, 2024)

The Australian Government undertook a trial of the Microsoft 365 Copilot Generative Artificial Intelligence across a 6-month period starting at the start of 2024. The trial was prompted by the rapid uptake and increase in availability of AI including Gen AI across Australia and the world. The trial included distributing over 5,765 Microsoft Copilot license to interested Australian Public Sector (APS) employees. Microsoft Copilot allowed the APS employees to use Generative AI in their job which integrated with the Microsoft Office suite (Word, Excel, OneNote etc). The Government claims that Copilot allows for safe and responsible use of Generative AI as the model is created and ran by Microsoft and all data is encrypted and suitably stored (Microsoft 2024).

A comprehensive report was compiled following the completion of the trial which was prepared by the Australian Government and Nous Group (Australian Government, 2024). This report included the results of a variety of quantitative and qualitative data with a focus on the following four objectives:

- Employee related outcomes
- Productivity
- Whole-of-government adoption of generative AI
- Unintended outcomes

Some limitations included evaluation fatigue, non-randomised participation and inconsistent rollouts.

The key learnings and outcomes from this trial and subsequent report are as follows:

- Participants reported an increase in work quality and efficiency across the following activities: summarisation of content, creating first drafts and information searches. An estimated 1 hour was gained in these activities. 70% of participants used Microsoft Teams and Word to summarise and create drafts of documents.
- 40% of survey respondents commented they were able to spend more time on other activities such as staff engagement, mentoring, organisational culture building, and building relationships with other stakeholders.
- The participants of the trial generally noted an increase in productivity, speed of work. A small increase in quality was noted, to a lesser extent of the previously mentioned benefits.

- Generally, participants were satisfied with Copilot trial and wished to continue using it in their work. 86% of the participants wanted to continue using Copilot following the conclusion of the trial.
- Some issues were noted with the adoption of the program and training was required to educate the workforce. Some cultural issues were noted and some participants felt uncomfortable using the software. Generally, participants felt they needed clear guidance and information on the use of the software to use it responsibly.

The following barriers to adoption of Gen AI were identified:

- Technical Barriers – including technical problems, security and information management issues.
- Capability – including finding efficient ways to use Copilot, understanding the requirements of the program and engineering.
- Legal – including legislative requirements of freedom of information, the need to disclose the use of the software and accountability.
- Cultural – regarding ethical issues and a degree of anti-AI sentiment across the workforce.
- Governance – including planning and rolling out of the program,

The report also found several unintended outcomes including inclusivity and accessibility improvements to employees with disabilities and to help employees who may struggle with writing. Other unintended outcomes included concern around impact to employment, bias towards western norms, environmental impact and a general loss of skill from the use of the program.

The following recommendations were presented following the trial:

- 1- **Product Selection** – Consider if a Gen AI solution is appropriate for the environment and use case.
- 2- **System Configuration** – Organisation should ensure their systems and IT security allow for gen AI products to safely operate and be utilised.
- 3- **Specialised Training** – Training should be offered to develop capabilities and should be tailored to the use case of the Gen AI product.
- 4- **Change Management** – Organisations should support the adoption of the products, highlighting benefits and supporting adoption.
- 5- **Clear Guidance** – Clear guidance should be provided outlining disclaimers, risks and accountabilities.
- 6- **Workflow Analysis** – Workflows should be interrogated to identify Gen AI use cases to assist in adoption.
- 7- **Use Case Sharing** – Organisations should share the ways they are using Gen AI to assist in adoption and reduce risk.
- 8- **Impact Monitoring** – The impact of Gen AI on employees and workplace and should be closely monitored to manage risks and identify issues ahead of time.

Generative AI Adoption at Lambda Corporation

Lambda Corporation, a Nordic multinational retail organisation, has taken a thoughtful and incremental approach to adopting generative AI within its departments to improve operational efficiency and decision-making. Their methodology begins with identifying specific use cases where generative AI can enhance productivity, such as content creation in marketing, customer interaction in service, and operational planning in logistics. The company prioritises tools that support and augment employee tasks rather than replace them.

Pilot programs are conducted in selected departments to test the effectiveness of generative AI applications, with an emphasis on iterative learning and refinement. Employees are encouraged to explore AI tools in their daily workflows, fostering an AI-use culture within the organisation.

Technically, Lambda ensures effective deployment by integrating generative AI tools seamlessly with existing IT systems. Data management is a key focus, with efforts made to use accurate, relevant, and unbiased datasets for training and validation. AI applications are chosen for their adaptability to departmental needs and their potential to deliver measurable benefits. Ethical considerations, including data privacy and fairness, are integral to the adoption process, ensuring the AI systems align with the company's values and regulatory requirements.

This systematic and adaptive approach enables Lambda Corporation to harness the potential of generative AI responsibly, enhancing productivity, decision-making, and workforce efficiency across its operations.

1. Human Resources (HR)

The HR team uses generative AI to make hiring and employee engagement easier. AI helps write job ads, screen candidates, and analyse employee feedback. While people still make the key decisions, AI speeds up tasks and provides helpful insights. This has led to faster hiring processes and better employee communication.

2. Purchasing Department

In the purchasing department, AI is used to predict demand, compare prices, and draft contracts. It helps analyse market trends and handle large amounts of data, making it easier to make smart buying decisions. This saves time and makes the department more efficient.

3. Customer Service Department

Generative AI improves customer service by supporting staff with tools like chatbots that handle simple questions. AI also analyses customer feedback to find areas for improvement. This has resulted in faster responses and happier customers, especially for new team members who get extra support from AI.

4. Logistics Department

The logistics team uses AI to plan delivery routes, manage inventory, and improve warehouse operations. AI tools help reduce waste and make the team more flexible when solving problems. This has made the supply chain smoother and more reliable.

5. Marketing Department

In marketing, AI helps create content like ads, social media posts, and email templates. It also analyses customer data to target the right audience for campaigns. This has made marketing efforts more creative, productive, and effective while saving time for the team.

Overall, generative AI is helping each department at Lambda work faster, smarter, and more creatively, making the company more efficient and innovative.

Department	Individual AI Use Perspective	Department Level AI Use Perspective	AI Use Scope	AI Use Effects on Departmental Performance	Department's AI Use Culture
HR	Varies among individuals; used to augment decision-making in both simple and complex tasks.	Relatively active; used to augment various tasks. AI use supported departmental decision-making.	From non-existent to relatively active, augmentative use in work and decision-making processes.	Enhances efficiency and quality of decision outcomes in both individual and teamwork.	Highly encouraged by leadership, active culture of AI use, shared learning. Positive sentiment.
Purchasing	Limited; mainly for creating product related content, translations and campaign ideas.	Limited; no evidence of use in significant department-level decisions.	Limited scope. Only a few members used from time to time.	Contributes to efficiency in creating textual content and equalising skills of staff, contributing to task allocation flexibility.	Lack of widespread encouragement and awareness of potential applications. Overall sentiment was curious.
Customer Service	Minimal; isolated instances of generative AI use in customer communication.	Notable absence of generative AI utilisation for work and decision-making.	Minimal, with potential for broader application in customer communication.	None reported. Potential to improve operational adaptability and resource allocation.	Low adoption rate, influenced by low encouragement to use AI. Mix of hesitant and curious sentiment.
Logistics	Used to enhance efficiency in technical tasks and solution development.	Clear difference in the application of machine learning AI for automation and generative AI for individual problem-solving.	Diverse, from enhancing digital workflow efficiency to aiding in technical tasks.	Significant operational benefits in automation and individual problem solving, possible broad departmental impact.	Digital and technical, with generative AI becoming embedded in certain teams' work processes, but not all.
Marketing	Extensive; used for creative ideation, technical tasks, and overcoming work challenges.	Widespread use of generative AI, changing operational methods and possibly influencing productivity.	Extensive, incorporated in various intensities for creative and technical work.	Positively affects creativity and technical capabilities, contributing to managing work with a compact team.	Active culture of AI use, encouraged exploration, and innovation with AI tools.

Figure 2: Use Case Table




Figure 3: Engagement Against Use Case

Utilisation of Gen AI in NSW Government

The New South Wales (NSW) Government has taken a structured approach to adopting Artificial Intelligence (AI), integrating it into various public services to improve efficiency and service delivery.

Their implementation process follows a clear framework, starting with identifying areas where AI can add value, such as streamlining administrative processes or enhancing customer service. The government then decides whether to develop AI tools in-house or procure solutions from external vendors. Pilot programs are run to test AI capabilities, with continuous iteration and refinement based on feedback (Nsw.gov.au, 2022).

On the technical side, the NSW Government ensured **robust data management**, selecting high-quality, unbiased datasets to train AI models effectively. They use **tailored AI algorithms** suited for specific applications, ensuring models are trained and validated rigorously.

Seamless **system integration** is a priority, with AI solutions being embedded into existing IT infrastructure to maintain efficiency and scalability. Ethical considerations, including **privacy protections and transparency**, are also key, ensuring AI systems align with public trust and **legal compliance**. This structured approach enables NSW to **leverage AI responsibly**, driving digital transformation while maintaining accountability (Digital.gov.au, 2023).

The role of Gen AI chatbots in local government context: Risks and real-world cases (Nickodem, 2024)

In recent years, Gen AI chatbots based on large language models (LLMs) have become increasingly popular in local government settings due to their ability to provide automated, real-time responses to public queries. Notwithstanding, their implementation presents notable risks, as demonstrated by several real-world cases. This analysis highlights key issues and contextualise them for CoA, emphasising the need for careful risk management should Gen AI chatbots be adopted.

Key risks and real-world applications

1. Inaccurate and misleading information: New York City Council

New York City's Gen AI chatbot, built using Microsoft's Azure AI services, demonstrated significant flaws when it provided incorrect and potentially harmful legal advice. For example, it claimed that employers could legally fire employees for complaining about sexual harassment and that restaurants could serve cheese that had been bitten by rats. Such errors illustrate the risk of "hallucination" - a common issue with LLM-based chatbots, where they generate responses that are made-up or factually incorrect. If a similar chatbot were deployed in Adelaide, inaccurate guidance on local regulations could lead to severe public confusion or legal disputes.

2. Legal liability for misrepresentation: Air Canada

In a case involving Air Canada, a customer relied on incorrect advice from the airline's chatbot about its bereavement travel policy, leading to financial loss. The British Columbia Civil Resolution Tribunal found Air Canada liable for negligent misrepresentation, ruling that the company was responsible for the chatbot's statements, as it was part of their website. A similar scenario in Adelaide could involve residents relying on faulty chatbot advice for matters such as building permits or zoning regulations. If incorrect guidance led to financial or legal consequences, the city could face claims of misrepresentation, damaging public trust and incurring costs.

3. Potential for adversarial attacks via "jailbreaking"

Research has shown that Gen AI chatbots are vulnerable to adversarial attacks. By using specific prompts, bad actors can "jailbreak" the chatbot and manipulate it into generating harmful content, such as instructions for illegal activities. If CoA chatbot was exploited in this manner, it could provide damaging or dangerous advice to users, creating public safety concerns and reputational harm for the council.

Gen AI chatbots can improve public services by offering instant responses and easing workloads. However, as seen in several real-world cases, they come with risks like providing incorrect or harmful information. Before adopting such a tool, CoA should carefully vet the chatbot by following these key recommendations:

Recommendations to mitigate risks

1. Thorough testing

CoA should test the chatbot's ability to give accurate answers to common questions and check how it responds to unusual or harmful prompts. Special attention should be given to questions about city laws and regulations, ensuring the chatbot does not give misleading legal advice. It is also important to test whether it provides consistent answers when asked the same question repeatedly.

2. Choosing the right chatbot model

CoA must decide whether to use a simple question-and-answer model or a more advanced conversational model. While the advanced model allows for more interaction, it also increases the chance of errors. A simpler model may be easier to manage and more reliable for users.

3. Reviewing vendor capabilities

CoA staff must work closely with the vendor to understand how the chatbot is trained, which AI model it uses and how it ensures accurate answers. They should also assess any safeguards the vendor provides to prevent errors or harmful content.

4. Cost and maintenance

Beyond the initial cost, maintaining and updating the chatbot will require ongoing resources. This includes retraining the AI to keep it accurate as new policies and information emerge.

5. Disclaimers and transparency

Since chatbots can sometimes produce incorrect responses, it is important to include clear warnings advising users to confirm important information with a city staff member. This approach, used by New York City Council, helps manage public expectations and reduces potential legal risks.

As can be seen, there are critical risks associated with the adoption of Gen AI chatbots. As such, local government bodies like CoA need to carefully vet whether the benefits outweigh the costs. Should such chatbots be endorsed, it is important that careful planning, regular communication and maintenance are in place to keep the chatbot in check and ensure that the services it provides enhance public services while maintaining trust.

3. ETHICS AND ENVIRONMENT

Ethical risks from a local government perspective

Integrating generative AI tools into City of Adelaide processes could improve organisational efficiency and help achieve the visions set out by the Strategic Plan. However, a substantial body of literature identifies ethical risks of this technology across a multitude of sectors which must be considered when designing and implementing such programs. Those most pertinent to a local government context are identified below.

Data privacy and trust

City of Adelaide handles a significant volume of personal sensitive information of its employees, residents and businesses every day such as contact details, expiations, economic figures and proposed developments. During onboarding, new staff are required to complete an exercise on record management given the potential risks of leaking data of this nature.

Input data used to train and utilise generative AI models are typically collected and stored within the program and has the potential to be leaked or stolen (Hiter, 2023). Unauthorised use of sensitive information poses risks to privacy, legal consequences and erodes user trust, which can have compounding effects on the rest of the organisation (Dey, 2024). When implementing a generative AI program, City of Adelaide must ensure that any identifying data is anonymised, robust cybersecurity measures are in place and that all employees are aware of the risks of uploading sensitive information to the model.

Workforce impact

Generative AI's capability of automating processes has led to growing concerns of workforce impact worldwide (Luckett, 2023). Although reports predict mass job creation through artificial intelligence, the same technology could also lead to mass job displacement, particularly in those sectors and industries involving routine and predictable tasks (Kshetri, 2024). Unaddressed, the prospect of job displacement may have a negative psychosocial on employees if perceptions of their employment security are affected.

City of Adelaide should proactively invest in upskilling and retraining employees most affected by the adoption of generative AI, alongside developing appropriate action plans and policies to bolster trust throughout the organisation.

Transparency and bias

Although generative AI models are capable of mimicking human behaviour and dialogs based on the information they are trained with, they lack certain critical thinking and decision-making processes to warrant being left unchecked. The black-box nature of Generative AI models presents multifaceted issues regarding transparency such as data collection, sources of output data, copyright infringement and accountability (Chugh, 2023). Without proper documentation, the assignment of responsibility for models' outputs are undeterminable, bringing into question the developer, user and any other body involved in the creation and deployment of generative AI programs.

These also models risk perpetuating biases and stereotypes within their input data that can have cascading effects throughout the organisation (Esposito and Tse, 2024). Biases embedded in training data can be reinforced by the model, leading to misconceptions being presented as fact.

Appropriately skilled and authorised officers should be appointed to monitor the generative AI model throughout its life cycle, clearly documenting processes and ensuring information shared attributes AI involvement. City of Adelaide must also seek that the team involved is reasonably representative of a range of demographics to improve identification of bias in the input and output data.

Environmental impacts

There are significant energy, raw material and water costs involved in the manufacturing of hardware and development and deployment of software associated with generative AI (Dolby, 2023). In turn, there has been an increase in e-waste as outdated hardware such as GPUs are replaced by new devices (Kuever, 2023). These costs are expected to escalate considerably with the growing demand for generative AI globally, however the environmental stresses are inequitably distributed such that already vulnerable areas are more impacted (Ren and Wierman, 2024).

Mirroring the opacity of generative AI processes, the figures for emissions caused by the technology are not readily available so exact numbers cannot be analysed. However, given the forecasted trajectory of consumption in the coming years efforts are required to minimise the environmental costs of using generative AI.

Strategies to reduce the environmental impact of generative AI technologies include promoting the use of energy-efficient data centres, encouraging the adoption of renewable energy sources, and advocating for the development of water-efficient AI infrastructure. Furthermore, the city could support initiatives aimed at reducing the carbon footprint of AI systems, such as using recycled water in data centres and improving the efficiency of cooling systems. Encouraging collaboration between industry, government, and environmental experts will be key to ensuring that AI technologies are developed and deployed in a way that aligns with sustainability goals.

4. RISK VS OPPORTUNITY FOR CITY OF ADELAIDE

Implementing Gen AI into an organisation introduces a wide range of risks and benefits for the organisation. For CoA to successfully utilise Gen AI software to its full capability and to experience the full range of benefits these risks must be properly managed and mitigated (Australian Government, 2024). Associated risks relating to Gen AI can vary markedly, ranging from typical Information Technology risks to emotional harm and discrimination. (National Institute of Standards and Technology, 2024).

Risks

A range of risks are discussed below which CoA may be susceptible to should Gen AI be utilised at the organisation:

Ethics & Environment, as discussed in Section 3.

Data Security & Privacy: Gen AI models are required to have access to large databases of information, some of which may be sensitive, confidential or personal. This information needs to be managed to prevent leaks to outside parties or the users of the software (PwC, 2024). Gen AI can also be targeted by cyber-attacks in a wide range of ways. One such example of a cyber-attack includes outside parties influencing the content generated by the model leading to undesirable outputs and other potentially harmful consequences. Furthermore, entering personal information into Gen AI models makes that information at risk of leaking along with the associated issues with transparency and consent of that information (National Institute of Standards and Technology, 2024). Security and Privacy issues will be much more significant when the AI model is external to CoA and not endorsed or hosted by a reputable and secure organisation (Such as Microsoft's Copilot model). When an internal AI model can be used, these risks will be mostly mitigated or significantly reduced.

Ownership of Information: Content generated by Gen AI may infringe on IP rights / Copyright, sometime due to the material the Gen AI was trained on. Furthermore, it is possible that an organisation utilising an external Gen AI model may need to research the IP rights with the Gen AI model owner (National Institute of Standards and Technology, 2024).

Legal: A range of legal considerations need to be made related to ethics, data security, various legislation, terms of use, copyright etc. Penalties could apply to the organisation and/or its employees if an infringement is reported (PwC, 2024). Decisions should not be made by Gen AI directly; however, Gen AI can be used to responsibly help a staff member made decisions, while ensuring compliance with Council Delegation of Authority. AI bias can lead to discrimination against groups or ideas which can be considered illegal in Australia due to the DDA Act 1992. There is a degree of uncertainty in the legal risk of Gen AI as it is a new technology, and a detailed assessment of the legal risks should be undertaken.

Confabulations (Hallucinations): Gen AI models can be prone to generating and presenting incorrect information in a confident demeanour. This can also include responses which do not relate to the prompt or inputs of the user and may even contradict previously generated responses. This expression is also known as 'hallucinations' or 'fabrications' (National Institute of Standards and Technology, 2024).

Bias & Fairness

Gen AI models can show a bias influenced by the content it was trained on. This bias can include discrimination, favouring certain topics or answers and may be unfair in a variety of ways (PwC, 2024). This bias could clash and be against the CoA Strategic Plan, CoA values and the intent of CoA employees and may harm the organisations operations.

Other Risks:

- Unknown Risks – Gen AI is an emerging technology, and new risks are probable to be discovered or introduced as the technology and research around it develops (National Institute of Standards and Technology, 2024).

Preventative Measures

The South Australian Government (2024) published a Guideline on the use of Gen AI which outlined some key risks and methods to mitigate them. These preventative measures along with other recommended preventative measures are outlined below:

Legal Due Diligence

Prior to implementing Gen AI at CoA, legal advice should be sought to comply with any relevant legislation. Considerations may need to be made with regards to the Local Government Act (1999), CoA Strategic Plan(s), Code of Conduct for CoA employees along with all other relevant documents, acts and laws.

Policy and Procedure

Adopting a CoA Generative AI policy will allow the organisation to take a stance on the use of the technology. In addition, an operating guideline on the proper use and guiding principles will allow the administration to be confident in its use and the procedure around its use. These guidelines are needed to guide users of Gen AI to mitigate the risks shown above.

Validation

All output from Generative AI should be checked and verified by a suitably experienced CoA staff member to ensure the output is correct. Generated content should not be blindly copy and pasted. This process will need to be built into the operating guideline, training and policy.

Training and Education

Prior to adoption of a CoA Gen AI model, the administration should be educated on the proper use and the potential risks of Gen AI. This should also be a part of employee onboarding and be regularly repeated to ensure all CoA members are properly trained on Gen AI.

Vetting AI Software

Prior to selecting any given Gen AI model for use at CoA, it must be vetted by the IM team to ensure the model can provide suitable security, privacy and other technical requirements. This will reduce the risk of issues being encountered. This vetting process should include any testing and validation processes prior to the model being used.

Transparency

The use of Gen AI at CoA will need to be transparent to avoid reputational issues and complaints. Residents and councillors should be aware of the use of Gen AI at CoA and the relevant policies which guide its use.

Monitoring

The performance of Gen AI at CoA should be constantly monitored to ensure the model is performing as expected and is not exposing CoA to any undue risk. Similarly, CoA policy and procedures should be constantly monitored and updated to mitigate any identified issues

Internal CoA AI Model

The CoA should adopt an internal Gen AI model which is made for and only used by the City of Adelaide to reduce outside interference and to improve monitoring capabilities.

Table 1 below shows the list of risks and the respective measures which will help to reduce the risks.

Table 1: Risk vs Mitigating Action

Risk	Mitigating Action
Data Security & Privacy	Using vetted and approved AI models for CoA purposes, appropriate policies and procedures, training, monitoring, transparency, suitably vetting and adopting internal AI model, transparent policy.
Ownership of Information	Appropriate Policy and Procedure, Internal COA AI Model, human checking of outputs to identify potential IP infringements (validation).
Legal	Policy and procedure, legal due diligence, training and education, validating AI outputs.
Confabulations	Validation, Vetting AI model, monitoring.
Bias and Fairness	Legal due diligence, vetting AI model, monitoring and checking AI outputs, training, policy and procedure.

Most risks associated with the use of Generative AI at the City of Adelaide can be mitigated through proper preparation, research, suitable policies and procedures. This allows the organisation to realise the full benefits of the technology while minimising risk.

Benefits

Generative AI introduces a wide range of benefits and potential opportunity and useful application, especially in an administrative organisation such as the City of Adelaide. McKinsey and Company (2025) discuss a range of benefits which are included below along with some other identifiable benefits:

Increased Productivity

Gen AI can boost productivity in a wide range of ways including automating repetitive tasks, brainstorming, summarising, drafting etc. This allows the employee to free up time for other work.

Assist in Decision Making

While Gen AI must not be used to make decisions due to legal consequences and a variety of other obligations (Delegations of power cannot be given to an AI), it can be used to assist in decision making. This could include analysing data, summarising feedback and more. CoA must be careful however, to ensure that staff members do not become dependant or over reliant on Gen AI.

Reduced Operating Costs

Due to the wide range of benefits related to Gen AI, cost savings could be expected due to better managed resources, staff member productivity and many other benefits.

Brainstorming

Due to Gen AI's ability to innovative and be creative, it is a great tool to be used for brainstorming ideas. This can ensure that council have considered all possible options before committing to a solution.

Improved Data Analysis

Gen AI can effectively analyse large amounts of data quickly and accurately. This can help reduce the time it takes CoA staff to complete analytical tasks or when working with large datasets. The results can be summarised and reviewed and presented in a wide range of ways.

Identify Risks

Gen AI can analyse situations and identify potential risks and vulnerabilities. This can greatly assist an organisation risk management and supplement the existing methods. Large issues could be identified before they have consequences.

Improved Accessibility / Collaboration

Staff members who have English as a second language or have intellectual disabilities can utilise Gen AI to help them drafting emails, spell checking and to make it more readable. This can boost communication across the organisation and improve collaboration.

Automation

Gen AI can automate complex and time-consuming tasks to reduce the number of hours CoA staff need to spend.

If the risks associated with the adoption of Gen AI at CoA can be properly managed, Gen AI provide a wide range of benefits to help the City of Adelaide perform at a high level and deliver better results for its businesses, residents and city users. CoA can use Generative AI in specific ways to help certain teams and sections of the organisation, this is discussed further in Section 6.

City of Adelaide Information Management Perspective

From the City of Adelaide's Information Management (IM) viewpoint, planning for Gen AI involves a balance between innovation with responsible data governance and integration considerations. A few key areas highlighted below stand out:

1. Technical Challenges and Infrastructure Requirements

- **Data Storage and Security:** Storing large volumes of data and model outputs securely is essential. Ensuring compliance with privacy and data protection regulations is vital.
- **Computing Resources:** Gen AI models often require significant processing power. Some organisations may opt for cloud-based solutions to manage scalability, whilst others may consider on-premises infrastructure for tighter data control.
- **Model Training and Maintenance:** Gen AI models need ongoing tuning and updates. Adequate version control, monitoring, and retraining processes should be established.
- **Governance and Ethics:** Policies around appropriate use, content moderation, and intellectual property rights must be in place to guide staff and protect the community.

2. Integration with Existing Systems

- **Records Management Systems:** Outputs from Gen AI must be captured and archived according to existing records management policies.
- **Customer Relationship Management (CRM):** AI-generated responses or content for community-facing platforms should seamlessly feed into or out of current CRM tools.

- **Collaboration Platforms:** Chatbots or virtual assistants built on Gen AI must integrate with the City's internal collaboration suite to ensure consistent and up-to-date information sharing.
- **Data Repositories and APIs:** AI solutions may need to draw on internal data sources (e.g., geographic information systems) and external APIs to generate relevant, context-rich content.

5. COA STAKEHOLDER ENGAGEMENT

Survey explanation and intent

A survey was created to assess the CoA Administration sentiment and understanding of Gen AI. This survey was sent out via The Next Edition (CoA internal newsletter) as a voluntary survey to all employees of CoA. The questions included in the survey are presented in *Appendix B*.

The intent of the survey was to acquire quantitative and qualitative data on the following subjects:

- The administrations understanding of Gen AI, its risks and its opportunities,
- The administration sentiment towards Gen AI, and
- How the administration currently uses Gen AI,
- How different demographics of the administration feel about Gen AI.

The individual results were kept confidential, and it was made clear to the participants that their responses were to help the CoA Graduates.

Key findings

Consensus from the organisational survey

There is **strong interest in using Gen AI**, with employees recognising its potential to improve efficiency and productivity. However, there is **low trust and risk tolerance**, particularly regarding data privacy and security. Staff are **keen for structured training and clear policies** to ensure responsible adoption. While employees see the value of Gen AI, usage remains relatively low both at work and in their personal lives.

To drive adoption, it is essential to address trust issues, provide training, and establish clear guidelines on ethical use, while also emphasising efficiency gains.

Insights from the organisational survey (n=67)

Majority are individual contributors – 70% of respondents are individual contributors, meaning Gen AI adoption will mostly affect frontline staff rather than leadership.

Diverse Age Demographics – Respondents span multiple age groups, with the highest percentage in 35-44 (31%) and therefore, adoption strategies should consider different levels of familiarity with technology.

Understanding of Gen AI is moderate – 69% either agree or strongly agree that they understand Gen AI, but 19% are neutral or disagree, **showing there is still a need for education and awareness**.

Low Trust in Gen AI – **Only 23% trust Gen AI** whilst 47% are neutral or disagree. This hesitation could slow down adoption unless concerns are properly addressed.

Risk awareness exists, but confidence is low – Whilst **53% understand the risks of Gen AI**, only 27% feel comfortable with those risks, suggesting that **many people don't yet feel prepared to use it safely**.

Willingness to use Gen AI at Work – 76% believe incorporating Gen AI would support a high-performing organisation, showing **strong potential for uptake** if the right conditions are in place.

Desire for Training and Clear Guidelines – 88% want CoA-specific training on risks, safe use, and best practice, whilst 94% want clear guidance on responsible usage. There is a strong **appetite for structured learning before rolling out Gen AI at scale**.

External Use of Gen AI is Limited – 30% have never used Gen AI outside of work, and only 11% use it daily. This suggests many employees have limited exposure to the technology, **reinforcing the need for introductory training**.

Efficiency is the Biggest Perceived Benefit, Data Privacy is the Biggest Risk – Word cloud analysis shows "efficiency," "timesaving," and "tasks" as the main benefits, while "data privacy," "security," and "confidentiality" are the biggest concerns, highlighting a trade-off between productivity gains and data protection risks.

Demographic insights

Survey respondents present a diverse cross-section of the organisation, with every role, age group, program being represented.

In terms of role distribution, the Individual Contributor formed the largest group of respondents at over 70 percent, followed by Team Leaders at over 16 percent and Senior Leadership Team at over 7 percent.

The age bracket between 35-44 years old reflect the highest response rate at nearly 31 percent, followed by 25-34 years old and 45-54 years old, both at nearly 28 percent.

Finally, a broad range of programs are represented, with notable participations from People team at close to 20 percent, followed by City Culture at nearly 14 percent, then Customer & Marketing as well as Infrastructure at over 12 percent respectively. No participation was noted in the Office of the CEO, which may hinder insights into perspective from C-suite level.

Overall, this is a fair representation of the role demographics within CoA, with the main demographic being Individual Contributor, aged between 25-54 years old, spanning across various programs.

General sentiment towards Gen AI

To gauge the CoA employees' sentiment towards Gen AI, a sentiment analysis was conducted in question 4 using a Likert scale, with five response options from Strongly Disagree to Strongly Agree for clear statements, along with open ended questions to provide qualitative insights into the matter, which will be represented in form of word clouds.

Insights from CoA employee survey

Quantitative analysis – Likert scale

The findings reveal that a majority of CoA employees have a clear understanding of Gen AI (nearly 69 percent Strongly Agree and Agree) and its associated risks (over 53 percent Strongly Agree and Agree). However, it is also noted that there is a considerable amount of contrasting response for Strongly Disagree and Disagree, at over 19 percent and nearly 28 percent respectively, and overall low neutrality on both statements. This indicates that CoA employees are divided into two categories, one group with a clear understanding of Gen AI and the other with limited understanding of Gen AI.

All in all, the results from the Likert scale reveal that most CoA employees have a clear understanding of Gen AI and its associated risks, causing lower level of trust in Gen AI.

To gain a deeper understanding of CoA employees' perceived risks and benefits towards Gen AI, which is a key factor in determining their sentiment and level of trust, open ended questions were asked in question 8 and 9, and answers were presented in form of word clouds.

20

Usage of Gen AI

Usage of Gen AI within the workplace versus outside of work

In question 6 and 7 of the survey, participants were asked on the frequency of usage of Gen AI at work, followed by outside of work. Results are similar for both questions, with contrasting level of usage. Participants have either high level of usage (daily, weekly at nearly 36 percent for “at work” and “outside of work”) or low level of usage (rarely, never at 58 percent for “at work” and over 48 percent for “outside of work”), and low level of neutrality overall (see figure 3 and 4). This indicates that CoA employees are generally more comfortable with using Gen AI outside of work, potentially due to the perceived risks of using Gen AI in the workplace as discussed in the earlier section.

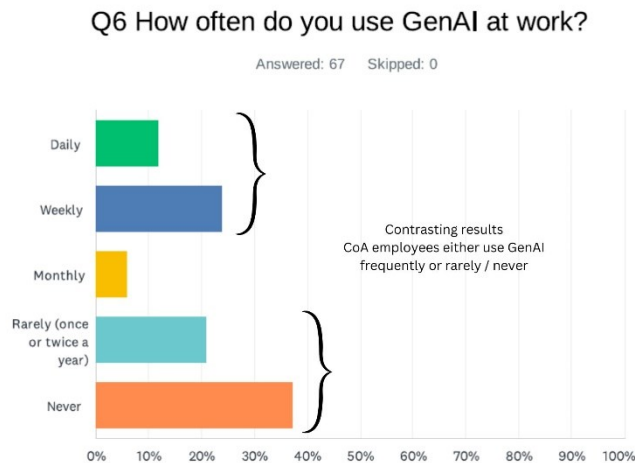


Figure 6: Question 6 | How often do you use Gen AI at work?




Figure 7: Question 7 | How often do you use Gen AI outside of work?

Sentiment towards usage of Gen AI within CoA

To gain a more tailored insight into CoA employees’ sentiment towards the usage of Gen AI within CoA, another sentiment analysis using a Likert scale was established in question 5, with five response options from Strongly Disagree to Strongly Agree against statements that are specific to CoA (see appendix C).

From the results, it can be inferred that CoA staff are generally open to the use / adoption of Gen AI within CoA, with over 80 percent of employees indicating Strongly Agree and Agree,

The findings also reveal the importance of a clear stance from CoA on the adoption of Gen AI and highlight a strong need for Gen AI training and clear operating guidelines. Over 61 percent of staff agree or strongly agree that they would be open to use Gen AI if CoA formally endorse its usage. Additionally, almost all employees indicate strong interest in having clear guidance on responsible usage (more than 94 percent) and Gen AI training (nearly 88 percent). This can also be seen in the open-ended question 10, whereby CoA were asked on how CoA can support safe usage of Gen AI – keywords highlighted were “*training, workshops, clear policies, guidelines, informing of potential risks*” (see figure 5).



The final section of the survey explores possible use cases of Gen AI within CoA from staff perspective. The results from this will guide the scope of the Gen AI adoption and implementation plan which will be further developed in the subsequent section. From the responses, there is a strong interest for the adoption of Gen AI within customer-facing roles, potentially the adoption of “*chatbot*” within the customer centre space. Content generation within the marketing space is another area with vested interests in Gen AI usage. Finally, keywords such as “*data, information, research*” are also highlighted, indicating high interest in using Gen AI for data mining, analysis and research purposes (see figure 6).



Discussion

As discussed, survey respondents come from a diverse range of roles, age group and program, giving a fair representation of the demographics within CoA, with the main demographic being Individual Contributor, aged between 25-54 years old. All program options were represented, except for contributor from Office of the CEO, which may mean limited insights into C-suite level perspective on Gen AI. However, this is countered by the relatively high participation rate from SLTs and Executives which give reasonable insights into leaderships' perspective on Gen AI.

In terms of sentiment analysis, CoA employees generally display a positive outlook on the usage and adoption of Gen AI within CoA, with even stronger level of confidence if CoA formally endorse its usage and provide clear operating guidelines as well as training. This will also aid in the contrasting results for understanding of Gen AI, whereby staff are either highly knowledgeable on Gen AI or have limited knowledge on it, with low level of neutrality.

Finally, potential use cases of Gen AI as presented by CoA staff include, but are not limited to, customer centre (potentially chatbot), content generation for marketing purposes and data analysis, data mining or research purposes.

6. COA USE CASES

There are many possible use cases for Gen AI at CoA and as the technology develops, more use cases will appear and be able to be capitalised on. Two potential use cases are discussed below including general administrative tasks and AI chatbots in the Customer Centre.

General Administration

One of the largest potential use cases of Gen AI at CoA would be using a Gen AI model to assist in general administrative tasks. Through integrating a Gen AI model into day-to-day tasks, the organisation could possibly realise a range of benefits ranging from improved efficiency to improved communication, as discussed in Section 4. A trial could be implemented, similar to the APS trial of Microsoft Copilot (2024). This would enable the administration to:

- Generate draft documents / policies / reports / etc, ensuring that they are verified and edited. This would save time and could increase consistency across the organisation if the model uses a generalised CoA writing style.
- Gen AI models can edit text to make it grammatically correct, formal or informal and check for consistency.
- Summarise long documents, emails or other forms of text. This allows key points to be conveyed to the reader directly and effectively and allows more time for the employee to spend their time on other tasks.
- Save time on research. The model could find relevant resources and information based on input from the user. The user can specify the type or resource they need, put limitations on date and country of origin to get exactly what they need without researching themselves.
- Certain AI models may also have additional capabilities such as managing calendars, booking meetings and organising reminders.
- Effectively analyse data, identify trends, insights and create reports. This could help inform policy and decision making across the organisation.

AI chatbots in City of Adelaide customer service

Gen AI chatbots can enhance CoA's customer service by providing instant responses to inquiries about council services, regulations and events. They can improve accessibility, reduce wait times, and assist both residents and businesses. This Gen AI chatbot can be integrated into CoA's website as a widget that follows the user throughout their on-page user journey as an "AI personal assistant". ([See real-life Gen AI chatbot example from City of New York Council](#)). However, there are key risks involved that must be managed to ensure accuracy, reliability and public trust.

Potential integration opportunities:

- Rates and payments: Real-time info on due dates and payment options.
- Permit and licensing assistance: Guidance for applications (e.g. building, events, business licenses).
- Waste and recycling enquiries: Bin collection schedules, recycling guidelines.
- Event and venue bookings: Help with reservations and event details.
- Parking and transport information: Updates on zones, fees, public transport links.
- General council services: FAQs on libraries, community programs, recreation.

Key risks involved:

- Inaccurate information: Incorrect details may cause confusion and complaints.
- Legal liability: Misleading advice on approvals or fines could create legal issues, as seen in cases like Air Canada’s chatbot misrepresentation.
- Manipulation risks: Adversarial attacks could prompt inappropriate responses, damaging CoA’s reputation.

Risk mitigation strategies:

- Thorough testing: Ensure accuracy in responses, particularly for regulatory queries.
- Chatbot model selection: A rule-based chatbot may be safer than an advanced AI model. ([See real-life rule-based chatbot example from City of Los Angeles Council](#)).
- Ongoing Monitoring: Regular updates and oversight to maintain accuracy.
- Transparency and disclaimers: Clearly inform users of chatbot limitations and advise verification with staff.

Overall, a well-managed chatbot can improve CoA’s customer service efficiency. However, ongoing efforts to safeguard trust through careful implementation and oversight is essential.

7. RECOMMENDATIONS

The findings from this report demonstrate that Generative Artificial Intelligence (Gen AI) can significantly benefit the City of Adelaide (CoA), enhancing productivity, decision-making, innovation, and operational efficiency. However, successful adoption requires careful management of several critical areas including ethical considerations, privacy, security, environmental sustainability, and workforce impact.

Based on our extensive research, stakeholder feedback, and insights gained from both governmental trials and best industry practices (including AWS guidelines), we propose the following strategic recommendations:

1. Strategic Adoption Approach

- Adopting a phased implementation approach beginning with low-risk pilot projects to build organisational readiness and familiarity.
- Using an **Ideation Canvas** and prioritisation frameworks to identify initial use cases, ensuring alignment with CoA's strategic goals.

2. Governance and Policy

- Developing and maintaining a comprehensive CoA-specific Gen AI policy that clearly outlines acceptable uses, privacy guidelines, data protection, accountability, and ethical standards as outlined in the Generative AI Operating Guideline.
- Establishing an internal governance structure (e.g., AI Centre of Excellence) to oversee model deployment, risk management, and compliance.

3. Infrastructure and Technical Considerations

- Prioritising the adoption of an internal AI model tailored exclusively for CoA, reducing external risks and enhancing security and control.
- Ensuring rigorous vetting of any Gen AI tools and robust cybersecurity measures, aligned with existing information management policies and infrastructure requirements.

4. Training and Education

- Implementing structured training programs tailored to different staff demographics, addressing both technical usage and ethical considerations.
- Providing ongoing education and clearly documented guidelines to enhance staff confidence, mitigate risks, and encourage responsible use.

5. Transparency and Communication

- Maintaining transparency regarding the use of Gen AI with staff, residents, councillors, and stakeholders, clearly communicating purposes, benefits, and risks.
- Continuously engaging stakeholders through regular updates and feedback mechanisms, ensuring public trust and internal buy-in.

6. Legal and Ethical Compliance

- Conduct detailed legal due diligence, ensuring compliance with local government legislation, the Privacy Act 1988, and relevant OAIC guidelines, as outlined in the Operating Guideline.
- Regularly monitor and evaluate models to detect biases, inaccuracies (hallucinations), and ethical breaches, taking immediate corrective actions as necessary.

7. Monitoring and Continuous Improvement

- Regularly update Gen AI policies and procedures based on new insights, technological advancements, or emerging risks.

8. CONCLUSION

Generative AI represents an unprecedented opportunity for the City of Adelaide to position itself as a forward-thinking, innovative, and high-performing local government body.

It aligns closely with the City of Adelaide's Strategic Plan 2024–2028, reinforcing commitments to economic growth, operational efficiency, innovation, and enhanced community engagement. Whilst the associated risks including data security, privacy concerns and ethical challenges are noteworthy, they are manageable with adept policies, targeted training, transparent governance, and careful technology selection.

The organisational survey clearly indicates strong internal interest and support for Gen AI, tempered by legitimate concerns about privacy, security, and ethical issues. Addressing these concerns proactively through structured training, clear guidelines, and transparent governance frameworks is essential for successful adoption.

The lessons learned from the Australian Government's Copilot trial, along with best practices advised by industry leaders like AWS, provide valuable roadmaps for a phased, sustainable approach to implementation. By carefully navigating these challenges through responsible stewardship and proactive risk management, the City of Adelaide can unlock substantial benefits from Generative AI, ensuring it remains competitive, efficient, innovative, and responsive to the needs of its community.

This report aims to serve as a foundational resource, guiding the responsible adoption and integration of Generative AI technologies across the City of Adelaide, laying a strong foundation for future policy development, innovation, and continuous improvement.

REFERENCES

1. Australian Government. (2024), *Australian Government Trial of Microsoft 365 Copilot*. Available at: <https://www.digital.gov.au/initiatives/copilot-trial> (Accessed: 18 November 2024)
2. Borji, A. (2023) *A Categorical Archive of ChatGPT Failures*. Available at: <https://arxiv.org/pdf/2302.03494> (Accessed: 9 January 2025).
3. Chugh, V. (2023). *Ethics in Generative AI*. Available at: <https://www.datacamp.com/tutorial/ethics-in-generative-ai> (Accessed: 28 August 2024)
4. Dey, S. (2024). *Council Post: Which Ethical Implications Of Generative AI Should Companies Focus On?* Forbes. Available at: <https://www.forbes.com/councils/forbestechcouncil/2023/10/17/which-ethical-implications-of-generative-ai-should-companies-focus-on/> (Accessed: 28 August 2024)
5. Dolby, N. (2023). *Artificial Intelligence Can Make Companies Greener, but It Also Guzzles Energy*. Wall Street Journal. Available at: <https://www.wsj.com/articles/artificial-intelligence-can-make-companies-greener-but-it-also-guzzles-energy-7c7b678>. (Accessed: 28 August 2024)
6. Digital.gov.au. (2023). *Interim guidance on government use of public generative AI tools - November 2023* | aga. [online] Available at: https://architecture.digital.gov.au/guidance-generative-ai?utm_source=chatgpt.com (Accessed 3 Apr. 2025).
7. Esposito, M., and Tse. T. (2024). Mitigating the Risks of Generative AI in Government through Algorithmic Governance. *Association for Computing Machinery*. <https://doi.org/10.1145/3657054.3657124> (Accessed: 9 January 2025)
8. George, A.S., George, A.S.H., and Martin, A.S.G. (2023) The Environmental Impact of AI: A Case Study of Water Consumption by Chat GPT. *Partners Universal International Innovation Journal (PUIIJ)*, 1(2), March-April. Available at: <https://doi.org/10.5281/zenodo.7855594> (Accessed: 9 January 2025).
9. Government of South Australia. (2024) *Guideline for the use of Generative Artificial Intelligence and Large Language Model Tools*. Available at: https://www.dpc.sa.gov.au/data/assets/pdf_file/0007/936745/Guideline-13.1-Use-of-Large-Language-Model-AI-Tools-Utilities.pdf (Accessed: 10 September 2024)
10. Hiter, S. (2023). *Generative AI Ethics: Concerns and Possible Solutions*. eWEEK. Available at: <https://www.eweek.com/artificial-intelligence/generative-ai-ethics/> (Accessed: 28 August 2024)
11. Huang, S. and Siddarth, D. (2023) *Generative AI and the Digital Commons*. Available at: <https://arxiv.org/pdf/2303.11074> (Accessed: 9 January 2025).
12. Jaffe, S., Shah, N.P., Butler, J., Farach, A., Cambon, A., Hecht, B., Schwarz, M. and Teevan, J. eds. (2024) *Generative AI in Real-World Workplaces: The Second Microsoft Report on AI and Productivity Research*. Available at: <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf> (Accessed: 16 December 2024).
13. Kuever, N. (2023). *Generative AI and its potential environmental impact*. Available at: <https://blog.bosch-digital.com/generative-ai-and-its-potential-environmental-impact/> (Accessed: 28 August 2024)
14. Kshetri, N. (2024) *Economics of Artificial Intelligence Governance*. *Computer*, 57, 113–118.

15. Li, P., Yang, J., Islam, M.A., and Ren, S. (2023) Making AI Less “Thirsty”: Uncovering and Addressing the Secret Water Footprint of AI Models. Available at: <https://arxiv.org/pdf/2304.03271> (Accessed: 9 January 2025)
16. Luckett, J. (2023) *Regulating Generative AI: A Pathway to Ethical and Responsible Implementation*, Journal of Computing Sciences in Colleges, 39(3), 47–65.
17. McKinsey & Company. (2025). *Superagency in the workplace: Empowering people to unlock AI’s full potential*. Available at: <https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/superagency-in-the-workplace-empowering-people-to-unlock-ais-full-potential-at-work> (Accessed: 17 March 2025).
18. Microsoft. (2024) *Introducing Microsoft 365 Copilot – your copilot for work*. Available at: <https://blogs.microsoft.com/blog/2023/03/16/introducing-microsoft-365-copilot-your-copilot-for-work/> (Accessed: 18 November 2024)
19. Microsoft. (2024) *Generative AI in Real-World Workplaces*. Available at: <https://www.microsoft.com/en-us/research/uploads/prod/2024/07/Generative-AI-in-Real-World-Workplaces.pdf> (Accessed: 20 December 2024)
20. National Institute of Standards and Technology (NIST USA). (2024) *Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile*. NIST.AI.600-1. Available at: <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf> (Accessed 11 December 2025).
21. Nickodem, K. (2024) *Unpacking the Potential Risks of Generative AI Chatbots on Local Government Websites*. Available at: canons.sog.unc.edu/2024/05/unpacking-the-potential-risks-of-generative-ai-chatbots-on-local-government-websites/ (Accessed 9 January 2025).
22. PwC Australia. (n.d.) *Generative AI: Mitigating Risks and Adopting Secure Practices*. Available at: <https://www.pwc.com.au/digitalpulse/gen-ai-risks-and-secure-adoption.html> (Accessed 9 January 2025)
23. Ren, S. and Wierman, A. (2024). *The Uneven Distribution of AI’s Environmental Impacts*. Harvard Business Review. Available at: <https://hbr.org/2024/07/the-uneven-distribution-of-ais-environmental-impacts>. (Accessed 28 September 2024)
24. Saenko, K. (2023) *Is generative AI bad for the environment? A computer scientist explains the carbon footprint of ChatGPT and its cousins*. Available at: <https://theconversation.com/is-generative-ai-bad-for-the-environment-a-computer-scientist-explains-the-carbon-footprint-of-chatgpt-and-its-cousins-204096> (Accessed: 9 January 2025).
25. Stahl, B.C. and Eke, D. (2024) *The ethics of ChatGPT – Exploring the ethical issues of an emerging technology*. International Journal of Information Management, 74, 102700. Available at: https://www.sciencedirect.com/science/article/pii/S0268401223000816?ref=pdf_download&fr=RR-2&rr=8ff1284d79b3e682 (Accessed: 9 January 2025).
26. nsw.gov.au. (2022). *Case Studies on the use of AI in NSW Government | Digital NSW*. [online] Available at: https://www.digital.nsw.gov.au/article/case-studies-ai-nsw-government?utm_source=chatgpt.com (Accessed 3 Apr. 2025).

APPENDIX A: PROPOSED OPERATING GUIDELINE

APPENDIX B: SURVEY QUESTIONS

Demographic Questions

- What is your role at CoA? (Individual Contributor, Team Leader (inc WGL), Frontline Leader, Program Manager, Senior Leadership Team (SLT))
- What age group are you? (15-24, 25-34, 35-44, 45-54, 55-64, 65+)
- Which Program do you work in
 - Adelaide Central Market Authority
 - AEDA
 - City Culture
 - City Operations
 - Customer & Marketing
 - Finance & Procurement
 - Governance
 - Information Management
 - Infrastructure
 - Office of Lord Mayor
 - Office of the CEO
 - Park Lands Policy & Sustainability
 - People
 - Regulatory Services
 - Strategic Property & Commercial
 - Strategy Insights & Performance

General Questions

Scale from **Strongly Disagree** to **Strongly Agree** (alternative questions shown in the level 2 dot points)

General sentiment

- I have a good understanding of what Generative AI is.
- I trust Gen AI.
- I have a good understanding of the risks associated with using Gen AI.
- I am comfortable with the risks of using Gen AI.

In the workplace

- I would use Gen AI more if CoA endorsed its usage.
- I am open to the use of Gen AI at CoA.
- I think incorporating Gen AI in my work would support a high performing organisation performance.
- I think it is important that CoA provides clear guidance on the appropriate and responsible use of Gen AI.
- I would want to undertake CoA specific training on Gen AI (risks, safe use and best practice)

Custom format Questions

- How often do you use Gen AI at work (daily/weekly/monthly/rarely (once or twice a year)/never).
 - Please provide some examples of how you have used GenAI in the workplace. (Custom answer format)
- How often do you use Gen AI outside of work (daily/weekly/monthly/rarely (once or twice a year)/never).

General thoughts (open text, 100 words)

- What are the potential benefits of using Gen AI at the City of Adelaide?
- What are the potential risks of using Gen AI at the City of Adelaide?
- How could CoA support you in using Gen AI safely?
- Do you see any potential use cases for Gen AI at the City of Adelaide?

APPENDIX C: SURVEY RESULTS