LONDON BOROUGH OF HAMMERSMITH & FULHAM

Report to: Policy and Oversight Board

Date: 29/04/2024

Subject: Generative AI: Opportunities, Challenges and Risks

Report author: Anthony King, Enterprise Architect

Responsible Director: Nicola Ellis, Strategic Director Corporate Services

---

SUMMARY

This report provides a brief overview of the current state and future implications of generative AI technologies.

RECOMMENDATIONS

1. For the Policy and Oversight Board to note and comment on the report.

Wards Affected: None

Our Values	**Summary of how this report aligns to the H&F Values**
Building shared prosperity	
Creating a compassionate council	Generative AI can help to improve and widen the accessibility of our digital and information services to residents
Doing things with local residents, not to them	Resident-facing digital and information products, including those employing Generative AI, are co-produced with Residents via the Digital Accessibility Group (DAG)
Being ruthlessly financially efficient	Generative AI, if implemented as part of business efficiency and improvement initiatives, could help us to achieve significant savings
Taking pride in H&F	Progressing our work in this area will help to ensure that H&F is not left behind in this significant shift which will have major impacts on our society.
Rising to the challenge of the climate and ecological emergency	

Background Papers Used in Preparing This Report

None

Introduction

Artificial intelligence (AI) is revolutionising how we approach various tasks and challenges. In the context of H&F, it presents opportunities to enhance service delivery, streamline operations, and improve decision-making.

What is generative AI?

Generative AI is a branch of artificial intelligence that focuses on creating new content or data from existing data. Generative AI can produce realistic images, texts, sounds, videos, and other types of media that are indistinguishable from human-made ones. Generative AI can also generate novel and creative content that does not exist in the real world, such as faces, landscapes, artwork, music, and stories.

The distinction between classic AI and the large language model (LLM)[1] – Generative AI

It is important to draw a clear distinction in the evolution of AI technology. Using the example of waste collection within a local authority such as H&F, in the classic AI approach, imagine a virtual assistant or chatbot designed to handle queries about waste collection. This chatbot relies on a predefined set of questions and answers, much like a table. Below is a simplified representation of what this table might look like, when representing how classic AI works:

User's Question	Chatbot Response (Classic AI)
"What day is my pickup"	"Your pickup is scheduled for Thursday 9am
"How do I recycle"	"To recycle glass, place it in the blue bin"

Note that if a resident asks a question not explicitly listed in this table – for instance or goes off topic, the chatbot will struggle to understand the query and therefore Classic AI is only as effective as predefined rules, scripts and prepopulated tables to make decisions or perform tasks. The classic AI approach is like having a predetermined script that the chatbot must strictly follow.

[1] Describes the ability of artificial intelligence applications or machines to understand and interpret natural language akin to a human being, and also includes the process of ingesting a large amount of data for training of these tools.

Generative AI – illustrative example

Now, let us imagine generative AI as the subject matter expert (SME) that has read every book, manual, and article in the organisation about a topic, so much so that it can now create new, helpful advice on the spot and is not encumbered by a pre-populated script of questions and answers.
This SME can dynamically reference data for the latest information as it is updated at source and handle unrelated questions.

Classic AI in the new era: navigating its crucial role in achieving business excellence.

It is essential to recognise that classic AI continues to play a major role in supporting business outcomes in organisations with diverse and complex business processes and systems.
The decision to leverage any kind of AI capability will be intricately tied to specific use cases. AI is like having a toolbox with a variety of tools – choosing the most optimal one depends on the task at hand. By acknowledging the strengths of classic AI, we ensure a strategic and nuanced approach, applying appropriate solutions to deliver the desired outcomes in the most optimal and efficient ways.

Navigating risks in the era of generative AI and large language models (LLMs): A heightened awareness

The advent of generative AI has amplified public awareness regarding the risks and challenges linked to deploying such capabilities in the pursuit of business outcomes. Concerns arise when considering the integration of generative AI, especially LLM's, into decision-making processes. These concerns span critical areas such as information security, privacy, accuracy, ethics, bias, and the potential for hallucinations,[2] heightening the need for a thorough understanding and a strategic approach to adoption.

Item	Risks, challenges and concerns	Description
1	Information Security	If not securely implemented, sensitive information might be inadvertently disclosed during interactions potentially leading to unauthorised access or data breaches.
2	Privacy concerns	The ability of generative AI to generate contextually relevant responses raises concerns about user privacy, as personal or sensitive information might be inferred or unintentionally shared.
3	Accuracy challenges	Despite advanced capabilities, (we are now on GPT-4), generative AI can still produce inaccurate or misleading responses, especially in complex and ambiguous situations,

[2] A 'hallucination' is when a Generative AI tool produces results or outputs that are deemed to be beyond a reasonable level of expectations. These are due to a variety of reasons such as a misunderstanding of the prompt or request from the user, misinterpretation of the data, incomplete data, or the intricacies of the underlying algorithm.

		impacting decision-making processes.

Item	Risks, challenges and concerns, continued	Description
4	Biased outputs	Large language models (LLM's) learn from vast datasets, which may unintentionally include biases present in the data. This can result in AI systems generating responses that reflect or amplify existing biases in society.
5	Hallucinations and unintended creativity	LLM's have been known to generate creative content whose sources cannot be verified, nor do they align with factual information.
6	Ethical Concerns	The ethical implications of AI decision-making, such as bias or unfair treatment, may arise, sparking debates about ethical responsibility around the deployment and utilisation of generative AI capabilities.
7	Lack of explainability	LLMs often operate as complex "black box" systems that are opaque to the end user, making it challenging to explain how they arrived at a particular decision or generated a specific response. This lack of transparency raises concerns about accountability and trust.

Unlocking the Power of Generative AI: capabilities and applications

Before we navigate the landscape of risks and challenges in Generative AI, let us begin by exploring its high-level capabilities and applications in a bit more detail. The following diagram vividly illustrates how Generative AI unleashes its transformative powers – summarising, expanding, creating intelligent chatbots, inferring insights, transforming developing content and much more – and provides a solid foundation for understanding the breadth of Generative AI's potential applications. The adoption of Generative AI could be transformative, and is akin to the industrial revolution, or the dawn of the computer age.

Generative AI: overview of capabilities and applications

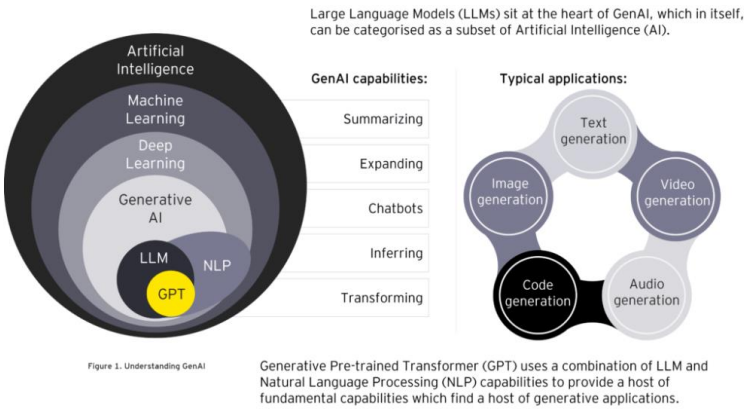

Source – The GenAI Framework for HMG

Generative AI in Action: potential H&F organisational initiatives

Below are examples of how this transformative technology could be harnessed in a local authority context in the short term, and examples use cases we are actively exploring:

- content generation and/or productivity support
- contact centre support/automation.
- code generation.

There are a range of adoption and implementation challenges facing any adoption:

Item	Challenges	Description
1	Ensuring data is AI ready.	Data Fragmentation held across multiple areas. Implications and Impact: Difficulties in creating a centralised and comprehensive dataset for training generative AI models (single source of truth), hindering the ability of the tool to generate contextually relevant responses. Mitigations: Adoption a data governance framework and practices The establishment of a robust data governance framework to ensure standardisation and consistency, including data quality and control measures (e.g. master data management, data stewardship). Digital and Data Adoption: Investment in support for staff to help them develop the digital and data skills to become not just 'AI ready' but also 'digital, data and technology ready.'
2	Developing our AI capabilities	Manifestation: The field of Generative AI and the attendant opportunities in deriving business outcomes are still very new. Therefore, there will always be a gap in ensuring that an organisation's people, processes, and technologies are leveraged

Item	Challenges	Description
		to exploit these capabilities in an optimal manner. One such gap is in prompt writing,[3] in the era of Generative AI and its transformative nature, this a critical skillset. Training on prompt writing and other AI capabilities are core to any adoption or implementation programme.

[3] Prompt writing refers to the ability of users to engage with large language models (LLM's) such as Microsoft co -pilot, ChatGPT and any other Generative AI tools, it involves formulating specific instructions using appropriate syntax, natural language, and sentences to elicit accurate and effective responses from the Generative AI tools.

Item	Challenges, continued	Description
3	Securing the necessary investment	Leveraging generative (and other forms of) AI at scale will take significant investment – not just financially, but also in terms of time to safely realise a return on investment without unintended consequences. Mitigations: Define a clear (generative) AI ambition with dedicated funding and a realistic understanding of the timescales to see a return on investment and realise benefits. Be clear that this involves investment in significant culture change, and that not every investment will realise a measurable return or benefit.
4	Safeguarding privacy and confidential information	Where AI systems utilise data from across the council, safeguards will need to be installed to ensure that confidential information, and information relating to individuals, is protected. It is important to note that as we embark on this journey, human led decision making and final decisions will continue to be core to the use of AI, to ensure accountability and ethical oversight.

Embracing Generative AI at H&F

Our Generative AI journey has involved several steps to date.

- We organised envisioning sessions with Microsoft and other strategic partners which included representatives from Digital services, REAP and other services, to equip ourselves with essential knowledge on generative AI. We also sought insights from Gartner and other thought leaders on the topic.

- A cross departmental working group was established to coordinate the adoption of generative AI across H&F and includes technical oversight, governance (privacy and security), adoption, delivery, etc.

- We have also aligned ourselves with the Generative AI Collaboration Forum established by Microsoft for Local Authorities, and with LOTI.

- We engaged with services across the council to elicit potential use cases that could leverage Generative AI capabilities to support business outcomes.

Navigating the new world: recommendations for H&F's preparedness

1. Develop a policy and governance framework for the ethical adoption of AI.

2. Continuous Monitoring and evaluation of the fast-changing GenAI landscape.

3. Skills development initiatives, including career transition support for creative industries and other areas that will be disrupted by LLM's.

4. Continued co-production of AI developments across the council.

5. Strategic partnerships with industry, educational institutions, and other stakeholders.

In conclusion, it is important to recognise that realising the greatest benefits from AI necessitates adopting hybrid approach that combines classic AI and generative AI to support business outcomes effectively.

There is a sense of inevitability in the adoption of Generative AI capabilities, given that over 80% of our line of business applications are vendor managed and are provided as Software as a Service (SaaS) and, as outlined earlier, majority of these vendors will start embedding GenAI capabilities within their applications. We therefore need to anticipate these and be prepared for these integrations.

As a forward-looking organisation, it is important to that we explore opportunities associated with this capability, positioning ourselves perhaps not at the bleeding edge of innovation, but at the leading edge, and therefore, working closely with other local authorities and the wider public sector on the adoption of Generative AI capabilities.

LIST OF APPENDICES

Appendix A – Link to a short explanatory video on Generative AI and what it is: https://www.youtube.com/watch?v=rwF-X5STYks